

MemXCT: Design, Optimization, Scaling, and Reproducibility of X-Ray Tomography Imaging

Mert Hidayetoğlu¹, Tekin Biçer², *Senior Member, IEEE*, Simon Garcia de Gonzalo, Bin Ren, Doğa Gürsoy³, Rajkumar Kettimuthu, *Senior Member, IEEE*, Ian T. Foster⁴, *Fellow, IEEE*, and Wen-Mei W. Hwu⁵, *Fellow, IEEE*

Abstract—This work extends our previous research entitled “MemXCT: Memory-centric X-ray CT Reconstruction with Massive Parallelization” that was originally published at SC19 conference (Hidayetoğlu *et al.*, 2019) with reproducibility of the computational imaging performance. X-ray computed tomography (XCT) is regularly used at synchrotron light sources to study the internal morphology of materials at high resolution. However, experimental constraints, such as radiation sensitivity, can result in noisy or undersampled measurements. Further, depending on the resolution, sample size and data acquisition rates, the resulting noisy dataset can be in the order of terabytes. Advanced iterative reconstruction techniques can produce high-quality images from noisy measurements, but their computational requirements have made their use an exception rather than the rule. We propose a novel memory-centric approach that avoids redundant computations at the expense of additional memory complexity. We develop a memory-centric iterative reconstruction system, MemXCT, that uses an optimized SpMV implementation with two-level pseudo-Hilbert ordering and multi-stage input buffering. We evaluate MemXCT on various supercomputer architectures involving KNL and GPU. MemXCT can reconstruct a large (11K × 11K) mouse brain tomogram in 10 seconds using 4096 KNL nodes (256K cores). The results presented in our original article at the SC19 were based on large-scale supercomputing resources. The MemXCT application was selected for the Student Cluster Competition (SCC) Reproducibility Challenge and evaluated on a variety of cloud computing resources by universities around the world in the SC20 conference. We summarize the results of the top-ranked SCC Reproducibility Challenge teams and identify the most pertinent measures for ensuring the reproducibility of our experiments in this article.

Index Terms—X-ray tomography, space-filling curves, cache utilization, knights landing, GPU computing, SpMV, reproducibility

1 INTRODUCTION

X-RAY computed tomography (XCT) is a widely used non-destructive 3D imaging technique for observing and understanding the internal morphology of samples. Synchrotron light sources, such as the Advanced Photon Source (APS), can provide high-brilliance X-rays that enable tomographic imaging of centimeter sized samples at sub-micrometer (μm) spatial resolution. Such experiments can generate from a few GBs to TBs of data volumes in a short time period with the typical pixelated detectors that can run at 16 GB/s [2]. However, the quality of data collected from CT experiments depends heavily on factors such as radiation exposure time (dose) and target spatial resolution. Much effort has been devoted to develop and implement advanced reconstruction

algorithms to improve image quality when collected data are noisy or imperfect (e.g., due to limited dose).

Modern XCT reconstruction approaches involve either (i) direct solvers based on analytical inversion or (ii) iterative solvers that can accurately model the experimental conditions and/or constrain the solution based on prior knowledge of the sample. Analytical methods such as the filtered backprojection (FBP) algorithm are computationally efficient, but reconstruction quality is often poor when measurements are noisy or undersampled. Iterative methods, on the other hand, can use advanced optimization and regularization techniques to handle inherent noise in X-ray measurements [4], [5], [6], [7], [8], [9], [10]. However, these methods are computationally more demanding than the analytical methods, since forward and backprojection operations require many (irregular) accesses and computation at each iteration. Thus, efficient implementation and parallelization are crucial for high-quality large-scale image reconstructions.

Most state-of-the-art reconstruction software and libraries can perform parallel reconstruction to some extent. However, these implementations mostly rely on data replication [11], [12] and/or redundant computations [13], [14], each of which significantly limits the runtime performance of the system due to repeated computation of intermediate indices of ray tracing and irregular data accesses.

MemXCT, our advanced reconstruction system, uses a memory-centric approach to avoid redundant computation by removing on-the-fly operations at the expense of greater memory complexity. Specifically, it utilizes an optimized

- Mert Hidayetoğlu, Simon Garcia de Gonzalo, and Wen-Mei W. Hwu are with the University of Illinois at Urbana-Champaign, Champaign, IL 61810 USA. E-mail: {hidayet2, grcdgnz2, w-hwu}@illinois.edu.
- Tekin Biçer, Doğa Gürsoy, Rajkumar Kettimuthu, and Ian T. Foster are with the Argonne National Laboratory, Lemont, IL 60439 USA. E-mail: {tbicer, dgursoy, kettimuthu, foster}@anl.gov.
- Bin Ren is with the College of William & Mary, Williamsburg, VA 23185 USA. E-mail: bren@cs.wvm.edu.

Manuscript received 27 Aug. 2021; revised 26 Oct. 2021; accepted 8 Nov. 2021. Date of publication 23 Nov. 2021; date of current version 24 Jan. 2022.

This work was partially supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research and Basic Energy Sciences, under Grant DE-AC02-06CH11357.

(Corresponding author: Mert Hidayetoğlu.)

Recommended for acceptance by S. L. Harrell and S. A. Michael.

Digital Object Identifier no. 10.1109/TPDS.2021.3128032




Fig. 1. Multiple zooms on a single 11293×11293 2D slice of a mouse brain image generated by reconstruction, with the methods described here, from a single sinogram with size 4501×11293 . Data were collected by Dyer *et al.* [3] at the APS. Single-slice reconstruction with 30 CG iterations on 4096 KNL nodes takes 10 seconds. Memory footprint of the application is 10.2 TB. The full mouse brain consists of 11293 slices.

sparse matrix-vector multiplication (SpMV) implementation with *multi-stage buffering* and *two-level pseudo-Hilbert ordering* to optimize data communication, partitioning, and accesses at different analysis levels and on both measurement (sinogram) and reconstructed (tomogram) data.

MemXCT memoizes on-the-fly operations and therefore requires more memory than alternatives. However, its overall per-node memory footprint decreases linearly with increasing computing resources, promoting scalability with massive parallelization and enabling iterative reconstruction of extremely large datasets with efficient resource use. Further, MemXCT stores intermediate data structures in compressed format and locality-enhancing layout order to minimize memory footprint and data movement.

To illustrate the practical implications of MemXCT optimizations, we show in Fig. 1a 2D tomogram from a 3D mouse brain image with total size $11293 \times 11293 \times 11293$, obtained by reconstructing a sinogram with 4501×11293 dimensions. This single-slice reconstruction was generated with 30 conjugate gradient (CG) iterations in about 10 seconds on 4096 KNL nodes (256K cores) using MemXCT. This high speed capability makes it feasible, for the first time, to apply advanced iterative reconstruction algorithms to extremely large CT datasets. The multi-scale nature of the reconstruction is presented by zooming progressively to the brain arteries, where great detail can be seen. High reconstruction quality at such scales is important as subsequent data analysis steps, such as segmentation of blood vessels and myelinated axon tracts, are highly dependent on the quality of the reconstruction.

The MemXCT work was originally published at the SC19 conference[1]. In 2020, the MemXCT was selected as the benchmark application for the Student Cluster Competition (SCC) Reproducibility Challenge and evaluated on a variety of cloud computing resources by 16 university teams. Students were tasked to reconstruct several images using MemXCT and compared their measured performance results with the evaluations reported in the SC19 paper as well as in this paper. We summarize the SCC Reproducibility Challenge results of the top-ranked SCC teams and present key insights from their results at the end of the paper.

The principal contributions of this paper are as follows:

- We propose a memory-centric approach that uses an *extended SpMV* to address race conditions due to scatter operations. Our implementation replaces scatter operations on sinogram and tomogram domains with gather operations.
- We implement a *two-level pseudo-Hilbert ordering* to organize communications and data accesses, further improving communication performance, utilization of different levels of memory hierarchy, and cache utilization.
- We introduce a *multi-stage input buffering*. Domain partitioning improves process-level data communication and locality; multi-stage buffering enables data reuse in first-level caches and minimizes memory latency.
- We extensively evaluate *MemXCT* performance on both KNL and GPU systems up to 4096 nodes, and show architecture-specific considerations and optimizations.
- We present the key observations from student teams who successfully reproduced the SC19 paper results and identify the key measures in the MemXCT repository that are the most helpful to their efforts.

2 BACKGROUND AND MOTIVATION

We first explain XCT experiments and their data acquisition process. Then, we describe the steps involved in iterative image reconstruction. Finally, we explain computational bottlenecks in iterative reconstruction approaches and introduce our solutions.

2.1 Data Acquisition and Measurement Process

In an XCT experiment, the sample is placed on a rotation stage and illuminated by an X-ray source, while collecting 2D images through a detector as the sample is being rotated: see Fig. 2. The mathematical model of the measurement process is based on Beer's law [15] that describes interaction between X-rays and matter: $I_\theta(s) = I_0(s)\exp[-p_\theta(s)]$, where $I_0(s)$ is the incident X-ray illumination on the sample and $I_\theta(s)$ are the collected measurements at a number of different θ angles as a result of a tomographic scan.

In Fig. 2, a set of collected projections from a sample is shown with p_{θ^i} . Considering parallel beam geometry, the sinogram $p_\theta(s)$ is a cross section of projections (shown in blue grid on p_{θ^*}) that consists of I_θ measurements from f . The goal of a tomographic reconstruction algorithm is to recover 2D image slice (tomogram) from its corresponding sinogram $p_\theta(s)$.

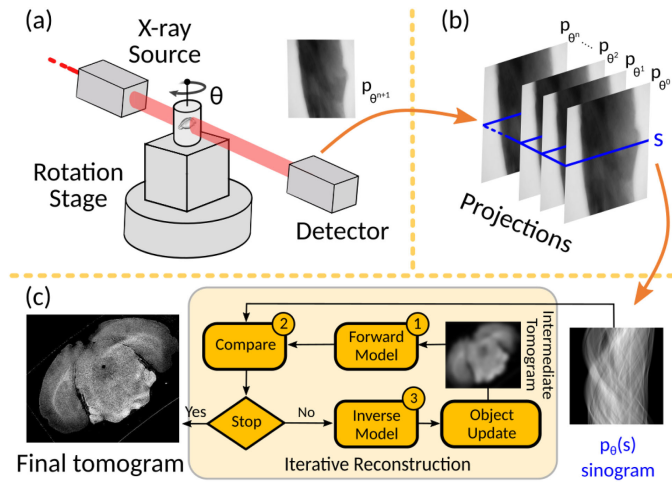


Fig. 2. (a) Illustration of the experimental setup where a target object is illuminated with a synchrotron light source. Photons are attenuated by different amounts according to the object's attenuation coefficient profile and the photons' travel distance within the object. (b) Presentation of how the detector measurements are organized and stored as projections (from different angles). Once the experiment is completed, the sinograms are extracted from the projections, and the images corresponding to the object are iteratively reconstructed as shown in (c).

2.2 Iterative Formulation

Tomographic reconstruction solves the problem

$$\hat{x} = \underset{x \in C}{\operatorname{argmin}} \|y - Ax\|^2 + R(x), \quad (1)$$

where \hat{x} is the reconstructed tomogram, A is the forward model, y is the sinogram, $R(x)$ is a regularizer functional, x is the search variable, and C is a constraint on x .

Almost all iterative solvers based on gradient descent perform three common steps in each iteration as depicted in Fig. 2. Specifically, for iteration i : first, the residual $r_i = y - Ax_i$ is found through *forward projection*; second, the gradient is found through *backprojection* as $\nabla \|y - Ax_i\|^2 = A^T r_i$; and finally, the candidate solution is updated in the negative-gradient direction as $x_{i+1} = x_i - \alpha \nabla \|y - Ax_i\|^2$, where α is the step size. Iterative approaches can also involve additional updates due to regularizer $R(x)$ and constraint C . However, this paper focuses on the common computational costs involving aforementioned steps.

