

syris: a flexible and efficient framework for X-ray imaging experiments simulation

Tomáš Faragó,^{a,b*} Petr Mikulík,^{c,d} Alexey Ershov,^{a,b} Matthias Vogelgesang,^e
Daniel Hänschke^{a,b} and Tilo Baumbach^{a,b}

Received 3 February 2017

Accepted 23 August 2017

Edited by Y. Amemiya, University of Tokyo, Japan

Keywords: simulation; high-speed imaging; parallelization; free-space propagation; coherence; X-ray imaging; synchrotron radiation.


^aInstitute for Photon Science and Synchrotron Radiation, Karlsruhe Institute of Technology (KIT), Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany, ^bLaboratory for Applications of Synchrotron Radiation, Karlsruhe Institute of Technology (KIT), D-76131 Karlsruhe, Germany, ^cDepartment of Condensed Matter Physics, Faculty of Science, Masaryk University (MU), Kotlářská 2, 611 37 Brno, Czech Republic, ^dCEITEC – Central European Institute of Technology, Masaryk University (MU), Kamenice 753/5, 625 00 Brno, Czech Republic, and ^eInstitute for Data Processing and Electronics, Karlsruhe Institute of Technology (KIT), Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany. *Correspondence e-mail: tomas.farago@kit.edu

An open-source framework for conducting a broad range of virtual X-ray imaging experiments, *syris*, is presented. The simulated wavefield created by a source propagates through an arbitrary number of objects until it reaches a detector. The objects in the light path and the source are time-dependent, which enables simulations of dynamic experiments, *e.g.* four-dimensional time-resolved tomography and laminography. The high-level interface of *syris* is written in Python and its modularity makes the framework very flexible. The computationally demanding parts behind this interface are implemented in OpenCL, which enables fast calculations on modern graphics processing units. The combination of flexibility and speed opens new possibilities for studying novel imaging methods and systematic search of optimal combinations of measurement conditions and data processing parameters. This can help to increase the success rates and efficiency of valuable synchrotron beam time. To demonstrate the capabilities of the framework, various experiments have been simulated and compared with real data. To show the use case of measurement and data processing parameter optimization based on simulation, a virtual counterpart of a high-speed radiography experiment was created and the simulated data were used to select a suitable motion estimation algorithm; one of its parameters was optimized in order to achieve the best motion estimation accuracy when applied on the real data. *syris* was also used to simulate tomographic data sets under various imaging conditions which impact the tomographic reconstruction accuracy, and it is shown how the accuracy may guide the selection of imaging conditions for particular use cases.

1. Introduction

The increasing complexity of novel X-ray imaging methods and data processing algorithms makes it more and more difficult to understand the dependency of the data analysis accuracy on the vast amount of experimental and data processing parameters. However, the optimal combination of these parameters is crucial and can substantially increase experiment success rates and efficiency, particularly when available already before the actual measurement.

State-of-the-art analysis of X-ray image data involves a number of mutually dependent steps, *e.g.* advanced normalization techniques (Van Nieuwenhove *et al.*, 2015), phase retrieval (Paganin *et al.*, 2002; Moosmann *et al.*, 2010), three-dimensional (3D) volume reconstruction (Thompson *et al.*, 1984) and segmentation (Pham *et al.*, 2000), which together form an image processing *pipeline*. The performance of such



a pipeline depends on the optimal combination of experimental conditions, algorithms inside it and their parameters. For example, high-speed experiments tend to produce noisy images with motion blur and substantial sample shape change (see Fig. 1). If we want to analyze the motion in such images, we need to select an exposure time such that the combination of the noise level and the amount of motion blur still allows a motion estimation algorithm to follow our sample. Moreover, the addition of a de-noising algorithm or another image enhancement technique into the pipeline may dramatically change the accuracy of the data analysis. Another example may be the lack of signal in tomographic projection regions where the sample is too thick to be penetrated by X-rays, known from limited angle tomography (Davison, 1983). Suitable imaging methods [e.g. laminography (Helfen *et al.*, 2005)] and algorithms had to be found so that we could obtain sufficiently accurate sample reconstruction even in the case of a partially missing signal.

The design of new X-ray instrumentation, investigation of novel imaging approaches and benchmarking of data processing pipelines can strongly benefit from prior simulations. A number of specialized simulation programs for specific parts of the image formation process exist, e.g. the beamline optics ray-tracing simulators *SHADOW* (Sanchez del Rio *et al.*, 2011) and *xrt* (Klementiev & Chernikov, 2014). *McXtrace* (Bergbäck Knudsen *et al.*, 2013) is a ray-tracing package which enables simulations of X-ray imaging experiments, like tomography. Propagation-based calculation of synchrotron radiation from various sources was implemented in *SRW* (Chubar & Elleaume, 1998). Propagation-based tools which can be employed to simulate X-ray imaging experiments are, for example, *SRCLsim*, which was applied to laminography (Helfen *et al.*, 2005) and laboratory-based tomography (Zápražný *et al.*, 2013), and *pyXSF* (Malecki *et al.*, 2012) for grating interferometry simulations.

Complementary to these approaches, *syris* aims for simulation of a broad range of X-ray imaging experiments, from time-resolved radiography (Myagotin *et al.*, 2012; Zabler *et al.*, 2013) to four-dimensional measurements (Moosmann *et al.*, 2013; dos Santos Rolo *et al.*, 2014; Walker *et al.*, 2014). This is possible by modeling the complete image formation process from X-ray source to detection and by including dynamics of the optical elements in the light path.

syris can produce the *ground truth* (e.g. a 3D sample structure, motion vectors, *etc.*) and the correspondent virtual data sets with various experimental parameters. The combination of the ground truth and simulated data sets can be used to benchmark complex data analysis pipelines, categorize them and subsequently select the one suitable for a specific use case. Moreover, such an integral approach enables auto-

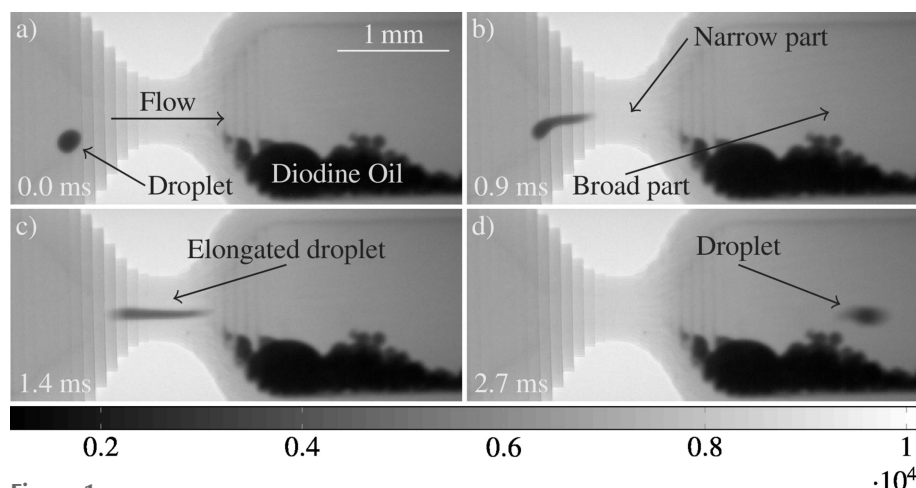


Figure 1 X-ray projections (grayscale in detector counts) of a capillary with varying diameter and an iodine droplet traveling through it. The droplet enters the narrow capillary part in (a), thus it accelerates and changes its shape in (b). The shape is completely different in (c) and as the iodine enters the capillary part with broader diameter in (d) it slows down and forms a droplet again.

matic optimization of complete experiments including sample, imaging and data processing parameters prior to the actual measurement.

In this paper, we will first define the scope of *syris* in §1.1. In §2 we will describe the image formation principles considered in the current framework implementation. Based on that, we will explain the design, implementation and parallelization of *syris* in §3. To show that the physical models and their implementation can produce realistic data sets, we will compare various simulations with real data in §4. In §5 we will show two use cases which demonstrate how the framework can be used to optimize data processing and measurement parameters.

First, we will create a virtual counterpart of a high-speed radiography experiment and use the simulated data together with ground truth to compare the performance of motion estimation algorithms, select the best performing one and optimize one of its parameters. Second, we will use *syris* to simulate tomographic data sets under different imaging conditions which lead to various reconstruction accuracies of the filtered back projection algorithm (Kak & Slaney, 1988). Finally, we will conclude our work and provide an outlook in §6.

