

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

SHARP: a distributed GPU-based ptychographic solver

### Permalink

<https://escholarship.org/uc/item/9qn4m1zh>

### Journal

Journal of Applied Crystallography, 49(4)

### ISSN

0021-8898

### Authors

Marchesini, Stefano  
Krishnan, Hari  
Daurer, Benedikt J  
[et al.](#)

### Publication Date

2016-08-01

### DOI

10.1107/s1600576716008074

Peer reviewed

# SHARP: a distributed, GPU-based ptychographic solver

Stefano Marchesini,<sup>1</sup> Hari Krishnan,<sup>1</sup> Benedikt J. Daurer,<sup>2</sup> David A. Shapiro,<sup>1</sup> Talita Perciano,<sup>1</sup> James A. Sethian,<sup>1</sup> and Filipe R.N.C. Maia<sup>2</sup>

<sup>1</sup>*Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

<sup>2</sup>*Uppsala University, Uppsala, Sweden*

(Dated: June 21, 2016)

Ever brighter light sources, fast parallel detectors, and advances in phase retrieval methods, have made ptychography a practical and popular imaging technique. Compared to previous techniques, ptychography provides superior robustness and resolution at the expense of more advanced and time consuming data analysis. By taking advantage of massively parallel architectures, high-throughput processing can expedite this analysis and provide microscopists with immediate feedback. These advances allow real-time imaging at wavelength limited resolution, coupled with a large field of view. Here, we introduce a set of algorithmic and computational methodologies used at the Advanced Light Source, and DOE light sources packaged as a CUDA based software environment named SHARP (<http://camera.lbl.gov/sharp>), aimed at providing state-of-the-art high-throughput ptychography reconstructions for the coming era of diffraction limited light sources.

## I. INTRODUCTION

Reconstructing the 3D map of the scattering potential of a sample from measurements of its far-field scattering patterns is an important problem. It arises in a variety of fields, including optics [1, 2], astronomy [3], X-ray crystallography [4], tomography [5], holography [6, 7] and electron microscopy [8]. As such it has been a subject of study for applied mathematicians for over a century. The fundamental problem consists of finding the correct phases that go along with the measured intensities, such that together they can be Fourier transformed into the real-space image of the sample. To help recover the correct phases from intensity measurements a range of experimental techniques have been proposed along the years, such as interferometry/holography [6], random phase masks [7, 9, 10], gratings [11]. A variety of numerical techniques have also been recently developed, for example by approximating the problem as a matrix completion problem [12], or by other convex relaxations [13] tractable by semidefinite programming.

Since its first demonstration [14], progress has been made in solving the phase problem for a single diffraction pattern recorded from a non-periodic object, including the dynamic update of the support [15] and a variety of projection algorithms [16–18]. Such methods, referred to as coherent diffractive imaging (CDI), attempt to recover the complete complex-valued scattering potential or electron density, and the complex exit wavefront scattered from the object, providing phase contrast as well as a way to overcome depth-of-focus limitations of regular optical systems.

Ptychography, a relatively recent technique, provides the unprecedented capability of imaging macroscopic specimens in 3D and attain wavelength limited resolution along with chemical specificity [19]. Ptychography was proposed in 1969 [20, 21], and later experimentally demonstrated [22, 23], with the aim of improving the resolution in x-ray and electron microscopy. Since then

it has been used in a large array of applications, and shown to be a remarkably robust technique for the characterization of nano materials. A few software implementations of the reconstruction algorithm exist such as ptypy (<http://ptycho.github.io/ptypy/>) and PtychoLib [24], and a repository for sharing experimental data has been established [25].

Ptychography can be used to obtain large high-resolution images. It combines the large field of view of a scanning transmission microscope with the resolution of scattering measurements. In a scanning transmission microscope, operated in transmission mode, a focused beam is rastered across a sample, and the total transmitted intensity is recorded for each beam position. The pixel positions of the image obtained correspond to the beam positions used during the scan, and the value of the pixel to the intensity transmitted at that position. This limits the resolution of the image to the size of the impinging beam, which is typically limited by the quality of focusing optics and work distance constraints. In ptychography, instead of only using the total transmitted intensity, one typically records the distribution of that intensity in the far-field, i.e. the scattering pattern produced by the interaction of the illumination with the sample. The diffracted signal contains information about features much smaller than the size of the x-ray beam, making it possible to achieve higher resolutions than with the scanning techniques. The downside of having to use the intensities is that one now has to retrieve the corresponding phases to be able to reconstruct an image of the sample, which is made even more challenging by the presence of noise, experimental uncertainties, and perturbations of the experimental geometry. While it is a difficult problem, it is usually tractable by making use of the redundancy inherent in obtaining diffraction patterns from overlapping regions of the sample. This redundancy also permits the technique to overcome the lack of several experimental parameters and measurement uncertainties. For example, there are methods to recover unknown illuminations [26–29]. As a testament to their success these

methods are even used as a way of characterizing high quality x-ray optics [30–32], the wavefront of x-ray lasers [33] and EUV lithography tools [34].

Ptychographical phasing is a non-linear optimization problem [35] still containing many open questions [36]. Several strategies, such as Alternating Directions [37], projections, gradient [35], conjugate gradient, Newton [38–40], spectral methods [36, 41] and Monte-carlo [42], have been proposed to handle situations when both sample and positions [35, 41–43], are unknown parameters in high dimensions, and to handle experimental situations such as accounting for noise variance [39, 44], partial coherence [3, 45–47] [41, 46, 48], background [38, 41, 49, 50] or vibrations.