2.3 Forward Projection and Backprojection

Forward and backprojections are the two most computationally demanding kernels in iterative tomographic reconstruction approaches. Many reconstruction libraries, such as Trace and TomoPy [11], [14], implement Siddon's algorithm [16] to perform ray tracing on tomogram domain so that the exact length and weight information on each voxel can be computed for each intersecting X-ray. For forward model, this information is used for computing the residuals, whereas for backprojection they are used for gradient and update operations. Since storing voxel indices and length information for all rays and voxels require significant memory, most state-of-the-art approaches perform ray tracing computations on-the-fly. For instance, a sinogram with dimensions 1501×2048 requires 56 GB memory to store intermediate data structures. Listing 1 illustrates the high level implementation of iterative reconstruction and

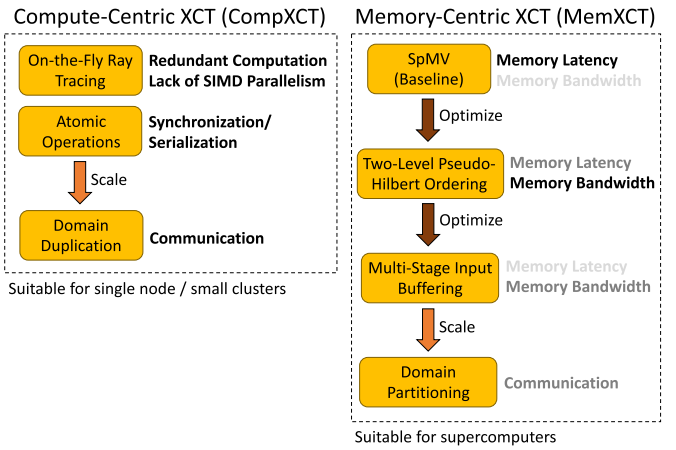


Fig. 3. Memory-centric and compute-centric approaches for XCT and their areas for potential optimizations. Shades of black color indicate higher contribution to performance bottleneck.

repetitive computations of indices and lengths arrays. We term reconstruction algorithms that compute on-the-fly ray tracing information as *compute-centric XCT* (CompXCT).

Listing 1. High Level Implementation of CompXCT

```

1  for (int i = 0; i < num_iters; ++i)
   { //iterations
2  for (int j = 0; j < sinogram_rows; ++j) {
3  float theta = rotations(j); //  $\theta$  for row j
4  for (int k = 0; k < sinogram_cols; ++k) {
5  int **indices =
6      intersecting_voxels(j, k, theta,
7                          tomogram);
8  float **lengths =
9      compute_lengths(j, k, theta, tomogram);
10 float m = sinogram[j, k]; // Data for simulated
11     ray
12     float residual =
13     forward_model(m, indices, lengths,
14                  tomogram);
15     // Apply inverse model and then update
16     backprojection(residual, indices, lengths,
17                  tomogram);
18 }
19 }

```

2.4 Computational Bottlenecks and Overview of Our Solution: MemXCT

Fig. 3 illustrates the techniques used for traditional CompXCT and proposed MemXCT approaches in yellow boxes. The existing bottlenecks and their influence are given next to the areas for potential optimizations with different shades of black, in which darker shades indicate higher contribution to performance bottleneck.

CompXCT eliminates the need for storing intermediate data structures; however, since these data structures are recalculated for each iteration, they introduce additional computational complexity to the execution. Further, on-the-fly calculation of indices limits the optimizations for

vectorization as repeated sorting and padding operations nullify any performance benefit. As a remedy, MemXCT embraces a memory-centric approach that memoizes ray-tracing operations and performs efficient SpMV using compressed (vectorizable) data structures on both sinogram and tomogram domains. Further, the irregular data access patterns are normalized using pseudo-Hilbert ordering and first level cache utilization are improved using multi-stage buffering. These optimizations result in high-performance computations and converts performance bottlenecks from computation to memory.

The parallelization of CompXCT approaches is typically based on X-rays, where each measured ray can independently be projected on tomogram domain. This type of parallelization performs well for forward projection kernel (residual computation) since the main operation is reading the intersected voxel values from tomogram domain, i.e., rays perform *gather operations* during forward projection. However, subsequent backprojection and update steps require synchronization among parallelized rays since there can be many updates from different rays on the same voxel, i.e., rays perform *scatter operations* during backprojection and updates. Recent work deals with race conditions by either applying atomic operations [17] or duplicating the pixel domain across threads/processes and then performing a reduction [11], [12]. Unfortunately, the performance of atomic operations hinges on hardware implementations, which differ substantially across architectures [18], and typically results in significant performance degradation on massively parallel architectures. Domain duplication is also impractical for GPU-like architectures, since each parallel unit requires a replica of the domain and the total memory footprint almost always exceed available resources. Further, distributed memory parallelization using duplication can result in redundant communication cost [11].

In contrast, MemXCT partitions both sinogram and tomogram domains among parallel units and transforms scatter operations to gathers. Note that this transformation can result in irregular data accesses on both domains, however our pseudo-Hilbert index ordering coupled with multi-level buffering amortize performance penalty due to irregular data accesses. MemXCT also exploits Hilbert ordering (hence named two-level pseudo-Hilbert ordering) for process-level domain partitioning, which results in efficient communication and better connectivity between processes [19], [20]. We describe these optimization in detail in the following section.

3 MEMXCT OPTIMIZATIONS

In this section, we provide detailed information about MemXCT. First, we explain the *baseline implementation*, in which we focus on explicit SpMV operations. Then, we present our system optimizations that use *two-level pseudo-Hilbert ordering* and *multi-stage input buffer*. We also explain our distributed memory parallelization considerations with MPI.

3.1 Baseline Implementation

MemXCT performs forward projection and backprojection operations as explicit SpMV operations to remove redundant and inefficient computations. Listing 2 shows the

baseline compressed sparse row (CSR) SpMV kernel used for both forward and back projection. In forward projection, x and y correspond to tomogram and sinogram data, respectively, and vice-versa in backprojection. In either case, ind and val data correspond to precomputed pixel-ray intersection indices and lengths that are reused to avoid redundant computation.

Listing 2. Baseline MemXCT Kernel

```

1  #pragma omp parallel for schedule(dynamic,
    partsize)
2  for (int i = 0; i < numrow; ++i) {
3    float acc = 0;
4    //vectorize
5    for (int j = displ[i]; j < displ[i+1]; ++j)
6      acc += x[ind[j]]*val[j];
7    y[i] = acc;
8  }
```

3.1.1 Regular and Irregular Accesses

Each fused multiply-add (FMA) operation in Listing 2 performs three important memory accesses: ind and val are *regular*, and x is *irregular*. The *regular* accesses are sequential and hence exhibit low memory latency. The *irregular* accesses to x exhibit long latency due to high L2 miss rates. MemXCT addresses this problem by improving cache reuse through novel data layout techniques that increase locality, as illustrated in Section 3.2.

3.1.2 Row Partitioning & Parallelization

Listing 2 involves gather operations only and hence is suitable for massive parallelization. MemXCT parallelizes the outer loop among *row partitions*. On KNL, partitions are distributed across dynamically scheduled OpenMP threads. On GPU, each partition corresponds to a CUDA thread block. Each OpenMP thread processes many row partitions, whereas each CUDA thread processes a single row in a partition.

3.1.3 Vectorization on KNL

MemXCT enables efficient instruction-level parallelization through vectorization of the inner loop in Listing 2. Each KNL vector-processing unit (VPU) can load multiple regular and irregular data, and multiply them in one step with an AVX-512 instruction. Then, they perform parallel reductions on partial data, and add results to the accumulator. We provide necessary data alignments for efficient vectorization.

3.1.4 Coalesced Memory Access on GPU

Our GPU implementation modifies Listing 2 to use the ELL (column-major) storage format instead of CSR. Transposed ELL data structures provide coalesced memory access through consecutive threads accessing consecutive memory locations. We minimize redundant computations and memory accesses due to ELL format by performing zero padding on thread-block (i.e., partition) level, rather than on matrix level.

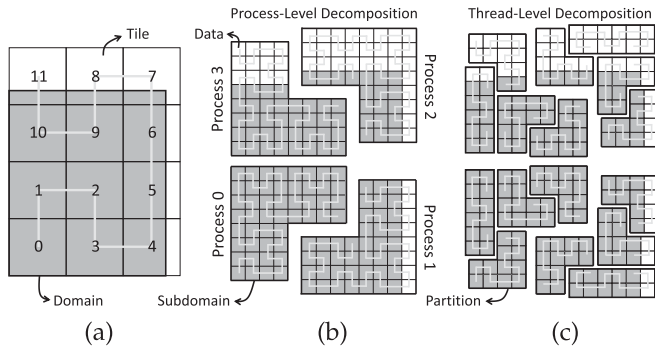


Fig. 4. Two-level pseudo-Hilbert ordering and domain decomposition at (a) tile level, (b) process level, and (c) thread level.

3.2 Two-Level Pseudo-Hilbert Ordering

MemXCT implements a two-level pseudo-Hilbert ordering on both tomogram and sinogram data of arbitrary size. Fig. 4 shows an example for a 13×11 domain. First, the domain is tiled with a minimum number of equi-sized square tiles with dimension a power of two. In the figure, tile size is 4×4 and 12 tiles are used to cover the 13×11 domain. These tiles are indexed with a Hilbert ordering for rectangular domains [21], as shown in Fig. 4a. A second-level Hilbert ordering is then applied to the data within each tile. Necessary rotations are performed to provide data connectivity among tiles, so as to achieve both data locality and connectivity for domain decomposition at the process and thread levels, as discussed next. The process-level domain decomposition is essential for MPI parallelization, as discussed in Section 3.4.

3.2.1 Irregular Data Access Patterns

As shown in Fig. 5, the processing of a single sinogram or tomogram results in a linear or a sinusoidal memory access footprint, respectively. In this example, they perform 25 and 30 accesses on respective domains. With row-major (naive) ordering of 2D data and 64 B cache line, each row in Fig. 5 would correspond to a single cache line. In this case, both tomogram and sinogram data would have 16 cache misses, yielding miss rates of 64% and 53%, respectively. Row-major (or column-major) ordering of 2D domains is very inefficient for XCT because of its poor cache locality. That is, a single cache line does not provide enough data reuse except on a few instances.

3.2.2 Cache Locality

In Fig. 5, Hilbert ordering maps each cache line to a 4×4 block in a 2D domain, increasing cache data reuse. As a result, cache misses are reduced to six and seven, yielding rates of 24% and 23% for forward projection and backprojection, respectively. Hilbert ordering is portable to different cache-line sizes thanks to its recursive nature. The cache locality provided by Hilbert ordering eliminates the memory latency bound of MemXCT to a great extent, as results show in Section 4.2.2.

3.2.3 Partition Locality

Hilbert ordering not only provides cache locality but also partition locality. That is, the outer loop of SpMV in Listing 2




Fig. 5. Data access patterns on 2D domains.

is partitioned with respect to the index of y , as described in Section 3.1.2. In MemXCT, each partition remains connected at both the process and thread levels, thanks to the connectivity provided by two-level pseudo Hilbert ordering. In contrast, a Morton ordering would yield disconnected partitions, since it does not guarantee adjacent memory locations to have adjacent locations in the 2D domain. Partition locality is essential for our further optimizations.

3.3 Multi-Stage Input Buffering

MemXCT implements a multi-stage buffering mechanism for: (1) further reducing L2 cache miss rates of irregular data by explicit staging of corresponding accesses on an L1 buffer, and (2) reducing memory bandwidth consumed for index data by using two-byte addressing. Listing 3 shows SpMV kernel with input buffering, where `stagedispl` and `stagenz` arrays correspond to starting points of (multi-stage) buffers and number of (nonzero) elements in buffers, respectively.

