



# A Monte Carlo ray-tracing simulation of coherent X-ray diffractive imaging

Giovanni Fevola,<sup>a</sup> Erik Bergbäck Knudsen,<sup>b</sup> Tiago Ramos,<sup>a</sup> Dina Carbone<sup>c</sup> and Jens Wenzel Andreassen<sup>a\*</sup>

<sup>a</sup>Department of Energy Conversion and Storage, Technical University of Denmark, Frederiksborgvej 399, Roskilde 4000, Denmark, <sup>b</sup>Department of Physics, Technical University of Denmark, Fysikvej 311, Kgs Lyngby 2800, Denmark, and <sup>c</sup>MAX IV Laboratory, Lund University, 22484 Lund, Sweden. \*Correspondence e-mail: jewa@dtu.dk

Received 21 May 2019

Accepted 22 October 2019

Edited by V. Favre-Nicolin, CEA and Université Joseph Fourier, France

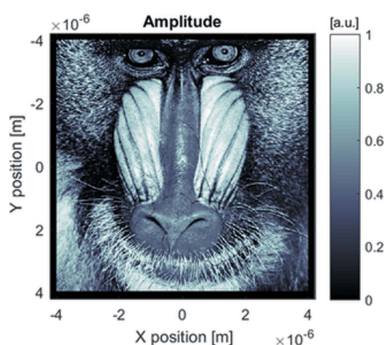
**Keywords:** X-ray microscopy; Monte Carlo simulations; ray tracing; coherent diffractive imaging; ptychography.

**Supporting information:** this article has supporting information at journals.iucr.org/s

Coherent diffractive imaging (CDI) experiments are adequately simulated assuming the thin sample approximation and using a Fresnel or Fraunhofer wavefront propagator to obtain the diffraction pattern. Although this method is used in wave-based or hybrid X-ray simulators, here the applicability and effectiveness of an alternative approach that is based solely on ray tracing of Huygens wavelets are investigated. It is shown that diffraction fringes of a grating-like source are accurately predicted and that diffraction patterns of a ptychography dataset from an experiment with realistic parameters can be sampled well enough to be retrieved by a standard phase-retrieval algorithm. Potentials and limits of this approach are highlighted. It is suggested that it could be applied to study imperfect or non-standard CDI configurations lacking a satisfactory theoretical formulation. The considerable computational effort required by this method is justified by the great flexibility provided for easy simulation of a large-parameter space.

## 1. Introduction

Coherent diffractive imaging (CDI) techniques have become increasingly popular in the X-ray community over the last two decades, to the point that most synchrotrons have at least one beamline dedicated to them (Miao *et al.*, 2015). CDI techniques are based on the inversion of coherent X-ray intensity data to obtain an image of the diffracting sample. This inversion is performed through algorithms that utilize boundary conditions in both real and reciprocal space, *e.g.* sample size and measured data, respectively, for the recovery of amplitude and phase of the sample-diffracted field (Veen & Pfeiffer, 2004). Thus, the recovered images are obtained without the use of lenses (Chapman & Nugent, 2010), which represents the most appealing feature of CDI, implying that the theoretical limit for the image resolution is determined only by the extent of reciprocal space that can be recorded with statistical significance (Noh *et al.*, 2016). The actual resolution, however, is affected by other factors related both to the physics of the experiment and to the processing of the collected data. Non-ideal conditions, such as partial coherence, mechanical instabilities and electronic noise, all contribute to the degradation of resolution by decreasing data quality. Although these can be accounted for in the reconstruction algorithms, and several strategies to mitigate their detrimental effects have been proposed in the literature (Thibault & Guizar-Sicairos, 2012; Godard *et al.*, 2012; Chang *et al.*, 2019), they still represent a strong limitation for the ultimate achievable resolution. This is demonstrated by the fact that, although remarkable for a non-destructive technique, the current



record resolutions for tomographic reconstructions based on full-field CDI [ $\sim 10$  nm (Takahashi *et al.*, 2010)] or ptychographic projections [14.6 nm (Holler *et al.*, 2017), 13.6 nm (Ramos *et al.*, 2017)] are probably an order of magnitude higher than the theoretical limit (Marchesini *et al.*, 2003; Dietze & Shpyrko, 2015), at least for inorganic materials (Howells *et al.*, 2009).

Due to the limited availability of coherent X-ray sources, it becomes of primary importance to be able to perform an offline optimisation of the experiment, prepare the data acquisition and predict the quality of its outcome. Therefore, reliable tools for the simulation of CDI experiments become an essential ingredient for the success of the experiment. The ability to reproduce data acquired under realistic conditions is essential to test and optimize retrieval algorithms and, furthermore, it helps to identify factors that could potentially limit the resolution and enables exploration of novel experimental configurations.

In the case of CDI, an inverse problem needs to be solved to retrieve the phase of the wavefield measured only through its intensities. What defines this inverse problem is the model that relates the object to the measurement: this model describes the interaction of X-rays with the sample and the propagation of the produced 'exit' wavefield to the detector, both of which depend on the chosen technique, *e.g.* measurement geometry, detector distance and X-ray wavelength. A typical example of such interaction is the attenuation and phase delay experienced by the X-ray beam when transmitted through a sample in the so-called forward scattering geometry. In this case, the propagation is well described by the Fresnel or the Fraunhofer propagator depending on whether the patterns are collected in the near- or far-field regime, respectively. These models can be used to simulate an experiment, and noise can be added to the simulated data by scaling the wavefields according to the expected flux and extracting values for shot noise from a Poisson distribution. Electronic noise can be added as well using, for example, a model of Gaussian distribution.

However, the described approach uses, for the inversion of data, the same model used to produce them [inverse crime (Colton & Kress, 1992)]; moreover, all noise is based on models and great care has to be taken to ensure that it is propagated correctly, however accurate they might be. Finally, other physical variables that might affect the experimental results, such as optical aberrations or the presence of other optical elements on the beam path, *e.g.* pinholes or beam-defining slits (Mastropietro *et al.*, 2011), are difficult to account for unless a specific and accurate model can be created to describe analytically their impact, which can in itself be a relatively complex task. We introduce here an alternative geometrical means for the simulation of coherent X-ray diffraction data, entirely based on ray tracing. This very general framework offers the possibility to customize the model of sample interaction (each intercepted element of the sample can modify any parameter of the ray), to include elements of the setup that affect data quality (optical elements, slits, pinholes, windows) and, for diffraction, relies on a very general and simple propagator based on Huygens

wavelets. This approach has the double advantage of producing diffraction data that are *intrinsically* noisy due to the Monte Carlo method used, and totally *independent* of the forward model on which reconstruction algorithms might rely, ensuring prevention of the inverse crime. The scheme for diffraction is not radically different from the already reported scheme for diffraction in *McXtrace* (Bergbäck Knudsen *et al.*, 2013), but it contains an important modification which is detailed in the following and, importantly, is applied for comparison to a realistic test case for which solutions from other approaches are available.

With respect to other tools available in the existing X-ray diffraction simulation software portfolio, our approach provides purely ray-tracing-based simulations for fully coherent X-ray diffraction experiments, which are still absent in the existing literature. While some notable tools such as *ShadowOUI* (Rebuffi & Sánchez del Río, 2016) and *OASYS* (Sanchez del Rio & Rebuffi, 2019) rely on hybrid models using both approaches, available software tools are mainly divided into those based on wave-propagation methods [*e.g.* *XRT* (Klementiev & Chernikov, 2014), *SRW* (Chubar *et al.*, 2013), *PHASE/REDUCE* (Bahrtdt, 2007)] and those based on ray tracing [*e.g.* *SHADOW* (Lai & Cerrina, 1986), *RAY* (Schäfers, 2008)]. The former exploit the laws of physical optics and Fourier integrals whereas the latter exploit the laws of geometrical optics. *McXtrace* (Bergbäck Knudsen *et al.*, 2013) falls into this second category and is, furthermore, Monte Carlo based. The ray-tracing approach might not appear the most suitable to describe diffraction, which is typically viewed as an undulatory phenomenon and is qualitatively explained by Huygens' principle. On the other hand, when complemented with a Monte Carlo based approach, ray tracing provides all degrees of freedom when it comes to the difficult task of describing non-ideal conditions from first principles. For instance, partial coherence, extremely relevant for our study, has been suitably addressed in this context (Prodi *et al.*, 2011; Cipiccia *et al.*, 2014), and beam monochromaticity, divergence and polarization effects are particularly easy to simulate. Furthermore, this method does not rely on paraxial approximation and can be particularly effective for simulation of CDI datasets that extend far in the reciprocal space (Shapiro *et al.*, 2014). This will be relevant for experiments carried out with high flux in the soft X-ray regime that aim at extreme resolutions as will certainly be the case at the new diffraction-limited sources (Hettel, 2014).