Here, we describe an algorithm approach and software environment SHARP (Scalable Heterogeneous Adaptive Real-time Ptychography) that enables high throughput streaming analysis using computationally efficient phase retrieval algorithms. The high performance computational back-end written in C/CUDA and implemented for NVIDIA GPU architectures is hidden from the microscopist, but can be accessed and adapted to particular needs by using a python interface or by modifying the source code.

Using SHARP we have built an intuitive graphical user interface that provides visual feedback, of both the recorded diffraction data as well as the reconstructed images, throughout the data acquisition and reconstruction processes at the Advanced Light Source (ALS).

We use a mathematical formulation of ptychography which was first introduced in [36–38, 40, 41].

## II. SHARP SOFTWARE ENVIRONMENT

### A. Forward model

In a ptychography experiment (see Fig. 1), one performs a series of diffraction measurement as a sample is rastered across an x-ray, electron or visible light beam. The illumination is formed by an x-ray optic such as a zone-plate. The measurement is performed by briefly exposing an area detector such as a CCD which records the scattered photons.

In a discrete setting, a two-dimensional small beam with distribution  $w(\mathbf{r})$  of dimension  $m_x \times m_y$  illuminates a subregion positioned at  $\mathbf{x}_{(i)}$  (referred to as frame) of an unknown object of interest  $\psi(\mathbf{r})$  of dimension  $n_x \times n_y$ . Here  $0 < m < n$ ,  $i = 1, \dots, K$  and  $K$  is the total number of frames (also referred to as “views” in the literature). For simplicity we consider square matrices. Generalization to non-square matrices is straightforward but requires more indices and complicates notation.

The pixel coordinates on a detector placed at a distance  $z_D$  from the sample are described as  $\mathbf{p} = (p_x, p_y, z_D)$ . Under far-field and paraxial approximations the pixel coordinates are related to reciprocal space co-




FIG. 1. Experimental geometry in ptychography: an unknown sample with transmission  $\psi(\mathbf{r})$  is rastered through an illuminating beam  $w(\mathbf{r})$ , and a sequence of diffraction measurements  $\mathcal{I}_{(i)} = |a_{(i)}(\mathbf{q})|^2$  are recorded on an area detector with pixel coordinates  $\mathbf{p}$  at a distance  $z_D$  from the sample.

ordinates

$$\begin{aligned} \mathbf{q} &= \mathbf{k}_{\text{out}} - \mathbf{k}_{\text{in}} \\ &= \frac{1}{\lambda} \left( \frac{(p_x, p_y, z_D)}{\sqrt{p_x^2 + p_y^2 + z_D^2}} - (0, 0, 1) \right) \\ &\simeq \frac{1}{\lambda z_D} (p_x, p_y, 0) \end{aligned}$$

where  $\mathbf{k}_{\text{in}} = (0, 0, k)$  and  $\mathbf{k}_{\text{out}} = k \frac{\mathbf{p}}{|\mathbf{p}|}$  are the incident and scattered wave vectors that satisfy  $|\mathbf{k}_{\text{in}}| = |\mathbf{k}_{\text{out}}| = k = 1/\lambda$ , and  $\lambda$  is the wavelength. With a distance  $p_m$  from the center to the edge of the detector, the diffraction limited resolution (half-period) of the microscope is given by the lengthscale  $r = \frac{\lambda z_D}{2p_m}$ . As a consequence, the coordinates in reciprocal and real space are defined as

$$\mathbf{q} = \left( \frac{\mu}{mr}, \frac{\nu}{mr} \right), \quad \mu, \nu \in \{0, \dots, m-1\}$$

and

$$\begin{aligned} \mathbf{r} &= (r\mu, r\nu), \quad \mu, \nu \in \{0, \dots, m-1\}, \\ \mathbf{x}_{(i)} &= (r\mu', r\nu'), \quad \mu', \nu' \in \{0, \dots, n-m\}. \end{aligned}$$

While  $\mathbf{x}_{(i)}$  is typically rastered on a coarser grid,  $\mathbf{r} + \mathbf{x}_{(i)}$  spans a finer grid of dimension  $n \times n$ .

In other words, we assume that a sequence of  $K$  diffraction intensity patterns  $\mathcal{I}_{(i)}(\mathbf{q})$  are collected as the position of the object is rastered on the position  $\mathbf{x}_{(i)}$ . The simple transform  $a_{(i)} = \sqrt{\mathcal{I}_{(i)}(\mathbf{q})}$  is a variance stabilizing transform for Poisson noise [51, 52]. The relationship among the amplitude  $a_{(i)}(\mathbf{q})$ , the illumination function  $w(\mathbf{r})$  and an unknown object  $\psi(\mathbf{r})$  to be estimated can be expressed as follows:

$$a_{(i)}(\mathbf{q}) = \left| \mathcal{F}w(\mathbf{r})\psi(\mathbf{r} + \mathbf{x}_{(i)}) \right| \quad (1)$$

and  $\mathcal{F}$  is the two-dimensional discrete Fourier transform,

$$(\mathcal{F}f)(\mathbf{q}) = \frac{1}{\sqrt{m^2}} \sum_{\mathbf{r}} e^{2\pi i \mathbf{q} \cdot \mathbf{r}} f(\mathbf{r}). \quad (2)$$

where the sum over  $\mathbf{r}$  is given on all the indices  $m \times m$  of  $\mathbf{r}$ . We define an operator  $T_{(i)}$ , that extracts a frame out of an image  $\psi$ , and build the illumination operator  $\mathbf{Q}_{(i)}$ , which scales the extracted frame point-wise by the illumination function  $w$ :

$$\begin{aligned}\mathbf{Q}_{(i)}[\psi](\mathbf{r}) &= w(\mathbf{r})\psi(\mathbf{r} + \mathbf{x}_{(i)}), \\ &= w(\mathbf{r})T_{(i)}[\psi](\mathbf{r}), \\ &= z_{(i)}(\mathbf{r}).\end{aligned}$$