Listing 3. Optimized MemXCT Kernel

```

1  #pragma omp parallel for schedule(dynamic)
2  for (int part = 0; part < numparts; ++part) {
3      float output[partsize] = {0};
4      float input[bufsize];
5      for (int stage = partdispl[part];
6           stage < partdispl[part+1]; +
7           +stage) {
8          int start = stagedispl[stage];
9          //vectorize
10         for (int i = 0; i < stagenz[stage]; ++i)
11             input[i] = x[map[start+i]];
12         for (int j = 0; j < partsize; ++j) {
13             int start = stage*partsize;
14             //vectorize
15             for (int i = displ[start+j];
16                  i < displ[start+j+1]; ++i)
17                 output[j] += input[ind[i]]*val[i];
18         }
19     }
20     int start = part*partsize;
21     //vectorize
22     for (int i = 0; i < partsize; ++i)
23         if (start+i < numRows)
24             y[start+i] += output[i];

```

3.3.1 Data Reuse From Input Buffer

Fig. 6a shows memory access footprints for processing a 64×64 tomogram and sinogram partition in a 256×256

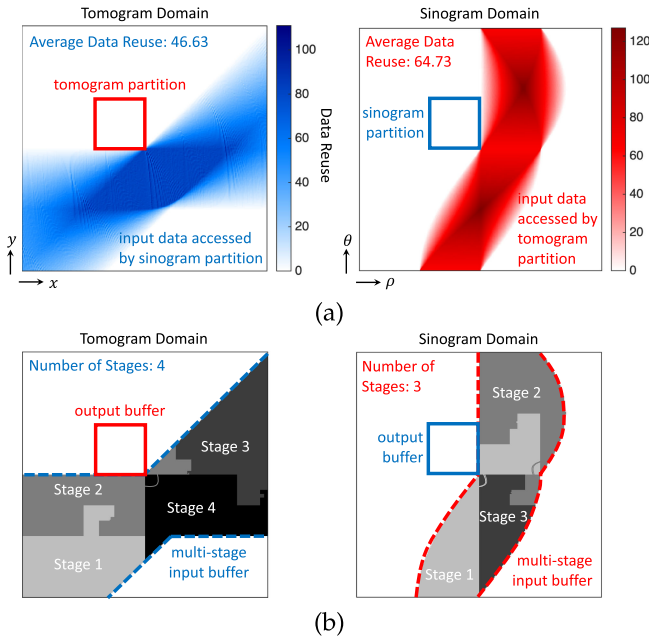


Fig. 6. (a) Tomogram and sinogram partitions and respective data access footprints, and (b) multi-stage buffer shapes.

tomogram and sinogram domain, respectively. The tomogram partition reads from the sinogram domain and the sinogram partition from the tomogram domain. Data access footprints are shown in tomogram and sinogram domains. Darker shades represent higher data reuse. For processing each partition, MemXCT explicitly moves required data from memory to an L1 buffer through `map` at line 9 of Listing 3. Then it reuses buffered data for irregular accesses and hence performs less memory access.

3.3.2 Multi-Staging

In a single-stage (naive) buffering, the required buffer size grows with the problem geometry. If the buffer is too large, it cancels the buffering benefit since it leaks into L2 or even memory [22]. As a remedy, MemXCT implements a multi-stage buffering strategy with a constant buffer size. In this case, irregular data are accessed through multiple buffer stagings. Stages are determined with respect to Hilbert ordering, ensuring data locality. As an example, Fig. 6b shows mapping of buffer stages on 2D tomogram and sinogram domains. With a buffer size of 32 KB, MemXCT performs irregular accesses through four and three stages for projection and backprojection, respectively. On KNLs, buffer size should be smaller than 32 KB L1 cache to avoid spill to L2 cache. On GPUs, the buffer is allocated through CUDA shared memory, guaranteeing reuse from L1 cache.

3.3.3 Staging Overhead

Input buffering involves a staging overhead. On KNL, an OpenMP thread stays idle while waiting staging to be completed. Similarly on GPU, CUDA threads across warps stall for synchronization before and after each staging. Input buffering also consumes some additional memory bandwidth for reading `map` data. The next two subsections describe two techniques for hiding buffering overhead and also to save some bandwidth.

3.3.4 Overlapping Staging and FMAs

MemXCT multi-stage buffering lends itself well to be used by the underlined hardware to hide staging overhead. On KNL, this is done through SMT (simultaneous multithreading): when there are multiple threads running on a single core, the hardware scheduler finds overlapping opportunities across buffer stagings and FMAs among threads. As a result, MemXCT enables effective SMT utilization. Similarly, GPU hardware also takes advantage of overlapping through block scheduling on SMs (streaming multiprocessors).

3.3.5 Saving Memory Bandwidth

Hilbert ordering and input buffering eliminate most of the memory latency due to irregular data accesses, and therefore MemXCT is bounded by memory-bandwidth consumed by regular data. To alleviate this bottleneck, we use 16-bit addressing to access input buffer (see `ind` at line 14 on Listing 3), rather than 32-bit addressing (as in `ind` at line 6 on Listing 2). 16-bit addressing can address buffer sizes up to 256 KB. This saves 25% of total bandwidth consumption of regular data, and provides additional speedup (see Section 4.2.3 for results.)

3.4 MPI Parallelization

Traditionally, one domain (either tomogram or sinogram) is partitioned while the other is duplicated across all processes. MemXCT partitions both tomogram and sinogram domains by distributing tiles evenly across MPI processes, as seen in Fig 4b. Each process is responsible for a single *tomogram subdomain* and a single *sinogram subdomain*. Each subdomain consists of a single or several tiles, e.g., subdomain 0 consists of tiles 0–2. While processes are not perfectly load balanced, it can be improved by finer tile granularity at the cost of more preprocessing.

3.4.1 Sparse Communications

MemXCT communicates only necessary data through `MPI_Alltoallv`. To explain, Fig. 7b shows communication footprint of two 256×256 tomogram subdomains shown in Fig. 7a. For example, tomogram subdomain 7 interacts only with sinogram subdomains 1, 2, 8–11, 13, and 14. Fig. 7c shows the corresponding communication matrix, where each entry represents communication between two processes. Communication size between each pair depends on the interaction footprint, e.g., process 7 sends more data to process 1 than 14, as seen in Fig. 7d. Partition locality provided by two-level pseudo-Hilbert ordering (described in Section 3.2) minimizes the footprint, and hence increase data reuse and reduce communications.

3.4.2 Overlapped Interactions

MemXCT communicates sinogram data rather than tomogram data because: (1) sinogram data are smaller in many applications, and (2) it yields more data reuse as compared to tomogram data. In forward projection, tomogram subdomains send partial sinogram data to corresponding sinogram subdomains where overlapped data are reduced, e.g., processes 8, 9, and 11 reduce partial sinogram data received

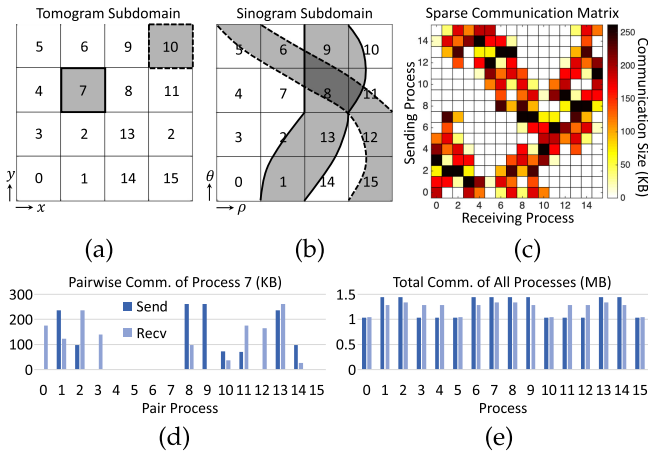


Fig. 7. Communication footprint of two (a) tomogram subdomains on (b) sinogram subdomains. (c) Resulting communication matrix for all processes. (d) The amount of outgoing/incoming data from/to process 7 to/from others. (e) Total incoming and outgoing data for all processes.

from processes 7 and 10. In backprojection, processes duplicate overlapped sinogram data and send them to interacted processes which perform backprojections on their respective tomogram subdomains. The communication matrix for backprojection is the transpose of the one in Fig. 7c. The total amount of communications and load balancing for all processes is shown in Fig. 7e.

3.4.3 Parallelization Overhead

MemXCT parallelization of forward projection mathematically corresponds to a factorization of the projection matrix as $A = RCA_p$. As a result, forward projection can be modeled as three fundamental steps: A_p , C , and R which correspond to partial forward projection, communication, and reduction operations. Backprojection can also be seen as simply $A^T = A_p^T C^T R^T$. The multiplication costs of A_p and original A are the same because they involve the same number of $\mathcal{O}(MN^2)$ nonzeros. Thus, communication C and reduction R can be seen as the MemXCT parallelization overhead. Both C and R have $\mathcal{O}(MN\sqrt{P})$ nonzeros, where P is the number of processes. It is because when P quadruples, total communication footprint on sinogram domain doubles. For communications, there is an extra $\mathcal{O}(\sqrt{P})$ term comes due to the handshake overhead between processes. On the other hand, compute-centric approach involves race conditions and hence it duplicates the tomogram domain across all processes [11]. In that case, duplicated domains has to be reduced through `MPI_Allreduce` at the end of each backprojection, yielding $\mathcal{O}(N^2 \log P)$ parallelization overhead. Resulting computational complexities are shown in Table 1.

3.5 Other Details

MemXCT requires an extra preprocessing step for avoiding redundant and inefficient computations as opposed to a compute-centric approach. The preprocessing involve: (1) Hilbert ordering and domain decomposition, (2) ray tracing for constructing forward projection matrix, (3) sparse transposition for constructing backprojection matrix, and (4) row partitioning and building corresponding buffer data

TABLE 1
Computational Complexities

	Sequential	Trace	MemXCT
Memory	$MN + N^2$	$MN/P + N^2$	$MN^2/P + MN/\sqrt{P}$
Comput.	MN^2	MN^2/P	$MN^2/P + MN/\sqrt{P}$
Comm.	N/A	$N^2 \log P$	$MN/\sqrt{P} + \sqrt{P}$

M : # of projections, N : # of channels P : # of processes.

structures. Preprocessing is MPI+OpenMP parallel and is performed on CPU.

3.5.1 Preserving Data Locality

MemXCT preserves data locality through all matrix manipulations, which is essential for our performance optimizations. For example, constructing backprojection matrix requires a sparse matrix transposition of the forward projection matrix. MemXCT performs this through a scan-based matrix transposition [23] which preserves data ordering, rather than an atomic-based transposition which randomizes data ordering.

3.5.2 Iterative Solution

There is a plethora of solution schemes for iterative XCT. To mention a few, recent work implements SIRT [11], SGD [17], and ICD [24] iterations. Any of them can be implemented for our proposed memory-centric approach in a plug-and-play manner with minor modifications. Nevertheless, MemXCT implements CG iterations [25], which has a faster convergence rate than any of them (at a higher per-iteration cost). It is simply because: (1) a full gradient is found rather than a partial gradient, (2) optimal step size is found analytically via an additional forward projection, and (3) visiting redundant directions is prevented via three-term recursion. We use a heuristic early termination of iterations to prevent *overfitting*, which is practically considered as a regularization method. Convergence results are shown in Section 4.1.2.

4 NUMERICAL RESULTS

All experiments in this paper are performed on ALCF Theta [26], NCSA Blue Waters [27], ALCF Cooley [28], an IBM Minsky node [29], and an Nvidia DGX-1 node [30]. Table 2 characterizes these machines. The fifth column (B/W) is the theoretical device memory bandwidth, assuming that ECC (error correcting code) degrades theoretical bandwidth of K20X and K80 by 15% [31]. Link describes interface between host and device.

We used six datasets, as shown in Table 3, in order to evaluate application performance. The first four are artificially created for performance evaluation and the latter two are from real synchrotron experiments at APS. Measurement sinograms are given for a single slice. The artificial datasets follow parallel raster scan geometry just as the real datasets. The irregular and regular data footprints, defined in Section 3.1.1, are given for all datasets in Table 3. The first/second entry for memory footprints are accessed in forward/backprojection, respectively. RDS1 involves a

TABLE 2
Key Features of Machines Used for Experiments

Machine	Nodes	Accel.	Device Mem.	B/W	Node Mem.	Link
Theta	4392	KNL	16 GB MCDRAM	400 GB/s	192 GB	90 GB/s
Blue W.	4228	K20X	6 GB DDR5	121.5 GB/s	32 GB	PCIe
Cooley	126	2×K80	12 GB DDR5	204 GB/s	384 GB	PCIe
Minsky	1	4×P100	16 GB HBM2	720 GB/s	128 GB	NVLink
DGX-1	1	8×V100	16 GB HBM2	900 GB/s	512 GB	NVLink

TABLE 3
Dataset Details and Memory Footprints

Name	Sinogram ($M \times N$)	Sample	Irregular Data	Regular Data
ADS1	360×256	Artificial	256/360 KB	215/215 MB
ADS2	750×512	Artificial	1.0/1.5 MB	1.8/1.8 GB
ADS3	1500×1024	Artificial	4.0/6.0 MB	14/14 GB
ADS4	2400×2048	Artificial	16/19 MB	90/90 GB
RDS1	1501×2048	Shale Rock	16/12 MB	56/56 GB
RDS2	4501×11283	Mouse Brain	500/198 MB	5.1/5.1 TB

TABLE 4
Comparison With Compute-Centric Approach

		Preproc.	Reconst.	Per-Iter.	Speedup
ADS2	Trace	N/A	26.05 s	579 ms	1×
	MemXCT	4.00 s	0.53 s	11.8 ms	49.2×
RDS1	Trace	N/A	425.3 s	9.45 s	1×
	MemXCT	146.3 s	62.00 s	1.37 s	6.86×

shale sample [32] and RDS2 involves a brain sample. RDS1 is available open source [33], and RDS2 is proprietary.