Outside of the X-ray context, the use of a geometrical approach for simulating diffraction dates back to the work of Young & Keller (see Kumar & Ranganath, 1991) and ray-tracing simulation of diffraction is widely explored: Mahan *et al.* (2018) established the conditions under which a Fresnel-Huygens principle is well described; Andreas *et al.* (2015) reported a vectorial ray tracer applied to laser interferometry in non-paraxial cases; Mout *et al.* (2016, 2018) have described the theory that makes Monte Carlo ray-tracing consistent, demonstrating diffraction of a multiple-component microscope in this framework. These works bear strong similarities with the ray-tracing scheme we have used here to simulate

CDI, which will be described. In this manuscript, we demonstrate the efficacy of our approach by focusing on a couple of simple cases of CDI measured in forward geometry and under far-field conditions. For demonstrative purposes, we choose examples that can be solved with existing approaches for comparison. Our approach, however, can be extended to other more complex cases where other methods fail and its impact is potentially the highest. These include simulation of novel approaches such as near-field CDI, grazing-incidence ptychography, or where the use of standard inversion algorithms fails to converge. This is the case of CDI under the Bragg condition and in the presence of strongly distorted phase fields. In this case the knowledge obtained from the simulations could help to reduce the space of possible solutions with a hybrid approach of inversion and fitting procedures (Mastropietro *et al.*, 2013). Other cases include the simulation of wavefronts from imperfect optics, where it would be useful to disentangle contributions from sample and wavefield. Simulations from low-contrast samples would help to identify the minimum dose necessary for inversion of data, and to optimize beforehand the setup of the experiments.

## 2. Methods – ray tracing of CDI

We focus on the simulation of CDI measured under far-field conditions. The implementation of a proper simulation, in this case, is achieved by ensuring that the two main requirements for CDI are satisfied. The first proviso of CDI is the thin sample assumption (Rodenburg, 2008), which allows the sample to be described with a 2D object function [the conditions for meeting this assumption are discussed by Dierolf *et al.* (2010)]. That withstanding, the multiplication approximation also holds, *i.e.*  $\psi$  (the wavefield immediately after the sample) can be expressed as the product of  $P$ , the wavefield immediately before the sample (also referred to as the probe), and  $O$ , the object function. Namely,

$$\psi(x, y) = P(x, y) O(x, y), \quad (1)$$

where  $(x, y)$  indicates the position vector in the sample plane. The second condition is that the intensity measured in the Fraunhofer, or far-field, regime, has a Fourier transform relation with the wavefield  $\psi$ , *i.e.* the intensity  $I$  in the far field should respect the following relation,

$$I(X, Y) = |\Psi(X, Y)|^2 = |\mathcal{F}[P(x, y) O(x, y)]|^2, \quad (2)$$

where  $\mathcal{F}$  denotes the Fourier transform and  $(X, Y)$  is the position vector in the detector plane. In the following, we discuss how the multiplication approximation and the transformation into the far field were implemented and tested in *McXtrace*. Before entering into the details, we need to recall some key concepts regarding ray tracing, *i.e.* the Monte Carlo approach and the computational flow. Further details can be found in the work by Bergbäck Knudsen *et al.* (2013). In *McXtrace*, a ray is described by 12 variables, including its position  $(x, y, z)$ , propagation vector  $\mathbf{k} = (k_x, k_y, k_z)$ , weight  $p$  and phase  $\phi$ . It can conceptually be viewed as an ensemble of photons sharing the same position in space, the same propa-

gation vector and the same phase. The weight of a ray accounts for the number of photons that it represents. We neglect the remaining four variables for a ray, which relate to time and polarization, as they are not relevant for our discussion. Polarization has little effect for small scattering angles of the beam in both the horizontal and vertical directions. The time accounts for the longitudinal coherence, which in the thin object approximation is not relevant. Synthetically, we can write a set of rays as

$$r_i = (x, y, z, k_x, k_y, k_z, p, \phi), \quad i = 1, \dots, N, \quad (3)$$

where  $N$  defines the size of such set and is an input parameter. We define the ‘trace’ as the set of values that a ray assumes at all possible abscissae. If  $z$  is defined as the propagation direction, we define the trace  $T_i$  of the  $i$ th ray as

$$T_i = \{(x, y, z, k_x, k_y, k_z, p, \phi), \quad z_0 \leq z \leq z_{\text{end}}\}, \quad (4)$$

where  $z_0$  and  $z_{\text{end}}$  denote the boundaries of the region of interest. In this description, each component of an X-ray instrument can be introduced sequentially by placing it at specific positions ( $z_c$ ) along the beam path, with its own customizable parameters, effectively translating any (relevant) element of a beamline into software code. Formally, the instrument defines a rule  $BL$  to compute the trace of a ray,

$$BL: r_i \rightarrow T_i, \quad i = 1, \dots, N. \quad (5)$$


$BL$  describes the interaction of every specific component with the ray, and how it affects the propagation of the rays as well as of propagation in free space. A trace is updated as a result of simple propagation in space or as a result of a particular interaction with a component. As with all Monte Carlo methods, a single trace does not provide relevant information per se, but cumulative statistics of all traces do. In *McXtrace* the computational flow of a simulation is, in general, sequential and static, *i.e.* a pre-defined number of independent traces are computed in a sequential way. In some limited cases, the sequence may be purposely broken at discrete logical points in the simulation, but we have no need for this feature in our discussion.

Fig. 1 illustrates the components of a setup for a CDI measurement with a focused X-ray beam. The simulation of this measurement with *McXtrace* follows the path of the rays through the four components. For each ray the following operations are performed in order:

(i) A ray is instanced by the source according to constraints input by the user. For a source satisfying ideal plane wave conditions, all rays have identical  $\mathbf{k}$ -vectors (monochromatic and parallel beam), phases and weights.

(ii) The ray interacts with the focusing optics via diffraction: effectively it takes a new direction, mimicking a Huygens’ wavelet (this step is described in detail later).

(iii) The ray interacts with the sample: after being propagated from the exit of the focusing optics to the sample section, the exit ray is computed according to equation (1). If, instead of its transfer function, the sample is specified in terms of refractive indices, the amplitude and phase of the object function are computed using  $\exp(-k\beta\tau)$  and  $\exp(jk\delta\tau)$ , in



**Figure 1** Illustration of a general setup for simulation of CDI in *McXtrace*. The beam, focus, sample and detector are components in the graphical user interface, each consisting of a set of lines of code modifying parameters of the rays in the simulation. The dice illustrate components where a Monte Carlo process takes place. We refer to  $\psi$  as the wavefield immediately after the sample, and to  $\Psi$  as the wavefield on the detector.

which  $k$  and  $\tau$  indicate the wavenumber and the thickness of the sample, respectively, and  $\delta$  and  $\beta$  are the real and imaginary parts of the refractive index ( $j$  denotes the imaginary unit).