With the operator  $\mathbf{Q}$ , eq. (1) can be represented compactly as:

$$\mathbf{a} = |\mathbf{F}\mathbf{Q}\psi^\vee|, \text{ or } \begin{cases} \mathbf{a} = |\mathbf{F}\mathbf{z}|, \\ \mathbf{z} = \mathbf{Q}\psi^\vee, \end{cases} \quad (3)$$

where the superscript  $\psi^\vee$  denotes the linearized version of the image (the superscript will be omitted for simplicity), and more explicitly as:

$$\begin{array}{c} \mathbf{a} \in \mathbb{R}^{Km^2} \\ \left[ \begin{array}{c} a_{(1)} \\ \vdots \\ a_{(K)} \end{array} \right] \end{array} = \begin{array}{c} \mathbf{F} \in \mathbb{C}^{Km^2 \times Km^2} \quad \mathbf{z} \in \mathbb{C}^{Km^2} \\ \left[ \begin{array}{ccc} F & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & F \end{array} \right] \left[ \begin{array}{c} z_{(1)} \\ \vdots \\ z_{(K)} \end{array} \right] \end{array}, \quad (4)$$

$$\begin{array}{c} \mathbf{z} \in \mathbb{C}^{Km^2} \\ \left[ \begin{array}{c} z_{(1)} \\ \vdots \\ z_{(K)} \end{array} \right] \end{array} = \begin{array}{c} \mathbf{Q} \in \mathbb{C}^{Km^2 \times n^2}, \quad \psi \in \mathbb{C}^{n^2} \\ \left[ \begin{array}{c} \text{diag}(w)T_{(1)} \\ \vdots \\ \text{diag}(w)T_{(K)} \end{array} \right] \left[ \begin{array}{c} \psi_1 \\ \vdots \\ \psi_{n^2} \end{array} \right] \end{array}. \quad (5)$$

where  $\mathbf{z}$  are  $K$  frames extracted from the object  $\psi$  and multiplied by the illumination function  $w$ , and  $\mathbf{F}$  is the associated 2D DFT matrix when we write everything in the stacked form [41]. When both the sample and the illumination are unknown, we can express the relationship (Eq. 5) between the image  $\psi$ , the illumination  $w$ , and the frames  $\mathbf{z}$  in two forms:

$$\mathbf{z} = \mathbf{Q}\psi = \text{diag}(\mathbf{S}w)\mathbf{T}\psi = \text{diag}(\mathbf{T}\psi)\mathbf{S}w \quad (6)$$

where  $\mathbf{S} \in \mathbb{R}^{Km^2 \times m^2}$  denotes the operator that replicates the illumination  $w$  into a stack of  $K$  frames, since  $\mathbf{Q}\psi = \text{diag}(\mathbf{S}w)\mathbf{T}\psi$  is the entry-wise product of  $\mathbf{T}\psi$  and  $\mathbf{S}w$ . Eq. (6) can be used to find  $\psi$  or  $w$  from  $\mathbf{z}$  and the other variable.

The Fourier transform relationship used in equations (1), (3) and (4) is valid under far-field and paraxial approximation, which is the focus of the current release of SHARP. For experimental geometries such as Near Field [53], Fresnel [54], Fourier ptychography [55], through-focus [56] partially coherent multiplexed geometries [48, 57, 58], under-sampling conditions [59] and to account for noise variance [41], one can substitute the simple Fourier transform with the appropriate propagator and variance stabilization [38].

## B. Phase retrieval

Projection operators form the basis of every iterative projection and projected gradient algorithms are implemented in SHARP and accessible through a library. The projection  $P_{\mathbf{a}}$  ensures that the frames  $\mathbf{z}$  match the experiment, that is, they satisfy Eq. (4), and is referred to as data projector:

$$P_{\mathbf{a}}\mathbf{z} = \mathbf{F}^* \frac{\mathbf{F}\mathbf{z}}{|\mathbf{F}\mathbf{z}|} \mathbf{a} \quad (7)$$

while the projection  $P_{\mathbf{Q}}$  onto the range of  $\mathbf{Q}$  (see Fig. 2):

$$P_{\mathbf{Q}} = \mathbf{Q}(\mathbf{Q}^*\mathbf{Q})^{-1}\mathbf{Q}^* \quad (8)$$

ensures that overlapping frames  $\mathbf{z}$  are consistent with each other and satisfy Eq. (5).

The projector  $P_{\mathbf{a}}$  is relatively robust to Poisson noise [51], but weighting factors to account for noisy pixels can be easily added [40].

