JOURNAL OF COMPUTATIONAL BIOLOGY
Volume 25, Number 3, 2018

Pp. 270–281
DOI: 10.1089/cmb.2017.0151

ICON-MIC: Implementing a CPU/MIC Collaboration Parallel Framework for ICON on Tianhe-2 Supercomputer

ZIHAO WANG,[1,2,6] YU CHEN,[1,2,6] JINGRONG ZHANG,[1,2] LUN LI,[1,3]
XIAOHUA WAN,[1] ZHIYONG LIU,[1,*] FEI SUN,[2,4,5] and FA ZHANG[1,*]

ABSTRACT

Electron tomography (ET) is an important technique for studying the three-dimensional structures of the biological ultrastructure. Recently, ET has reached sub-nanometer resolution for investigating the native and conformational dynamics of macromolecular complexes by combining with the sub-tomogram averaging approach. Due to the limited sampling angles, ET reconstruction typically suffers from the "missing wedge" problem. Using a validation procedure, iterative compressed-sensing optimized nonuniform fast Fourier transform (NUFFT) reconstruction (ICON) demonstrates its power in restoring validated missing information for a low-signal-to-noise ratio biological ET dataset. However, the huge computational demand has become a bottleneck for the application of ICON. In this work, we implemented a parallel acceleration technology ICON-many integrated core (MIC) on Xeon Phi cards to address the huge computational demand of ICON. During this step, we parallelize the element-wise matrix operations and use the efficient summation of a matrix to reduce the cost of matrix computation. We also developed parallel versions of NUFFT on MIC to achieve a high acceleration of ICON by using more efficient fast Fourier transform (FFT) calculation. We then proposed a hybrid task allocation strategy (two-level load balancing) to improve the overall performance of ICON-MIC by making full use of the idle resources on Tianhe-2 supercomputer. Experimental results using two different datasets show that ICON-MIC has high accuracy in biological specimens under different noise levels and a significant acceleration, up to 13.3×, compared with the CPU version. Further, ICON-MIC has good scalability efficiency and overall performance on Tianhe-2 supercomputer.

Keywords: electron tomography, hybrid task allocation strategy, ICON, MIC acceleration, parallel NUFFT, Tianhe-2 supercomputer.

[1]High Performance Computer Research Center, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China.
[2]University of Chinese Academy of Sciences, Beijing, China.
[3]School of Mathematical Sciences, University of Chinese Academy of Sciences, Beijing, China.
[4]National Key Laboratory of Biomacromolecules, CAS Center for Excellence in Biomacromolecules, Institute of Biophysics, Chinese Academy of Sciences, Beijing, China.
[5]Center for Biological Imaging, Institute of Biophysics, Chinese Academy of Sciences, Beijing, China.
[6]These authors contributed equally to this work.
*Corresponding authors.

1. INTRODUCTION

Electron tomography (ET) is an important technique for studying the three-dimensional (3D) structures of the biological ultrastructure (Fridman et al., 2012; Lučić et al., 2013). Recently, to address the low signal-to-noise ratio (SNR) of biological samples, a new route called sub-tomogram averaging is proposed (Castaño-Dez et al., 2012). Combined with the sub-tomogram averaging approach, ET has reached sub-nanometer resolution for investigating the native and conformational dynamics of macromolecular complexes (Bharat et al., 2015). In ET, a series of two-dimensional (2D) projection micrographs (tilt series) are taken in different orientations by rotating the sample around a single fixed axis (Y-axis). Then, the micrographs are used to reconstruct the 3D density of the ultrastructure based on the projection slice theorem (Mersereau and Oppenheim, 1974). The tilt angle of ET is limited within a range of $-70°$ to $70°$ to ensure that a reasonable number of electrons can pass through and form reliable images. The absence of the high tilt angles will cause the ''missing wedge'' problem (Penczek et al., 1995). Traditional ET reconstruction methods, such as weighted back projection (WBP) (Radermacher, 2007), SIRT (Gilbert, 1972), and INFR (Chen and Förster, 2014), often suffer from the ''missing wedge'' problem and reconstruct 3D tomograms with ray artifacts. Such ''missing wedge'' artifacts will seriously hinder the reconstruction interpretation, especially for in situ specimens that are embedded in crowded environments (Kovacik et al., 2014).

In recent years, the topic of solving the ''missing wedge'' problem in ET has been widely discussed, and many algorithms have been proposed. The dual-axis and conical tomography (Lanzavecchia et al., 2005; Arslan et al., 2006) have been proposed to compensate for the ''missing wedge'' artifacts, which rotate the sample along two vertical or different tilt axes to collect multiple tomography data. These approaches cannot prevent the radiation damage problem and leave a conical area in which information cannot be measured. Other algorithms try to apply prior constraints to the reconstructed tomogram to restore the ''missing'' information, such as FIRT (Chen et al., 2016) and DART (Batenburg and Sijbers, 2011). The prior constraints include density smoothness and density localness. In addition, compressed sensing (CS) ET has been used to solve the reconstruction problem as an underdetermined problem based on a theoretical framework called ''compressed sensing'' (Donoho, 2006), and it demonstrated some success for data with a high SNR [e.g., material science data (Goris et al., 2012; Leary et al., 2013; Saghi et al., 2016)]. To address a case with a low SNR (e.g., biological cryo-ET data), Deng et al. (2016) proposed iterative compressed-sensing optimized nonuniform fast Fourier transform (NUFFT) reconstruction (ICON) by combining CS and NUFFT. Experimental results of ICON for different datasets showed that the algorithm can restore the missing information and measure the fidelity of the information restoration by using a validation procedure.

Although ICON has demonstrated its power in restoring validated missing information for a low SNR biological ET dataset, the huge computational demand becomes a bottleneck for its wider application. As high-performance computing platforms are becoming increasingly popular, to improve the efficiency of the program and shorten its running time, many ET reconstruction algorithms have been ported to a heterogeneous system containing acceleration units [e.g., graphics processing units (Palenstijn et al., 2011), many integrated core (MIC) units (Dahmen et al., 2016)]. Further, to realize the real-time ET reconstruction process, it is essential to make full use of super computing resources such as Tianhe-2 (Liao et al., 2014), which is one of the top 5 supercomputers in the world.