1.1. Scope of *syris*

syris is a *framework* which orchestrates virtual X-ray imaging experiments by connecting *components* responsible for particular aspects of the virtual experiment, e.g. sample shape description, X-ray source representation, and so on. The components interact with each other by a clear application programming interface (API), which makes them easy to change or extend.

Out of the box, the framework supports a broad range of virtual X-ray imaging experiments, because it ships with implementations of the components required for sample creation, motion and image formation description. The

current implementation models the wavefield of the X-ray beam, its interaction with the sample and free-space propagation along the optical axis z based on the scalar diffraction theory (Goodman, 2005). This theory has proven to be a sufficiently accurate model for the contrast formation of typical full-field X-ray imaging applications, where the beam is scattered mostly in the forward direction (e.g. Bragg diffraction is not considered).

The simulation of X-ray imaging experiments can be computationally very expensive. For this reason the physical and mathematical models are selected in such a way that they enable both realistic simulations and fast implementation at the same time. Moreover, we leverage the power of modern graphics processing units (GPUs) by implementing the computationally demanding parts of *syris* in OpenCL (Munshi, 2009).

2. Physical principles considered in *syris*

In this section we will explain the physical models behind the image formation components included in *syris*.

First, we apply a fully deterministic approach to obtain the intensity pattern at the virtual imaging plane. A wavefield emitted by a source may pass through an arbitrary number of objects described by their transmission functions. Its evolution between the objects is modeled by free-space propagation. We then retroactively incorporate the stochastic nature of the photon emission and the detection process by considering partial lateral coherence, polychromaticity, shot noise and electronic noise in an X-ray imaging detector system.

2.1. Profile of the incident beam

In the deterministic part of the calculations, we first assume perfect coherence properties, namely a point source emitting a monochromatic wavefield $u_0(\mathbf{x}, z_1)$ with wavelength λ . Here, \mathbf{x} denotes the two-dimensional (2D) coordinates in the plane perpendicular to the optical axis z . z_1 is the position of the first penetrated object in the light path. The wavefield may have spatially varying intensity distribution $I_0(\mathbf{x}, z_1) = |u_0(\mathbf{x}, z_1)|^2$. In the case of a spherical incident wave, large z_1 and small spatial extent across the direction perpendicular to z , we may use the parabolic approximation of the spherical phase profile (Cloetens, 1999) and write

$$u_0^{\text{parabolic}}(\mathbf{x}, z_1) = [I_0(\mathbf{x}, z_1)]^{1/2} \exp(jkz_1) \exp[jk/(2z_1)\mathbf{x}^2], \quad (1)$$

where $k = 2\pi/\lambda$ is the wavenumber. If the source of the spherical wave is sufficiently far away, which is typically the

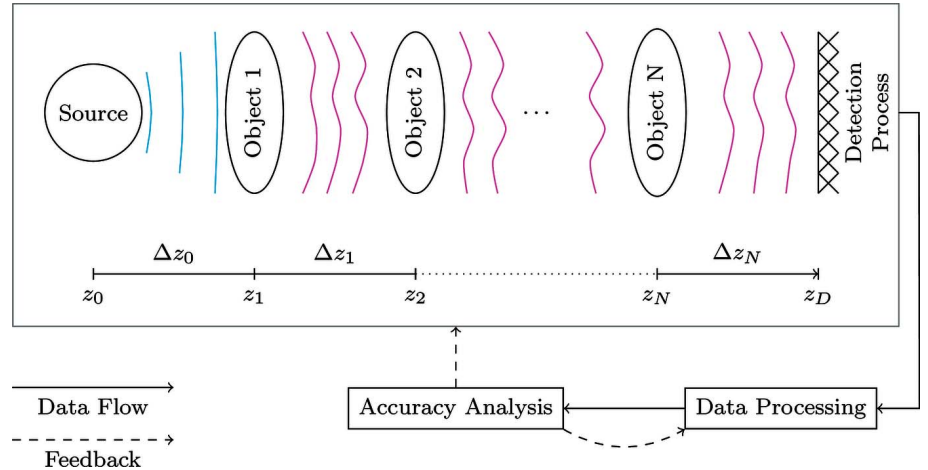


Figure 2

X-ray path which can be simulated in *syris* and a possible use case. Simulated data may be passed to ‘Data Processing’ which extracts the desired information about the sample (e.g. a 3D volume). ‘Accuracy Analysis’ evaluates the data processing performance based on the ground truth provided by the simulation. The estimation may then be used to adjust the virtual or real experimental and data processing parameters.

case at synchrotrons, even the parabolic approximation may be relaxed and a constant phase profile may be used instead,

$$u_0^{\text{const}}(\mathbf{x}, z_1) = [I_0(\mathbf{x}, z_1)]^{1/2} \exp(jkz_1). \quad (2)$$

2.2. Transmission functions

An arbitrary number of time-dependent objects can be placed in the X-ray light path, as depicted in Fig. 2. Their positions are given by z_i ($z_i > z_j$ for $i, j \in \mathbb{N}$, $j > 0$ and $i > j$) along the X-ray beam between the source at z_0 and the detection plane $z_d > z_N$; their separation is expressed by $\Delta z_i = z_{i+1} - z_i$.

Every object i in the light path is described by the 3D complex refractive index,

$$n_i(\mathbf{x}, z, t) = 1 - \delta_i(\mathbf{x}, z, t) + j\beta_i(\mathbf{x}, z, t). \quad (3)$$

We will omit the time-dependence in the following text for brevity until §3.2. The transmission function integrates the refractive index along z and it is defined as (Born & Wolf, 1999)

$$T_i(\mathbf{x}) = \exp[jk \int n_i(\mathbf{x}, z) dz] = \exp\{-k[B_i(\mathbf{x}) - j\varphi_i(\mathbf{x})]\}, \quad (4)$$

where $B_i(\mathbf{x}) = \int \beta_i(\mathbf{x}, z) dz$ and $\varphi_i(\mathbf{x}) = \int [1 - \delta_i(\mathbf{x}, z)] dz$ correspond to the local absorption and phase shift in the i th object’s exit plane.

The relation between the wavefield $u_{i-1}(\mathbf{x}, z_i)$ in the i th object’s entrance plane and $u_i(\mathbf{x}, z_i)$ in its exit plane can be expressed as

$$u_i(\mathbf{x}, z_i) = T_i(\mathbf{x}) u_{i-1}(\mathbf{x}, z_i). \quad (5)$$

2.3. Free-space wavefield propagation between objects

Exact within the scalar theory, the relation of the lateral wave profile on two parallel planes with separation Δz can be described by means of the angular spectrum formalism (Goodman, 2005). With $\tilde{u}(\xi) = \mathcal{F}[u(\mathbf{x})]$ being the 2D Fourier transform of a wavefield with 2D spatial frequencies ξ , we can write

$$\tilde{u}(\xi, z + \Delta z) = \tilde{P}(\xi, \Delta z) \tilde{u}(\xi, z), \quad (6)$$

where

$$\tilde{P}(\xi, \Delta z) = \exp\left\{jk\Delta z[1 - (\lambda\xi)^2]^{1/2}\right\} \quad (7)$$

is the so-called propagator. The wavefield at the distance Δz behind the i th object can thus be computed by the recursive relation

$$u_i(\mathbf{x}, z_i + \Delta z) = \mathcal{F}^{-1}\{\tilde{P}(\xi, \Delta z) \mathcal{F}[u_{i-1}(\mathbf{x}, z_i) T_i(\mathbf{x})]\}. \quad (8)$$

The Fresnel approximation of the propagator in (7) is based on the parabolic approximation of the spherical waves and it is well known to be sufficiently accurate for most X-ray imaging applications. It is defined as

$$\tilde{P}_F(\xi, \Delta z) = \exp(jk\Delta z) \exp(-j\pi\lambda\Delta z\xi^2). \quad (9)$$