Using relationship (6), we can update the image  $\psi$  from  $w$  and frames  $\mathbf{z}$ :

$$\psi \leftarrow \frac{\mathbf{Q}^*\mathbf{z}}{\mathbf{Q}^*\mathbf{Q}} \quad (9)$$

or the illumination  $w$  from an image  $\psi$  and frames  $\mathbf{z}$  [26, 27] multiplying (Eq. 6) on the left by  $\text{diag}\mathbf{T}\bar{\psi}$  and solving for  $w$ :

$$w \leftarrow \frac{\mathbf{S}^* \text{diag}(\mathbf{T}\bar{\psi})\mathbf{z}}{\mathbf{S}^*\mathbf{T}|\psi|^2}, \quad (10)$$

$\mathbf{S}$  where  $\bar{\psi}$  denotes the complex conjugate of  $\psi$ . See [29] for alternative updates, and [28] for convergence theory behind a similar blockwise optimization strategy. Several possible pathologies need to be accounted for when updating both  $\psi$  and  $w$ :

- Combined drift of the illumination and the image in real space. Drift is eliminated by keeping the illumination in the center of the frame by computing the center of mass and correcting for drifts after every update of the illumination.
- Fourier space drifts and grid pathologies are suppressed by enforcing either the absolute value  $a_w = |\mathcal{F}w_0|$  or support  $m_w$  of the Fourier transform of the unknown illumination  $w_0$ .
- A possible global phase factor between the solution and the reconstruction is taken into account in the error calculation.

A typical reconstruction with SHARP uses the following sequence:

1. Input data  $\mathcal{I}(\mathbf{q})$ , translations  $\mathbf{x}$ . Optional inputs: initial image  $\psi^{(0)}$ , illumination  $w^{(0)}$ , illumination Fourier mask  $m_w$  and illumination Fourier amplitudes  $a_w$ .

2. If  $w^{(0)}$  is not provided, initialize illumination by setting  $w^{(0)}$  to the inverse Fourier transform of the square root of the average frame.
3. If  $\psi^{(0)}$  is not provided, initialize the image by filling  $\psi^{(0)}$  with random numbers uniformly drawn from  $[0, 1)$ .
4. Build up  $\mathbf{Q}$ ,  $\mathbf{Q}^*$ , and  $(\mathbf{Q}^*\mathbf{Q})^{-1}$ , and frames  $\mathbf{z}^{(0)} = \mathbf{Q}\psi^{(0)}$ ;
5. Update the frames  $\mathbf{z}$  according to [60] using projector operators defined in (Eqs. (7,8)) below:

$$\mathbf{z}^{(l)} := [2\beta P_{\mathbf{Q}}P_{\mathbf{a}} + (1 - 2\beta)P_{\mathbf{a}} + \beta(P_{\mathbf{Q}} - I)]\mathbf{z}^{(l-1)},$$

where  $\beta \in (0.5, 1]$  is a scalar factor set by the user (set to 0.75 by default, which works in most cases).

6. Update image  $\psi^{(\ell)}$  using using Eq. 9.
7. If desired, compute a new illumination  $w$  using Eq. 10. If  $m_w$  is given apply the illumination Fourier mask constraint:

$$w^{(\ell)} := \mathcal{F}^{-1}\{(\mathcal{F}w)m_w\},$$

else if  $w_I$  is given apply the illumination Fourier intensities constraint:

$$w^{(\ell)} := \mathcal{F}^{-1}\left\{\frac{\mathcal{F}w}{|\mathcal{F}w|}a_w\right\},$$

else simply keep the unconstrained illumination  $w^{(\ell)} := w$ .

Now compute center of mass of  $w^{(\ell)}$  and shift it to fix the translation of the object.

8. If desired do background retrieval, that is, estimate static background and remove it in the iteration as described in [41] (p.7, Eq. 30).
9. Iterate from 5 until one of the metrics from Eqs. 11,12,14 drops below a user defined level or until a maximum iteration for time-critical applications, and return  $\psi^{(\ell)}$  and  $w$ .

The metrics  $\varepsilon_{\Delta}$ ,  $\varepsilon_a$ ,  $\varepsilon_Q$ ,  $\varepsilon_{\Delta}$  used to monitor progress and stagnation are the *normalized mean square root error* (nmse) from the corresponding projections of  $\mathbf{z}$ :

$$\varepsilon_{\mathbf{a}}(\mathbf{z}) = \frac{\| [P_{\mathbf{a}} - I]\mathbf{z} \|}{\|\mathbf{a}\|}, \quad (11)$$

$$\varepsilon_Q(\mathbf{z}) = \frac{\| [P_Q - I]\mathbf{z} \|}{\|\mathbf{a}\|}, \quad (12)$$

$$\varepsilon_{\Delta}(\mathbf{z}^{(l)}, \mathbf{z}^{(l-1)}) = \frac{\| \mathbf{z}^{(l)} - \mathbf{z}^{(l-1)} \|}{\|\mathbf{a}\|} \quad (13)$$

where  $I$  is the identity operator, and  $\mathbf{z}^{(0)} = 0$ ,

For benchmarking purposes, when using a simulation from a known solution  $\psi_0$ , the following metric can also be used:

$$\varepsilon_0(\mathbf{z}) = \frac{1}{\|\mathbf{Q}^*\mathbf{z}_0\|} \min_{\varphi} \|\mathbf{Q}^*(e^{i\varphi}\mathbf{z} - \mathbf{z}_0)\|, \quad (14)$$

where  $\varphi$  is an arbitrary global phase factor, and  $\mathbf{z}_0 = \mathbf{Q}\psi_0$ . Notice the additional scaling factor  $\mathbf{Q}^*$  used in  $\varepsilon_0$ .