In this work, we implemented a parallel acceleration technology ICON-MIC on Xeon Phi cards to address the huge computational demand of ICON. During this step, we parallelize the element-wise matrix operations and use the efficient summation of a matrix to reduce the matrix computation cost. By using a more efficient MKL fast Fourier transform (FFT) (Wang et al., 2014) calculation, we developed parallel versions of NUFFT to achieve high ICON acceleration. The NUFFT library has not previously existed on MIC. We then proposed a hybrid task allocation strategy, two-level load balancing (TLLB), to improve the overall performance of ICON-MIC by making full use of the idle resources on the Tianhe-2 supercomputer. We used two different ET datasets, including a resin embedded dataset and a cryo-ET dataset, to test ICON-MIC. The experimental results show that ICON-MIC has high reconstruction accuracy and can well restore missing information under different noise levels. Further, ICON-MIC exhibits significant acceleration factors compared with the CPU version ICON (ICON-CPU) and has good weak and strong scalability. In addition, ICON-MIC has good overall performance and can make good use of the supercomputer resources by using the TLLB. ICON-MIC has also been developed into software packages, which can be downloaded from the homepage of the authors as follows: http://ear.ict.ac.cn.

2. RELATED WORK

2.1. Iterative compressed-sensing optimized NUFFT reconstruction

ICON is an iterative reconstruction algorithm based on the theoretical framework of ''compressed sensing,'' and its complete workflow can be divided into four steps: ''Pre-processing,'' ''Gray value adjustment,'' ''Reconstruction and pseudo-missing-validation,'' and ''Verification filtering'' (Deng et al., 2016). A series of tests showed that ''Reconstruction and pseudo-missing-validation'' accounts for at least 95% of the execution time of ICON. Thus, the major task for accelerating ICON is paralleling this step effectively on MIC. The parallelization of ''reconstruction'' and ''pseudo-missing-validation'' is similar, and only ''reconstruction'' is discussed in this article. The major steps of ICON ''reconstruction'' can be briefly described as follows.

Step 1. Fidelity preservation step using the steepest descent method (Goldstein, 1965).

$$r = A^h W A x^k - A^h W f \tag{1}$$

$$\alpha = \frac{r^T r}{r^T A^h W A r} \tag{2}$$

$$y^{k+1} = x^k + \alpha * r \tag{3}$$

where x^k is the 2D reconstructed image of the kth iteration. A is the projection operation, and is defined as a nonuniform Fourier sampling matrix, that performs a Fourier transform on the noninteger grid points. A^h stands for the conjugate transpose of A. W follows the description of INFR (Chen and Förster, 2014) and contains the weights that account for the nonuniform sampling in the Fourier space (similar to the ramp filtering in WBP). f is the Fourier transform of the acquired projections. r is the residual. α is the coefficient used to control the updating step. y^{k+1} is the intermediate updating result of the $(k+1)th$ iteration.

Step 2. Prior sparsity restriction step.

$$x^{k+1} = H(y^{k+1}) = \begin{cases} 0, & if \quad y^{k+1} < 0 \\ y^{k+1}, & if \quad y^{k+1} \geq 0 \end{cases} \tag{4}$$

where y^{k+1} is the intermediate updating result of the $(k+1)th$ iteration. $H(\cdot)$ is a logic function. $x^{k+1}$is the 2D reconstructed slice of the $(k+1)th$ iteration.

The operations of these two steps are classified into the following three types: a. element-wise matrix operations; b. matrix summation; and c. the NUFFT and the adjoint NUFFT. For each type of operation, a parallelization strategy is proposed in section 3.

2.2. Nonuniform FFT

First, we give a brief description of NUFFT. Given the Fourier coefficients $\hat{f}_k \in \mathbb{C}, \boldsymbol{k} \in I_N$ and $I_N = \{\boldsymbol{k} = (k_t)_{t=0,\ldots,d-1} \in \mathbb{Z}^d : -\frac{N_t}{2} \leq k_t < \frac{N_t}{2},\ t = 0, \ldots, d-1\}$ as input, NUFFT tries to evaluate the following trigonometric polynomial efficiently at the reciprocal points $x_t \in [-\frac{1}{2}, \frac{1}{2})$, $t = 0, \ldots, M-1$:

$$f_j = f(x_j) = \sum_{k \in I_N} \hat{f}_k e^{-2\pi i k x_j},\ j = 0, \ldots, M-1 \tag{5}$$

Correspondingly, the adjoint NUFFT tries to evaluate Equation (6) at the frequencies $\boldsymbol{k}$:

$$\hat{h}_{\boldsymbol{k}} = \sum_{j-0}^{M-1} f_i e^{2\pi i \boldsymbol{k} x_j} \tag{6}$$

NFFT3.0 (Keiner et al., 2009) is a successful and widely used open source C library for NUFFT and adjoint NUFFT. However, to our knowledge, no corresponding library on MIC is available yet. Thus, we paralleled the NUFFT and the adjoint NUFFT based on the algorithms described in NFFT3.0 and the algorithm of 2D NUFFT is displayed in Algorithm 1 for deep analysis.

Algorithm 1: NUFFT

Input: $M, N=\{N_1, N_2\}, \sigma=\{\sigma_1, \sigma_2\}, m, x_j \in \left[-\frac{1}{2}, \frac{1}{2}\right)^2,\ j=0, \ldots, M-1, \hat{f}_k \in \mathbb{C}, \boldsymbol{k} \in I_N$
$\boldsymbol{n}=\boldsymbol{\sigma N}=\{n_1, n_2\}=\{\sigma_1 N_1, \sigma_2 N_2\}$

1: For $\boldsymbol{k} \in I_N$ compute
$\hat{g}_k=\frac{\hat{f}_k}{|I_n| c_k(\tilde{\varphi})}$
$c_k(\tilde{\varphi})=\hat{\varphi}(k_1)\hat{\varphi}(k_2)$
2: For $\boldsymbol{I} \in I_n$ compute by 2-variate FFT
$g_l=\sum_{\boldsymbol{k} \in I_N} \hat{g}_{\boldsymbol{k}} e^{-2\pi i \boldsymbol{k}(\boldsymbol{n}^{-1} \odot \boldsymbol{l})}$
3: For j = 0, . . . ,M-1 compute
$f_l=\sum_{l \in I_{n,m}(x_j)} g_l \tilde{\psi}\left(x_j-\boldsymbol{n}^{-1} \odot \boldsymbol{l}\right)$
$I_{n,m}(x_j)=\left\{\boldsymbol{l} \in I_n : \boldsymbol{n} \odot x_j - m\mathbf{1} \le \boldsymbol{l} \le \boldsymbol{n} \odot x_j + m\mathbf{1}\right\}$
$\tilde{\psi}(x)=\varphi(x_1)\varphi(x_2)$