2.3.1. Parabolic incident wavefield. We will now take into account the parabolic incident wave from (1). Let us first separate $u'_i(\mathbf{x}, z_i)$ from the parabolic phase profile of the wavefield,

$$u_i(\mathbf{x}, z_i) = u'_i(\mathbf{x}, z_i) \exp\left(\frac{jk}{2z_i} \mathbf{x}^2\right), \quad (10)$$

and then apply the Fresnel diffraction integral in real-space,

$$u_i(\mathbf{x}, z_i + \Delta z) = \frac{\exp(jk\Delta z)}{j\lambda\Delta z} \int u_i(\boldsymbol{\eta}, z_i) \times \exp\left[j\frac{k}{2\Delta z}(\mathbf{x} - \boldsymbol{\eta})^2\right] d\boldsymbol{\eta}. \quad (11)$$

Here, $\boldsymbol{\eta}$ are the 2D spatial coordinates in the plane z_i and \mathbf{x} those in the plane $z_i + \Delta z$. Insertion of (10) into (11) leads to (Cloetens, 1999)

$$u_i(\mathbf{x}, z_i + \Delta z) = \exp\frac{jk}{2(z_i + \Delta z)} \mathbf{x}^2 \frac{\exp(j\Delta\Psi_i) \exp(jk\Delta D)}{M j\lambda\Delta D} \times \int u'_i(\boldsymbol{\eta}, z_i) \exp\left[j\frac{k}{2\Delta D}\left(\frac{\mathbf{x}}{M} - \boldsymbol{\eta}\right)^2\right] d\boldsymbol{\eta}, \quad (12)$$

where the magnification $M = (z_i + \Delta z)/z_i$, the de-focusing distance $\Delta D = z_i\Delta z/(z_i + \Delta z)$ and the global phase shift $\Delta\Psi_i = 2\pi(\Delta z)^2/[\lambda(z_i + \Delta z)]$ has been introduced. Remarkably, this is again an expression of the general form in (10) and comparison with (11) reveals that the corresponding $u'_i(\mathbf{x}, z_i + \Delta z)$ is equivalent to propagating $u'_i(\mathbf{x}, z_i)$ by a distance ΔD combined with spatial magnification M , amplitude reduction $1/M$ and multiplication with a pure phase

factor $\exp(j\Delta\Psi_i)$. We can express the computation of a wavefield after the i th object recursively again,

$$u'_i(M\mathbf{x}, z_i + \Delta z) = \frac{\exp(i\Delta\Psi_i)}{M} \times \mathcal{F}^{-1}\{\tilde{P}(\xi, \Delta D) \mathcal{F}[u'_{i-1}(\mathbf{x}, z_i) T_i(\mathbf{x})]\}. \quad (13)$$

If we are interested in the detected intensity, we may omit the parabolic phase profiles and work with adjusted propagation distances and magnifications instead, which simplifies the numerical computation.

2.4. Partial coherence and polychromaticity of the incident wavefield

The previous formalism assumes a monochromatic laterally coherent incident wavefield $u_0(\mathbf{x}, z_1)$. In contrast, the commonly used real X-ray sources (X-ray tubes, synchrotrons) are based on the widely uncorrelated emission of photons with various energies and are thus by nature considered as intrinsically incoherent. The observed degree of partial lateral coherence results only from free-space propagation.

To account for the corresponding partial lateral coherence, we employ the van Cittert–Zernike theorem (Born & Wolf, 1999) which is a sufficiently good approximation for large ratios z_1/σ , where z_1 is the distance between the source and the first transmitted object and σ is the characteristic source size. When we obtain an intensity pattern based on calculations with a perfectly coherent source and assume no magnification (parallel geometry), the lateral extension of the source can be accounted for by convolving such an intensity pattern with the source intensity distribution $S(\mathbf{x})$, geometrically rescaled through the sample to the detection plane. For the case of a single object and its distance z_1 from the source, the intensity at $z_1 + \Delta z$ is

$$I_{\text{incoh}}(\mathbf{x}, z_1 + \Delta z) = \mathcal{F}^{-1}\left\{\tilde{S}\left(\frac{\Delta z}{z_1} \xi\right) \mathcal{F}[I(\mathbf{x}, z_1 + \Delta z)]\right\}. \quad (14)$$

For more than one transmitted object in the beam path, this approximation is valid for $\Delta z \gg \sum_{i=1}^{N-1} \Delta z_i$.

Finally, we account for a polychromatic incident wavefield by incoherent superposition of monochromatic intensities, which enables us to calculate the intensity pattern at the detector plane by

$$I_{\text{poly}}(\mathbf{x}, z_d) = \int I_{\text{incoh}}(\mathbf{x}, z_d, \lambda) d\lambda. \quad (15)$$

2.5. Detection of the intensity profile

Indirect detectors are commonly used for X-ray imaging experiments. They use a thin scintillator crystal which converts X-ray photons to visible light. Such visible light patterns are subsequently magnified by an optical system and detected by a digital camera.

First, we address the X-ray to visible light conversion process (Douissard *et al.*, 2010) by

$$I_v(\mathbf{x}, z_d) = \int I(\mathbf{x}, z_d, \lambda) \{1 - \exp[-\mu(\lambda)p]\} \frac{hc}{\lambda} L(\lambda) d\lambda, \quad (16)$$

where $\mu(\lambda)$ is the linear attenuation coefficient of the scintillator's material, p its thickness and $L(\lambda)$ the light yield.

The number of visible light photons emitted by the scintillator, transmitted through the optical system and converted to electrons in the camera sensor further depends on the emission spectrum of the scintillator, optics collection efficiency and the camera quantum efficiency, which we combine into a single *detector conversion factor* $A(\lambda_v)$, where λ_v is the wavelength of the visible light. The detection process also causes blurring of the recorded image, which we account for by convolving the image with the convolution kernel $R(\mathbf{x}, \lambda_v)$. If e_d^- is the mean number of electrons in the camera sensor present without incident light, the total number of electrons which are read out by camera electronics is

$$e^-(\mathbf{x}) = e_d^- + I_v(\mathbf{x}, z_d) \int A(\lambda_v) R(\mathbf{x}, \lambda_v) d\lambda_v. \quad (17)$$

The signal recorded by a camera is amplified by the *overall system gain* K and contains signal-dependent Poisson-distributed shot noise with variance $\sigma_e^2(\mathbf{x})$, signal-independent normally distributed electronic noise with variance σ_d^2 and uniformly distributed quantization noise with variance σ_q^2 . The total variance of the recorded image can be written as (Jähne, 2010)

$$\sigma^2(\mathbf{x}) = K^2[\sigma_d^2 + \sigma_e^2(\mathbf{x})] + \sigma_q^2. \quad (18)$$

We illustrate the most important aspects of the image formation process based on the physical principles described above in Fig. 3.




Figure 3 Simulated intensity patterns of a plane wave which passes through a PMMA sphere and propagates 1 m in free space to the imaging plane. (a) Monochromatic incident wavefield with energy 15 keV; (b) polychromatic incident wavefield with energy range 10–30 keV; (c) partial spatial coherence due to an extended source applied on (b) (vertically five times larger coherence than horizontally); (d) short exposure time (the sphere moves horizontally by one pixel per camera exposure time) and the associated strong shot noise applied on (c); (e) also shot noise application on (c) but with 50-times longer exposure time than in (d), thus the SNR is larger but there is motion blur by 50 pixels. Grayscale is adjusted for every image separately.

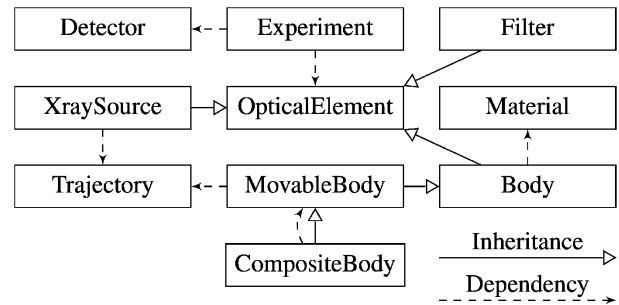


Figure 4 Simplified class diagram of *syris*. *Experiment* orchestrates a virtual experiment and outputs image sequences. It uses various classes which create, modify or detect the virtual wavefield.

3. Implementation

syris is an open-source project with a Python (<https://www.python.org>) user-level API currently running on Python 2. The source code with many examples is freely available on Github (<https://github.com/ufo-kit/syris>) and the documentation is located on Read the Docs (<https://syris.readthedocs.io/en/latest/>).

The functionality is split across several packages, modules and classes. All computationally demanding parts are implemented in OpenCL and can be executed on GPUs. The core classes are shown in Fig. 4.