4.1 Evaluating Overall Performance

4.1.1 Comparison With Compute-Centric Approach

We compare our memory-centric MemXCT with compute-centric Trace [11], an open-source high-performance implementation that employs SIRT iterations. To enable a one-to-one comparison, we implement SIRT, and run 45 iterations with both codes on a single KNL for the ADS2 and RDS1 datasets (see Table 3). Table 4 reports solution times and corresponding speedups. In the best case, where MemXCT

TABLE 5
Reconstruction on Various Nodes-Machines

Nodes-Machine	Preproc. Speed.	Recon. Speed.	All Slices
1-Theta (1 KNL)	139 s	1×	63.3 s 1×
8-Theta (8 KNL)	16.5 s	8.42×	3.33 s 19.0×
8-Cooley (16 K80)	25.5 s	5.45×	2.89 s 21.9×
32-Blue W. (32 K20X)	14.6 s	9.52×	1.82 s 34.8×
32-Theta (32 KNL)	4.54 s	30.6×	1.37 s 46.2×
32-Cooley (64 K80)	6.31 s	22.0×	1.22 s 51.9×

* Calculated based on single-slice execution time.

memory footprint fits within MCDRAM, it performs each iteration $49.2 \times$ faster than Trace. In the worst case, where MemXCT is bounded by slow DRAM bandwidth due to its large memory footprint, it still performs each iteration $6.86 \times$ faster. Trace memory footprint fits within MCDRAM in both cases, but it has to perform redundant and inefficient computation.

Table 4 reports preprocessing overhead of MemXCT. Although the preprocessing step appears to be a significant fraction of overall time, when it comes to many-slice reconstruction, the preprocessing cost is paid only once for the first slice. It is then reused for all the remaining slices, as shown in Table 5.

4.1.2 Iterative Convergence

MemXCT uses CG for iterative solutions, as discussed in Section 3.5.2, as opposed to SIRT used by Trace. Fig. 8a presents convergence properties by comparing L-curves of CG and SIRT up to 500 iterations: horizontal and vertical axes represents residual and reconstruction norms, respectively (see Section 2.2 for iterative formulation). As L-curve suggests, CG solution experiences overfitting soon after 30 iterations where the image does not further improve, but

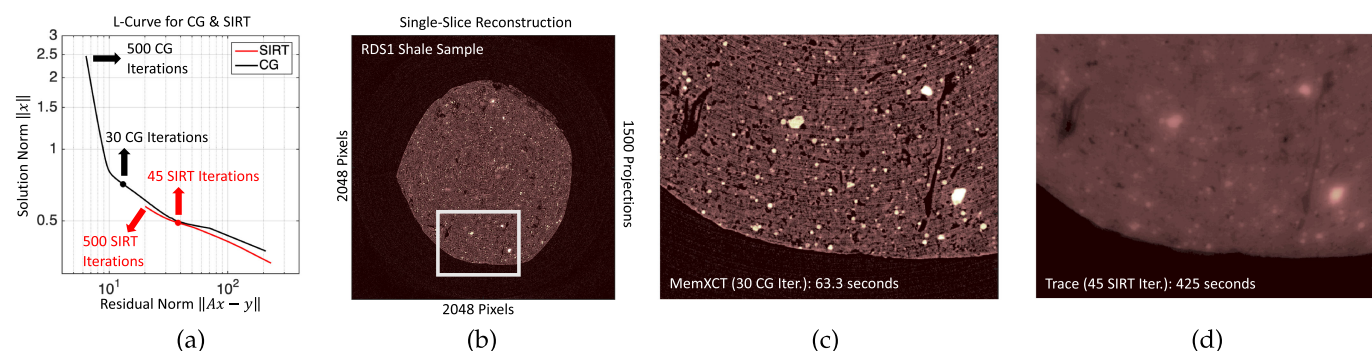


Fig. 8. For RDS1: (a) L-curves for CG and SIRT iterations. (b) Single-slice reconstruction with (c) CG and (d) SIRT iterations.

instead starts to be polluted by noise. Therefore we terminate CG solution after 30 iterations. On the other hand, SIRT does not converge even with 500 iterations. Fig. 8b shows a single-slice reconstruction from RDS1, and (c) and (d) compare image details after 30 CG iterations with MemXCT and 45 SIRT iterations with Trace, respectively.

4.1.3 Machine-Specific Considerations

We reconstruct RDS1 on small number of Theta, Cooley, and Blue Waters nodes. Machines were described briefly earlier. Unfortunately, reconstruction does not fit within DGX-1 or on less than eight nodes of Cooley or 32 nodes of Blue Waters, due to limited memory of their respective GPUs and the large memory complexity of MemXCT. Nevertheless, reconstruction fits well into a single Theta node thanks to its large DRAM capacity.

Table 5 reports RDS1 preprocessing and reconstruction times on various numbers of nodes-machines. Results show that 32 nodes of all machines enjoy comparable time to solution: reconstruction of all slices reduces from 1.44 days down to about or less than an hour. Preprocessing time also scales well and, in fact, it is insignificant compared to reconstruction time of all slices. Table 5 also shows the super-linear speedup property of MemXCT on KNLs: 8-Theta is $19\times$ faster than 1-Theta. This is a result of shrinking per-node memory footprint and extra memory bandwidth gain when it fits within 16 GB MCDRAM capacity. The super-linear speedup also demonstrates effective domain decomposition and reduced communications described in Section 3.

4.2 Performance Optimizations

This subsection presents single-device performance gains due to the optimizations described in Section 3. Fig. 9 shows forward and backprojection performance metrics for ADS1 through ADS4 on KNL and GPU. ADS3 and ADS4 are too large to fit in a single GPU (see Table 3). Hilbert ordering and input buffering are applied to the baseline in order because the first optimization enables the second. Since performance is dependent on both dataset and device, we tune all results (including the baseline) independently for maximum GFLOPS as described in Section 4.2.4.

Since there are two FLOPs per non-zero element in the projection matrix (one multiplication and one addition per FMA), we calculate the GFLOPS metric as $2N_{nz}/t$, where t is the time measured for a single forward/back-projection. Similarly, we calculate average memory bandwidth utilization for regular data only as $N_{nz} \times B_{reg}/t$, where B_{reg} is regular data (in bytes) read from memory per FMA, respectively. L2 miss rates are measured by Intel VTune profiler. The performance metrics are aggregated over many iterations.

4.2.1 The Baseline

GFLOPS on KNL drops as dataset size increases, regardless of whether or not it fits into high-bandwidth MCDRAM. The reason for this behaviour is that the baseline is bounded by memory latency rather than memory bandwidth, due to the high L2 miss rates of irregular access. As a result, larger datasets suffer more performance degradation due to their higher L2 miss rates, as seen in Fig. 9b. In contrast, GPU

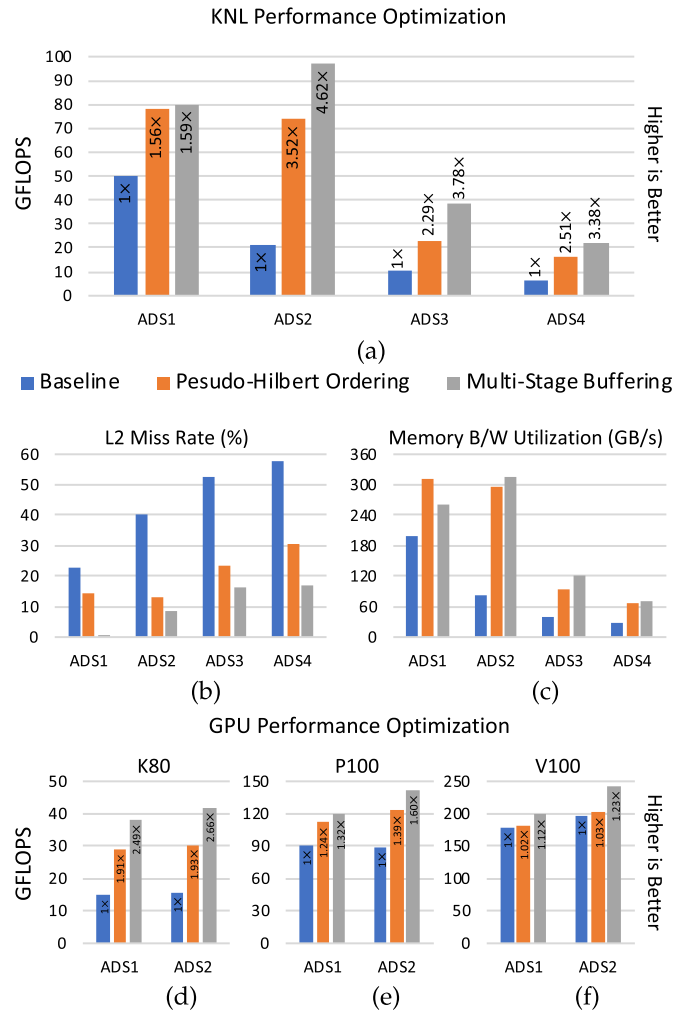


Fig. 9. KNL performance: (a) GFLOPS, (b) L2 miss rate, and (c) memory bandwidth utilization. GPU GFLOPS performance: (d) K80, (e) P100, and (f) V100.

performance improves slightly with larger datasets, as more parallelism hides the latency.

4.2.2 Pseudo-Hilbert Ordering

provides data locality and reduces the L2 miss rates for all datasets as seen in Fig. 9b. ADS1 does not benefit from Hilbert ordering as much as other datasets due to its small size. Nevertheless, as opposed to the baseline, the performance of Hilbert ordering is bounded by memory bandwidth consumed by regular data rather than memory latency. On KNL, regular data smaller than 16 GBs (ADS1 and ADS2) fits completely into MCDRAM and performs better whereas large regular data (ADS3 and ADS4) does not and is bounded by low DRAM memory bandwidth. On GPUs, the improvement is less noticeable with larger L2 cache because irregular data are better cached. As a result, pseudo-Hilbert ordering speeds up K80, P100, and V100 baselines by about $1.93\times$, $1.39\times$, and $1.03\times$, respectively.

This analysis shows that ADS1 and ADS2 use at least 78% and 74% of the theoretical MCDRAM bandwidth, respectively. The ADS3 regular data (28 GB) is larger than MCDRAM (16 GB) where it is cached partially. We cannot say the same for ADS4 because it has too large regular data




Fig. 10. Tuning input buffer and block size for (a)-(c) various SMT/core on KNL and (d) for V100.

to be significantly cached at MCRAM. As a result, we deduce that ADS4 uses at least 73% of the theoretical DRAM bandwidth. These utilizations agree well with STREAM benchmarks in previous work [34]. Similarly, K80, P100 and V100 use 78%, 69%, and 92% of their theoretical HBM bandwidths, respectively. These numbers agree well with benchmark measurements in previous work [35].

4.2.3 Multi-Stage Buffering

as we described in Section 3.3, comes with an staging overhead. On KNL, the gain amortizes the overhead when dataset is large, as seen in ADS2 and up. ADS1 is not large enough to show any further performance improvement with input buffering. The bandwidth utilization for input buffering in Fig. 9c and Figs. 9d, 9e, and 9f is adjusted with respect to additional memory-to-buffer mappings as well as reduced bytes for buffer-address indices.

It is worth to comment that reduced L2 miss rate saves significant memory bandwidth on K80 since its utilization due to regular data increases to at least 67% of theoretical peak as seen in Fig. 9d. The respective utilizations on P100 and V100 drop slightly, if not remain the same, because their L2 miss rates are already low and bandwidth utilizations are already high thanks to Hilbert ordering and their large L2 capacity. As a result, we can deduce that GFLOPS improvements on P100 and V100 are solely provided by the bandwidth saving due to reduced number of bytes needed for addressing shared-memory.

4.2.4 Tuning

The baseline and Hilbert ordering are relatively simple on both KNL and GPU architectures. We perform an exhaustive search and find out that blocks size of 128 scheduled dynamically among 128 threads (2 SMT/core) provides good single KNL performance for ADS1 through ADS4. Similarly, block size of 32 or 64 provides good single GPU performance.