$\varphi(x)$ and $\hat{\varphi}(k)$ are the window functions. In this work, the (dilated) Gaussian window functions [Eqs. (7) and (8)] are used.

$$\varphi\ (x)=(\pi b)^{-\frac{1}{2}} e^{-\frac{(nx)^2}{b}} \left(b=\frac{2\sigma}{2\sigma-1}\frac{m}{\pi}\right) \tag{7}$$

$$\hat{\varphi}\ (k)=\frac{1}{n} e^{-b\left(\frac{\pi k}{n}\right)^2} \tag{8}$$

where x is a component of the reciprocal points $\boldsymbol{x}$. k is a component of the frequencies $\boldsymbol{k}$. σ is a component of the oversampling factors $\boldsymbol{\sigma}$ with $\sigma > 1$. n is one component of $\boldsymbol{n}=\boldsymbol{\sigma N}$. $m \in \mathbb{N}$ and $m \ll n$. In this work, $\sigma=2$ and $m=6$.

NUFFT operations can be classified into the following three types: (1) element-wise matrix operations, (2) 2D FFT, and (3) window function calculations.

3. ACCELERATION OF ICON USING MIC COPROCESSORS

3.1. Parallel element-wise matrix operations

In ET reconstruction, the matrix size is usually larger than the number of processing units on MIC, so an appropriate number of processing units (threads) must be selected to balance the control and computing resources. According to experiments on the Xeon Phi 31SP card, if the thread number is close or equal to 228 (57 cores with 4 hardware threads in each core), having all of the threads participate in the computation will cause performance degradation. Thus, 200 threads are used when parallelizing the element-wise operations, and a matrix is divided into 200 parts, in which each part is assigned to one thread for calculation. Experiments show that the matrices constituting large arrays in ICON have high allocation costs when using 4 KB pages. To reduce the allocation cost, 2 MB pages are used in offload mode, which also reduces the TLB misses and page faults. The 512-bit vector processing unit (VPU) (Duran and Klemm, 2012) is used on each core, which means that 16 single-precision or 8 double-precision operations can be executed in one instance to achieve a high computational throughput for element-wise matrix operations.

3.2. NUFFT and adjoint NUFFT parallelization

As mentioned in section 2.2, there are three classes of operations in NUFFT and adjoint NUFFT. The parallelization strategy of the element-wise matrix operations in NUFFT and adjoint NUFFT is the same as the strategy described in section 3.1.

The experimental results show that the calculation of FFT operations accounts for $\sim 80\%$ of the total operating time for ICON, so Intel MKL FFT is used to achieve significant acceleration. Intel MKL FFT is a specially designed FFT library for MIC, and the following three methods are used for Intel MKL FFT: The first method uses the FFTW3 interface, which has good portability; the second method uses the FFT interface by Intel for memory alignment; and the third method also uses the FFT interface by Intel but changes the multidimensional data layout in the coprocessor memory (More, 2013). The performance of these three methods is tested by using a 2D forward and backward complex to complex FFT by repeating

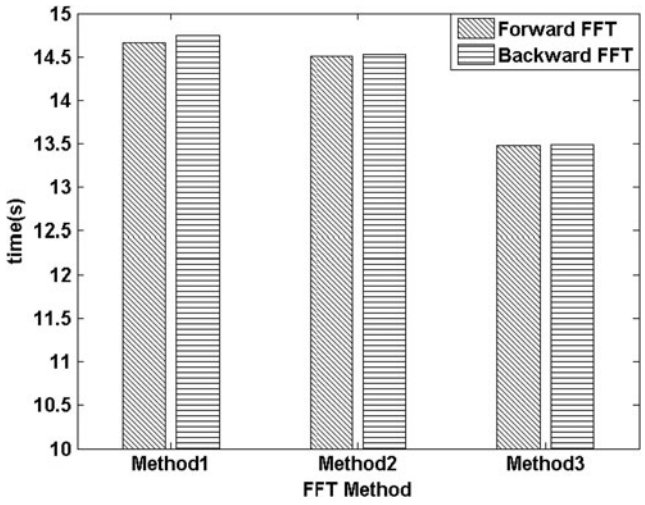

FIG. 1. Performance comparison of main FFT methods. FFT, fast Fourier transform.

the test 100 times on a Xeon Phi 31SP coprocessor of Tianhe-2. In the following test, the compiler-assisted offload is used to manage the functions and the data that are offloaded. The test data scale is 8192×8192, and the data are stored in complex form, in which each element takes 8 bytes. The performance comparisons are shown in Figure 1.

The first method uses the FFTW3 interface directly, and the performance is normal. The second method uses the FFT by Intel, which is more suitable for the SIMD architecture of Xeon Phi cards and uses memory alignment to improve performance. The cache line size on Xeon Phi cards is 64 bytes, so the algorithm-performing column FFT will have poor cache associativity through cache re-usage when the matrix element distance from one row to the next is a multiple of 64. Unfortunately, because of the camera characteristics in ET, the scale of the input image using ICON is always a multiple of 64. Thus, the data layout in the coprocessor memory must be changed by using the third FFT method to obtain better performance when the NUFFT is running in parallel.