The topmost class for conducting virtual experiments is the *Experiment* class. It uses various classes which create, modify and detect the simulated wavefield. It orchestrates the virtual experiment and outputs the detected images as a function of time.

3.1. Optical elements

The most important class for the implementation of the X-ray image formation is the *time-dependent* *OpticalElement*. It is a cornerstone for other classes which create or modify an X-ray wavefield. One of its subclasses is the *XraySource* class, which produces the incident wavefield and is currently implemented for bending magnet or wiggler (Williams, 2001) sources.

The *Body* class represents a 3D object that uses the *Material* class to compute the object's transmission function. The object shape in the *Body* class is currently based on two models. Metaballs (Blinn, 1982) describe an isosurface as a combination of falloff curves and are suitable for modeling objects with smooth transitions. Alternatively, triangular meshes allow us to create arbitrary object shapes (see Fig. 5b).

The *Material* class provides the complex refractive index, which can currently be looked up in the CXRO database (Henke *et al.*, 1993), the X_{0h} database (<http://x-server.gmca.aps.anl.gov/x0h.html>) and the *pmasf* program, which combines optical constants from Henke (1993, 1997) and NIST tables for the energy range 30–433 keV with fine sampling around absorption edges.




Figure 5 Illustration of different object shape creation techniques. (a) Projected thickness of two metaballs computed by *syris*, (b) projected thickness of a biological screw joint (van de Kamp *et al.*, 2011) represented as a 3D triangular mesh and projected by *syris*. For illustration, (c) shows a rendering of the mesh used to compute (b).

3.2. Time-dependence

Wavefield creation and modification by an `OpticalElement` is time-dependent, which enables simulations of, for example, beam drift and sample motion. Time-dependence is currently implemented in `MovableBody` and `XraySource` by rigid-body motion. These objects can have a `Trajectory` with a velocity profile. Trajectories are implemented as B-splines (De Boor *et al.*, 1978) $\mathbf{B}(t)$ with time parameter t . This parametrization enables us to place the optical element at any point of the spline at any time, *i.e.* we can create arbitrary velocity profiles. Moreover, `MovableBody` has a defined direction vector. Its position and orientation at a time t is given by placing it at $\mathbf{B}(t)$ and aligning its direction vector with the spline derivative $\mathbf{B}'(t)$. For instance, a metaball with radius $20\ \mu\text{m}$ moving along the x -axis between $0\ \text{mm}$ and $1\ \text{mm}$ with constant velocity $5\ \mu\text{m s}^{-1}$ can be created like this:

```
# Create a linear trajectory
x = np.linspace(0, 1, num=128)
y = z = np.zeros_like(x)
# Create spline control points
control_points = zip(x, y, z) * q.um
# Trajectory with a constant velocity
trajectory = Trajectory(control_points,
                        velocity=5 * q.um / q.s)
# MetaBall with radius 20 micrometer
body = MetaBall(trajectory, 20 * q.um)
```

A `CompositeBody` encompasses a number of movable bodies to describe complex motion patterns. All the sub-bodies follow the global motion of the `CompositeBody` and additionally their own relative trajectories.

3.3. Image formation

The calculation of the image formation described in §2 up to the point of adding noise is realised by series of pixel-wise multiplications in real and Fourier space. We use the fast Fourier transform (FFT) executed on GPUs (<http://github.com/Manticore-attic/pyfft>) to quickly convert one space to the other.

We would like to point out that the transmission function and free-space propagation can suffer from aliasing artifacts

caused by the discretization of high spatial frequency components (Sypek *et al.*, 2003; Kirkland, 2010) if the sampling theorem (Shannon, 1949) is not satisfied. If an object's transmission function or a propagator would cause phase shift between two adjacent pixels larger than $\pm\pi$ (thus violate the sampling theorem), *syris* outputs a warning.

Periodic convolution caused by the finite extent of the computational grid may cause artifacts as well. This can be avoided in *syris* by using a twice as large computational grid centered around that of the desired size.

The following code snippet calculates polychromatic free-space propagation to $1\ \text{m}$ of a wavefield transmitted through a sphere with radius $256\ \mu\text{m}$:

```
# A dimensionless 2D pixel grid
grid = (1024, 1024)
energies = range(15, 30) * q.keV
pixel_size = 1 * q.um
material = make_pmasf('PMMA', energies)
sphere = make_sphere(grid, 256 * q.um,
                    pixel_size=pixel_size,
                    material=material)

# Propagate to 1 m
result = propagate([sphere], grid,
                  energies, 1 * q.m,
                  pixel_size)
```

3.4. Parallelization

The pixel-wise calculations are independent of the neighboring pixels, thus they can be parallelized by using OpenCL. For a comparison of CPU (central processing unit) and GPU performance see Table 1, which shows the execution times of polychromatic wavefield propagation on various platforms.

In addition to pixel-wise parallelization, we also provide `qmap`, a function which distributes images across multiple GPUs for further speedup.

4. Simulation comparison with real experiments

In this section we will compare simulations with real data for various experiments conducted at the KIT synchrotron

Table 1

CPU *versus* GPU performance of a 2D polychromatic wavefield propagation calculation with different number of pixels ($N \times N$) and 100 energy points.

CPU is the Intel® Xeon® E5-2680 v2@2.8 GHz, GPU is the NVIDIA® GeForce® GTX Titan. The same code executed on a corresponding OpenCL platform is used.

	N				
	512	1024	2048	4096	8192
CPU	2.38 s	3.90 s	8.08 s	29.63 s	189.39 s
GPU	1.95 s	1.96 s	2.20 s	4.08 s	9.00 s

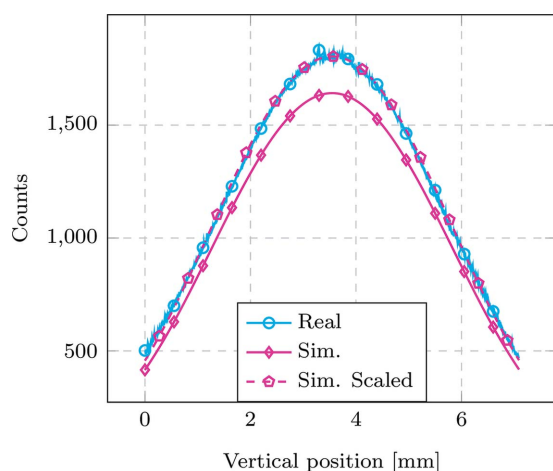
radiation facility to show the validity of the image formation model described in §2 and its implementation described in §3.

4.1. Bending-magnet beam profile

The aim of this example is to show that *syris* can simulate filtered white-beam intensity profiles for a bending-magnet source and that they can be converted into an image with the number of counts given by an indirect detector.

Here, *syris* used the energy range from 10 keV to 40 keV in 1 keV steps to compute intensity profiles of the beam at the imaging plane 30 m from the source directly (Thompson *et al.*, 2001), without free-space propagation. A simulated indirect detector computed the number of visible-light photons emitted by a scintillator for every profile according to (16) and converted them to detector counts by (17) and by using the camera gain (see §2.5). Noise was not simulated in this case. The full set of parameters are given in Table 3, Appendix A.

A comparison between the real and the simulated data is shown in Fig. 6 (as a consequence of broad horizontal divergence, the horizontal beam profile is almost constant and not depicted here). The difference in the absolute number of detector counts between real data and simulation is caused by the fact that we worked with theoretical values of some

**Figure 6**

Vertical white-beam profile of a bending magnet. The cyan curve depicts real detected counts, the magenta curve represents simulated detected counts and the dashed magenta curve shows simulated counts multiplied by a global factor to show that the virtual profile is qualitatively in good agreement with the real data.

imaging parameters, *e.g.* the light yield and emission spectrum of the scintillator, which can vary between scintillators of the same crystal. To show that the beam profile is qualitatively correct, we applied a global multiplicative factor to the simulated counts, depicted by dashed magenta lines in Fig. 6.

4.2. Absorption and edge enhancement

In this experiment we imaged an acrylic glass cuboid of size $0.42 \text{ mm} \times 0.42 \text{ mm} \times 10 \text{ mm}$, positioned 942 mm from the detector in order to confirm correct absorption and edge enhancement simulation.