(iv) The ray is propagated to a numerical detector, referred to in the following as Phase Sensitive Detector (*PSD*), where it keeps its phase information and is accumulated coherently on the target pixel. Namely,

$$PSD_{\phi}(m, n) = \sum_{i \in T_{m,n}} c_i p_i \exp(j\phi_i), \quad (6)$$

where  $T_{m,n}$  is the subset of traces hitting a pixel with indices  $m, n$ . Rays with a weight  $p$  and phase  $\phi$  accumulate, with a scale factor  $c$ , over the coherent sum  $PSD_{\phi}$ ;  $c$  is dependent on the distance between the pixel and ray origin. Interference among rays is produced at this step.

Operations (ii) and (iv) deserve more space for an accurate description because they contain the core model of diffraction that we use to simulate a realistic coherent X-ray experiment properly. For simplicity, in the following we consider that the sample in step (iii) does not induce any further modification to the rays, *i.e.* it is  $O(x, y) = 1$  in equation (1). In conventional ‘incoherent’ ray-tracing simulations, rays travel in a straight line and constructive interference between different rays is not accounted for, *i.e.* all rays hitting the area of the detector pixel are summed incoherently. Here, instead we are interested in adding rays coherently to predict the correct interference diffraction pattern. The diffraction of the rays in our model is described by Huygens’ principle, which phenomenologically explains diffraction by modelling a plane wave as a sum of elemental spherical waves. Although it is possible to split one ray into many and rely on the law of large numbers and Monte Carlo sampling to generate the wavefront on a semicircle, in practice the procedure is impractical due to the number of

rays required for sufficient sampling of the wavefront. Instead, we have developed an approach that considers diffraction from the detector perspective, in which all detector pixels collect diffraction from every ray. Furthermore, diffraction is collected only on a sub-pixel level. This may be done efficiently as it amounts to locally substituting a ray with an associated spherical wavefront, and sampling that wavefront at the sites of detector sub-pixels. An additional upshot of this reverse procedure is that it inherently avoids sampling points outside the detector.

Besides computational efficiency, the need to define an appropriate sub-pixel size stems from the following argument. Let  $\Psi(X, Y)$  be the wavefield in the detector plane. Then the recorded value  $PSD_{\phi}$  on the pixel with indices  $(m, n)$  can be written as

$$PSD_{\phi}(m, n) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Psi(X, Y) \Lambda_W(mW - X, nW - Y) dX dY \\ = (\Psi \star \Lambda_W)(mW, nW), \quad (7)$$

where  $\star$  denotes the convolution operator,  $(X, Y)$  are the coordinates in the detector plane and  $\Lambda_W$  is the 2D top-hat function of size  $W$ ,

$$\Lambda_W(X, Y) = \begin{cases} 1, & \text{if } |XY| < 1, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

Equation (7) states that the matrix of the recorded values is given by the convolution of the actual wavefield function with the pixel area, which causes an effective blurring of the wavefield at the detector plane that strongly depends on the value of  $W$  (see the example in Section S1 of the supporting information). One way to prevent blurring is to impose  $W \ll a$ , where  $a$  is the smallest feature in the object or illumination. In this way, the top-hat function can be approximated by a Dirac delta and its effect in the convolution becomes negligible. Another option to circumvent this problem is to translate all rays hitting a pixel to its centre ( $W = 0$ ). This is followed in the work by Mout *et al.* (2018), and has also been tested by us, yielding the same results. We emphasize that this requirement, and the way it is addressed in both cases, only relates to the way interference is evaluated in a ray-tracing context and should not be confused with the Nyquist sampling requirement of CDI.

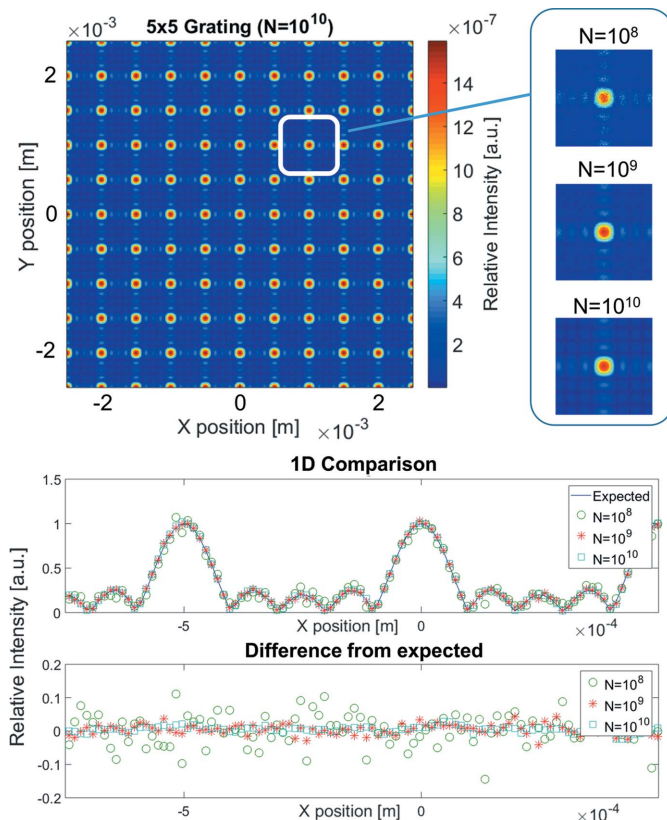
Concerning the propagation mentioned in step (iv), we point out that no paraxial approximation is made. Instead, the distance between wavelet origin and detector pixel is determined on a pixel-by-pixel and ray-by-ray basis, and thus the exact distance is taken into account. The factor  $c_i$  in equation (6) scales the weight accordingly, *i.e.*  $c_i \propto 1/r$ , where  $r$  is the distance between the ray origin and (sub)pixel. This term can be held constant in the far-field approximation and set to  $1/z$  for each ray, whereas it must be evaluated for each ray in the near-field. However, this is carried out at no extra cost in a ray-tracing context as the distances ( $r$ ) are computed regardless.

### 3. Results and discussion

We describe three study cases we have examined to show the correctness and the applicability of our approach: diffraction from a grating, full-field CDI from a complex object illuminated by a uniform beam with flat wavefront, and ptychography of a complex object with an experimental probe function (unpublished data). For all examples, the wavelength  $\lambda$  and the sub-pixel size  $W$  are set to 1 Å and the sample-to-detector distance  $z$  is set to 5 m. The diffraction patterns, when noted, are reported on a logarithmic scale, whereas they are always normalized with respect to the number of simulated rays  $N$  to allow a comparison. Details of the analysis of the more basic case of diffraction from a single slit are reported in Appendix A and Appendix B.

#### 3.1. Grating-like source

As a first test, we show the simulation of a  $5 \times 5$  grating in Fig. 2. The grating consists of a square grid of square apertures of size  $w = 1$  nm separated by a distance  $d = 1$  μm. Each aperture is an elementary source where a ray is emitted from a randomly chosen position within the region defined by the



**Figure 2** Simulated diffraction of a  $5 \times 5$  grating. The distance between the apertures is 1 μm, and an  $800 \times 800$  pixel detector of 0.5 cm size is used. Intensities are on a linear scale, and the ray weight is normalized to  $N$ , thus allowing the comparison for simulations with varying  $N$ . The three cases  $N = 10^8$ ,  $10^9$  and  $10^{10}$  are considered for comparison in the inset and the 1D plot. The diffraction patterns show slight blur and non-zero background for the first case, whereas they become clear in the others. The profile comparison (bottom plot) shows in detail the fluctuations for the three cases.