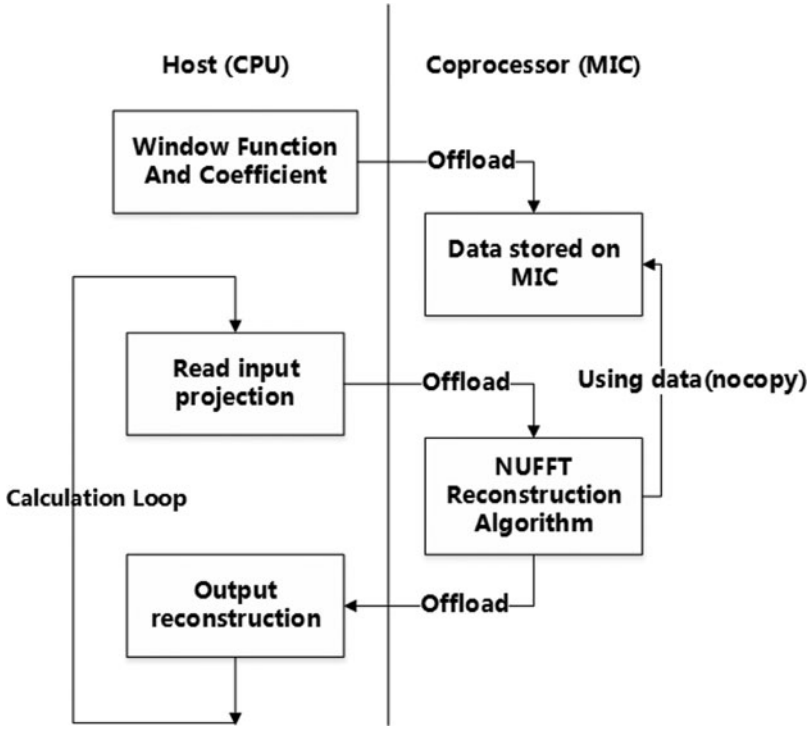

FIG. 2. NFFT precomputing using data persistence. MIC, many integrated core; NUFFT, nonuniform fast Fourier transform.

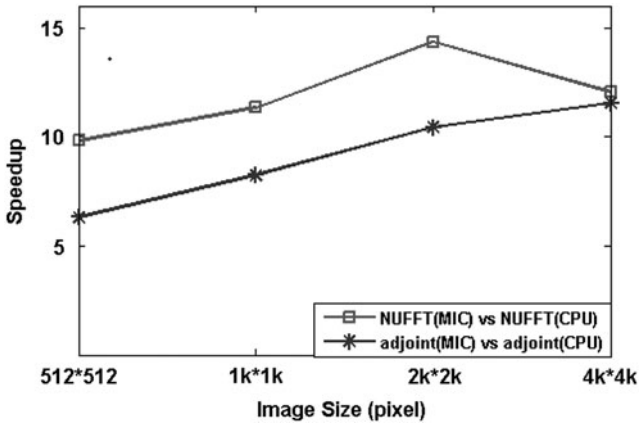

FIG. 3. The speedups of parallel NUFFTs compared with NFFT3.0.

When calculating the window functions, since ICON is an iterative algorithm, NUFFT and adjoint NUFFT will be repeated many times. To decrease the calculation and memory transfer time, the data persistence technology shown in Figure 2 is used, in which the window functions are precomputed and stored in the device memory.

A resin-embedded ET dataset with an MDCK cell section (see section 4.1 for detail) is used to test the performance of parallel NUFFTs on MIC compared with NFFT3.0. NFFT3.0 ran on one core (thread) of an Intel® Xeon™ CPU E5-2620 v2 @ 2.1 GHz (6 cores per CPU), and parallel NUFFTs ran on a Xeon Phi 31SP coprocessor of Tianhe-2. The test datasets include image sizes of 512×512, $1k\times1k$, $2k\times2k$, and $4k\times4k$. The experimental results show that parallel NUFFTs are 10 times faster than the library NFFT3.0, as illustrated in Figure 3.

3.3. Efficient summation of a matrix

Commonly, CPU programs will execute the summation by using a single thread. However, for MIC, the computational capability of one thread is too weak to sum an entire matrix in a reasonable time.

An OpenMP reduction clause is usually used to avoid executing the summation on one core of MIC. However, it cannot take advantage of the 512-bit vector-processing unit on MIC. An array notation is used as part of Intel Cilk Plus (Robison, 2012) to help the compiler with vectorization and achieve the efficient utilization of all available processing resources. The three summation strategies mentioned earlier were compared, and the result is shown in Figure 4. The strategy using Intel Cilk Plus reduction with VPU is the most efficient.

3.4. Extend ICON-MIC on multiple Xeon Phi cards on Tianhe-2

To further satisfy the large amount of computational requirements, the ICON-MIC is extended to multiple Xeon Phi cards on Tianhe-2. To make ICON-MIC compatible with the architecture of Tianhe-2, we proposed a hybrid task allocation model named TLLB that takes advantage of the message passing interface. Considering the exclusivity of one node on Tianhe-2 and the fact that idle CPU cores also have good computing capability, the idle CPU cores involved in the computation are utilized to make full use of the idle resources on one node. The TTLB, which is a heterogeneous computing framework, combines the static allocation (for the level on the Xeon Phi cards and idle CPU cores) with dynamic allocation (for the level on the CPU nodes), as described in Figure 5.

In ET, the reconstruction of a 3D volume can be divided into a series of similar tasks. Each node on Tianhe-2 has three Xeon Phi cards and idle CPU cores, so we perform precomputation to estimate the computing capability and assign different weights to the Xeon Phi cards and CPU cores to achieve load balancing. Then, we separate all tasks into a series of task subsets. Each subset contains a series of similar tasks, and the exact number of tasks depends on the weight and the entire number of tasks. During reconstruction, each node will dynamically request one task subset after the previous task subset is finished. Within one node, each task will be statically assigned to Xeon Phi cards and idle CPU cores according to their weight.

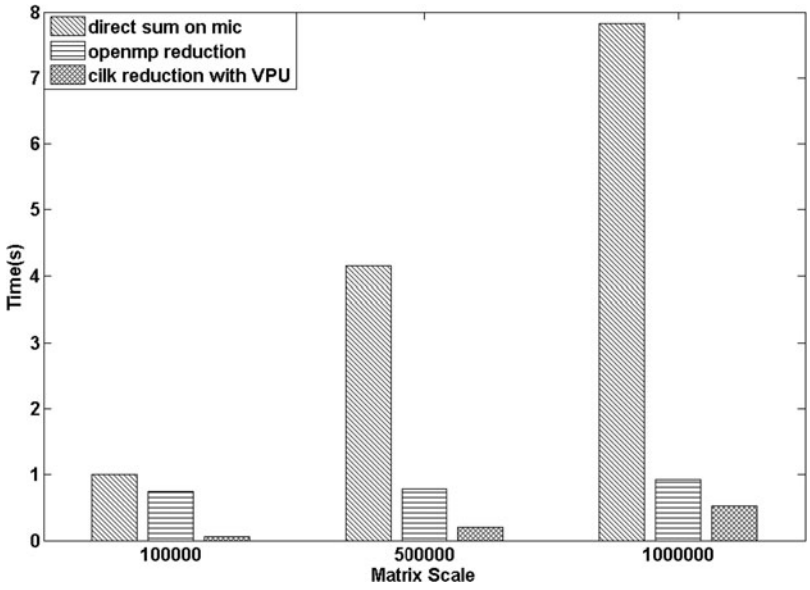

FIG. 4. Summation of a matrix. VPU, vector processing unit.