The initial values for the input data and translations can either be loaded from file or set by a python interface. The starting “zero-th” initial image is loaded from file, set to a random image, or taken as a constant image.

### C. Computational Methodology

SHARP was developed to achieve the highest performance, taking advantage of the algorithm described earlier and using a distributed computational backend. The ptychographic reconstruction algorithm requires one to compute the product of several linear operators ( $\mathbf{Q}$ ,  $\mathbf{Q}^*$ ,  $\mathbf{F}$ ,  $\mathbf{F}^*$ ,  $\mathbf{S}$ ,  $\mathbf{S}^*$ ) on a set of frames  $\mathbf{z}$ , an image  $\psi$  and an illumination  $w$  several times. We use a distributed GPU architecture across multiple nodes for this task (Fig. 2).

To implement fast operators, a set of GPU kernels and MPI communication are necessary. The split ( $\mathbf{Q}\psi$ ) and overlap ( $\mathbf{Q}^*\mathbf{z}$ ) kernels are among the most bandwidth demanding kernels and play an important role in the process.

The strategy used to implement those kernels impacts directly the overall performance of the reconstruction algorithm. To divide the problem among multiple nodes, SHARP initially determines the size of the final image based on the list of translations, frames size, and resolution. It subsequently assigns a list of translations to every node and loads the corresponding frames onto GPUs.

The split ( $\mathbf{Q}\psi$ ) and FFT ( $\mathbf{F}$ ) operations are easily parallelized because of the framewise intrinsic independence. Summing the frames onto an image ( $\mathbf{Q}^*\mathbf{z}$ ) requires a reduction for every image pixel across neighboring MPI nodes. Within each GPU the image is divided into blocks and we first determine which frames contribute to each block. The contributing frames are summed and then the resulting image is summed across all MPI nodes. We use shared memory or constant memory, depending on GPU compute capability, to store frame translations, and we use kernel fusion to reduce access to global memory. The last step of summing across all MPI nodes does not necessarily have to be done at every iteration, at the cost of slower convergence [61], but that is the default.

Timing to compute the overlap at each iteration depends on the size of the image and number of frames on top of each pixel, i.e. the density but not the size of the frames.

In addition to the high performance ptychographic algorithm, the SHARP software environment provides a flexible and modular framework which can be changed and adapted to different needs. Furthermore, the user has control of several options for the reconstruction algorithm, which can be used to guarantee a balance between performance and quality of the results. These include the choice of illumination Fourier mask, illumination Fourier intensities and the  $\beta$  parameter, as well as how often to do different operations such as illumination retrieval, background retrieval, and synchronization of the different GPUs. For more details we refer the reader to the

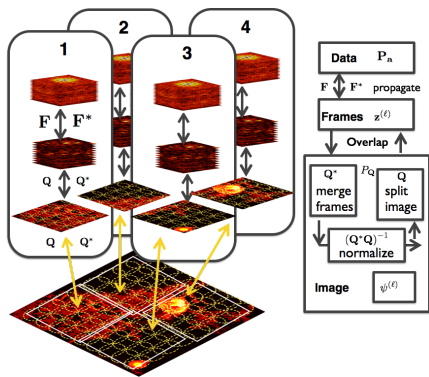


FIG. 2. Schematic of the ptychographic reconstruction algorithm implemented in SHARP. The iterative reconstruction scheme is shown on the right. To achieve the highest possible throughput and scalability one has to parallelize across multiple GPUs as shown on the left for the case of 4 GPUs. As most ptychographic scans use a constant density of scan point across the object, we expect to be able to achieve a very even division, resulting in good load balancing. SHARP enforces an overlap constraint between the images produced by each of the GPUs, and also enforces that the illumination recovered on each GPU agree with each other. This is done by default at every iteration.

documentation (<http://www.camera.lbl.gov/sharp>).

### III. APPLICATIONS AND PERFORMANCE

SHARP enables high-throughput streaming analysis using computationally efficient phase retrieval algorithms. In this section we describe a typical dataset and sample that can be collected in less than 1 minute at the ALS, and the computational backend to provide fast feedback to the microscopist.

To characterize our performance, we use both simulations and experimental data. We use simulations to compare the convergence of the reconstruction algorithm to the “true solution” and characterize the effect of different light sources, contrast, scale, noise, detectors or samples for which no data exists yet.

Experimental data from ALS used to characterize battery materials, green cement, magnetic materials, at different wavelengths and orientation has been successfully reconstructed [62–66] using the software described in this article.

We also describe a streaming example in which a front-end that operates very close to the actual experiment sends the data to the reconstruction backend that runs remotely on a GPU/CPU cluster. Further details about the streaming front-end and processing back-end pipeline will be published in an upcoming paper by our group.




FIG. 3. Convergence rate (top) per iteration and timing (bottom) to process 10,000 frames of dimension  $128 \times 128$  extracted from an image of size  $1000 \times 1000$  as a function of the number of nodes. All residuals decrease rapidly; numerical precision limits the (weighted) comparison with the known solution  $\varepsilon'(z)$ . Reconstruction is achieved ( $\varepsilon_0 < 5e-4$ ) in under 2 seconds using a cluster with 4 compute nodes with 4 GTX Titan GPU per node (16 total, 43000 cores), 96 GBytes GPU memory, 1 TByte RAM, and 24 TBytes storage, infiniband. Timing contributions for corresponding computational kernels are  $(\mathbf{Q}^* \mathbf{Q})^{-1} \mathbf{Q}^*$  30 %,  $\mathbf{F}, \mathbf{F}^*$  20 %,  $\mathbf{Q}$  20 %,  $\mathbf{S}^{ast}$  5 %, elementwise operations 20 %, and residual calculation 5 %. No illumination retrieval was done, as the exact illumination was given. The simulation was done using periodic boundary conditions to avoid edge effects.