For input buffering optimization, parameters should be re-tuned along with buffer size for effective SMT utilization

TABLE 6
Comparison With MKL and cuSPARSE for ADS2

	KNL	K80	P100	V100
MKL/cuSPARSE	1×	1×	1×	1×
MemXCT Baseline	1.42×	0.52×	1.39×	1.79×
Pseudo-Hilbert Ordering	4.99×	1.13×	1.93×	1.84×
Multi-Stage Buffering	6.55×	1.56×	2.23×	2.11×

on KNL. Figs. 10a, 10b, and 10c shows GFLOPS heat maps for ADS2 with various block and buffer sizes, and different numbers of SMT/core. Results show that input buffering can use more SMTs per core because it finds opportunity to overlap buffer stagings and accesses among SMTs. The following factors need to be considered when using SMTs in this way: A larger block size increases the per-block memory footprint and thus data reuse from the buffer, but also increases the number of buffer stagings per block, which comes with an overhead. In that case, it is better to increase buffer size for limiting the number of stagings. However, too few stagings decreases the opportunity to overlap buffer stagings and accesses among SMTs. Also, too large a buffer size can result in leaks to L2.

On KNL, peak GFLOPS performance for ADS1 and ADS2 is achieved with 4 SMT/core and buffer size of 8 KB. Tuning for ADS3 and ADS4 is cumbersome since they are severely limited by DRAM bandwidth and SMT effect is not that apparent. Therefore we use 2 SMT/core, and buffer size of 16 KB for them. Block size of 128 provides good performance for all datasets.

On GPUs block size of 512 or 1024 and buffer size of 48 KB or 96 KB provides good single-GPU performance for all GPUs. It is worth to note that addressable shared-memory size is limited to 48 KB on P100 and K80 and therefore tuning space is half the size. We follow a similar tuning strategy for GPUs as explained for KNL, as seen in Fig. 10d.

4.2.5 Comparison With Existing Libraries

We compare MemXCT forward/backprojection kernel performance with SpMV functions of existing libraries. We use Intel MKL library on KNL, and Nvidia cuSPARSE library on GPU. Table 6 reports speedups of our MemXCT baseline implementation and optimizations compared to CSR (on KNL) and column-major ELL (on GPU) SpMV of respective libraries. Evidently, MemXCT baseline implementation is faster than the existing libraries, with the exception for K80. This is due to K80 small L2 cache compared to P100 and V100 GPUs. On P100 and V100, our baseline is 1.39× and 1.79× faster since cuSPARSE pads CSR matrix with -1 [36] which requires extra branches whereas we pad with 0 and perform redundant multiplication with 0 to avoid thread divergence on GPUs. Also, cuSPARSE pads CSR on a matrix level whereas we pad on a thread-block (partition) level, which saves some of the redundant operations. As a result, we outperform existing libraries thanks to our application-specific SpMV optimizations. Nevertheless, both MKL and cuSPARSE are optimized for general SpMV whereas our kernels are optimized for XCT application.

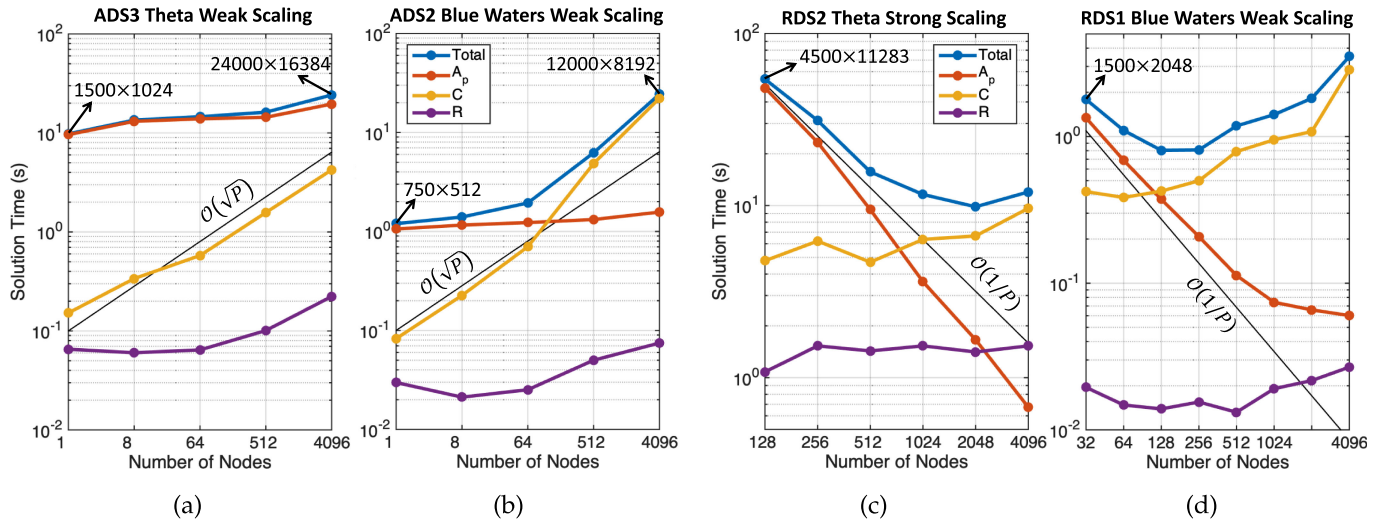


Fig. 11. Weak scaling: (a) ADS3/Theta, (b) ADS2/Blue Waters. Strong scaling: (c) RDS2/Theta, (d) ADS3/Blue Waters.

4.3 Scalability

For presenting scaling of MemXCT on a large number of nodes, we perform weak and strong scaling experiments on Theta and Blue Waters. For each experiment, we perform a full solution with 30 CG iterations, and report total reconstruction times as well as A_p , C , and R kernel times as they are defined in Section 3.4. Kernel times involve both forward and backprojection. In order to accurately measure each individual kernel, we placed the necessary artificial barriers, though they slightly increase total execution time. Among the kernels, C involves inter-node communication therefore we include its host-device communication times. For demonstrating scaling, $\mathcal{O}(\sqrt{P})$ and $\mathcal{O}(1/P)$ curves are included in Fig. 11 as well as the sinogram data dimensions.

4.3.1 Weak Scaling

Each weak scaling experiment is performed by starting from the reconstruction of a root dataset, and then doubling the number of channels and projections in a sinogram on each step. As computational cost increases eight times per step, the number of nodes is increased eight times accordingly. Figs. 11a and 11b show weak scaling of root datasets ADS3 and ADS2 on Theta and Blue Waters, respectively. Both experiments exhibit good weak scaling except communication operations C because of its $\mathcal{O}(\sqrt{P})$ complexity, as explained in Table 1. On Blue Waters, weak scaling is bounded by communication on 512 nodes and up; on Theta, communication is not the bounding component, but rather A_p . The difference is due to architectural differences across machines.

4.3.2 Strong Scaling

For strong scaling, we reconstruct RDS2 (brain) and RDS1 (shale) samples on Theta and Blue Waters respectively. We increase numbers of nodes and the dataset sizes are fixed. The minimum number of nodes needed for RDS2 and RDS1 to fit well within their respective systems is 128 and 32 nodes respectively. Reconstructions are scaled up to 4096 nodes of each system. Figs. 11c and 11d show strong scaling results. Theta exhibits good scaling up to 2048 nodes where as Blue

TABLE 7
Comparison of Theta and Blue Waters

	RDS1	RDS2	12000×8192
BW	805 ms (128 K20X)	74 s (4096 K20X)	24.4 s (4096 K20X)
Theta	474 ms (128 KNL)	10 s (2048 KNL)	3.25 s (4096 KNL)

Waters scales up to 128 nodes. The main reason (apart from the difference in network bandwidth and topology) is that the RDS2 solution is $\sim 91 \times$ more costly than the RDS1, and therefore the former scales better than the latter. Also, A_p kernel exhibits super-linear speedup, demonstrating efficient domain partitioning and high-bandwidth memory utilization, as discussed in Section 4.1.3. RDS2 reconstruction results are in Fig. 1. In contrast, A_p performance saturates on 1024 nodes of Blue Waters (and up) since smaller per-node computation penalizes GPU performance.

4.3.3 Comparison of Theta and Blue Waters Systems

For cross-comparing Theta and Blue Waters systems, we pick the fastest reconstructions of both RDS datasets and their corresponding weak scaling versions when ran on 4096 nodes. Results are shown in Table 7. RDS1 (shale) sample reconstruction runs fastest on 128 nodes on both Theta and Blue Waters. In this case, Theta is about $1.7 \times$ faster than Blue Waters. RDS2 (brain) sample reconstruction runs fastest on 2048 nodes of Theta and requires at least 4096 nodes to fit well into Blue Waters. In this case, Theta is $7.4 \times$ faster than Blue Waters. Lastly, Theta is about $7.5 \times$ faster on reconstructing the 12000×8192 dataset.

5 RELATED WORK AND DISCUSSION

Iterative reconstruction algorithms provide superior image quality considering analytical approaches, however their usage has so far been limited with their computational demands [37], [38], [39].

The multicore parallelization of iterative reconstruction algorithms has long been studied [40], [41], [42], [43]. However, most prior approaches provide limited scalability and only consider optimizations in the object domain. Further,

they suffer from redundant computations when dealing with large datasets.

With the emergence of high-throughput many-core architectures and state-of-the-art reconstruction algorithms, the applicability of iterative techniques has become more feasible [13], [24], [44], [45]. Especially in medical imaging, iterative reconstruction approaches have been extensively used to provide high-quality 3D reconstructions [46], [47], [48] in order to limit patient radiation exposure [4], [49]. Most of these approaches rely on GPUs to meet computational requirements of these algorithms [50], [51], [52]. Although GPUs can provide sufficient computational throughput, their limited memory can accommodate only small- to medium-scale datasets. For large-scale dataset host-device communication is known to be dominating overhead in GPUs [53]. Further, most of these approaches rely on on-the-fly computation of ray tracing, i.e., CompCXT type of data access pattern, which shows sub-optimal performance compared to MemCXT.

Li *et al.* developed cuMBIR [17], a framework for model-based iterative reconstruction on GPUs. They implement two solvers ICD and SGD, and propose several optimizations including unified thread mapping. Their method is also based on on-the-fly computation of A matrix, and performs redundant computations. Further, their solution focuses on small to medium scale datasets that can fit into single node GPUs. In comparison, MemCXT is optimized for very large tomographic datasets that require not only single node performance but also efficient large-scale execution.

Wang *et al.* [44] use KNL many-core capabilities to reconstruct 1024^2 tomograms (or 1024^3 3D volumes). They introduce the *non-uniform parallel super-voxel* to exploit sinusoidal bands in sinograms and optimize data access patterns. Our approach shares similarities with their work, but we consider parallelization and performance optimizations in both the tomogram and sinogram domains. Moreover, we further improve data access patterns and communication using two-level pseudo Hilbert ordering and multi-stage input buffering on both domains. Our evaluation also extends to extremely large objects (11293^3 rather than 1024^2), showing the usability of our approach for very large datasets.

In a more recent work by our group, we further extended the optimizations proposed in this paper and improved the reconstruction performance with high-end GPU clusters [54]. Specifically, we applied hierarchical communications and reductions to minimize inter-node communication overhead, improved the memory access performance with 3D batch parallelism, and used mixed-precision types to maximize throughput and bandwidth utilization. These optimizations enabled reconstruction throughput to reach 34% of Summit's peak performance (or 65 PFLOPS).

The two-level pseudo Hilbert ordering technique used in this work is a form of space-filling curves. To our knowledge, the first space-filling curve example was presented by Peano in 1890 [55]. Later, space-filling curves have been successfully applied to improve performance of many applications [56], [57], [58], [59], [60], [61], [62]. In [63], Mellor-Crummey *et al.* showed that data and computation reordering based on space-filling curves can improve the performance of irregular applications significantly. Moon *et al.*

analyzed clustering properties of Hilbert space-filling curve [64] on different query shapes [65]. Reissmann *et al.* showed the effects of Hilbert and Morton curves [66] on energy and locality [62]. In our work, we applied two-level pseudo Hilbert ordering to optimize the performance of tomographic reconstruction data access patterns and inter-process communication.