4. RESULTS AND DISCUSSION

4.1. Resin-embedded ET dataset

The resin-embedded ET dataset is an MDCK cell section. The tilt angles of the dataset originally ranged from −68° to +68°, with a 1° increment. To verify the ability of ICON to restore missing information, every two projections are extracted from the original dataset to generate a new tilt series with a 2° increment for the experiments. The tilt series are aligned using atomalign (Han et al., 2014). The original image size is 4k × 4k, with a pixel size of 0.72 nm.

4.2. Cryo-ET dataset

The Cryo-ET dataset contains the mitochondria of mice hepatic cells. The data set is recorded by the China National Key Laboratory of Biomacromolecules and was collected by the FEI company's production-Tecnai 20. The tilt angles of the dataset originally ranged from −52° to +59°, with a 1° increment. The tilt series are also aligned using atomalign (Han et al., 2014). The original image size is 2k × 2k, with a pixel size of 0.4 nm.

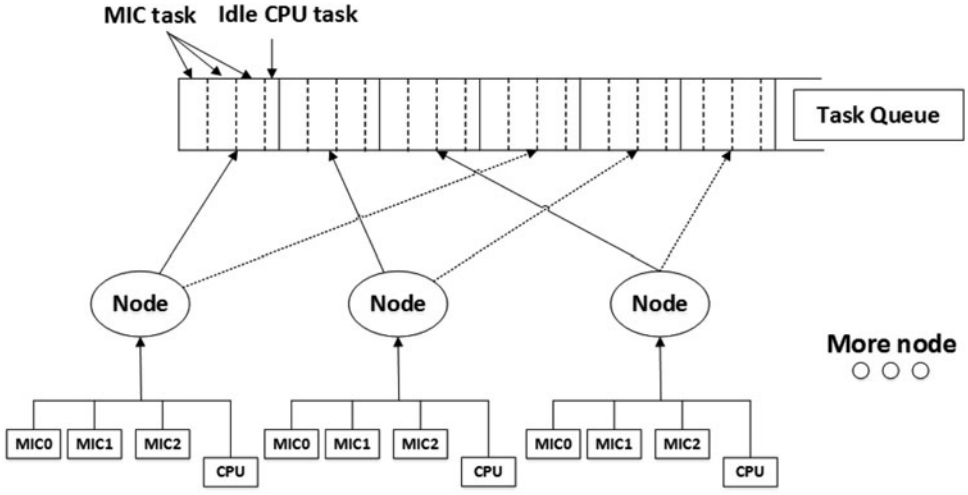

FIG. 5. TLLB for ICON on Tihanhe-2. TLLB, two-level load balancing.