aperture. The simulated diffraction pattern of the grating presents all the characteristics predicted by theory. Namely:  $n - 2$  subsidiary maxima between two main peaks can be observed when the illumination function is a grid of  $n \times n$  apertures; the distance  $y$  between the main peaks is related to the distance between the apertures  $d$  by the following equation,

$$d \sin \theta = d \frac{y}{z} = l \lambda, \quad (9)$$

in which  $l$  is an integer and  $z$  is the sample-to-detector distance; the envelope observable in the diffraction pattern depends on the size of the apertures and follows the expected trend of

$$\text{sinc}\left(\frac{w}{z\lambda} X\right) \text{sinc}\left(\frac{w}{z\lambda} Y\right), \quad (10)$$

where  $w$  is the size of the aperture and  $X, Y$  are coordinates in the detector plane. In particular, we show in Fig. 2 that, for a sufficiently high number of simulated rays ( $N = 10^{10}$  in this case), the simulated trend follows closely the one predicted by theory,


$$|\Psi(X, 0)| = \left| \Psi_0 \sin\left(\frac{N\pi d X}{\lambda z}\right) \sin\left(\frac{\pi d X}{\lambda z}\right) \right|. \quad (11)$$

The same results are obtained by using a Fourier transform as in equation (2) with  $P(r) = 1$ . This test is an extension of the basic diffraction experiments of single and double slits already reported by Bergbäck Knudsen *et al.* (2013). It demonstrates that interference of multiple point-like sources distributed on a grid placed at a given section is properly modelled and recorded with our approach and, therefore, represents a test-bench simulation. Moreover, it serves us to roughly estimate the number of rays needed for simulation of imaging cases, which we have learned is dependent on the density of rays at the source and at the detector. However, this is difficult to extend to more general cases. We find in this case that a choice of  $N = 10^8$  leads to poor cancellation of the zeros and fluctuations in the profile with standard deviation  $\sigma = 0.054$ , whereas  $N = 10^9$  already gives satisfactory results ( $\sigma = 0.035$ ), and  $N = 10^{10}$  can be regarded as an ideal noiseless case ( $\sigma = 0.033$ ). For an estimation of values and trend of the error as a function of  $N$ , see Appendix B, where we evaluate it in a simpler case.

This also allows us to assert that, despite the high number of phase evolutions, the machine numerical precision does not seem to interfere with the phasor accumulation, as we are eventually able to obtain accurate cancellations of zeros and no systematic errors on the profile. This we have tested and observed for several reasonable detector distances and sizes.

#### 3.2. Full-field CDI

In a full-field CDI (Miao *et al.*, 1999) experiment, the entire sample is illuminated by a coherent beam and the diffraction pattern is recorded in the Fraunhofer regime. As the information about the structure of the object is encoded in the phase and only intensity measurements are available, a phase



**Figure 3**

Full-field CDI. Columns, left to right: wavefield immediately after the object; far field, recorded on a 2 cm-size detector (second column) and zoom on the inner portion (third column); inverse Fourier transform of the diffraction pattern in column 2. Upper and lower figures display amplitude and phase of each recorded wavefield, respectively. Wavelength and sample-to-detector distances are 1 Å and 5 m, respectively. A uniform amplitude and constant phase have been selected for the illuminating beam to allow an immediate visual comparison of the results.

retrieval algorithm is required to recover the initial object. This algorithm exploits the setup configuration by imposing the support constraint, which provides the *a priori* knowledge that the object is limited in size and fully illuminated by the beam as well as the exact knowledge of the diffracted intensity data.

To illustrate the outcome of our tests we refer to Fig. 3. In this simulation, we assume the validity of the multiplicative assumption to verify the second requirement of CDI techniques — the correct propagation in the far-field zone. For this reason, we used a two-dimensional mask of complex refractive indices, with amplitude and phase mimicking popular test images (baboon and man) to demonstrate how the information of attenuation and phase shift of a thin sample is propagated by the simulator into reciprocal space. In this case, the setup is that of a full-field CDI experiment. The images are input with  $480 \times 480$  pixel square masks of  $8 \mu\text{m}$  lateral size; the object is fully covered by a slightly larger beam, with a uniform intensity profile and constant phase to allow for a direct visual comparison.<sup>1</sup> A  $400 \times 400$  pixel detector with 2 cm lateral size is set at a distance of 5 m from the sample. The object and detector sizes determine the sampling requirement for CDI (Spence *et al.*, 2004), which yields in this case  $\sim 30 \mu\text{m}$  as the minimum pixel size, lower than the  $50 \mu\text{m}$  pixel size of this detector. Therefore, this configuration produces the undersampled diffraction patterns represented

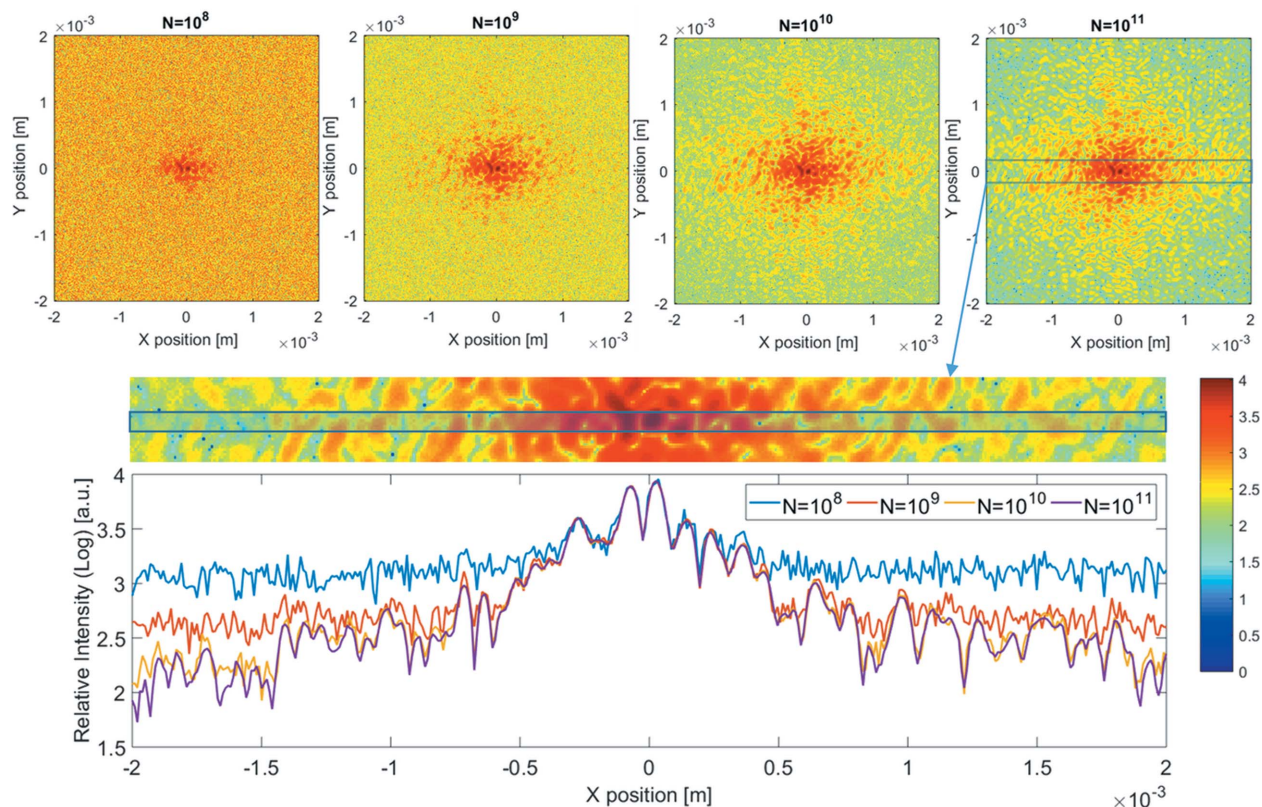
<sup>1</sup> The constant phase is not in general a requirement of a CDI experiment. Conceptually, it is not a useful addition to use a realistic probe in this case, as CDI phase retrieval algorithms do not typically retrieve the probe function (but only the wavefield  $\psi$ ) and because we are using a Fourier transform to invert the pattern.

in the second column of Fig. 3 (an example of an oversampled simulation follows below). Given the use of a constant phase beam, the validity of our simulation can be verified by simple inverse Fourier transform of the simulated patterns. The result, depicted in the last column of the same figure, with a reconstructed pixel size of 25 nm, proves that the information about the object has been properly encoded. The noise introduced by the Monte Carlo process produces a mixing of amplitude and phase of the retrieved object which can be observed for a lower number of traced rays.

