



Universiteit
Leiden
The Netherlands

Faster X-ray computed tomography in real-world dynamic applications

Graas, A.B.M.

Citation

Graas, A. B. M. (2026, February 4). *Faster X-ray computed tomography in real-world dynamic applications*. Retrieved from <https://hdl.handle.net/1887/4291923>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4291923>

Note: To cite this publication please use the final published version (if applicable).

Chapter 2

X-ray Tomography for Fully-3D Time-Resolved Reconstruction of Bubbling Fluidized Beds

This chapter is based on the article: Adriaan B.M. Graas, Evert C. Wagner, Tristan van Leeuwen, J. Ruud van Ommen, K. Joost Batenburg, Felix Lucka, and Luis M. Portela. “X-ray tomography for fully-3D time-resolved reconstruction of bubbling fluidized beds”. In: *Powder Technology* 434 (2024), p. 119269. ISSN: 0032-5910. DOI: [10.1016/j.powtec.2023.119269](https://doi.org/10.1016/j.powtec.2023.119269)

2.1 Introduction

In gas-solid fluidized beds, a mixture of gas and solid particles attains fluid-like characteristics. Fluidized beds are widely applied in chemical, pharmaceutical, and mineral industries, and laboratory-scale set-ups are used to investigate and validate fluidized beds in production applications. Gas is typically added from the bottom of a column or tank, and the *solids*, i.e., suspended particles, achieve different mixing regimes depending on particle size and gas velocity [47]. In bubbling regimes, a gas-solid fluidized bed has a homogeneous dense phase, and gas travels in dispersed voids, i.e., *bubbles*, upwards through the bed.

Experimental imaging of the gas-solid distribution of bubbling beds traditionally aims to find quantities such as sizes and shapes of bubbles, as well as the bubbles' solid contents. These are of paramount importance, e.g., to explain the catalytic behavior, or to describe the conversion of reactants. Since the solids mixture is opaque to visible light, several alternatives to optical imaging have been developed in the past decades [48]. Intrusive techniques, such as optical probes, use point measurements to arrive at local gas holdup, cord length and bubble velocity [49, 50]. Non-intrusive techniques, such as electrical capacitance tomography (ECT) [51], and X-ray computed tomography (CT), resolve the spatial distribution of the solids in a 2D slice or 3D volume. Positron emission particle tracking (PEPT) [52] measures the particle velocity in a volume. Of a more recent interest is also the study of the dynamic behavior of bubbles and solids. This, however, requires a technique for a fully spatio-temporally resolved gas-solid distribution. Bubble motion has a strong effect on the convection of particles, and, therefore, plays an important role in mass transport and advection of heat.

At Delft University of Technology, a set-up consisting of a triplet of conebeam sources (X-ray tubes) and 32-by-2 double-line detectors was introduced in 2010 [53]. The three sources and detectors were placed in an equilateral triangular geometry. This enabled tomographic reconstruction of bubbles in the horizontal plane of a 24 cm diameter column. Thanks to the double-line detectors, which operated at an effective 250 Hz, bubble velocities could be inferred. This subsequently allowed pseudo-3D bubble reconstruction through stacking the temporal evolution of the 2D slice [54]. Recently, the detectors have been upgraded to 1548-by-1524 pixel flat panels, and have enabled X-ray radiographic experiments to measure gas holdup in fluidized beds, cavitation, and bubble columns [55, 56, 57]. As our new detectors have the ability to capture the complete 3D volume in every frame, we are a step closer towards tomographic reconstructions that are both fully-3D and time-resolved.

In this work, we introduce and evaluate our new experimental method for tomographic reconstruction. This consists of (i) the stationary set-up, with three sources and detectors, (ii) a geometric calibration procedure and data processing method, and (iii) the Simultaneous Iterative Reconstruction Technique (SIRT). We examine its limitations, and present its capabilities. Since there exists a large amount of related techniques on tomography of multiphase flows, including fluidized beds, we start by giving a brief overview of work related to ours in Section 2.2. In Section 2.3, we introduce the set-up and calibration procedure. In Section 2.4, we describe a data processing method that is necessary to directly obtain tomographic reconstructions from the gas-solid distribution, and introduce SIRT. In Section 2.5,

we study reconstruction artifacts with numerical simulations, perform phantom experiments, and demonstrate its application in the acquisition of time-resolved reconstructions from Geldart B bubbling fluidized beds. In Section 2.6, we conclude with a reflection on the possibilities and limitations of the new technique.