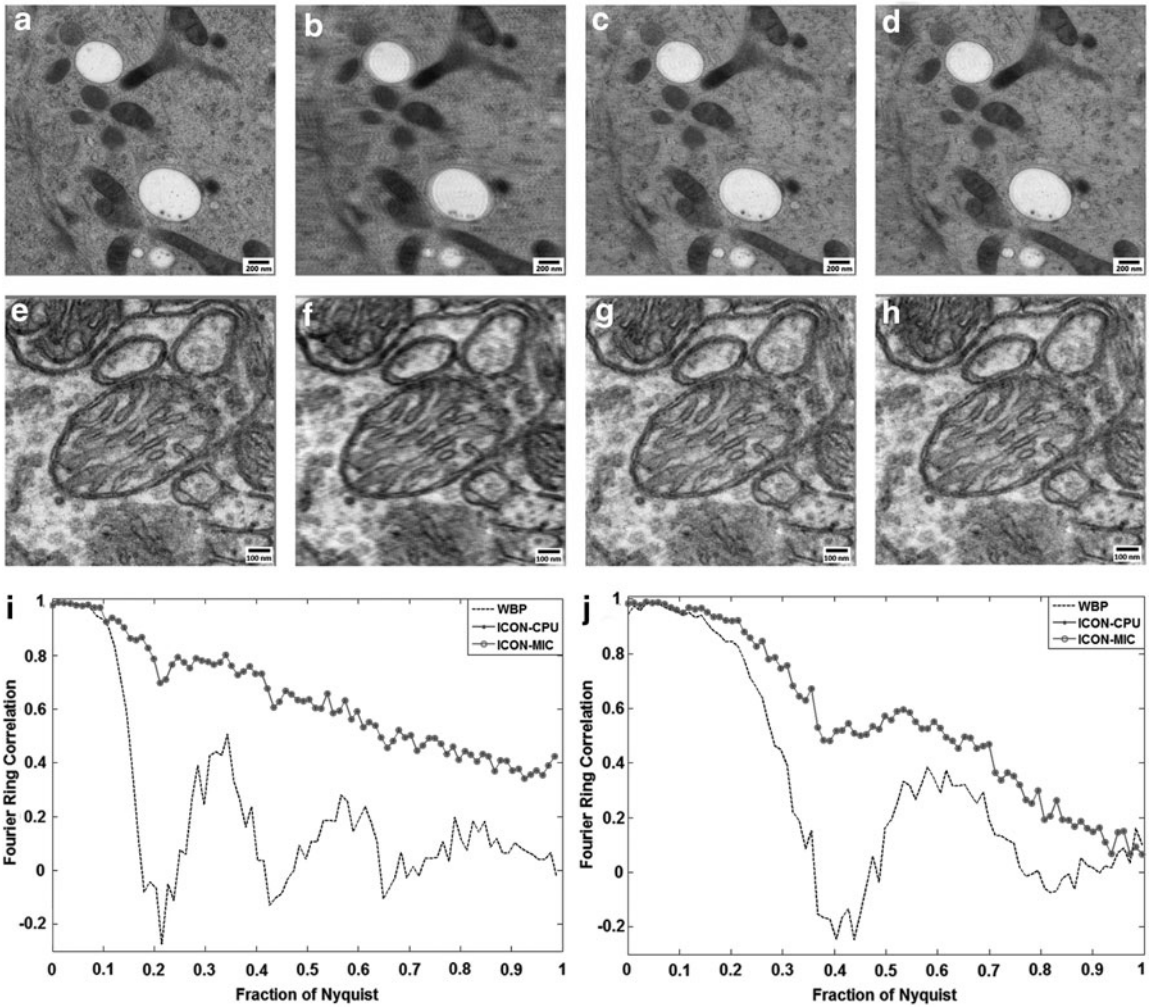

FIG. 6. Evaluate ICON-MIC by the pseudo-missing-validation procedure. **(a, e)** The omit projection (''Ground truth''); **(b–d, f–h)** The re-projections of the omit tomograms reconstructed by WBP, ICON-CPU, and ICON-MIC, respectively; **(i, j)** The pseudo-missing-validation FRCs of WBP, ICON-CPU, and ICON-MIC. FRC, Fourier ring correlation; WBP, weighted back projection.

4.3. Reconstruction precision

The reconstruction precision of ICON-MIC on the resin-embedded ET dataset is investigated by using the pseudo-missing-validation procedure (Deng et al., 2016). Here, the minimum tilt (the −0.29° tilt) projection was excluded as the omit projection (''ground truth''), as shown in Figure 6a. The reconstructed tomograms

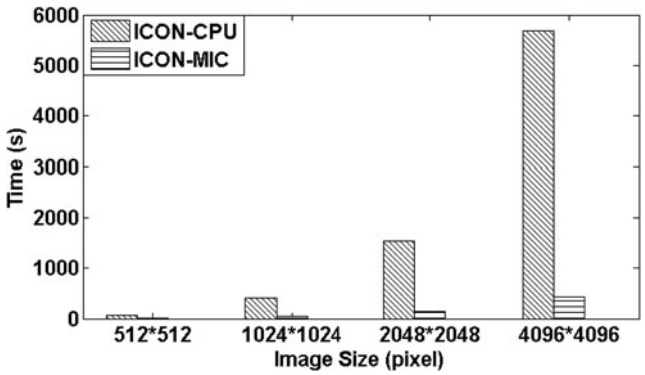

FIG. 7. The comparison of time consumption of ICON-CPU and ICON-MIC.

Table 1. The Speedups of ICON-MIC Compared with ICON-CPU

Image size	512×512	1024×1024	2048×2048	4096×4096
ICON-MIC	5.2×	9.4×	10.9×	13.3×

ICON, iterative compressed-sensing optimized NUFFT reconstruction; MIC, many integrated core.

are then projected again at –0.29°. The re-projections of ICONs (Fig. 6c, d) are identical to each other, and the NCCs between each are 1. The re-projections of ICONs present clearer, more detailed structures and are more similar to the ''ground truth'' compared with WBP (Fig. 6b). Such visual assessments are further verified quantitatively by comparing the Fourier ring correlation (FRC) curves between the re-projections and the ''ground truth.'' The FRCs of ICONs coincide with each other, and they are better than that of WBP (Fig. 6i). The coincident FRCs of ICONs further demonstrate the accuracy of ICON-MIC from the perspective of restoring missing information.

We further investigated ICON-MIC by using the noisier Cryo-ET dataset. The same pseudo-missing-validation procedure was conducted. We re-project the reconstructed tomograms at 0°. Even with a higher noise level, the improvement of ICON-MIC is also obvious (Fig. 6h, j), which demonstrates the applicability of ICON-MIC on different types of datasets under different noise levels.

4.4. Speedup

The acceleration of ICON-MIC is evaluated by comparing the running time required to reconstruct one slice for 200 iterations. A resin-embedded ET dataset of an MDCK cell section is constructed with image sizes of 512×512, 1k×1k, 2k×2k, and 4k×4k. ICON-CPU runs on one core (thread) of an Intel Xeon CPU E5-2620 v2 @ 2.1 GHz, ICON-MIC with one Xeon Phi 31SP coprocessor of Tianhe-2. The acceleration of ICON-MIC improves when the slice size increases (Fig. 7 and Table 1). The maximum speed of ICON-MIC is 13.3× for reconstructing a 4k×4k slice. With efficient acceleration, the reconstruction time of one 4k×4k slice is reduced from hours to minutes.

4.5. Scalability and overall performance on Tianhe-2 supercomputer

We first tested the weak and strong scalability of ICON-MIC on Tianhe-2 supercomputer. Then, we evaluated how the execution time varied with the number of processors.

In the weak scalability test, we fixed the number of tasks assigned to one processor. We tested the weak scalability with image sizes of 1024×1024 and 2048×2048. We gradually increased the number of Xeon Phi cards from 3 to 48. As the nodes gradually increased, the total number of images being processed also increased. The execution time shown in Figure 8 increased from 875 to 900 seconds, which indicated that the developed parallelization strategy is good for weak scalability.

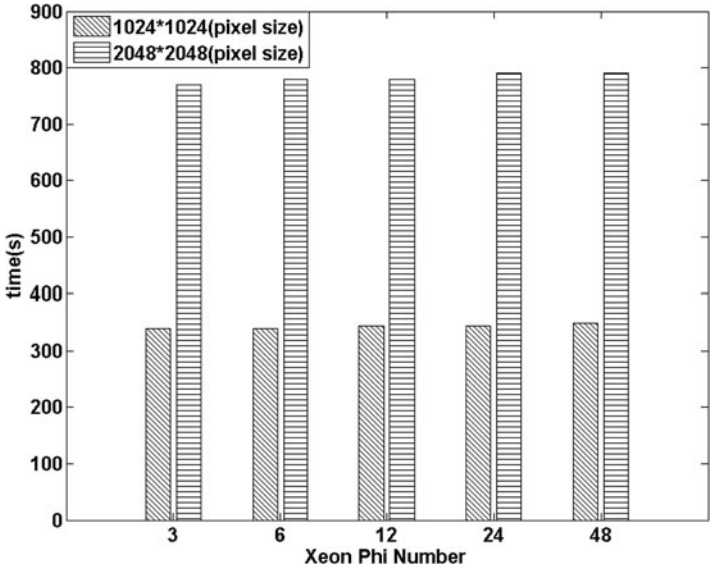

FIG. 8. Weak scalability results on Tianhe-2.