Our simulation tool allows us to assess how the Monte Carlo approach affects the quality of the patterns. To illustrate this, we use a detector with a pixel size of  $10 \mu\text{m}$ , which is hence oversampling and small enough to resolve the speckles produced by the largest dimensions of the sample. The speckle size is expected to be  $62.5 \mu\text{m}$  in this case. In Fig. 4 we illustrate how speckle visibility is affected by the number of simulated rays. The higher the number of simulated rays the better the speckles are resolved, particularly the least intense ones farthest out in reciprocal space. The effect of enhanced speckle visibility with increasing  $N$  on a retrieved image is reported in Section S2.2 of the supporting information, together with an estimation of the noise.

### 3.3. Ptychography

In ptychography (Rodenburg & Faulkner, 2004; Pfeiffer, 2018), the sample is illuminated by the beam one portion at a time and the object size is not limited by the size of the beam. A scan pattern is defined to illuminate the object at different



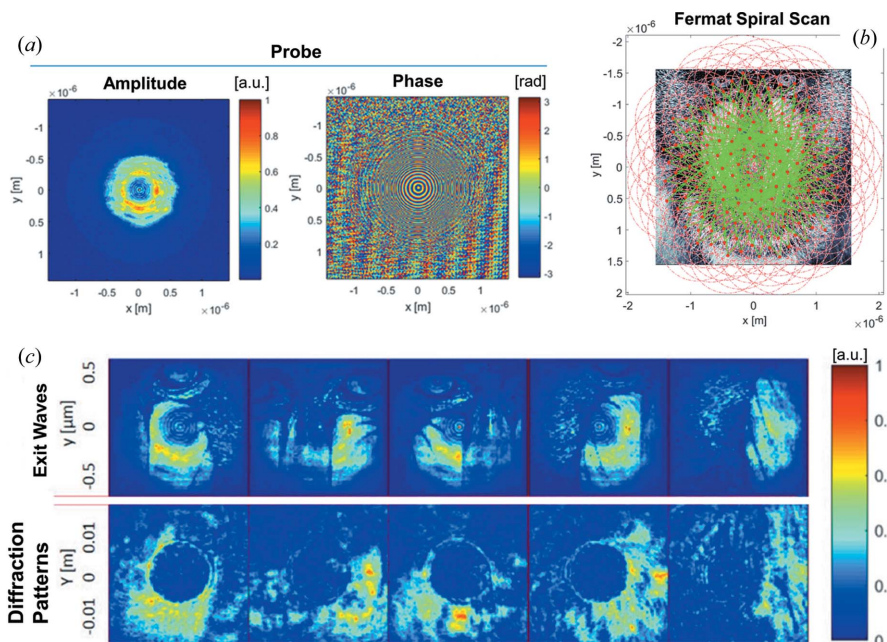
**Figure 4** Full-field CDI. Speckle resolution as a function of simulated rays. Three test cases are shown for a 4 mm-sized detector with 400 pixels set at a 5 m distance from the sample. All plots are on a logarithmic scale, with the same color bar. The profile is plotted along the  $X$  direction and is averaged on a length of seven pixels along the  $Y$  direction. Averaging helps to distinguish between speckles and fluctuations on a smaller scale due to noise. The speckles at higher spatial frequency are better resolved for increasing numbers of simulated rays  $N$ , whereas they are buried into background noise with fewer rays.

positions in overlapping areas and a set of diffraction patterns is collected. With respect to the definition of the setup, a sufficient amount of overlap between the illuminated portions of the sample must be ensured, and any symmetry of the scanning pattern that may cause raster-grid pathology (Thibault *et al.*, 2008) must be broken.

Here we discuss the implementation of our method for a ptychography scan applied to the same object of Fig. 3. In this case, we use the probe produced by Fresnel zone plate focusing optics and obtained with phase retrieval approaches from experimental data obtained in a previous experiment [Fig. 5(a)]. These data are used to describe the complex mask component that is inserted right before the sample in our instrument. The same probe could be produced by our simulation package by using a complete description of the focusing optics in Fig. 1. We prefer to use a real experimentally determined probe for our simulation. Some projections of the ptychography dataset generated in this test are reported in Fig. 5(c). The probe is shifted over the sample according to the Fermat spiral scheme depicted in Fig. 5(b), with steps sufficiently small to ensure 80% probe overlap [evaluated as in Huang *et al.* (2014)] and complete, uniform coverage (except for the corners). In this particular case, 160 steps are sufficient. We prefer the Fermat spiral scheme over a Cartesian grid or other geometries as it efficiently mitigates reconstruction artefacts. In this simulation, the specimen is  $\sim 3.1 \mu\text{m} \times 3.1 \mu\text{m}$

in size and its transfer function input is a  $480 \times 480$  matrix. A  $200 \times 200$  pixel detector of 3.5 cm lateral size [similar to Pilatus 100 K (Bech *et al.*, 2008)] is used. A total of  $10^{10}$  rays are used for each scan position. Each step runs for 12 min on a 16-core node. The phase and amplitude of the object are reconstructed via the inversion of the simulated diffraction patterns using cSAXS Matlab code. The result of the ptychography phase retrieval is shown in Fig. 6 and was obtained by running 300 iterations (with as many probe and object updates) of the algorithm ePIE (Maiden & Rodenburg, 2009). A good guess of the probe, its real part in our case, had to be provided to ensure convergence. The level of noise introduced by the simulation is acceptable for the pattern to return a fair reconstruction of the specimen and the probe (refer to Section S2 of the supporting information for a resolution assessment).

The reconstruction quality is comparable with that of the full-field case (see Section S2 of the supporting information for a resolution assessment). In this case, the single diffraction patterns are obtained with one-tenth of the rays used in Fig. 3. However, because of the multiple illuminations, a significantly larger total number of rays impinge on the sample. This results in estimates of the average signal-to-noise ratio being on the same order of magnitude (5.4 in CDI versus 3.3 in ptychography). We find that decreasing  $N$  by one order of magnitude in the ptychography case (*i.e.* going from  $10^{10}$  to  $10^9$ ) leads to



**Figure 5** Ptychography: (a) amplitude and phase of the probe used for illumination,  $500 \times 500$  pixels; (b) illumination scheme of Fermat spiral scan; (c) some results from the simulated ptychography dataset. The exit waves are in the top row, with each relative diffraction pattern below. The exit waves are recorded on a  $\sim 1 \mu\text{m}$ -wide detector covering most of the probe. Aside from high-frequency features, some near-field features are still visible in the diffraction patterns.

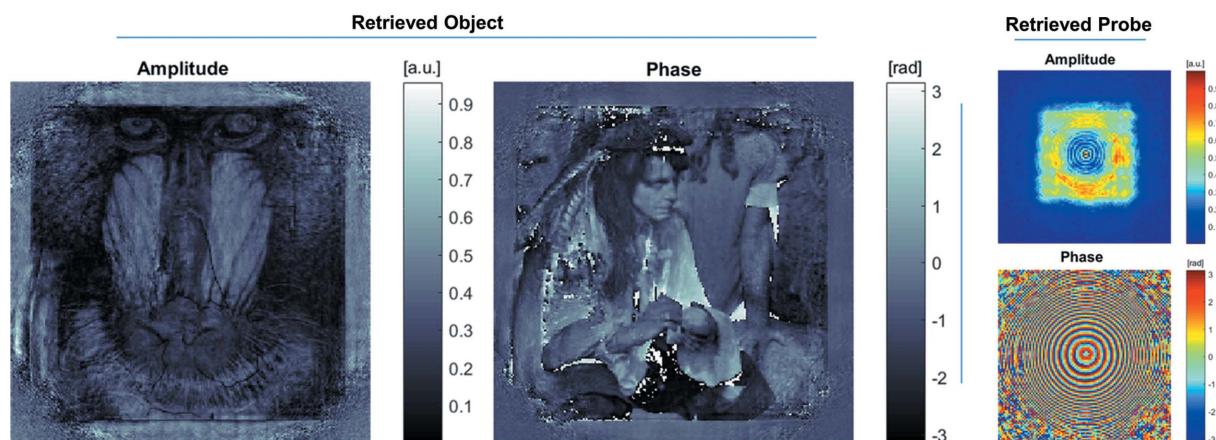
failure of the reconstruction, although the diffraction patterns do not suffer considerable loss in quality. As in Fig. 4, the difference between the cases  $10^{10}$  and  $10^{11}$  does not appear substantial, but can nonetheless mark the threshold for convergence in phase retrieval.