#### A. Simulations and performance

As a demonstration, we start from a sample that was composed of colloidal gold nanoparticles of 50 nm and 10 nm deposited on a transparent silicon nitride membrane. An experimental image was obtained by scanning electron microscopy, which provides high resolution and contrast but can only view the surface of the sample.

We simulate a complex transmission function by scaling the image amplitude from 0 to 50 nm thickness, and using the complex index of refraction of gold at 750 eV energy from [henke.lbl.gov]. The illumination is generated by a zone-plate with a diameter of 220 microns and 60 nm outer zone width, discretized into  $(128 \times 128)$  pixels in the far field.

#### B. Experimental example

Figure 4 shows ptychographic reconstructions of a dataset generated from a sample consisting of gold balls with diameters of 50 and 10 nm. The data were generated using 750 eV x-rays at beamline 5.3.2.1 of the Advanced Light Source, with high stability position control of the




FIG. 4. Reconstruction of a test sample consisting of gold balls with diameters of 50 and 10 nm. Detector pixel size 30 microns,  $1920 \times 960$  pixels 80 mm downstream from the sample, cropped and downsampled to 128, scan of  $(50 \times 50)$  points, illumination is generated by a zone-plate with a diameter of 220 microns and 60 nm outer zone width. A) Phase image generated by SHARP using the algorithm described in section II B applying the illumination Fourier mask constraint and turning on background retrieval. The red arrow points to a collection of 50 nm balls while the blue arrow points to a collection of 10 nm balls. The pixel size is 10 nm. B) Same as (A) except without enforcing the illumination Fourier mask. C) Same as (A) but without using the background retrieval algorithm.

soft x-ray scanning transmission microscope. Exposure time was 1 second and the dataset consists of a square scan grid with 40 nm spacing ( see [62] for details of the experimental setup). The reconstructions consisted of 300 iterations of the RAAR algorithm with a illumination retrieval and background retrieval step every other iteration. The initial illumination is generated by (1) computing the average of the measurements, (2) setting everything below a threshold to 0, and everything above a threshold to a constant average value (3) applying an inverse FFT. The image is initialized with complex independent identically distributed (i.i.d.) pixels, or a constant average value.

### C. Interface and Streaming

Common processing pipelines used for ptychographic experiments usually have a series of I/O operations and many different components involved. We have developed a streaming pipeline, to be deployed at the COSMIC beamline at the ALS, which allows users to monitor and quickly act upon changes along the experimental and computational pipeline.

The streaming pipeline is composed of a front-end and a back-end (Fig. 5). The front-end consists of a Graphical User Interface (see Fig. 5), a worker that grabs frames from the detector, and an interface that monitors network activity, experimental parameters (position, wavelength, exposure, etc...), and provides a live view of the ongoing reconstruction.

On the back-end side, the streaming infrastructure is composed of a communication handler and a collection of workers addressing different tasks such as dark calibration, detector correction, data reduction, ptychographic reconstruction and writing output to file.

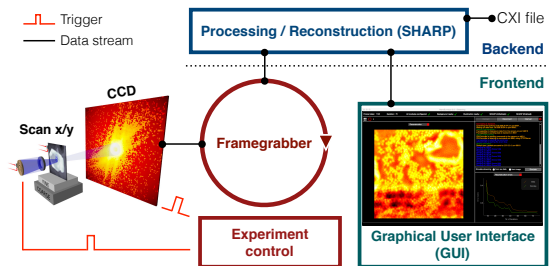


FIG. 5. Overview of the components involved in the software structure of the streaming pipeline. In order to maximize the performance of this streaming framework, the frontend operates very close to the actual experiment, while the backend runs remotely on a powerful GPU/CPU cluster. As soon as diffraction data is recorded by the CCD camera, a live view of the ptychographic reconstruction is transmitted to the Graphical User Interface, and the user is able to control and monitor (top panel) the current status of the data streams and analysis, (bottom right panel) .

This software architecture allows users an intuitive, flexible and responsive monitoring and control of their experiments. Such a tight integration between data acquisition and analysis is required to give users the feedback they expect from a STXM instrument.

## IV. CONCLUSIONS

In this paper we described SHARP, a high-performance software environment for ptychography reconstructions, and its application as part of quick feedback system used by the ptychographic microscopes installed at the Advanced Light Source.

Our software provides a modular interface to the high performance computational back-end and can be adapted to different needs. Its fast throughput provides near real time feedback to microscopists and this also makes it suitable as a corner stone for demanding higher dimensional analysis such as spectro-ptychography or tomography.

With the coming new generation light sources and faster detectors, the ability to quickly analyse vast amounts of data to obtain large high-dimensional images will be an enabling tool for science.