2.2 Related work

X-ray set-ups for imaging dynamic processes are common in engineering, chemical and medical sciences [58]. Single-source systems allow acquisition with a fast-rotating component, and may use a 4D (i.e., 3D + time) motion-corrected or sparse-angular technique for reconstruction. When measurements must be acquired in a shorter time than what is achievable by a gantry, often referred to as *ultrafast* imaging, multi-source multi-detector systems are considered [59, 60]. Example applications are combustion processes, multiphase pipe flows, and cardiography, and set-ups range from a dual source and detector pair on a gantry [61] to 29 pairs on a stationary circular track [62].

For the purpose of imaging dynamics in fluidized beds several techniques have been proposed. Multi-source X-ray set-ups, on the basis of line detectors are an established technique [53, 59]. In early gamma-ray CT and radioactive particle tracking [52], for example, this enabled a time-averaged gas distribution or time-averaged velocity field [63, 64]. In recent decades, several experimental methods have been developed to image dynamics in fully-3D [65], although for some techniques this limited the temporal resolution. Slow-moving solids, for instance, can be analyzed in 3D using X-ray radiography with Particle Image Velocimetry (PIV) [66]. To our knowledge, currently three set-ups have been used for fully-3D time-resolved tomography of fluidized beds: (i) an electron beam X-ray CT scanner at Helmholtz-Zentrum Dresden-Rossendorf, achieves a high spatial resolution with a 500 Hz temporal resolution in an 8 mm thick cross-sectional volume [67]; (ii) a fully-3D time-resolved MRI technique has been demonstrated with single bubble injection at 140 Hz [68]; (iii) electrical capacitance volume tomography (ECVT) has been applied in a $20 \times 20 \times 20$ -sized volume of size 5 cm width and 10 cm height at 80 Hz [69].

X-ray is a useful modality for imaging large scale multiphase flows in process engineering sciences, although it requires an environment for working with ionizing radiation [3, 5]. Magnetic Resonance Imaging (MRI) is a costly alternative, and can image particles that contain MR-sensitive nuclei in small volumes [70]. It delivers good contrast and recently achieved good temporal resolution [71]. Other modalities that have been successfully employed in 3D are the electrical resistance and capacity tomographic modalities, ERT and ECT/ECVT. They are inexpensive soft-field techniques that are able to reach a very high temporal resolution [72]. However, increasing their low spatial resolution is difficult due the limitations imposed by the placement of additional electrode pairs [51, 65, 70], and the unfavourable properties of the underlying mathematical image reconstruction problem [73]. The reader is referred to [65] for an overview and comparison of modalities.

In comparison to the existing 4D techniques, the main advantage of our set-up is that it enables imaging large scale flows at all heights simultaneously, with a high vertical resolution

and at a high framerate. It allows following individual bubbles along their entire paths upwards through a fluidized bed, and it therefore does not require statistical quantities for the analysis of bubble dynamics. While the readout bottleneck of CMOS detectors limits the attainable framerates, the most prominent limitation of our set-up is the effective resolution in the horizontal plane, which is surpassed by the previously mentioned electron-beam and MRI set-ups. Its potential resolution is 1548-by-1548 voxels, the number of pixels in a detector row, but the associated reconstruction problem is exceptionally difficult to solve: all n^2 voxels in the n -by- n horizontal plane must be resolved from only $\sim 3n$ values associated with their X-ray projections on the detectors. While the spatial localization of measured X-rays is very precise, compared to the ERT/ECVT and magnetic modalities, the quality of our set-up depends on the ability of the reconstruction algorithm to explain the measurement data. This ability will be tested experimentally in the forthcoming sections, and we will reflect on the development of new reconstruction algorithms in the conclusion.