#### 4. Outlook

Relevant features of the method reported by Mout *et al.* (2018) can be applied to our implementation for further testing of more complex and unsolved cases. Specifically, the

vectorial expression for the accumulation of rays at each step allowing for a cascaded system indicates a viable path for the implementation of thick samples through cascaded objects. Given the similarity of our framework, we presume that this procedure is applicable to our case, at least with respect to the forward imaging mode. Concerning the general 3D case, we note that CDI techniques can already be simulated on a 3D object with the support of the *ASTRA* toolbox (van Aarle *et al.*, 2015). The line integration of the refractive indices represents a computationally expensive operation that could in principle be performed in *McXtrace*, but not as efficiently as in *ASTRA*. If the sample is thin enough, the internal scattering is negligible and it is recommended to use *ASTRA*, otherwise a 3D sample can be represented by a stack of masks along the ray direction, using a multi-slice approach as in the work by Maiden *et al.* (2012). Simulation of 3D volumes would allow a comparison with experimental data and would provide definitive evidence of the validity of this tool.

Another interesting case is the grazing incidence setup, where the theoretical analytical expression has proved significantly distant from actual measurement in some cases and seems only applicable to describe the height function of truncated crystals (Zhu *et al.*, 2015). Imperfect temporal coherence is also easily simulated by assigning different wavelengths to the rays. In particular, it is possible to assign to the rays energy values distributed on a Gaussian centered on the main energy, with a deviation. In principle, non-paraxial cases and near-field cases can be tested but have not been explored in this work for imaging (Appendix A reports a



**Figure 6** Left to right: amplitude and phase of the retrieved object; plot of the retrieved probe, amplitude (top) and phase (bottom). The phase retrieval is performed with 300 iterations of ePIE after providing a good guess of the probe.



comparison with the Fraunhofer propagator in the case of wider angles with diffraction from a single slit). The model does not in fact assume paraxial propagation, and the spherical factor in the amplitude is accounted for each ray rather than set to the detector distance as in the far-field hypothesis.

Mout *et al.* (2018) also used a general expression for the weight of the rays, which could help to optimize the number of traced rays to achieve an adequate signal-to-noise ratio, reducing the computational cost of this method. We show in a basic example in Appendix B that not all portions of a detector are affected equally by Monte Carlo noise, which points to the study of specific, dynamical sampling strategies. The other main route to decrease computational cost is through parallelization. Formulated in this framework, any setup is easily parallelized, and GPU parallelized code can easily reduce processing time. Furthermore, the number of rays modern hardware can afford to simulate has increased exponentially over the years, making the cost of ray tracing for imaging within reach and with good prospects. While we routinely run simulations of billions of rays, the first *SHADOW* publication quotes 5000 as the maximum number of traceable rays, and as a result regards simulation of imaging experiments as ‘unsuitable’ (Lai & Cerrina, 1986). If one is only interested in a single portion of the reciprocal space, this method can already have a reasonable cost.

We finally note that future work should generally aim to: assess the signal-to-noise ratio of simulated diffraction patterns of a volume of refractive indices, representing realistic volumetric samples; simulate CDI experiments with complex sample-interaction mechanisms (Vartanyants *et al.*, 2007), and in combination with tomography and focusing optics. With respect to the simulated noise, it is of interest to characterize it statistically and assess whether it is Poisson-like as in real experiments and to relate the number of simulated rays to the simulated signal-to-noise ratio.

## 5. Conclusions

We reported a scheme for simulating CDI techniques in a ray-tracing framework. We have implemented it in the open-source simulator *McXtrace* and demonstrated its function for a novel problem in the configurations of full-field CDI andptychography in the far field and using a forward direction geometry. Possibly the main advantage of this method is to assist non-specialists in the simulation of coherent imaging experiments. All user-defined distances and parameters are real space quantities and do not need conversion to reciprocal coordinates, and allow for deviation from ideal test situations (translation or rotation of a component).

We acknowledge that like all Monte Carlo methods this one does not stand out for elegance (as it is not based on a physical insight of a setup) or efficiency and is hardly recommendable for standard cases where an analytical solution is available. However, this simple approach bears no attached hypothesis and is possibly extendable to less standard, unsolved cases; the scheme is suitable for GPU parallelization and therefore has wide room for improvement with regard to efficiency. The

rationale behind our endeavour is to eventually offload the complexity of a special setup to brute computational force. The main contribution of this paper lies in the definition of adequate constraints for simulation of diffraction, such as the sampling requirements and an estimate of the number of rays to trace to achieve a satisfactory signal-to-noise ratio. We believe that these are potentially applicable to the simulation of other untested CDI variants, although more tests are still to be made. In line with the open-source spirit of *McXtrace*, we trust that current and future users can exploit and help to develop the lines of inquiry mentioned in the outlook.

## 6. Related literature

The following references, not cited in the main body of the paper, have been cited in the supporting information: Nieuwenhuizen *et al.* (2013); Van Heel & Schatz (2005); Vila-Comamala *et al.* (2011); Huang *et al.* (2009).

## APPENDIX A Comparison with Fraunhofer propagator

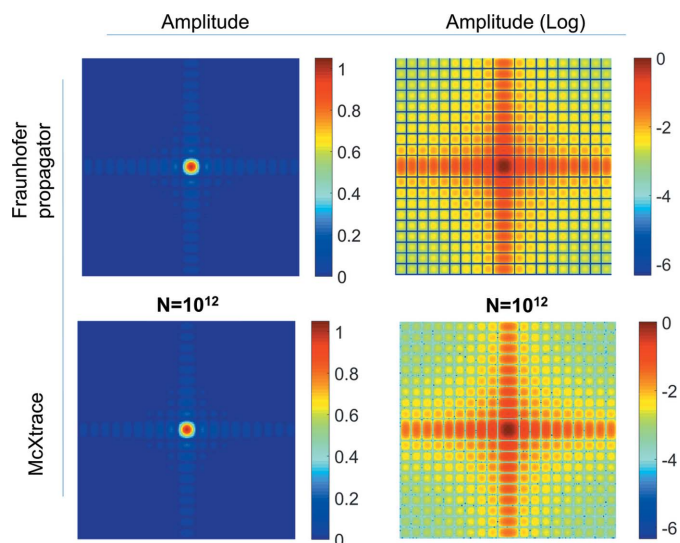
In this section, we compare the ray-tracing simulation of diffraction from a single slit with the solution obtained with a numerical implementation of the Fraunhofer propagator,

$$\Psi(X, Y) = -\frac{j \exp(jkz)}{\sqrt{N_x N_y}} \exp\left(jk \frac{X^2 + Y^2}{z}\right) F(\psi), \quad (12)$$


in which  $N_x$  and  $N_y$  are the number of pixels in the horizontal and vertical directions. The single slit has 100 nm width; the detector has  $200 \times 200$  square pixels of  $400 \mu\text{m}$  size, corresponding to a detector size of  $8 \text{ cm} \times 8 \text{ cm}$  at a distance  $z = 4 \text{ m}$  from the slits. These parameters produce a real space pixel size of 5 nm, which is relatively far in reciprocal space. The comparison of the two cases is illustrated in Fig. 7. An overall agreement is found at low frequencies, whereas a slightly faster decay of intensity is observed for the Fraunhofer propagator solution, and a worse cancellation of the zeros is observed for the ray tracing due to the stochastic way it is achieved. The faster decay is highlighted in Fig. 8. A third comparison is made with the sinc pattern derived from geometrical optics, here recalled,

$$\Psi(X, Y) = \text{sinc}\left(\frac{wX}{\lambda z}\right) \text{sinc}\left(\frac{wY}{\lambda z}\right), \quad (13)$$


where  $w$  denotes the slit width,  $\lambda$  the wavelength and  $z$  the detector distance. The ray-tracing simulation shows a closer resemblance with this last function and shows significant relative deviation from the Fraunhofer approach at high angles. The decay is similar when other line profiles are evaluated. It is interesting in this case to define a cut-off in detector space where the two approaches match. We find that for  $|X| < 2 \text{ cm}$  the relative difference between them is below 10%.