For preparing the virtual shape of the sample, we performed a computed tomography scan of the cuboid in a laboratory setup with pixel size of $6.1 \mu\text{m}$, so that it would fit the field of view of the detector. Then we performed 3D reconstruction, segmentation and converted voxel data to a 3D mesh, which was further processed to decrease the number of polygons. Then we used a synchrotron source and an effective pixel size of $1.22 \mu\text{m}$ to acquire high-resolution X-ray projections of the cuboid's edge in the white-beam mode. The cuboid's long edge was positioned vertically and the diagonal of the base was perpendicular to the beam. We closed the slit located at 4.727 m from a bending-magnet source to $200 \mu\text{m} \times 200 \mu\text{m}$ in order to limit the blurring caused by the source size.

To simulate the experiment, we used *syris* to compute intensities of the wavefield propagated from the cuboid's exit plane to the imaging plane. We used energies from 8 keV to 30 keV in 1 keV steps together with (15) to simulate white beam. The resulting intensity pattern was blurred by the rescaled source intensity distribution according to (14). To simulate the closed slit, we considered a virtual source located in the slit position and used the slit opening as the full width at half-maximum (FWHM) of the source's intensity distribution modeled by a Gaussian. Finally, *syris* converted the blurred intensity patterns into detector counts as in the previous example. The full experiment description can be found in Table 4, Appendix A.

Horizontal line averages of real and simulated flat-field corrected projection images are depicted in Fig. 7. The averages were computed by rotating the images to make the cuboid edge vertically aligned and by averaging them along the vertical direction to suppress noise from surface roughness. Flat-field corrected projection images are defined by $\log(I_0/I)$, where I is the intensity with the sample and I_0 is that without it.

4.3. Tomography simulation

We performed a tomographic scan and its simulation to show that *syris* can provide complex tomographic virtual data sets which can give rise to typical reconstruction artifacts, like streaks and rings.

We first conducted tomographic data acquisition of a wax attached to a steel needle located 8.2 cm in front of the detector in order to give rise to edge enhancement. After the data acquisition, we performed 3D reconstruction, segmented




Figure 7 Horizontal line averages (averages of flat-field corrected projection images along the vertical direction) of a real (cyan) and simulated (magenta) acrylic glass cuboid edge. On the left is the full field of view showing the correct absorption simulation. On the right is the cropped part around the cuboid edge showing that the simulation of the edge enhancement is in good agreement with the real data. The propagation distance between the cuboid and the detector was 942 mm.

the wax and the needle and converted the segmentations to 3D meshes.

Then we used these meshes to conduct virtual tomographic data acquisition. *syris* computed X-ray projections at various rotation angles with the sample positioned in the same way as in the real experiment. The wavefield from the sample was propagated to the detector in the white-beam mode in the same way as in the previous experiment but here no source blurring was applied. This is because the slit located at 4.727 m from the source was closed to $200\ \mu\text{m} \times 200\ \mu\text{m}$, which in combination with a slit-to-sample distance of $>25\ \text{m}$, sample-to-detector distance of 8.2 cm and effective pixel size of $1.22\ \mu\text{m}$ leads to source blurring by less than one pixel. In this case the noise simulation was turned on.

Comparison between real flat-field corrected projection and simulation is depicted in Fig. 8. Differences in subtle details (e.g. small features, surface roughness) are due to the segmentation procedure and 3D surface post-processing (surface smoothing and polygon reduction).

Next, we performed 3D reconstruction and show the comparison between the tomographic slices obtained from

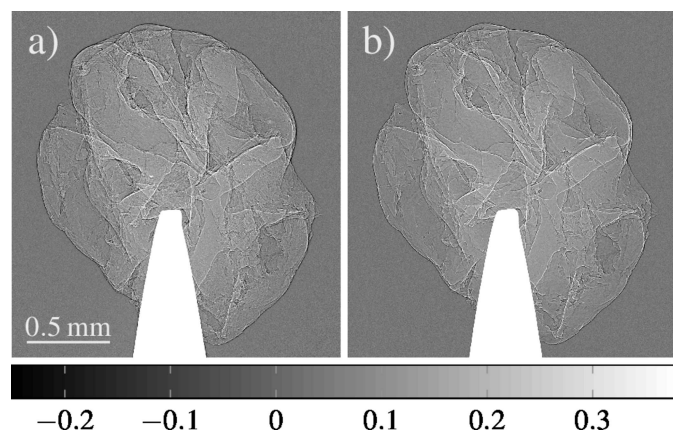


Figure 8 Flat-field corrected projection images from (a) the real data set and (b) the simulated one.

real and simulated projections in Fig. 9. Streak artifacts caused by edge enhancement can be observed near the notch with strong white–black transitions. Ring artifacts are the consequence of noise and can be observed in the top right region of the slices.

5. Examples of simulation-based experiment and data processing optimization

We will now demonstrate the capabilities of *syris* by conducting two virtual experiments. First, we will simulate a high-speed radiography experiment and show that high-fidelity simulations can help one to choose proper data processing

algorithms. Second, we will simulate tomographic data sets acquisition and vary the imaging conditions between them to compare the accuracy of the tomographic reconstruction. Such comparison helps one to choose the best imaging conditions for a particular real measurement.

5.1. Selection and optimization of a motion estimation algorithm for the analysis of a high-speed radiography measurement

In this section we will use *syris* to create a virtual counterpart of a high-speed radiography of a semi-solid aluminium alloy with moving particles inside (Zabler *et al.*, 2013), see Fig. 10. The purpose of the real experiment was to investigate the motion of the alloy and the particles, for which a motion estimation algorithm was used.

Here, we will compare the performance of different motion estimation algorithms, select the one which gives the most accurate motion vectors and optimize one of its parameters based on the ground truth data (2D motion vectors for every pixel in an image) available from the simulation. Since the

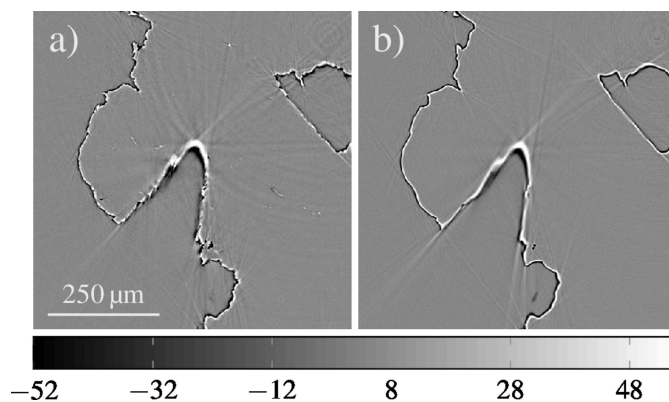


Figure 9 Reconstructed slice ROIs of linear attenuation coefficients (in cm^{-1}) from (a) real projections and (b) simulated projections. Streak artifacts are present near the notch and ring artifacts can be observed in the top right corners.




Figure 10
 (a) Flat-field corrected projection image of a semi-solid alloy with particles inside (Zabler *et al.*, 2013) and (b) its simulation based on the same experimental parameters.

simulation models many important aspects of image formation, one may use the found parameter to analyze the real data.

First, we re-create the original experiment in *syris* by using the wiggler source properties of the European Synchrotron Radiation Facility (ESRF) beamline ID15a. We use triangular meshes to create static parts of the sample and metaballs to create the particles inside the liquid. *syris* calculates the incident wavefield on the sample filtered by 20 mm of silicon, computes the transmission functions of the sample parts and applies them to the incident wavefield, which gives us the wavefield in the sample exit plane. *syris* further calculates the detected intensity pattern after the free-space propagation of this wavefield to the scintillator, located 10 cm after the sample. The comparison of the real and the virtual experiment can be found in Fig. 10.

Next, we compare the accuracy of four motion estimation algorithms (also called optical flow algorithms):

- (i) M_1 : Horn & Schunck (Horn & Schunck, 1981);
- (ii) M_2 : M_1 + robust flow-driven (Papenberg *et al.*, 2006);
- (iii) M_3 : M_2 + combined local–global approach (Bruhn *et al.*, 2005);
- (iv) M_4 : M_3 + intermediate flow filtering (Sun *et al.*, 2014).