6 REPRODUCIBILITY

Reproducibility of experimental results is essential to scientific research. To promote the reproducibility of experimental results in the HPC research community, the SC Student Cluster Reproducibility Committee has, since 2016, selected one of the papers from the past year's SC conference to serve as the benchmark for the Reproducibility Challenge [67]. Based on our paper's artifact descriptor and its suitability to the SCC, we were deeply honored that our SC19 paper [1] was selected to serve as the 2020 Reproducibility Challenge benchmark.¹ Nineteen teams, each of which is composed of six college students and their advisor who design, construct and utilize their own clusters to run computational experiments, were selected to participate SCC at SC20 and were tasked to replicate the findings of the publication.²

This section describes the pertinent features of the MemXCT codebase and collateral materials for the Reproducibility Challenge. Some of these features were already in our repository so that future papers can reproduce our results for comparison purposes. Some were added and/or enhanced in consultation with the Student Cluster Reproducibility Committee. We summarize the key findings of the nine top-scored teams and present the insights that can be potentially applicable to other HPC research experiments.

6.1 Datasets Used for Reproducibility

To support the reproducibility and for benchmarking, we provide *Test Datasets* and *Challenge Datasets* to the participating teams.

6.1.1 Test Datasets

The Test datasets are released two weeks in advance to the SCC kickoff. This allowed students to: 1) Compile and run the codes, 2) Verification by visualizing the results, and 3) Tuning system performance according to the teams' respective hardware. The Test datasets comprise ADS1, ADS2, ADS3, and ADS4, the artificial portion of the benchmark set that were used to generate the results presented in this paper. These benchmark datasets were chosen to represent very different points in the four dimensions (columns) shown in Table 3.

A major constraint in MemXCT is the memory footprint, which grows with respect to the memory complexity as formulated in Table 1. For example, ADS1 and ADS2 are small enough to fit into a single GPU whereas ADS3 and ADS4

1. Please see the blog post entitled "SC20 Student Cluster Reproducibility Committee Chooses Benchmark Wisely" by Scott Michael and Stephen Harrell for details.

2. Please see the blog post entitled "Global Lineup Will Compete at SC's First Virtual Student Cluster Competition" by Christine Baissac-Hayden for details.

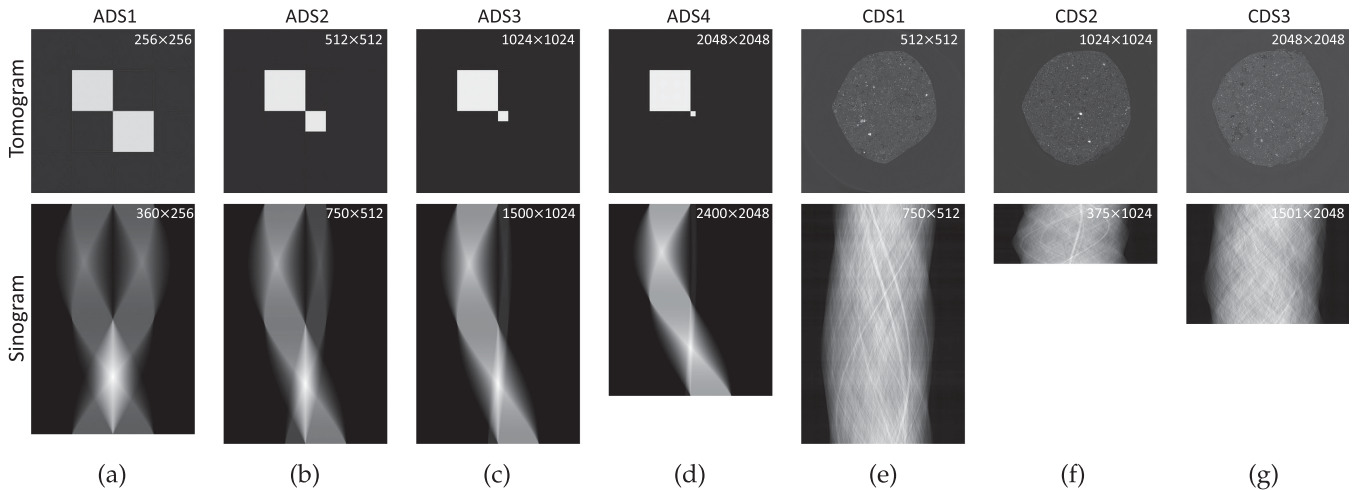


Fig. 12. (a)–(d) Artificial *Test Datasets* are used for performance benchmarks for the original paper and at the reproducibility challenge. (e)–(g) Real *Challenge Datasets* that are introduced for additional reproduction.

TABLE 8
Challenge Datasets Used for Reproducibility

Name	Sinogram ($M \times N$)	Sample	Irregular Data	Regular Data
CDS1	750×512	Shale Rock	1.0/1.5 MB	1.8/1.8 GB
CDS2	375×1024	Shale Rock	4.0/1.5 MB	3.5/3.5 GB
CDS3	1501×2048	Shale Rock	16/12 MB	56/56 GB

are larger. Therefore, ADS3 and ADS4 datasets are easier to accommodate in KNL nodes. Providing the exactly same benchmark datasets used in our original experiments gives students the opportunity to directly compare their performance results to those in the original paper.

For verification, we asked students to visualize their input (sinogram) and output (tomogram) datasets by using the Fiji software [68], although, students are free to use any other visualization tool. They import the raw output image files and inspect these to verify the execution of their code. The input (sinogram) and output (tomogram) Test Datasets are shown in Figs. 12a, 12b, 12c, and 12d.

6.1.2 Challenge Datasets

We revealed the Challenge Datasets by the time of the SCC. These datasets comprise CDS1, CDS2, and CDS3 of different dimensions and are listed in Table 8 along with their memory footprints. We generated these datasets by rescaling the sinogram of RDS1 in Table 3. Each scaling was applied to a different slice of the original 3D dataset and the output tomograms were not revealed to the students. We matched the dimensions of CDS1 and CDS3 with those of ADS2 and RDS1, respectively, to provide students a one-to-one benchmarking opportunity. We chose dimension of CDS2 large enough to require substantial computational throughput but small enough to still fit into a single GPU memory. The sinogram and tomogram of the Challenge Datasets are shown in Figs. 12e, 12f, and 12g. All datasets were provided through the repositories as explained in the next subsection.

Authorized licensed use limited to: Stanford University. Downloaded on August 28, 2024 at 06:10:46 UTC from IEEE Xplore. Restrictions apply.

6.2 Codebases and Repositories

The CPU and GPU codebases that were used for performing benchmarks in this paper and for the SCC reproducibility challenge were provided in respective GitHub repositories:

<https://github.com/merthidayetoglu/MemXCT-CPU>

<https://github.com/merthidayetoglu/MemXCT-GPU>

The MemXCT application requires MPI and C++ compiler (and CUDA compiler for GPU). Instructions for compiling, job execution, and preparing input file were provided to the participant teams. The repositories include persistent links to the Testing Datasets ADS1, ADS2, ADS3, and ADS4 and Challenge Datasets CDS1, CDS2, and CDS3. We also included the visualizations of the Test Datasets in the repository, for student teams to verify their results against any bugs.

6.3 Machines Used for Reproducibility

In a typical year, each of the SCC teams designs and assembles their own cluster on the SC show floor. This often results in a very diverse set of hardware and software configurations. However, due to the COVID-19 pandemic, the SC20 SCC was completely virtual. Student teams utilized virtual machines provided by Microsoft Azure cloud service. Although, the students had access to a wide variety of Azure instances, some of the more exotic hardware configurations that had been seen in the past with the physical competition were not available. Details of the virtual machine configurations and the teams that utilized them are shown in Table 9. Each team configured their own software environment, including OS, drivers, compilers, and MPI implementations. The details of the software configurations can be found in teams' critiques in the special issue herein.

6.4 Reproducibility Results

For reproducibility, we turned on all MemXCT optimizations on, i.e., pseudo-Hilbert ordering and multi-stage buffering, and asked students to first replicate the computational throughput (in GFLOPS) and bandwidth utilization (in GB/

TABLE 9
Machines on Azure Used for Reproducibility

Machine	CPU	Cores	GPU	Interconnect	Team(s)
NC24r Promo	Intel Xeon E5-2690 v3	2 x 12	4 x NVIDIA K80	N/A	Tsinghua University
NC24rs v2	Intel Xeon E5-2690 v4	2 x 12	4 x NVIDIA P100	ConnectX-3	ETH Zürich, Peking University, ShanghaiTech University, Tsinghua University
NC12s v2	Intel Xeon E5-2690 v4	1 x 12	2 x NVIDIA P100	N/A	Clemson, University of Texas
NC24rs v3	Intel Xeon E5-2690 v4	2 x 12	4 x NVIDIA V100	ConnectX-3	Nanyang Technological University, ETH Zürich, Georgia Institute of Technology, University of California San Diego, Tsinghua University
HC44rs	Intel Platinum 8168	2 x 21	N/A	ConnectX-6	ETH Zürich, University of Texas, Peking University
F16s v2	AMD EPYC 7551	1 x 32	N/A	ConnectX-4	Tsinghua University
HB60rs	AMD EPYC 7551	2 x 32	N/A	ConnectX-5	Nanyang Technological University, Tsinghua University
HB120rs v2	AMD EPYC 7V12	2 x 60	N/A	ConnectX-6	ETH Zürich, Georgia Institute of Technology, University of California San Diego, ShanghaiTech University, Tsinghua University

s) in Fig. 9 on a single CPU and single GPU. Then, we asked students to observe strong scaling from single node to four nodes.

6.4.1 Single-CPU Results

From sampling the student reports, we see that one of the challenges in reproducing the results is the resource utilization in the cloud that inherently includes virtualization overhead. Further, the cloud resources are much more limited, both in terms of computational throughput and bandwidth, compared to the (dedicated) high-end resources used in the original experiments. Nevertheless, the students seem to be able to verify that the application is indeed memory bound and the execution time is roughly proportional to the memory bandwidth. The top teams report that MemXCT can achieve high CPU utilization even though it is memory-bound, an important conclusion of the original experiments.

6.4.2 Single-GPU Results

In Fig. 13, we compare the top teams' results against the ones reported in the original paper. Since the teams didn't have access to KNL resources, we only show the GPU results. Specifically, the teams had access to the same

generations of GPUs, i.e., K80, P100, and V100, which were used in the original paper. Because of the limited cloud budget given to participating teams, not all teams reproduced results on all generations of GPUs except Tsinghua (the overall winner) and Clemson. Nevertheless, the top teams were able to reproduce the compute throughput and memory bandwidth results within the normal variation of experimental results with their available GPUs. This does not only confirm the integrity of the original measurements but also the stability of GPU execution speed in different deployment environments, i.e., supercomputers versus cloud clusters.

6.4.3 Scaling Results

As for scaling, since the teams didn't have access to large-scale supercomputers, such as Theta or Blue Waters, they scaled MemXCT only up to 4 nodes on the Azure cloud. Still, the majority of the teams seemed to reproduce strong-scaling results with the small number of nodes, and presented a similar trend shown in Fig. 11d. Since cloud resources provide weaker network resources compared to supercomputers, it was interesting to see that the scaling trends still hold. This suggests that the MemXCT application can be used in a modest-sized cloud computing cluster and still achieve desired scaling behavior.




Fig. 13. Reproduction of GPU throughput (in GFLOPS) and bandwidth utilization (in GB/s) through CDS1 (equivalent to ADS2) and CDS2 datasets. The reproduced benchmarkings on three generations of GPUs highly matches with the original results presented in Fig. 9c.

A specific performance trend that was not reproduced is the super-linear strong scaling as in Fig. 11c because it is specific to KNL architecture. That is, MemXCT performance is bounded by memory bandwidth and strong scaling have super-linear speedup property because when each partition fits into the 16-GB high-bandwidth MCDRAM on KNL, we observe an extra speedup. However, CPUs and GPUs that were available in the Azure cloud did not have high-bandwidth memory that would cache each partition, and therefore student teams could not report any super-linear strong scaling.

6.5 Further Insights on Reproducibility

In this section, we discuss a few challenges on different aspects of reproducibility of MemXCT and other applications in general.

6.5.1 Verification

Each application may require a different reproducibility approach. This mainly comes from the nature of the application, rather than the underlying computational workload. In MemXCT case, for example, the application reconstructs images, so the output is already visual and hence easy to verify by just looking at the screen, in a very specific way for MemXCT. The computational workload of the image reconstruction, however, is exactly the same as of any other iterative solution of a sparse linear system. Except, the technique proposed in this paper is highly optimized for the governing X-ray imaging problem. For verification, we initially thought about defining a distance metric between two images (or using an already-defined metric). However, a single number hardly provides any insight to the students about the correctness of the MemXCT system. Hence, we asked students to observe the iterative convergence of the result and visually verify its correctness by comparing it with the provided ground truth.