2.3 Material and methods

2.3.1 Experimental set-up

Our set-up, aimed at tomographic reconstruction, consists of three continuous X-ray sources and three CMOS detectors, arranged in an equilateral triangle (see Fig. 2.1). Fig. 2.2 shows a photograph of the set-up with a polymethyl methacrylate (PMMA) cylinder in the center. A rotation table is used for calibration, and air is supplied at the bottom of the cylinder through a bundle of needles. The sources and detectors can be repositioned to accommodate the experiment. Detectors are synchronized using an external trigger. With the detectors configured to operate in a region of interest (ROI), i.e., by disabling a segment of the detector rows, framerates up to 200 Hz at 1548×100 pixels can be achieved. Before preprocessing (Section 2.4.1), we subtracted darkfield images, i.e., measurement frames with the sources off, and performed dead-pixel corrections on our measurements [74]. Dead pixels are defected pixels that are (partially) non-responsive to X-ray radiation, and appear along certain detector rows and columns in our data. They are filled in by an unweighted average over their direct non-defected neighbours.

For the experiments in Section 2.5, we placed the detectors relatively close to the column, and image about 20 cm of bed height. We cropped to 1548×550 pixels, which resulted in a framerate of about 65 Hz, or, equivalently, an exposure time of 15 ms. The PMMA cylinder contained spherical polystyrene particles with a diameter of 560 μm (Geldart B type) [53]. The minimum fluidization velocity of this material has experimentally been determined to be 15 cm/s. Set-up specifications and further details are summarized in Table 2.1.

2.3.2 Calibration

For an accurate tomographic reconstruction (cf. Section 1.1), a precise and coherent prescription of the geometry is essential. The geometry consists of all positions of sources and detectors, as well as the orientation of each detector plane, in a single coordinate system. Per source-detector pair, this can be encoded in nine parameters: two 3D positions and three




Figure 2.1: Top view schematic of the X-ray set-up (not to scale).




Figure 2.2: Photograph of the X-ray set-up with filled PMMA cylinder. Red arrows point to detectors, the yellow arrow to a polystyrene phantom (Section 2.5.2), and the blue arrow to the gas inlet. Sources are not in view.

X-ray source	Yxlon Y.TU 160-D06
X-ray source voltage	Max. 150 kVp
X-ray detector	Teledyne Dalsa Xineos 3131
Detector framerate	22 to 200 Hz
Detector resolution	1548×1524 pixels, 0.2 mm/pixel
Detector surface area	307×302 mm
X-ray source voltage	120 kVp
X-ray anode current	1.5 mA
Detector ROI	1548×550 pixels
Detector framerate	65 Hz
Object-source	935, 947, 937 mm
Object-detector	273, 268, 294 mm
Column material	Polymethyl methacrylate
Column diameter	50 mm inner, 60 mm outer
Particle material	Polystyrene
Particle size	560 μm (Geldart type B)
Minimum fluidization	15 cm/s

Table 2.1: Hardware specifications (top part) and configuration for the experiments in this work (bottom part).

rotation angles. For brevity, we denote all geometry parameters together with the tuple $\psi \in \mathbb{R}^{3 \cdot 9}$. The object-source and object-detector distances, listed in Table 2.1, reflect the outcome of a careful manual measurement. They denote the distances between the column center and X-ray source, and column center and detector midpoint. To further increase the accuracy of ψ , we developed a workflow and methodology for a data-driven calibration procedure that allows for the simultaneous calibration of multiple sources and detectors. Compared to a manual measurement, this is both fast and flexible, and can accommodate both intentional and unexpected modifications to the set-up. In our experience, having an accessible calibration procedure significantly improves the quality of reconstructions.