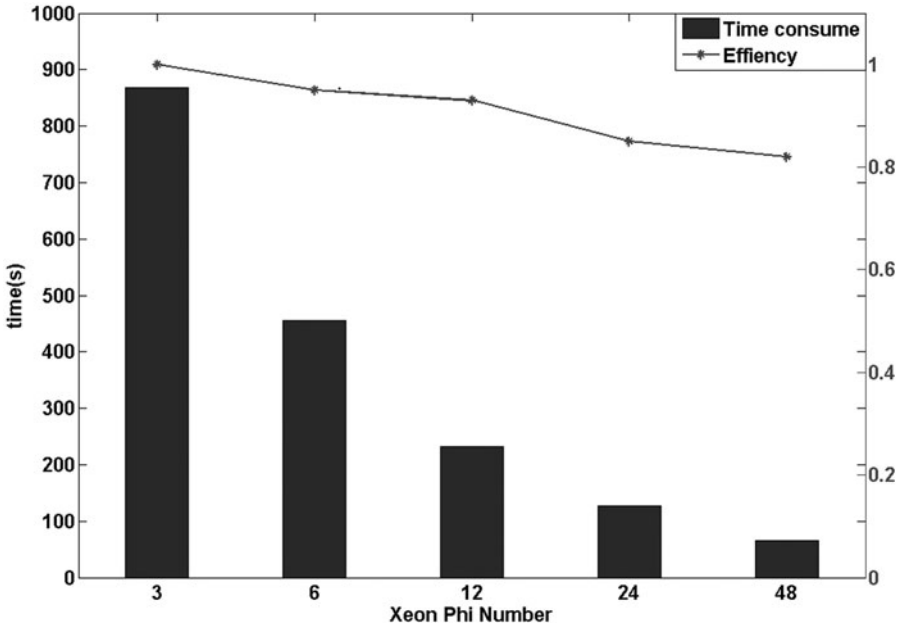

FIG. 9. Strong scalability results on Tianhe-2.

In the strong scalability test, we fixed the total number of tasks in all nodes. We tested the strong scalability with image sizes of 1024×1024 and the total number of reconstruct slice is 48, which was equal to the highest number of Xeon Phi cards. We only increased the number of Xeon Phi cards from 3 to 48. From Figure 9, we can observe that the parallel efficiency decreased to 92 percent when using 12 Xeon Phi cards and decreased further to 83 percent when using 48 Xeon Phi cards. The observed degradation of the strong scalability efficiency is acceptable.

We then tested the overall performance of the ICON-MIC on Tianhe-2 supercomputer. We use three groups of data, including the image sizes of 512×512, 1024×1024, and 2048×2048. The number of reconstruct slice is equal to 512. We use 10 nodes, which include 30 Xeon Phi cards, and we take advantage of TLLB task allocation by using 12 threads on one idle CPU of each node. From Figure 10, we can

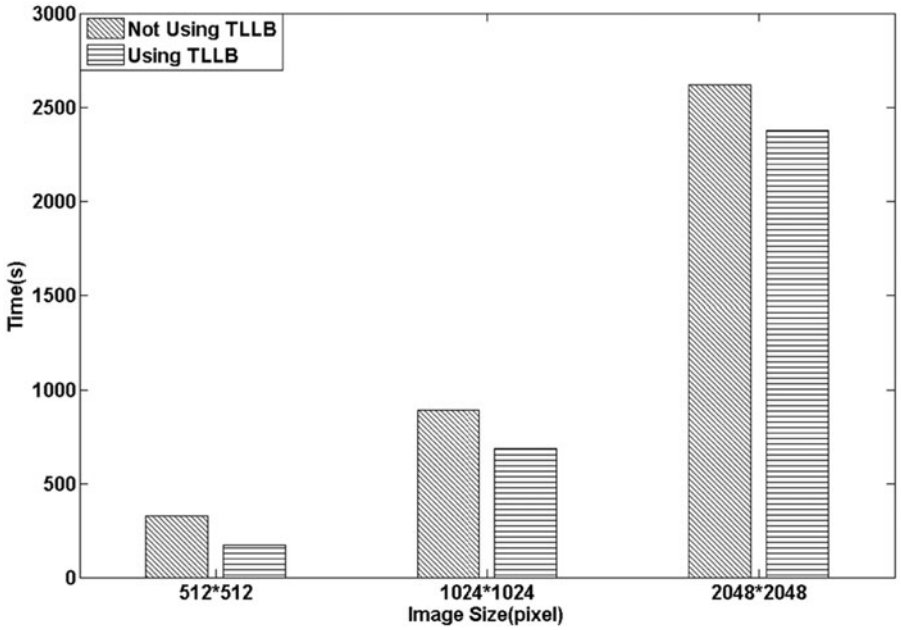

FIG. 10. Overall performance on Tianhe-2.

observe that after using the TLLB task allocation, we can further reduce the time consumption of ICON-MIC by making full use of the whole available computing resource on Tianhe-2.

5. CONCLUSION

In this work, we analyze the iterative framework of ICON and classify the operations of ICON's major steps into three types. Accordingly, we design parallelization strategies for ICON to generate a parallel program, ICON-MIC. We parallelize the element-wise matrix operations and use the efficient summation of a matrix to reduce the matrix computation cost. We also develop a parallel version of NUFFT and an adjoint NUFFT on MIC by using a more efficient calculation of FFT. To satisfy the vast computation requirements and consider the exclusivity of one node on the Tianhe-2 supercomputer, we propose a hybrid task allocation strategy to improve the scalability and overall performance.

We test ICON-MIC on two different datasets, including a resin-embedded ET dataset of an MDCK cell section and a cryo-ET dataset, which contains data on mitochondria of mice hepatic cells. By the pseudo-missing-validation procedure, the FRC curves show that ICON-MIC has reasonable numerical accuracy and can well restore miss information on both high and low SNR specimens. The experimental results also show that ICON-MIC has good acceleration, reaching $13.3\times$ for ICON-MIC in reconstructing one $4k\times4k$ slice, and ICON-MIC has good weak and strong scalability efficiency. The experimental results indicate that ICON-MIC can use the heterogeneous computational resources of the Tianhe-2 supercomputer effectively and has good overall performance.