**Figure 7**  
Comparison of diffraction patterns from the Fraunhofer propagator and ray-tracing simulations from a single slit.  $N = 10^{12}$  rays are traced in the *McXtrace* simulation.



**Figure 8**  
Line-cut of a single-slit diffraction pattern taken in the centre of the diffraction at  $y = 0$ . The ray-tracing simulation is compared with the solution obtained with the Fraunhofer propagator and with the sinc function from geometrical optics. (a) Profile across the entire detector, showing an overall agreement among the three solutions; (b) the ray-tracing simulation matches the sinc pattern but deviates significantly from the Fraunhofer propagator solution at high frequencies; (c) plot of the difference between the Fraunhofer solution and ray tracing for a single line-cut. The difference is below 10% of the Fraunhofer solution for  $|X| < 2$  cm.



**Figure 9**  
Diffraction patterns of a single slit for an increasing number of simulated rays ( $N = 10^8$ – $10^{11}$ ). The same colour axis and scale refers to the four log scale plots. The weight of the rays is normalized by  $N$  so that the sum of all weights equals unity.

## APPENDIX B

### Characterization of noise in Monte Carlo ray tracing

We evaluate noise in the above-mentioned case of the single slit where satisfactory results are obtained even with a relatively low number of rays and, additionally, an analytical solution exists. For the analysis, we focus solely on the simulated amplitude of the diffraction patterns. Error is evaluated as the difference from the true solution. For that we refer to two options: the simulation of the highest number of rays and the sinc pattern from equation (13). We test numbers of rays ranging from  $10^8$  to  $10^{11}$ , whereas the best guess of a solution has  $10^{12}$  rays.

Figs. 9 and 10 show the strong signal–noise dependence (*i.e.* the noise pattern reflects the single slit patterns in Fig. 9) with the largest relative error for low signal, which is also a characteristic of Poisson counting statistics and points to a useful modelling of noise. Mean and variance of this noise do not match in this case in general, as it should be for a Poisson-like noise, but they obviously match for the event (the number of rays hitting a pixel) numbers on the detector [Fig. 11(a)]. This

second fact implies that when the flux is set to match the number of rays, and there is fully constructive interference, the counts are Poisson distributed. Aside from this case, the variance is not observed to match the mean counts for all pixels.

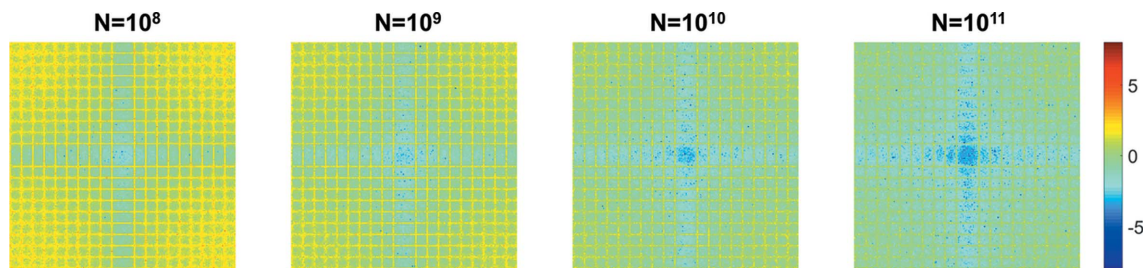
The norm of the error decreases as  $1/\sqrt{N}$  [Fig. 11(b)], as expected for a Monte Carlo method. The same conclusions apply whether the error is evaluated as the difference from the best guess with  $N = 10^{12}$ , or as the difference from the sinc pattern.

### Acknowledgements

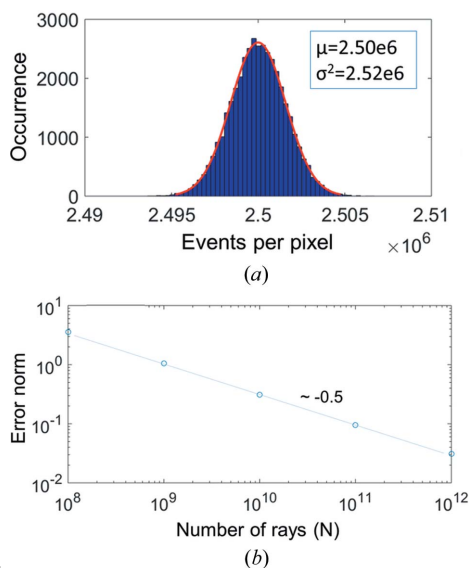
We gratefully acknowledge P. S. Jørgensen, A. F. Pedersen, H. Simons and P. K. Willendrup for fruitful scientific discussions.

### Funding information

Data processing (Fourier ring correlation and ptychography reconstructions)



**Figure 10** Relative error patterns. The same colour axis and scale refers to the four log scale plots. The simulated patterns from Fig. 9 are subtracted from and normalized by the simulation with the largest number of traced rays ( $N = 10^{12}$ ). The total error decreases with increasing  $N$  and is higher where the signal is lower.



**Figure 11** (a) Event statistics, which follow a Poisson distribution; the reported case refers to  $N = 10^{11}$ . (b) Error trend on logarithmic axes. The total error is evaluated as L2-norm of the difference with the sinc pattern. The points fit on a straight line with  $-0.5$  slope.

was carried out using the cSAXS ptychography MATLAB package developed by the Science IT and the coherent X-ray scattering (CXS) groups, Paul Scherrer Institut, Switzerland. This study was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (SEEWHI Consolidator grant No. ERC-2015-CoG-681881) and by the Ministry of Higher Education and Science (DANSCATT grant No. 7055-00007B).

**References**

Aarle, W. van, Palenstijn, W. J., De Beenhouwer, J., Altantzis, T., Bals, S., Batenburg, K. J. & Sijbers, J. (2015). *Ultramicroscopy*, **157**, 35–47.  
 Andreas, B., Mana, G. & Palmisano, C. (2015). *J. Opt. Soc. Am. A*, **32**, 1403.  
 Bahrtdt, J. (2007). *Phys. Rev. ST Accel. Beams*, **10**, 060701.  
 Bech, M., Bunk, O., David, C., Kraft, P., Brönnimann, C., Eikenberry, E. F. & Pfeiffer, F. (2008). *Appl. Radiat. Isot.* **66**, 474–478.  
 Bergbäck Knudsen, E., Prodi, A., Baltser, J., Thomsen, M., Kjær Willendrup, P., Sanchez del Rio, M., Ferrero, C., Farhi, E., Haldrup, K., Vickery, A., Feidenhans'l, R., Mortensen, K., Meedom Nielsen, M., Friis Poulsen, H., Schmidt, S. & Lefmann, K. (2013). *J. Appl. Cryst.* **46**, 679–696.