Each algorithm is applied on the same pair of consecutive simulated flat-field corrected projection images with low contrast-to-noise ratio, which imposes a significant challenge for motion estimation. The estimated 2D motion vectors for every pixel in the first image are compared with the ground truth motion vectors in terms of the *endpoint error* (EE), which determines the difference between the ground truth $\mathbf{x}_{GT} = (x_{GT}, y_{GT})$ and the estimated $\mathbf{x}_{est} = (x_{est}, y_{est})$ and is defined as

Table 2

Performance comparison of the four optical flow algorithms applied on a synthetic data set with noise.

Model	$\mu(\text{EE})$
$M_1 = \text{Horn \& Schunck}$	0.664
$M_2 = M_1 + \text{robust flow-driven}$	0.655
$M_3 = M_2 + \text{combined local–global}$	0.624
$M_4 = M_3 + \text{flow filtering}$	0.560

$$\text{EE} = [(x_{\text{GT}} - x_{\text{est}})^2 + (y_{\text{GT}} - y_{\text{est}})^2]^{1/2}. \quad (19)$$

Since the static background occupies a substantial part of the image, we measure the accuracy only in the vicinity of the moving particles. A performance comparison of the four optical flow methods is given in Table 2.

We select algorithm M_4 which produces the lowest average endpoint error with respect to the ground truth and we show how its accuracy depends on the settings of one of its parameters.

A variational optical flow model in its general form consists of two terms: a data term and a smoothness term (Sun *et al.*, 2014). To control the influence of both terms on the overall model, a special weighting parameter is used, the so-called smoothness parameter α . Optimizing this parameter for particular imaging conditions is crucial in order to obtain the best results.

We used the simulated data and the ground truth to compare the performance of M_4 with respect to α (see Fig. 11). One can see that increasing the smoothness parameter from $\alpha = 0$ improves the accuracy of the result and the best performance corresponds to $\alpha = 2.3$. Since the simulated data set closely resembles the real one, this parameter value may be selected for processing the real data.

5.2. Tomographic reconstruction accuracy dependence on different imaging conditions

In this section we will use *syris* to simulate tomographic data sets with various imaging conditions and compare the 3D reconstruction accuracy between them. Such comparison

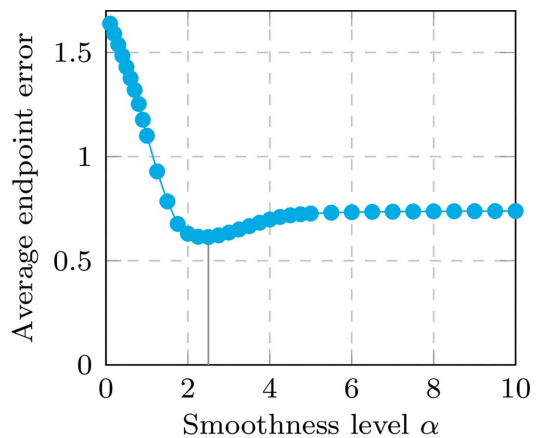


Figure 11
 Motion estimation accuracy of algorithm M_4 as a function of the smoothness parameter α .

may guide the selection of imaging conditions for real experiments.

Suppose we acquire N projections for one data set and the sample rotation speed is synchronized with the camera exposure time t , so that there is no rotational motion blur, *i.e.* the sample does not move by more than one pixel between two consecutive projections. If we increase the sample rotation speed, we can either reduce the exposure time t for a single projection or keep t constant and record fewer projections to cover the 180° angular range required for 3D reconstruction. In the former strategy, the angular sampling remains unchanged and only the signal-to-noise ratio (SNR) decreases in the reconstructed slices, whereas the latter strategy reduces both the SNR and the angular sampling. Moreover, it introduces rotational motion blur. We will show how these two reduction strategies impact the accuracy of the filtered back projection algorithm (FBP) and discuss when it is beneficial to reduce the amount of acquired projections N instead of the exposure time t .

First we use Blender (<https://www.blender.org>) to create the 3D phantom shape, which consists of multiple polygonal meshes and includes sharp transitions which induce streak artifacts in the tomographic reconstruction. *syris* assigns aluminium, calcium, scandium and titanium materials to various meshes, based on which it can compute the phantom's transmission function. A 2D slice of the phantom with the applied materials is shown in Fig. 12.

To obtain projections at various angles, *syris* rotates the phantom around the tomographic rotation axis. Projections are calculated by using a monochromatic plane incident wave with energy 20 keV, computing the intensity pattern in the exit plane of the phantom (*i.e.* free-space propagation is not considered) and applying Poisson-distributed shot noise. To create data sets with various reduction strategies, *syris* either decreases t or N by a *reduction factor* and increases the sample rotation speed accordingly, so that the projections are acquired over 180° .

The initial exposure time (reduction factor of 1) is selected in such a way that it does not lead to the rotational motion blur, and the reference projection (without sample) has a SNR $\mu/\sigma \simeq 358$, where μ is the signal mean and σ is the standard deviation. The initial number of projections is 1600, which satisfies the angular sampling limit for our phantom (Kak & Slaney, 1988). The ground truth slice is obtained from the reconstruction of a data set with reduction factor 1 and no noise. We use the mean squared error (MSE) for comparing the ground truth slice and the respective slices from the simulated data sets (see Fig. 13).




Figure 12 Linear attenuation coefficients (in cm^{-1}) for energy 20 keV in a 3D phantom slice perpendicular to the rotation axis. The rectangles mark regions used for the FBP accuracy analysis in this section.

The MSE for the reduced t mode grows linearly with the reduction factor, as shown in Fig. 14, which is expected due to the fact that the noise power spectrum grows linearly with the reduction of t (Riederer *et al.*, 1978).

Rotational motion blur plays a role in the quality of slice reconstructions for the reduced N case. However, when the ROI is close to the rotation axis, such as the one marked in cyan in Fig. 12, the MSE grows almost linearly as in the case of the reduced t , meaning that the error stemming from the reduced number of projections is small compared with the

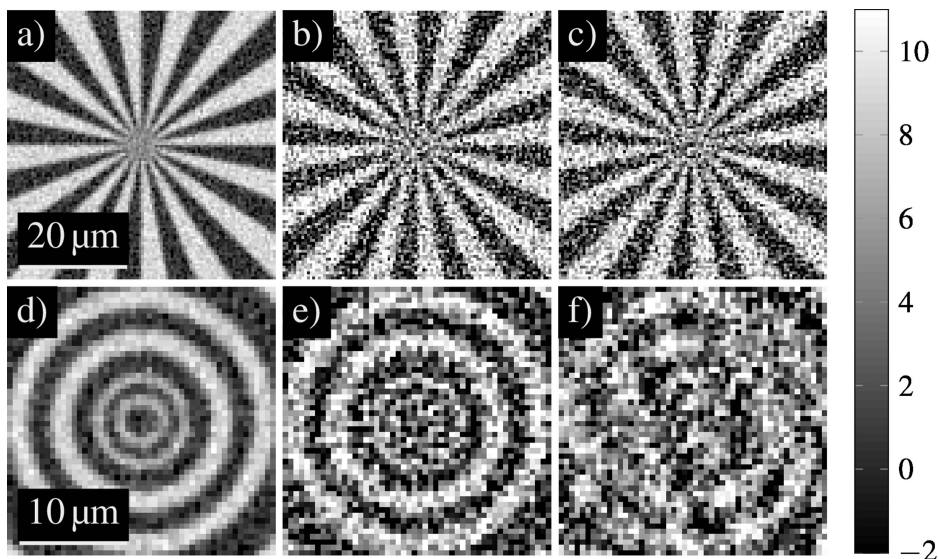


Figure 13 Reconstructions of linear attenuation coefficients (in cm^{-1}) for two phantom slice ROIs. The first row shows a ROI marked in cyan in Fig. 12, reconstructed from (a) 1600 projections, (b) the same number of projections as in (a) but with 16 times shorter t , (c) the same t as in (a) but 16 times less N . The second row shows a ROI further away from the rotation axis marked in magenta in Fig. 12 with the same t and N settings as in the first row.