6.5.2 System Variations and Scaling

An important challenge experienced by the SCC teams was to reproduce the results when the systems used by the teams are different from the ones used for the original experiments. Hardware variations are obvious. The teams were not able to match the numerical execution speed and memory bandwidth results for CPUs simply because the hardware they used were much weaker than the CPUs used in the original experiments. However, they were able to reproduce similar ratios and trends. In contrast, when the teams used one or more of the GPUs in the original experiments, their results matched closely. A more subtle problem is software variations. Different OS, library versions, and virtualization schemes can affect the achievable execution speed in a significant way.

In the area of scaling, we have scaled the MemXCT application up to 4096 nodes of Blue Waters and Theta. The scale of these systems (described in Table 2) are not available for reproducibility as mentioned before, and the teams attempted to reproduce scaling trends only with 4 nodes (low-node-count regions of Fig. 11.) The MemXCT application reports the times spent for partial matrix multiplication, communication, and reduction shown as A_p , C , and R ,

respectively, in Fig. 11. Times required for these operations show different properties when the number of nodes is increased. For example, while computation (A_p) efficiently scales with an increasing number of nodes, communication (C) does not. Most of the teams are able to explain these performance properties. Several teams were able to observe the same strong-scaling trend in spite of the weaker networks in the cloud environment. This is likely due to the fact that in these low-node-count settings, the amount of work assigned to each node is still sufficient to hide the communication overhead in spite of the weaker networks in the cloud environment.

6.5.3 Benchmarking Across Systems

Comparing performance across systems, e.g., CPU and GPU, is another challenge for benchmarking and reproducibility. First, the fundamental differences of CPU and GPU architecture make them incomparable. That is, GPUs are extremely good when applications show data parallelism. For example, students report up to $10 \times$ speedup of a single GPU over a single CPU. The reported KNL performance in this paper, however, has a comparable performance with GPUs, as shown in Fig. 9. Second, CPUs have much more memory on the host compared to the device memory on GPU. This sometimes prevents one-to-one comparison because a large benchmark does not fit into a single GPU. We solve this problem by benchmarking over four artificial datasets with various sizes, among which ADS1 and ADS2 fit into device memories of both GPU and KNL. Lastly, since GPUs (and other accelerators) have smaller device memory, strong scaling on the whole system may not be feasible when starting from a single device. Mainly because the single-device workload is extremely small when scaled to thousands of devices. That's why we use a large RDS2 benchmark (with 10.2 TB memory footprint) on 128 KNLs as the starting point for strong scaling on Theta. Similarly, we use a smaller RDS1 benchmark (112 GB memory footprint) on 32 K20x as the starting point for strong scaling on Blue Waters. In both cases, we scale up to 4096 nodes of both systems, as shown in Figs. 11c and 11d. The figure clearly contrasts two systems by their super-linear and sub-linear scaling properties.

In general, we can obviously not expect the teams to reproduce the same performance results when they use different hardware/software systems. However, we did expect that their results show similar characteristics and trends to the original paper. In this respect, the top teams were indeed able to observe the desired application performance behavior in their test systems. This gives us confidence that future users of MemXCT will experience these desired performance behavior in their own systems.

7 CONCLUSION

We have described a memory-centric approach to iterative XCT reconstruction. Our implementation, MemXCT, performs forward projection and backprojection operations as explicit SpMV's with no on-the-fly (redundant) computations or race conditions; these features allow MemXCT to achieve up to $50 \times$ speedup over a high-performance compute-centric approach. Although MemXCT has a high

memory complexity, its per-node memory footprint decreases linearly with increasing number of nodes, which favors large-scale resources. For solving the performance bottlenecks, we propose and implement two-level pseudo-Hilbert ordering and multi-stage input buffering techniques in MemXCT. We compare the performance of our computational kernels on KNL and several generations of GPUs, and demonstrate scaling up to 4096 nodes using ALCF Theta (KNL) and NCSA Blue Waters (GPU) systems. We show that MemXCT reconstructs a large-scale (11K × 11K) mouse brain tomogram in 10 seconds using 4096 KNL nodes. Our experimental results were reproduced successfully by the 2020 Student Cluster Competition teams. We present the pertinent features of the MemXCT artifact and collateral materials that enabled teams to verify their results and compare against our results within a strict time limit. Even though the reproducibility approach we take is specific to MemXCT, similar reproducibility design can be adopted by future HPC research artifacts to allow other researchers to reproduce the reported experimental results.

ACKNOWLEDGMENTS

The authors would like to thank Stephen Harrell and Scott Michael for their invaluable guidance through the entire Reproducibility Challenge as well as the development of this article. We gratefully acknowledge the computing resources provided and operated by the Argonne Leadership Computing Facility, which is a U.S. Department of Energy, Office of Science User Facility. We thank Vincent De Andrade from Advanced Photon Source for sharing the mouse brain dataset. This research is part of the Blue Waters sustained-petascale computing project, which was supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. This research was based in part upon work supported by: The Center for Applications Driving Architectures (ADA), a JUMP Center co-sponsored by SRC and DARPA. This material was based in part upon work supported by the XPACC Center for Exascale Simulation of Plasma-Coupled Combustion and Department of Energy, under Award Number DE-NA0002374.

REFERENCES

- [1] M. Hidayetoğlu *et al.*, "MemXCT: Memory-centric X-ray CT reconstruction with massive parallelization," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 2019, Art. no. 85.
- [2] APS Science 2018, Advanced Photon Source, Argonne Nat. Lab., Lemont, IL, USA, Tech. Rep. ANL-18/40, p. 15, Jan. 2019, ISSN 1931-5007.
- [3] E. L. Dyer *et al.*, "Quantifying mesoscale neuroanatomy using X-ray microtomography," *eNeuro*, vol. 4, 2017, Art. no. ENEURO-0195.
- [4] E. Y. Sidky, C.-M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," *J. X-Ray Sci. Technol.*, vol. 14, no. 2, pp. 119-139, 2006.
- [5] X. Han *et al.*, "Algorithm-enabled low-dose micro-CT imaging," *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 606-620, Mar. 2011.
- [6] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Global Conf. Signal Inf. Process.*, 2013, pp. 945-948.
- [7] K. A. Mohan, S. V. Venkatakrishnan, L. F. Drummy, J. Simmons, D. Y. Parkinson, and C. A. Bouman, "Model-based iterative reconstruction for synchrotron X-ray tomography," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2014, pp. 6909-6913.
- [8] D. J. Duke *et al.*, "Time-resolved X-ray tomography of gasoline direct injection sprays," *SAE Int. J. Engines*, vol. 9, no. 1, pp. 143-153, 2016.
- [9] D. Gürsoy, T. Biçer, A. Lanzirrotti, M. G. Newville, and F. D. Carlo, "Hyperspectral image reconstruction for X-ray fluorescence tomography," *Opt. Exp.*, vol. 23, pp. 9014-9023, Apr. 2015.
- [10] D. Gürsoy, T. Biçer, J. D. Almer, R. Kettimuthu, S. R. Stock, and F. De Carlo, "Maximum a posteriori estimation of crystallographic phases in X-Ray diffraction tomography," *Philosoph. Trans. Roy. Soc. London A: Math. Phys. Eng. Sci.*, vol. 373, no. 2043, 2015, Art. no. 20140392.
- [11] T. Bicer *et al.*, "Trace: A high-throughput tomographic reconstruction engine for large-scale datasets," *Adv. Structural Chem. Imag.*, vol. 3, Jan. 2017, Art. no. 6.
- [12] T. Bicer, D. Gürsoy, R. Kettimuthu, F. De Carlo, G. Agrawal, and I. T. Foster, "Rapid tomographic image reconstruction via large-scale parallelization," in *Proc. 21st Int. Conf. Parallel Distrib. Comput.*, 2015, pp. 289-302.
- [13] W. van Aarle *et al.*, "The ASTRA Toolbox: A platform for advanced algorithm development in electron tomography," *Ultra-microscopy*, vol. 157, pp. 35-47, 2015.
- [14] D. Gürsoy, F. De Carlo, X. Xiao, and C. Jacobsen, "TomoPy: A framework for the analysis of synchrotron tomographic data," *J. Synchrotron Radiat.*, vol. 21, pp. 1188-1193, Sep. 2014.
- [15] A. Beer, "Bestimmung der absorption des rothen Lichts in farbigen flüssigkeiten," *Ann. Physik*, vol. 162, pp. 78-88, 1852.
- [16] R. L. Siddon, "Fast calculation of the exact radiological path for a three-dimensional ct array," *Med. Phys.*, vol. 12, no. 2, pp. 252-255, 1985.
- [17] X. Li *et al.*, "cuMBIR: An efficient framework for low-dose X-ray CT image reconstruction on GPUs," in *Proc. Int. Conf. Supercomput.*, 2018, pp. 184-194.
- [18] S. G. D. Gonzalo, S. Hammond, S. Huang, O. Mutlu, J. Gómez-Luna, and W. Hwu, "Automatic generation of warp-level primitives and atomic instructions for fast and portable parallel reduction on GPUs," *IEEE/ACM Int. Symp. Code Gener. Optim.*, pp. 73-84, 2019, doi: 10.1109/CGO.2019.8661187.
- [19] P. M. Campbell, K. D. Devine, J. E. Flaherty, L. G. Gervasio, and J. D. Teresco, "Dynamic octree load balancing using space-filling curves," *Williams College Dept. Comput. Sci., Tech. Rep. CS-03-01*, 2003.
- [20] M. Parashar and J. C. Browne, "On partitioning dynamic adaptive grid hierarchies," in *Proc. 29th Hawaii Int. Conf. Syst. Sci.*, 1996, pp. 604-613.
- [21] J. Zhang, S.-I. Kamata, and Y. Ueshige, "A pseudo-hilbert scan algorithm for arbitrarily-sized rectangle region," in *Proc. Int. Workshop Intell. Comput. Pattern Anal. Synthesis*, 2006, pp. 290-299.
- [22] Y. Wang *et al.*, "A high-throughput X-ray microtomography system at the advanced photon source," *Rev. Sci. Instrum.*, vol. 72, no. 4, pp. 2062-2068, 2001.
- [23] H. Wang, W. Liu, K. Hou, and W.-C. Feng, "Parallel transposition of sparse data structures," in *Proc. Int. Conf. Supercomput.*, 2016, Art. no. 33.
- [24] X. Wang, A. Sabne, S. Kisner, A. Raghunathan, C. Bouman, and S. Midkiff, "High performance model based image reconstruction," *ACM SIGPLAN Notices*, vol. 51, 2016, Art. no. 2.
- [25] R. Barrett *et al.*, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, vol. 43. Philadelphia, PA, USA: SIAM, 1994.
- [26] ALCF, "Theta | Argonne leadership computing facility." Accessed: Jan. 16, 2019. [Online]. Available: <https://www.alcf.anl.gov/theta>
- [27] NCSA, "Blue waters user portal | home." Accessed: Jan. 16, 2019. [Online]. Available: <https://bluewaters.ncsa.illinois.edu/>
- [28] ALCF, "Cooley | Argonne leadership computing facility." Accessed: Jan. 16, 2019. [Online]. Available: <https://www.alcf.anl.gov/user-guides/cooley>
- [29] IBM, "IBM power system S822LC for high performance computing - overview." Accessed: Mar. 20, 2019. [Online]. Available: <https://www.ibm.com/us-en/marketplace/high-performance-computing>
- [30] Nvidia, "Deep learning server for AI research | NVIDIA CGX-1." Accessed: Mar. 20, 2019. [Online]. Available: <https://www.nvidia.com/en-us/data-center/dgx-1/>
- [31] OLCF, "Titan user guide - Oak ridge leadership computing facility." Accessed: Jan. 16, 2019. [Online]. Available: <https://www.olcf.ornl.gov/for-users/system-user-guides/titan/titan-user-guide/#nvidia-k20x-gpus>
- [32] W. Kaniitpanyacharoen *et al.*, "A comparative study of X-ray tomographic microscopy on shales at different synchrotron facilities: ALS, APS and SLS," *J. Synchrotron Radiat.*, vol. 20, no. 1, pp. 172-180, 2013.