Chang, H., Enfedaque, P., Zhang, J., Reinhardt, J., Enders, B., Yu, Y.-S., Shapiro, D., Schroer, C. G., Zeng, T. & Marchesini, S. (2019). *Opt. Express*, **27**, 10395–10418.  
 Chapman, H. N. & Nugent, K. A. (2010). *Nat. Photon.* **4**, 833–839.  
 Chubar, O., Fluerasu, A., Berman, L., Kaznatcheev, K. & Wiegart, L. (2013). *J. Phys. Conf. Ser.* **425**, 162001.  
 Cipiccia, S., Vittoria, F. A., Weikum, M., Olivo, A., Jaroszynski, D. A., Dino, A. & Jaroszynski, D. A. (2014). *Opt. Express*, **22**, 2395–2398.  
 Colton, D. & Kress, R. (1992). *Inverse Acoustic and Electromagnetic Scattering Theory*, p. 121; p. 289. Springer: Berlin.  
 Dierolf, M., Menzel, A., Thibault, P., Schneider, P., Kewish, C. M., Wepf, R., Bunk, O. & Pfeiffer, F. (2010). *Nature*, **467**, 436–439.  
 Dietze, S. H. & Shpyrko, O. G. (2015). *J. Synchrotron Rad.* **22**, 1498–1508.  
 Godard, P., Allain, M., Chamard, V. & Rodenburg, J. (2012). *Opt. Express*, **20**, 25914–25934.  
 Hettel, R. (2014). *J. Synchrotron Rad.* **21**, 843–855.  
 Holler, M., Guizar-Sicairos, M., Tsai, E. H. R., Dinapoli, R., Müller, E., Bunk, O., Raabe, J. & Aeppli, G. (2017). *Nature*, **543**, 402–406.  
 Howells, M. R., Beetz, T., Chapman, H. N., Cui, C., Holton, J. M., Jacobsen, C. J., Kirz, J., Lima, E., Marchesini, S., Miao, H., Sayre, D., Shapiro, D. A., Spence, J. C. H. & Starodub, D. (2009). *J. Electron Spectrosc. Relat. Phenom.* **170**, 4–12.  
 Huang, X., Miao, H., Steinbrener, J., Nelson, J., Shapiro, D., Stewart, A., Turner, J. & Jacobsen, C. (2009). *Opt. Express*, **17**, 13541–13553.  
 Huang, X., Yan, H., Harder, R., Hwu, Y., Robinson, I. K. & Chu, Y. S. (2014). *Opt. Express*, **22**, 12634–12644.  
 Klementiev, K. & Chernikov, R. (2014). *Proc. SPIE*, **9209**, 92090A.  
 Kumar, P. B. S. & Ranganath, G. S. (1991). *Pramana J. Phys.* **37**, 457–488.  
 Lai, B. & Cerrina, F. (1986). *Nucl. Instrum. Methods Phys. Res. A*, **246**, 337–341.  
 Mahan, J. R., Vinh, N. Q., Ho, V. X. & Munir, N. B. (2018). *Appl. Opt.* **57**, D56–D62.  
 Maiden, A. M., Humphry, M. J. & Rodenburg, J. M. (2012). *J. Opt. Soc. Am. A*, **29**, 1606.  
 Maiden, A. M. & Rodenburg, J. M. (2009). *Ultramicroscopy*, **109**, 1256–1262.  
 Marchesini, S., Chapman, H. N., Hau-Riege, S. P., London, R. A., Szoke, A., He, H., Howells, M. R., Padmore, H., Rosen, R., Spence, J. C. H. & Weierstall, U. (2003). *Opt. Express*, **11**, 2344.  
 Mastropietro, F., Carbone, D., Diaz, A., Eymery, J., Sentenac, A., Metzger, T. H., Chamard, V. & Favre-Nicolin, V. (2011). *Opt. Express*, **19**, 19223–19232.  
 Mastropietro, F., Eymery, J., Carbone, G., Baudot, S., Andrieu, F. & Favre-Nicolin, V. (2013). *Phys. Rev. Lett.* **111**, 215502.  
 Miao, J., Charalambous, P., Kirz, J. & Sayre, D. (1999). *Nature*, **400**, 342–344.  
 Miao, J., Ishikawa, T., Robinson, I. K. & Murnane, M. M. (2015). *Science*, **348**, 530–535.  
 Mout, M., Flesch, A., Wick, M., Bociort, F., Petschulat, J. & Urbach, P. (2018). *J. Opt. Soc. Am. A*, **35**, 1356.

- Mout, M., Wick, M., Bociort, F., Petschulat, J. & Urbach, P. (2016). *Appl. Opt.* **55**, 3847–3853.
- Nieuwenhuizen, R. P. J., Lidke, K. A., Bates, M., Puig, D. L., Grünwald, D., Stallinga, S. & Rieger, B. (2013). *Nat. Methods*, **10**, 557–562.
- Noh, D. Y., Kim, C., Kim, Y. & Song, C. (2016). *J. Phys. Condens. Matter*, **28**, 493001.
- Pfeiffer, F. (2018). *Nat. Photon.* **12**, 9–17.
- Prodi, A., Knudsen, E., Willendrup, P., Schmitt, S., Ferrero, C., Feidenhans'l, R. & Lefmann, K. (2011). *Proc. SPIE*, **8141**, 814108.
- Ramos, T., Jørgensen, J. S. & Andreasen, J. W. (2017). *J. Opt. Soc. Am. A*, **34**, 1830–1843.
- Rebuffi, L. & Sánchez del Río, M. (2016). *J. Synchrotron Rad.* **23**, 1357–1367.
- Sanchez del Rio, M. & Rebuffi, L. (2019). *AIP Conf. Proc.* **2054**, 060081.
- Rodenburg, J. M. (2008). *Advances in Imaging and Electron Physics*, Vol 150, pp. 87–184. San Diego: Academic Press.
- Rodenburg, J. M. & Faulkner, H. M. L. (2004). *Appl. Phys. Lett.* **85**, 4795–4797.
- Schäfers, F. (2008). *Modern Developments in X-ray and Neutron Optics*, edited by A. Erko, M. Idir, T. Krist & A. G. Michette, pp. 9–41. Berlin, Heidelberg: Springer.
- Shapiro, D. A., Yu, Y.-S., Tyliczszak, T., Cabana, J., Celestre, R., Chao, W., Kaznatcheev, K., Kilcoyne, A. L. D., Maia, F., Marchesini, S., Meng, Y. S., Warwick, T., Yang, L. L. & Padmore, H. A. (2014). *Nat. Photon.* **8**, 765–769.
- Spence, J. C. H., Weierstall, U. & Howells, M. (2004). *Ultramicroscopy*, **101**, 149–152.
- Takahashi, Y., Zettsu, N., Nishino, Y., Tsutsumi, R., Matsubara, E., Ishikawa, T. & Yamauchi, K. (2010). *Nano Lett.* **10**, 1922–1926.
- Thibault, P., Dierolf, M., Menzel, A., Bunk, O., David, C. & Pfeiffer, F. (2008). *Science*, **321**, 379–382.
- Thibault, P. & Guizar-Sicairos, M. (2012). *New J. Phys.* **14**, 063004.
- Van Heel, M. & Schatz, M. (2005). *J. Struct. Biol.* **151**, 250–262.
- Vartanyants, I. A., Grigoriev, D. & Zozulya, A. V. (2007). *Thin Solid Films*, **515**, 5546–5552.
- Veen, F. & Pfeiffer, F. (2004). *J. Phys. Condens. Matter*, **16**, 5003–5030.
- Vila-Comamala, J., Diaz, A., Guizar-Sicairos, M., Manton, A., Kewish, C. M., Menzel, A., Bunk, O. & David, C. (2011). *Opt. Express*, **19**, 21333–21344.
- Zhu, C., Harder, R., Diaz, A., Komanicky, V., Barbour, A., Xu, R., Huang, X., Liu, Y., Pierce, M. S., Menzel, A. & You, H. (2015). *Appl. Phys. Lett.* **106**, 101604.