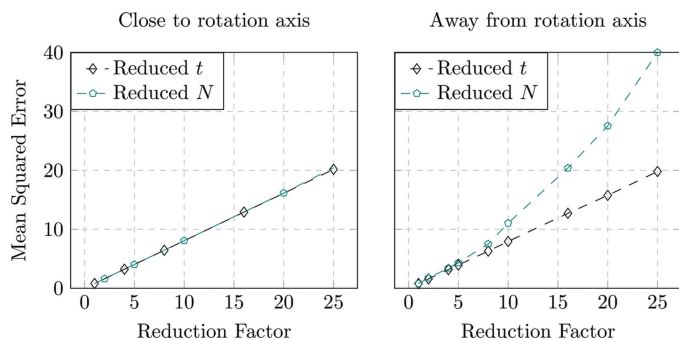


Figure 14

MSE of the reconstructions of linear attenuation coefficients (in cm^{-1}) in two slice ROIs as a function of either reducing the exposure time t (black) or the number of projections N (green).

reduced exposure time, as depicted on the left of Fig. 14. The MSE for a ROI positioned away from the rotation axis (magenta in Fig. 12) grows more rapidly and is shown on the right of Fig. 14.

To conclude, reducing N may lead to comparable reconstruction accuracy as by reducing t , especially when the ROI is positioned close to the rotation axis. Thus, reducing N may be an interesting option when the amount of acquired data needs to be minimized, *e.g.* for saving storage space or faster tomographic reconstruction.

6. Conclusion and outlook

We have developed a flexible and efficient framework, *syris*, for conducting virtual X-ray imaging experiments. It takes into account many important aspects of the X-ray image formation process. Currently implemented, enabling simulations of four-dimensional X-ray imaging experiments, are:

- (i) Bending magnet and wiggler sources.
- (ii) Metaballs and triangular meshes for shape description.
- (iii) Wavefield propagation through multiple objects.
- (iv) Free-space propagation of the wavefield between objects.
- (v) Polychromaticity and reduced lateral coherence.
- (vi) X-ray to visible light conversion and detection.
- (vii) X-ray source and object motion.

We demonstrated the capabilities of our framework by simulating various data sets and comparing them with the real ones acquired at a synchrotron imaging beamline. Furthermore, we simulated a high-speed radiography experiment and showed that realistic simulations may help to choose suitable image-processing algorithms and optimize their parameters in order to achieve good results when they are applied on real data. Second, we modeled a 3D phantom and used it to create virtual tomographic data sets under various imaging conditions (exposure times and numbers of projections). We further showed how they influence the accuracy of the FBP algorithm and that the accuracy may guide the selection of the conditions for particular use cases.

Our aim for the future is to create a database of testing data sets with given ground truth data, similar to a number of

popular databases for computer vision (Baker *et al.*, 2011; Martin *et al.*, 2001) and medical image analysis (van Ginneken *et al.*, 2010; Arganda-Carreras *et al.*, 2015; Hogeweg *et al.*, 2012). Such a database will be suitable for benchmarking and optimization of various algorithms which implement common X-ray imaging tasks, such as beam fluctuations compensation, noise reduction, motion blur treatment, phase retrieval, 3D reconstruction, segmentation and motion estimation.

Thanks to its parallelized implementation in OpenCL, *syris* can quickly compute new data sets based on changed virtual experimental parameters. This, combined with its flexibility, makes the framework highly suitable for investigating novel imaging approaches, creating data sets for the mentioned database and finding suitable imaging conditions prior to and during real experiments in order to make beam times more efficient.

APPENDIX A

Simulation parameters

Simulation parameters from §4 and §5 are listed here in Tables 3, 4, 5 and 6. Parameters which were not relevant for a particular simulation are marked by a dash.

Acknowledgements

We like to thank Marcus Zuber and Angelica Cecilia who supported us during the measurements.

Table 3

Simulation parameters for the KIT experiment in §4.1.

Category	Parameter	Value	
X-ray	Storage ring energy	2.5 GeV	
	Source type	Bending magnet	
	Source size (FWHM)	–	
	Magnetic field	1.5 T	
	Electric current	106 mA	
	Beam mode	White beam	
	Energies	10 keV to 40 keV, 1 keV step	
	Filters	0.7 mm aluminium	
	Slit position	–	
	Slit opening	–	
	Detector	Type	Indirect
		Effective pixel size	5.5 μm
Objective lens		Mitutoyo Plan Apo	
Magnification		2 \times	
Numerical aperture		0.055	
Scintillator		LuAg:Ce (Lu ₃ Al ₅ O ₁₂)	
Density		6.73 g cm^{-3}	
Thickness		50 μm	
Light yield		15 photons keV^{-1}	
Emission peak		525 nm	
Camera		PCO.dimax	
Sensor size		2016 \times 2016	
Sensor pixel size		11 μm	
A/D factor	10 e [−] per count		
Quantum efficiency	49% at 525 nm		
Dynamic range	1600:1		
Frame rate	100 frames s ^{−1}		

Table 4
Simulation parameters for the KIT experiment in §4.2.

Category	Parameter	Value	
X-ray	Storage ring energy	2.5 GeV	
	Source type	Bending magnet	
	Source size (FWHM)	0.89 mm × 0.22 mm	
	Magnetic field	1.5 T	
	Electric current	103 mA	
	Beam mode	White beam	
	Energies	8 keV to 30 keV, 1 keV step	
	Filters	0.2 mm aluminium	
	Slit position	4.727 m	
	Slit opening	200 μm × 200 μm	
	Detector	Type	Indirect
		Effective pixel size	1.22 μm
		Objective lens	Mitutoyo Plan Apo
		Magnification	10×
Numerical aperture		0.28	
Scintillator		LSO:Tb (Lu ₂ SiO ₅)	
Density		7.4 g cm ⁻³	
Thickness		12 μm	
Light yield		36 photons keV ⁻¹	
Emission peak		545 nm	
Camera		PCO.dimax	
Sensor size		2016 × 2016	
Sensor pixel size		11 μm	
A/D factor		10 e ⁻ per count	
Quantum efficiency	49% at 545 nm		
Dynamic range	1600:1		
Frame rate	25 frames s ⁻¹		
Sample	Material	PMMA (acrylic glass)	
	Dimensions	0.42 mm × 0.42 mm × 10 mm	
	Source distance	25273 mm	
	Detector distance	942 mm	
Aquisition	Experiment	Radiography	
	Data processing	Flat-field correction	

Table 5
Simulation parameters for the KIT experiment in §4.3.

Category	Parameter	Value	
X-ray	Storage ring energy	2.5 GeV	
	Source type	Bending magnet	
	Source size (FWHM)	–	
	Magnetic field	1.5 T	
	Electric current	106 mA	
	Beam mode	White beam	
	Energies	8 keV to 30 keV, 1 keV step	
	Filters	0.7 mm aluminium	
	Slit position	–	
	Slit opening	–	
	Detector	Type	Indirect
		Effective pixel size	1.22 μm
		Objective lens	Mitutoyo Plan Apo
		Magnification	10×
Numerical aperture		0.28	
Scintillator		LSO:Tb (Lu ₂ SiO ₅)	
Density		7.4 g cm ⁻³	
Thickness		12 μm	
Light yield		36 photons keV ⁻¹	
Emission peak		545 nm	
Camera		PCO.dimax	
Sensor size		2016 × 2016	
Sensor pixel size		11 μm	
A/D factor		10 e ⁻ per count	
Quantum efficiency	49% at 525 nm		
Dynamic range	1600:1		
Frame rate	25 frames s ⁻¹		
Sample	Material	Wax on a steel needle	
	Dimensions	1.5mm × 1.5 mm × 1.5 mm	
	Source distance	–	
	Detector distance	82 mm	
Aquisition	Experiment	Tomography	
	Number of projections	3000	
	Data processing	Flat-field correction	
	Reconstruction algorithm	Filtered back projection	

Funding information

The presented work was funded by the German Federal Ministry of Education and Research (BMBF) as UFO-1 and UFO-2 under the grants 05K10CKB and 05K10VKE.

References

Arganda-Carreras, I., Turaga, S. C., Berger, D. R., Ciresan, D., Giusti, A., Gambardella, L. M., Schmidhuber, J., Laptev, D., Dwivedi, S., Buhmann, J. M., Liu, T., Seyedhosseini, M., Tasdizen, T., Kamensky, L., Burget, R., Uher, V., Tan, X., Sun, C., Pham, T., Bas, E., Uzunbas, M. G., Cardona, A., Schindelin, J. & Seung, H. S. (2015). *Front. Neuroanat.* **9**, 142.

Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M. J. & Szeliski, R. (2011). *Int. J. Comput. Vis.* **92**, 1–31.