- [33] F. De Carlo *et al.*, "TomoBank: A tomographic data repository for computational X-Ray science," *Meas. Sci. Technol.*, vol. 29, no. 3, 2018, Art. no. 034004.
- [34] I. B. Peng, R. Gioiosa, G. Kestor, P. Cicotti, E. Laure, and S. Markidis, "Exploring the performance benefit of hybrid memory system on HPC environments," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. Workshops*, 2017, pp. 683–692.
- [35] Z. Jia, M. Maggioni, B. Staiger, and D. P. Scarpazza, "Dissecting the NVIDIA Volta GPU architecture via microbenchmarking," 2018. [Online]. Available: <https://arxiv.org/abs/1804.06826>
- [36] Nvidia, "CUSPARSE Library." Accessed: Jan. 16, 2019. [Online]. Available: https://developer.download.nvidia.com/compute/DevZone/docs/html/CUDALibraries/doc/CUSPARSE_Library.pdf
- [37] M. Beister, D. Kolditz, and W. A. Kalender, "Iterative reconstruction methods in X-Ray CT," *Physica Medica*, vol. 28, no. 2, pp. 94–108, 2012.
- [38] J. Hsieh, B. Nett, Z. Yu, K. Sauer, J.-B. Thibault, and C. A. Bouman, "Recent advances in CT image reconstruction," *Current Radiol. Rep.*, vol. 1, no. 1, pp. 39–51, 2013.
- [39] J. Amanatides and A. Woo, "A fast voxel traversal algorithm for ray tracing," in *Proc. Eur. Comput. Graph. Conf. Exhib.*, 1987, pp. 3–10.
- [40] J. Agulleiro and J.-J. Fernandez, "Fast tomographic reconstruction on multicore computers," *Bioinformatics*, vol. 27, no. 4, pp. 582–583, 2011.
- [41] J. Treibig, G. Hager, H. G. Hofmann, J. Hornegger, and G. Wellein, "Pushing the limits for medical image reconstruction on recent standard multicore processors," *Int. J. High Perform. Comput. Appl.*, vol. 27, pp. 162–177, 2012.
- [42] C. Johnson and A. Sofer, "A data-parallel algorithm for iterative tomographic image reconstruction," in *Proc. 7th Symp. Front. Massively Parallel Comput.*, 1999, pp. 126–137.
- [43] M. D. Jones, R. Yao, and C. P. Bhole, "Hybrid MPI-OpenMP programming for parallel OSEM PET reconstruction," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 5, pp. 2752–2758, Oct. 2006.
- [44] X. Wang, A. Sabne, P. Sakdhnagool, S. J. Kisner, C. A. Bouman, and S. P. Midkiff, "Massively parallel 3D image reconstruction," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 2017, Art. no. 3.
- [45] J. Qi and R. M. Leahy, "Iterative reconstruction techniques in emission computed tomography," *Phys. Med. Biol.*, vol. 51, no. 15, 2006, Art. no. R541.
- [46] C.-Y. Chou, Y.-Y. Chuo, Y. Hung, and W. Wang, "A fast forward projection using multithreads for multirays on GPUs in medical image reconstruction," *Med. Phys.*, vol. 38, no. 7, pp. 4052–4065, 2011.
- [47] G. Pratz, G. Chinn, P. D. Olcott, and C. S. Levin, "Fast, accurate and shift-varying line projections for iterative reconstruction using the GPU," *IEEE Trans. Med. Imag.*, vol. 28, no. 3, pp. 435–445, Mar. 2009.
- [48] D. Lee, I. Dinov, B. Dong, B. Gutman, I. Yanovsky, and A. W. Toga, "CUDA optimization strategies for compute-and memory-bound neuroimaging algorithms," *Comput. Methods Programs Biomed.*, vol. 106, no. 3, pp. 175–187, 2012.
- [49] B. Jang, D. Kaeli, S. Do, and H. Pien, "Multi GPU implementation of iterative tomographic reconstruction algorithms," in *Proc. IEEE Int. Symp. Biomed. Imag.: Nano Macro*, 2009, pp. 185–188.
- [50] A. Sabne, X. Wang, S. J. Kisner, C. A. Bouman, A. Raghunathan, and S. P. Midkiff, "Model-based iterative CT image reconstruction on GPUs," in *Proc. 22nd ACM SIGPLAN Symp. Princ. Practice Parallel Program.*, 2017, pp. 207–220.
- [51] S. S. Stone *et al.*, "Accelerating advanced MRI reconstructions on GPUs," *J. Parallel Distrib. Comput.*, vol. 68, no. 10, pp. 1307–1318, 2008.
- [52] F. Xu and K. Mueller, "Accelerating popular tomographic reconstruction algorithms on commodity PC graphics hardware," *IEEE Trans. Nucl. Sci.*, vol. 52, no. 3, pp. 654–663, Jun. 2005.
- [53] T. B. Jablin, P. Prabhu, J. A. Jablin, N. P. Johnson, S. R. Beard, and D. I. August, "Automatic CPU-GPU communication management and optimization," *ACM SIGPLAN Notices*, vol. 46, pp. 142–151, 2011.
- [54] M. Hidayetoğlu *et al.*, "Petascale XCT: 3D image reconstruction with hierarchical communications on multi-GPU nodes," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, 2020, pp. 1–13.
- [55] G. Peano, "Sur une courbe, qui remplit toute une aire plane," *Mathematische Annalen*, vol. 36, no. 1, pp. 157–160, 1890.
- [56] M. Bader, *Space-Filling Curves: An Introduction With Applications in Scientific Computing*, vol. 9. Berlin, Germany: Springer, 2012.
- [57] F. Günther, M. Mehl, M. Pögl, and C. Zenger, "A cache-aware algorithm for PDEs on hierarchical data structures based on space-filling curves," *SIAM J. Sci. Comput.*, vol. 28, no. 5, pp. 1634–1650, 2006.
- [58] H. Sagan, *Space-Filling Curves*. Berlin, Germany: Springer, 2012.
- [59] M. Bader and C. Zenger, "Cache oblivious matrix multiplication using an element ordering based on a Peano curve," *Linear Algebra Appl.*, vol. 417, no. 2/3, pp. 301–313, 2006.
- [60] Witten and Neal, "Using Peano curves for bilevel display of continuous-tone images," *IEEE Comput. Graphics Appl.*, vol. 2, no. 3, pp. 47–52, May 1982.
- [61] J. K. Lawder and P. J. H. King, "Querying multi-dimensional data indexed using the hilbert space-filling curve," *ACM SIGMOD Rec.*, vol. 30, no. 1, pp. 19–24, 2001.
- [62] N. Reissman, J. C. Meyer, and M. Jahre, "A study of energy and locality effects using space-filling curves," in *Proc. IEEE Int. Parallel Distrib. Process. Symp. Workshops*, 2014, pp. 815–822.
- [63] J. Mellor-Crummy, D. Whalley, and K. Kennedy, "Improving memory hierarchy performance for irregular applications using data and computation reorderings," *Int. J. Parallel Program.*, vol. 29, no. 3, pp. 217–247, 2001.
- [64] D. Hilbert, *Über die Stetige Abbildung einer Linie auf ein Flächenstück*. Berlin, Germany: Springer, 1935, pp. 1–2.
- [65] B. Moon, H. V. Jagadish, C. Faloutsos, and J. H. Saltz, "Analysis of the clustering properties of the hilbert space-filling curve," *IEEE Trans. Knowl. Data Eng.*, vol. 13, no. 1, pp. 124–141, Jan./Feb. 2001.
- [66] G. M. Morton, "A computer oriented geodetic data base and a new technique in file sequencing," 1966. [Online]. Available: <https://dominoweb.draco.res.ibm.com/reports/Morton1966.pdf>
- [67] M. A. Heroux and C. K. Garrett, "Special issue on SC16 student cluster competition reproducibility initiative," *Parallel Comput.*, vol. 70, pp. 3–4, 2017. [Online]. Available: <https://doi.org/10.1016/j.parco.2017.10.002>
- [68] J. Schindelin *et al.*, "Fiji: An open-source platform for biological-image analysis," *Nat. Methods*, vol. 9, no. 7, pp. 676–682, 2012.




Mert Hidayetoğlu is currently working toward the PhD degree at the University of Illinois at Urbana-Champaign, Champaign, Illinois with research at the intersection of large-scale applications, high-performance computing, and software systems. His research interests include unstructured data accesses, communications, and computations on supercomputers with multi-GPU node architecture. He was a Givens Fellow with Argonne National Laboratory and is currently a member of IBM-Illinois Center for Cognitive Computing Research. He is a recipient of the SC20 Best Paper Award and HPEC'20 Graph Challenge Champion Award, and 2021 ACM-IEEE CS George Michael Memorial HPC Fellowship.



Tekin Biçer (Senior Member, IEEE) received the master's and PhD degrees from Computer Science and Engineering Department, Ohio State University, Columbus, Ohio. He is currently a computer scientist with Data Science and Learning Division, Argonne National Laboratory. He also holds joint appointments in the X-Ray Science Division at Advanced Photon Source and Consortium for Advanced Science and Engineering, University of Chicago. His research interests include large-scale, high-performance, parallel and distributed computing systems. He is a senior member of the ACM.



Simon Garcia de Gonzalo received the PhD degree from Computer Science Department, University of Illinois at Urbana-Champaign, Champaign, Illinois. He is currently a postdoctoral researcher and STARS 2020 Fellow with Barcelona Supercomputing Center (BSC), Computer Sciences Department. He works within the Accelerators and Communications for HPC Team. His research interests include area of performance portability, runtime systems, and programming models for high-performance computing.



Bin Ren received the PhD degree from the Department of Computer Science and Engineering, Ohio State University, Columbus, Ohio. He is currently an assistant professor with the Department of Computer Science, College of William & Mary. His research interests include compilers, parallel computing, software systems, real-time machine learning, and machine learning systems. He has published more than 50 peer-reviewed papers in leading venues related to these areas, including ASPLOS, PLDI, PPOPP, SC, ICML,

AAAI, and *IEEE Transactions on Parallel and Distributed Systems*. He is the recipient of NSF CAREER Award, Jeffress Trust Award, CGO'13 and SC'20 Best Paper Awards, and the first place in Design Contest of ISLPED'20.



Doğa Gürsoy received the BS and MS degrees from Middle East Technical University, Ankara, Turkey, in 2004 and 2007, respectively, and the PhD degree in electrical engineering from the Graz University of Technology, Graz, Austria, in 2010. He is currently a computational scientist with the Advanced Photon Source at Argonne National Laboratory. He is also an adjunct associate professor with the Department of Electrical Engineering and Computer Science, Northwestern University, and a fellow of Northwestern

Argonne Institute of Science and Engineering. His research interests include developing computational methods for solving inverse problems in imaging and applying them to interesting problems in biological and materials sciences.



Rajkumar Kettimuthu (Senior Member, IEEE) is currently a computer scientist with Argonne National Laboratory, a senior scientist with the University of Chicago and a senior fellow with the Northwestern-Argonne Institute of Science and Engineering. His research interests include intelligent compute and data infrastructure for science, advanced wired and wireless communications for science, and quantum networks. He has coauthored more than 125 peer-reviewed articles, most of which appeared in premier journals and

conferences in computer science, and several of which won best paper awards. His work has featured in more than ten news articles. He has participated in more than 25 DOE and NSF workshops to set strategy and direction for networking and distributed computing research. He is a recipient of the prestigious R&D 100 award. He is a distinguished member of the ACM.



Ian T. Foster (Fellow, IEEE) received the BSc (hons I) degree in computer science from the University of Canterbury, Christchurch, New Zealand, and the PhD degree in computer science from Imperial College, London, U.K. He is currently the director of Argonne's Data Science and Learning Division, Argonne senior scientist and distinguished fellow and the Arthur Holly Compton distinguished service professor of computer science with the University of Chicago. His research interests include high-performance computing, distributed systems, and data-driven discovery. He has published hundreds of scientific papers and eight books on these and other topics. He is an elected fellow of the AAAS and ACM.



Wen-Mei W. Hwu (Fellow, IEEE) is currently a senior distinguished research scientist with NVIDIA. Prior to NVIDIA, he was a professor, Sanders-AMD endowed chair, acting department head and chief scientist of the Parallel Computing Institute, University of Illinois at Urbana-Champaign. He and his Illinois team developed the superblock compiler scheduling and optimization framework that has been adopted by virtually all modern vendor and open-source compilers today. For his research contributions, he received the ACM SigArch Maurice Wilkes Award, ACM Grace Murray Hopper Award, IEEE Computer Society Charles Babbage Award, ISCA Influential Paper Award, MICRO Test-of-Time Award, IEEE Computer Society B. R. Rau Award, CGO Test-of-Time Award and Distinguished Alumni Award in CS of the University of California, Berkeley. He has also won numerous best paper awards for major conferences. He is a fellow of the ACM.

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