The calibration is started by placing a marker object in the view of the three detectors. The marker object is sized similarly to a fluidized bed column, and contains a set of metal markers $\mathcal{I}$. The unknown coordinates of the markers are denoted by $\{m_i \in \mathbb{R}^3\}_{i \in \mathcal{I}}$. Using the rotation table, the marker object is then imaged from a small number of angles n_α . For three detectors, each marker $i \in \mathcal{I}$ is projected $3n_\alpha$ times, and thus yields a projection vector $p_i \in \mathbb{R}^{2 \cdot 3n_\alpha}$. The vector is obtained by manually annotating the positions of the markers in the measurements.

To solve for ψ , a nonlinear least-squares optimization objective, Eq. (2.1), is set up. Here, the mapping $P_\psi: m_i \mapsto p_i$ encodes the geometric point-projection for multiple rotation-table positions. The Levenberg-Marquardt optimizer (see, e.g., [75]) is then used to solve the objective for all markers and ψ simultaneously:

$$\arg \min_{\{m_i\}_{i \in \mathcal{I}}, \psi} \sum_{i \in \mathcal{I}} \|P_\psi(m_i) - p_i\|_2^2. \quad (2.1)$$

Note that, for this procedure to work, it is not necessary to know the marker coordinates ahead of time. Since three or more rotation angles provide an abundance of projections, the marker positions $\{m_i \in \mathbb{R}^3\}_{i \in \mathcal{I}}$ may be optimized jointly with the sought parameters ψ . This conveniently allows the construction of marker objects for different column dimensions or detector regions of interest.

2.3.2.1 Data-driven calibration Our marker-based calibration (Section 2.3.2) allows the inference of geometry parameters ψ from the projection of markers. The parameters are encoded using the 3D positions of the sources and detectors, and the extrinsic Euler angles of the orientations of the detectors planes. Before our data acquisition, we scanned a fabricated PVC pipe with nine uniquely identifiable metal markers points. The pipe was positioned on a rotation table, and after imaging a full rotation, three angles were selected for annotation. We picked the three angles such that they are far apart, and such that the markers on the images do not overlap. This led to 9 projections for each marker, and a total of 81 marker projections.

The subsequent nonlinear optimization problem (see Eq. 2.1) was solved with the Levenberg-Marquardt [75] implementation in *SciPy*. The gradient of P_ψ was computed numerically. Since the distances between markers were not given *a-priori*, the found geometry parameters were determined up to a scaling factor. To see this, note that stretching the geometry (moving sources and detectors closer or farther) leads to the same projections. The known inner diameter of the PVC pipe was used to correct the scaling. For this, we first reconstructed the pipe in high resolution, using a full-angular scan from a rotation table. An ellipse was subsequently fitted in the horizontal reconstruction plane to yield the reconstructed diameter of the pipe. The comparison between the reconstructed diameter and the known diameter was used to determine the scaling factor.

2.4 Theory

2.4.1 Measurement principle

Data of fluidized beds is described by the measurement principle explained in the introduction, section 1.1. The aim of a tomographic algorithm is typically to infer μ , using Eq. (1.4) for many detector pixels and additionally formulated assumptions on $\mu(\eta)$. For a continuous medium, without sharp variations in its composition, μ can be assumed as a continuous function of space and time. However, in a fluidized bed, there are sharp variations at the interface between the particles and the gas, and, also there can exist local sharp variations along a ray, depending on the local distribution of the particles. Therefore, μ should be seen as a local-averaged quantity, averaged over a small volume and short exposure time, associated with the space and time resolution of the measurements

In our case, the imaging object consists of a PMMA cylinder, and a bubbling fluidized bed – a mixture of suspended solids and gas. The quantity of interest is not μ but the distribution of particles and gas. Similarly to μ , this can be expressed in terms of a local-

averaged gas fraction α , i.e., the local percentage of the volume that is occupied by the gas-phase (see, e.g., [76] and [77] for a formal definition). Note that α can vary between a minimum value, greater than zero, corresponding to the densest packing of particles, and a maximum value of one, corresponding to a pure gas. Alternatively, α can be normalized by defining a bubble fraction, b , with $b = 0$ corresponding to the average gas fraction in the *solids* (homogeneous dense phase) and $b = 1$ corresponding to the gas fraction in the *bubbles* (pure gas). I.e., b is defined as

$$b = \frac{\alpha - \bar{\alpha}_{\text{solids}}}{1 - \bar{\alpha}_{\text{solids}}}, \quad (2.2)$$

where $\bar{\alpha}_{\text{solids}}$ denotes the average gas fraction in the solids. Assuming that the attenuation in air is negligible, the relation between b and μ is given by

$$b = 1 - \frac{\mu}{\bar{\mu}_{\text{solids}}}, \quad (2.3)$$

where $\bar{\mu}_{\text{solids}}$ denotes the average attenuation in the solids.