Bergbäck Knudsen, E., Prodi, A., Baltser, J., Thomsen, M., Kjaer 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.

Blinn, J. F. (1982). *ACM Trans. Graph.* **1**, 235–256.

Born, M. & Wolf, E. (1999). *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press.

Bruhn, A., Weickert, J. & Schnörr, C. (2005). *Int. J. Comput. Vis.* **61**, 1–21.

Chubar, O. & Elleaume, P. (1998). *Proceedings of the Sixth European Particle Accelerator Conference (EPAC'98)*, pp. 1177–1179.

Cloetens, P. (1999). PhD thesis, Vrije Universiteit Brussel, Belgium.

Davison, M. E. (1983). *SIAM J. Appl. Math.* **43**, 428–448.

De Boor, C., De Boor, C., De Boor, C. & De Boor, C. (1978). *A Practical Guide to Splines*, Vol. 27. New York: Springer-Verlag.

Douissard, P.-A., Cecilia, A., Martin, T., Chevalier, V., Couchaud, M., Baumbach, T., Dupré, K., Kühbacher, M. & Rack, A. (2010). *J. Synchrotron Rad.* **17**, 571–583.

Gunneken, B. van, S. G. A. III, de Hoop, B., van Amelsvoort-van de Vorst, S., Duindam, T., Niemeijer, M., Murphy, K., Schilham, A., Retico, A., Fantacci, M. E., Camarlinghi, N., Bagagli, F., Gori, I., Hara, T., Fujita, H., Gargano, G., Bellotti, R., Tangaro, S., Bolaos, L., Carlo, F. D., Cerello, P., Cheran, S. C., Torres, E. L. & Prokop, M. (2010). *Med. Image Anal.* **14**, 707–722.

Goodman, J. W. (2005). *Introduction to Fourier Optics*. San Francisco: Roberts.

Helfen, L., Baumbach, T., Mikulík, P., Kiel, D., Pernot, P., Cloetens, P. & Baruchel, J. (2005). *Appl. Phys. Lett.* **86**, 071915.

Henke, B. L., Gullikson, E. M. & Davis, J. C. (1993). *At. Data Nucl. Data Tables*, **54**, 181–342.

Hogeweg, L., Sánchez, C. I., de Jong, P. A., Maduskar, P. & Gunneken, B. (2012). *Med. Image Anal.* **16**, 1490–1502.

Horn, B. & Schunck, B. (1981). *Artif. Intell.* **17**, 185–203.

Jähne, B. (2010). *Opt. Photon.* **5**, 53–54.

Kak, A. C. & Slaney, M. (1988). *Principles of Computerized Tomographic Imaging*. IEEE Press.

Kamp, T. van de, Vagovič, P., Baumbach, T. & Riedel, A. (2011). *Science*, **333**, 52.

Table 6
Simulation parameters for the ESRF experiment in §5.1.

Category	Parameter	Value	
X-ray	Storage ring energy	6 GeV	
	Source type	Asymmetrical mutipole wiggler	
	Source size	–	
	Magnetic field	1.84 T	
	Electric current	100 mA	
	Beam mode	White beam	
	Energies	10 keV to 150 keV, 1 keV step	
	Filters	20 mm silicon	
	Detector	Type	Indirect
		Effective pixel size	5.5 μm
Objective lens		N/A	
Magnification		3.6 \times	
Numerical aperture		0.25	
Scintillator		YAG:Ce ($\text{Y}_3\text{Al}_5\text{O}_{12}$)	
Density		4.55 g cm^{-3}	
Thickness		100 μm	
Light yield		15 photons keV^{-1}	
Emission peak		550 nm	
Camera		Photron SA1	
Sensor size		1024 \times 1024	
Sensor pixel size		20 μm	
A/D factor		1 e^- per count	
Quantum efficiency		50% at 550 nm	
Dynamic range		800:1	
Frame rate		500 frames s^{-1}	
Sample		Sample holder	Two structured boron nitride plates
	Alloy	AlGe32 (wt%)	
	Particles sizes	Coarse globular particles of about 50 μm average diameter	
	Particles movement	Up to 3.9 mm s^{-1}	
	Particle clusters movement	1.6 mm s^{-1}	
	Source distance	–	
	Detector distance	10 cm	
Aquisition	Experiment	<i>In situ</i> radiography	
	Data processing	Flat-field correction	

Kirkland, E. J. (2010). *Advanced Computing in Electron Microscopy*. New York: Springer Science and Business Media.
 Klementiev, K. & Chernikov, R. (2014). *Proc. SPIE*, **9209**, 92090A.
 Malecki, A., Potdevin, G. & Pfeiffer, F. (2012). *Europhys. Lett.* **99**, 48001.

Martin, D., Fowlkes, C., Tal, D. & Malik, J. (2001). *Proceedings of the Eighth IEEE International Conference on Computer Vision (ICCV 2001)*, Vol. 2, pp. 416–423.
 Moosmann, J., Ershov, A., Altapova, V., Baumbach, T., Prasad, M. S., LaBonne, C., Xiao, X., Kashef, J. & Hofmann, R. (2013). *Nature (London)*, **497**, 374–377.
 Moosmann, J., Hofmann, R., Bronnikov, A. & Baumbach, T. (2010). *Opt. Express*, **18**, 25771–25785.
 Munshi, A. (2009). *2009 IEEE Hot Chips 21 Symposium (HCS)*, pp. 1–314. IEEE.
 Myagotin, A., Ershov, A., Helfen, L., Verdejo, R., Belyaev, A. & Baumbach, T. (2012). *J. Synchrotron Rad.* **19**, 483–491.
 Paganin, D., Mayo, S., Gureyev, T. E., Miller, P. R. & Wilkins, S. W. (2002). *J. Microsc.* **206**, 33–40.
 Papenberg, N., Bruhn, A., Brox, T., Didas, S. & Weickert, J. (2006). *Int. J. Comput. Vis.* **67**, 141–158.
 Pham, D. L., Xu, C. & Prince, J. L. (2000). *Annu. Rev. Biomed. Eng.* **2**, 315–337.
 Riederer, S. J., Pelc, N. J. & Chesler, D. A. (1978). *Phys. Med. Biol.* **23**, 446–454.
 Sanchez del Rio, M., Canestrari, N., Jiang, F. & Cerrina, F. (2011). *J. Synchrotron Rad.* **18**, 708–716.
 Santos Rolo, T. dos, Ershov, A., van de Kamp, T. & Baumbach, T. (2014). *Proc. Natl Acad. Sci.* **111**, 3921–3926.
 Shannon, C. E. (1949). *Proceedings of the IRE*, **37**, 10–21.
 Sun, D., Roth, S. & Black, M. (2014). *Int. J. Comput. Vis.* **106**, 115–137.
 Sypek, M. (2003). *Opt. Eng.* **42**, 3158–3164.
 Thompson, A. C., Attwood, D. T., Gullikson, E. M., Howells, M. R., Kortright, J. B., Robinson, A. L., Underwood, J. H., Kim, K.-J., Kirz, J., Lindau, I., Pianetta, P., Winick, H., Williams, G. P. & Scofield, J. H. (2001). *X-ray Data Booklet*, 2nd ed. University of California, Berkeley, CA, USA.
 Thompson, A., Llacer, J., Campbell Finman, L., Hughes, E., Otis, J., Wilson, S. & Zeman, H. (1984). *Nucl. Instrum. Methods Phys. Res.* **222**, 319–323.
 Van Nieuwenhove, V., De Beenhouwer, J., De Carlo, F., Mancini, L., Marone, F. & Sijbers, J. (2015). *Opt. Express*, **23**, 27975–27989.
 Walker, S. M., Schwyn, D. A., Mokso, R., Wicklein, M., Müller, T., Doube, M., Stampanoni, M., Krapp, H. G. & Taylor, G. K. (2014). *PLoS Biol.* **12**, e1001823.
 Williams, G. P. (2001). *X-ray Data Booklet*. University of California, Berkeley, CA, USA.
 Zabler, S., Ershov, A., Rack, A., Garcia-Moreno, F., Baumbach, T. & Banhart, J. (2013). *Acta Mater.* **61**, 1244–1253.
 Zápřažný, Z., Korytár, D., Mikulík, P. & Ác, V. (2013). *J. Appl. Cryst.* **46**, 933–938.