## V. ACKNOWLEDGMENTS

We acknowledge useful discussions with Chao Yang, H-T Wu, J. Qian and Z. Wen. This work was partially funded by the Center for Applied Mathematics for Energy Research Applications, a joint ASCR-BES funded project within the Office of Science, US Department of Energy, under contract number DOE-DE-AC03-76SF00098, by the Swedish Research Council and by the



- [1] J. R. Fienup, *Appl. Opt.* **21**, 2758 (1982).
- [2] D. R. Luke, J. V. Burke, and R. G. Lyon, *SIAM review* **44**, 169 (2002).
- [3] J. R. Fienup, J. C. Marron, T. J. Schulz, and J. H. Seldin, *Appl. Opt.* **32**, 1747 (1993).
- [4] M. Eckert, *Acta Crystallographica Section A* **68**, 30 (2012).
- [5] A. Momose, T. Takeda, Y. Itai, and K. Hirano, *Nature medicine* **2**, 473 (1996).
- [6] R. Collier, C. Burckhardt, and L. Lin, *Optical holography* (Academic Press, New York, 1971).
- [7] S. Marchesini, S. Boutet, and e. a. Sakdinawat, *Nature Photonics* **2**, 560 (2008).
- [8] P. W. Hawkes and J. C. H. Spence, eds., *Science of microscopy* (Springer, New York, 2007).
- [9] A. Fannjiang and W. Liao, *J. Opt. Soc. Am. A* **29**, 1847 (2012).
- [10] Y. Wang and Z. Xu, *ArXiv e-prints* **1310.0873** (2013).
- [11] F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David, *Nature physics* **2**, 258 (2006).
- [12] E. J. Candes, Y. C. Eldar, T. Strohmer, and V. Voroninski, *SIAM Journal on Imaging Sciences* **6**, 199 (2013).
- [13] I. Waldspurger, A. d'Aspremont, and S. Mallat, *Mathematical Programming* (2013).
- [14] J. Miao, P. Charalambous, J. Kirz, and D. Sayre, *Nature* **400**, 342 (1999).
- [15] S. Marchesini, H. He, H. Chapman, S. Hau-Riege, A. Noy, M. Howells, U. Weierstall, and J. Spence, *Phys. Rev. B* **68**, 140101 (2003).
- [16] H. H. Bauschke, P. L. Combettes, and D. R. Luke, *Journal of the Optical Society of America. A, Optics, image science, and vision* **19**, 1334 (2002).
- [17] S. Marchesini, *Rev Sci Instrum* **78**, 011301 (2007), [arXiv:physics/0603201](https://arxiv.org/abs/physics/0603201).
- [18] S. Marchesini, *JOSA A* **24**, 3289 (2007).
- [19] J. M. Rodenburg, *Advances in Imaging and Electron Physics* **150**, 87 (2008).
- [20] W. Hoppe, *Acta Crystallographica Section A* **25**, 495 (1969).
- [21] R. Hegerl and W. Hoppe, *Berichte Der Bunsen-Gesellschaft Fur Physikalische Chemie* **74**, 1148 (1970).
- [22] P. D. Nellist, B. C. McCallum, and J. M. Rodenburg, *Nature* **374**, 630 (1995).
- [23] H. N. Chapman, *Ultramicroscopy* **66**, 153 (1996).
- [24] Y. S. G. Nashed, D. J. Vine, T. Peterka, J. Deng, R. Ross, and C. Jacobsen, *Opt. Express* **22**, 32082 (2014).
- [25] F. R. Maia, *Nature methods* **9**, 854 (2012).
- [26] P. Thibault, M. Dierolf, A. Menzel, O. Bunk, C. David, and F. Pfeiffer, *Science* **321**, 379 (2008).
- [27] P. Thibault, M. Dierolf, O. Bunk, A. Menzel, and F. Pfeiffer, *Ultramicroscopy* **109**, 338 (2009).
- [28] R. Hesse, D. R. Luke, S. Sabach, and M. K. Tam, *SIAM Journal on Imaging Sciences* **8**, 426 (2015).
- [29] S. Marchesini and H.-T. Wu, *Rank-1 accelerated illumination recovery in scanning diffraction*, Tech. Rep. LBNL-6734E, [arXiv:1408.1922](https://arxiv.org/abs/1408.1922) (Lawrence Berkeley National Laboratory, 2014) [arXiv:1408.1922](https://arxiv.org/abs/1408.1922) [math.OC].
- [30] C. Kewish, P. Thibault, M. Dierolf, O. Bunk, A. Menzel, J. Vila-Comamala, K. Jefimovs, and F. Pfeiffer, *Ultramicroscopy* **110**, 325 (2010).
- [31] S. Hönig, R. Hoppe, J. Patommel, A. Schropp, S. Stephan, S. Schöder, M. Burghammer, and C. G. Schroer, *Opt. Express* **19**, 16324 (2011).
- [32] M. Guizar-Sicairos, S. Narayanan, A. Stein, M. Metzler, A. R. Sandy, J. R. Fienup, and K. Evans-Lutterodt, *Applied Physics Letters* **98**, 111108 (2011).
- [33] A. Schropp, R. Hoppe, V. Meier, J. Patommel, F. Seiboth, H. J. Lee, B. Nagler, E. C. Galtier, B. Arnold, U. Zastra, et al., *Scientific reports* **3**, 01633 (2013).
- [34] A. Wojdyla, R. Miyakawa, and P. Naulleau, in *SPIE Advanced Lithography* (International Society for Optics and Photonics, 2014) pp. 904839–904839.
- [35] M. Guizar-Sicairos and J. R. Fienup, *Opt. Express* **16**, 7264 (2008).
- [36] S. Marchesini, Y.-C. Tu, and H.-t. Wu, *Applied and Computational Harmonic Analysis*, 10.1016/j.acha.2015.06.005 (2015).
- [37] Z. Wen, C. Yang, X. Liu, and S. Marchesini, *Inverse Problems* **28**, 115010 (2012).
- [38] C. Yang, J. Qian, A. Schirotzek, F. Maia, and S. Marchesini, *Iterative Algorithms for Ptychographic Phase Retrieval*, Tech. Rep. 4598E, [arXiv:1105.5628](https://arxiv.org/abs/1105.5628) (Lawrence Berkeley National Laboratory, 2011).
- [39] P. Thibault and M. Guizar-Sicairos, *New Journal of Physics* **14**, 063004 (2012).
- [40] J. Qian, C. Yang, A. Schirotzek, F. Maia, and S. Marchesini, *Inverse Problems and Applications, Contemp. Math* **615**, 261 (2014).
- [41] S. Marchesini, A. Schirotzek, C. Yang, H.-t. Wu, and F. Maia, *Inverse Problems* **29**, 115009 (2013).
- [42] A. Maiden, M. Humphry, M. Sarahan, B. Kraus, and J. Rodenburg, *Ultramicroscopy* **120**, 64 (2012).
- [43] M. Beckers, T. Senkbeil, T. Gorniak, K. Giewekemeyer, T. Salditt, and A. Rosenhahn, *Ultramicroscopy* **126**, 44 (2013).
- [44] P. Godard, M. Allain, V. Chamard, and J. Rodenburg, *Opt. Express* **20**, 25914 (2012).
- [45] B. Abbey, K. A. Nugent, G. J. Williams, J. N. Clark, A. G. Peele, M. A. Pfeifer, M. de Jonge, and I. McNulty, *Nature Physics* **4**, 394 (2008).
- [46] N. C. Jesse and G. P. Andrew, *Applied Physics Letters* **99**, 154103 (2011).
- [47] L. W. Whitehead, G. J. Williams, H. M. Quiney, D. J. Vine, R. A. Dilanian, S. Flewett, K. A. Nugent, A. G. Peele, E. Balaur, and I. McNulty, *Phys. Rev. Lett.* **103**, 243902 (2009).
- [48] L. Tian, X. Li, K. Ramchandran, and L. Waller, *Biomedical Optics Express* **5**, 2376 (2014).
- [49] S. T. Thurman and J. R. Fienup, *JOSA A* **26**, 1008 (2009).
- [50] M. Guizar-Sicairos and J. R. Fienup, *Opt. Express* **17**, 2670 (2009).
- [51] F. J. Anscombe, *Biometrika* **35**, 246 (1948).