ACKNOWLEDGMENTS

This research is supported by the National Key Research and Development Program of China (2017YFA0504702), the NSFC projects Grant No. U1611263, U1611261, 61232001, 61472397, 61502455, 61672493 and Special Program for Applied Research on Super Computation of the NSFC-Guangdong Joint Fund (the second phase). The authors would like to thank Prof. Wanzhong He (NIBS, Beijing) for providing the resin-embedded ET dataset. All intensive computations were performed on Tianhe-2 supercomputer at the National Supercomputer Center in Guangzhou (NSCC-GZ), China.

AUTHOR DISCLOSURE STATEMENT

No competing financial interests exist.

REFERENCES

Arslan, I., Tong, J.R., and Midgley, P.A. 2006. Reducing the missing wedge: High-resolution dual axis tomography of inorganic materials. *Ultramicroscopy* 106, 994–1000.

Batenburg, K.J., and Sijbers, J. 2011. Dart: A practical reconstruction algorithm for dis- crete tomography. *IEEE Trans. Image Process*. 20, 2542–2553.

Bharat, T.A., Russo, C.J., Löwe, J., et al. 2015. Advances in single-particle electron cryomicroscopy structure determination applied to sub- tomogram averaging. *Structure* 23, 1743–1753.

Castaño-Dez, D., Kudryashev, M., Arheit, M., et al. 2012. Dynamo: A flexible, user-friendly development tool for subtomogram averaging of cryo-EM data in high-performance computing environments. *J. Struct. Biol.* 178, 139–151.

Chen, Y., and Förster, F. 2014. Iterative reconstruction of cryo-electron tomograms using nonuniform fast Fourier transforms. *J. Struct. Biol*. 185, 309–316.

Chen, Y., Zhang, Y., Zhang, K., Deng, Y., et al. 2016. Firt: Filtered iterative reconstruction technique with information restoration. *J. Struct. Biol*. 195, 49–61.

Dahmen, T., Marsalek, L., Marniok, N., et al. 2016. The ettention software package. *Ultramicroscopy* 161, 110–118.

Deng, Y., Chen, Y., Zhang, Y., et al. 2016. Icon: 3D reconstruction with missing-information restoration in biological electron tomography. *J. Struct. Biol*. 195, 100–112.

Donoho, D.L. 2006. Compressed sensing. *IEEE Trans. Inf. Theor*. 52, 1289–1306.

Duran, A., and Klemm, M. 2012. The Intel® many integrated core architecture, 365–366. In *2012 International Conference on High Performance Computing and Simulation (HPCS)*. IEEE, Madrid, Spain.

Fridman, K., Mader, A., Zwerger, M., et al. 2012. Advances in tomography: Probing the molecular architecture of cells. *Nat. Rev. Mol. Cell Biol*. 13, 736–742.

Gilbert, P. 1972. Iterative methods for the three-dimensional reconstruction of an object from projections. *J. Theor. Biol*. 36, 105–117.

Goldstein, A.A. 1965. On steepest descent. *J. Soc. Indus. Appl. Math Series A Control* 3, 147–151.

Goris, B., Van den Broek, W., Batenburg, et al. 2012. Electron tomography based on a total variation minimization reconstruction technique. *Ultramicroscopy* 113, 120–130.

Han, R., Zhang, F., Wan, X., et al. 2014. A marker-free automatic alignment method based on scale-invariant features. *J. Struct. Biol*. 186, 167–180.

Keiner, J., Kunis, S., and Potts, D. 2009. Using NFFT 3A software library for various nonequispaced fast Fourier transforms. *ACM Trans. Math. Softw*. 36, 19.

Kovacik, L., Kereche, S., Hoog, J.L., et al. 2014. A simple Fourier filter for suppression of the missing wedge ray artefacts in single-axis electron tomographic reconstructions. *J. Struct Biol*. 186, 141–152.

Lanzavecchia, S., Cantele, F., Bellon, P., et al. 2005. Conical tomography of freeze-fracture replicas: A method for the study of integral membrane proteins inserted in phospholipid bilayers. *J. Struct. Biol*. 149, 87–98.

Leary, R., Saghi, Z., Midgley, P.A., et al. 2013. Compressed sensing electron tomography. *Ultramicroscopy* 131, 70–91.

Liao, X., Xiao, L., Yang, C., et al. 2014. Milkyway-2 supercomputer: System and application. *Front. Comput. Sci*. 8, 345–356.

Lučić, V., Rigort, A., and Baumeister, W. 2013. Cryo-electron tomography: The challenge of doing structural biology in situ. *J. Cell. Biol.* 202, 407–419.

Mersereau, R.M., and Oppenheim, A.V. 1974. Digital reconstruction of multidimensional signals from their projections. *Proc. IEEE* 62, 1319–1338.

More, A. 2013. Intel Xeon Phi coprocessor high performance programming. *J. Comput. Sci. Technol*. 13.

Palenstijn, W., Batenburg, K., and Sijbers, J. 2011. Performance improvements for iterative electron tomography reconstruction using graphics processing units (gpus). *J. Struct. Biol*. 176, 250–253.

Penczek, P., Marko, M., Buttle, K., et al. 1995. Double-tilt electron tomography. *Ultramicroscopy* 60, 393–410.

Radermacher, M. 2007. Weighted back-projection methods, 245–273. In: Frank, J. (ed). *Electron Tomography*. Springer, New York, New York.

Robison, A.D. 2012. Cilk plus: Language support for thread and vector parallelism. *Talk at HP-CAST* 18, 25.

Saghi, Z., Divitini, G., Winter, B., et al. 2016. Compressed sensing electron tomography of needle-shaped biological specimens—Potential for improved reconstruction fidelity with reduced dose. *Ultramicroscopy* 160, 230–238.

Wang, E., Zhang, Q., Shen, B., et al. 2014. High-performance computing on the Intel Xeon Phi. *Springer* 5, 2.

Address correspondence to:
Prof. Zhiyong Liu
High Performance Computer Research Center
Institute of Computing Technology
Chinese Academy of Sciences
Beijing
China 100101

E-mail: zyliu@ict.ac.cn

Prof. Fa Zhang
High Performance Computer Research Center
Institute of Computing Technology
Chinese Academy of Sciences
Beijing
China 100101

E-mail: zhangfa@ict.ac.cn