- [52] M. Mäkitalo and A. Foi, *IEEE Transactions on Image Processing* **22**, 91 (2013).
- [53] M. Stockmar, P. Cloetens, I. Zanette, B. Enders, M. Dierolf, F. Pfeiffer, and P. Thibault, *Scientific reports* **3**, 01927 (2013).
- [54] D. Vine, G. Williams, B. Abbey, M. Pfeifer, J. Clark, M. De Jonge, I. McNulty, A. Peele, and K. Nugent, *Physical Review A* **80**, 063823 (2009).
- [55] G. Zheng, R. Horstmeyer, and C. Yang, *Nature Photonics* **7**, 739 (2013).
- [56] J. Marrison, L. Rätty, P. Marriott, and P. O'Toole, *Scientific reports* **3**, 02369 (2013).
- [57] D. J. Batey, D. Claus, and J. M. Rodenburg, *Ultramicroscopy* **138**, 13 (2014).
- [58] S. Dong, R. Shiradkar, P. Nanda, and G. Zheng, *Biomedical Optics Express* **5**, 1757 (2014).
- [59] T. B. Edo, D. J. Batey, A. M. Maiden, C. Rau, U. Wagner, Z. D. Pešić, T. A. Waigh, and J. M. Rodenburg, *Phys. Rev. A* **87**, 053850 (2013).
- [60] R. Luke, *Inverse Problems* **21**, 37 (2005).
- [61] J. Liu, S. J. Wright, C. Ré, V. Bittorf, and S. Sridhar, *The Journal of Machine Learning Research* **16**, 285 (2015).
- [62] D. A. Shapiro, Y.-S. Yu, T. Tyliczszak, J. Cabana, R. Celestre, W. Chao, K. Kaznatcheev, A. Kilcoyne, F. Maia, S. Marchesini, *et al.*, *Nature Photonics* (2014).
- [63] Y.-S. Yu, C. Kim, D. A. Shapiro, M. Farmand, D. Qian, T. Tyliczszak, A. D. Kilcoyne, R. Celestre, S. Marchesini, J. Joseph, *et al.*, *Nano letters* **15**, 4282 (2015).
- [64] S. Bae, R. Taylor, D. Shapiro, P. Denes, J. Joseph, R. Celestre, S. Marchesini, H. Padmore, T. Tyliczszak, T. Warwick, *et al.*, *Journal of the American Ceramic Society* **98**, 4090 (2015).
- [65] Y. Li, S. Meyer, J. Lim, S. C. Lee, W. E. Gent, S. Marchesini, H. Krishnan, T. Tyliczszak, D. Shapiro, A. L. D. Kilcoyne, *et al.*, *Advanced Materials* **27**, 6590 (2015).
- [66] X. Shi, P. Fischer, V. Neu, D. Elefant, J. Lee, D. Shapiro, M. Farmand, T. Tyliczszak, H.-W. Shiu, S. Marchesini, *et al.*, *Applied Physics Letters* **108**, 094103 (2016).