Referencing method A referencing method is a technique to remove certain features, e.g., the PMMA cylinder, from a projection (Eq. (1.4)). When μ_{ref} denotes the attenuation of the unwanted features in the object, the desired projection should be from $\mu - \mu_{\text{ref}}$. Substitution in Eq. (1.4) gives

$$\begin{aligned} \int_l [\mu(\eta) - \mu_{\text{ref}}(\eta)] d\eta &= \int_l \mu(\eta) d\eta - \int_l \mu_{\text{ref}}(\eta) d\eta \\ &= \log(I_0/I) - \log(I_0/I_{\text{ref}}) \\ &= \log\left(\frac{I_{\text{ref}}}{I}\right). \end{aligned} \quad (2.4)$$

This shows that the flatfield, I_0 in Eq. (1.4), can be replaced by a reference measurement I_{ref} , and that this removes the unwanted features. In Eq. (2.4), Beer-Lambert's law was applied twice, and I_0 dropped out of the equations under the assumption that I and I_{ref} were obtained in similar experimental conditions.

Projections of the bubble fraction Using the referencing method, we now derive a linear relation between b and projections of b , which we denote with y . A tomographic algorithm working with y therefore reconstructs the bubble fraction. Let I_{full} be a reference of the PMMA cylinder filled entirely with solids, i.e., $b = 0$ everywhere in the column, and

let I denote a measurement frame of an inhomogeneous bubbling fluidized bed. Then

$$\begin{aligned}
 y &= -\frac{1}{\bar{\mu}_{\text{solids}}} \log \left(\frac{I_{\text{full}}}{I} \right) \\
 &= -\frac{1}{\bar{\mu}_{\text{solids}}} \int_l [\mu(\eta) - \mu_{\text{full}}(\eta)] d\eta \\
 &= -\frac{1}{\bar{\mu}_{\text{solids}}} \int_l [(1 - b(\eta)) \bar{\mu}_{\text{solids}} - \bar{\mu}_{\text{solids}}] d\eta \\
 &= \int_l b(\eta) d\eta.
 \end{aligned} \tag{2.5}$$

Since referencing removes the contributions of the PMMA cylinder, we have omitted it from the equation. Fig. 2.3 displays an example of I and I_{full} , and panel (a) of Fig. 2.9 shows the corresponding outcome, y . In the following, we explain how $\bar{\mu}_{\text{solids}}$ and I_{full} can be obtained experimentally.

Average attenuation in the solids Let ℓ_{int} be the intersection of a ray ℓ with the interior of a homogeneously filled PMMA cylinder. Then

$$\frac{1}{|\ell_{\text{int}}|} \int_{\ell_{\text{int}}} \mu_{\text{solids}}(\eta) d\eta \tag{2.6}$$

describes the average attenuation along ℓ_{int} , and is, therefore, an estimate of $\bar{\mu}_{\text{solids}}$. In our set-up, we use the measurements of I_{full} , with a measurement of the empty PMMA cylinder, I_{empty} , acquired in the same experimental conditions (see Fig. 2.3). Applying the referencing procedure gives

$$\begin{aligned}
 \log \left(\frac{I_{\text{empty}}}{I_{\text{full}}} \right) &= \int_l [\mu_{\text{full}}(\eta) - \mu_{\text{empty}}(\eta)] d\eta \\
 &= \int_{\ell_{\text{int}}} \mu_{\text{solids}}(\eta) d\eta \approx |\ell_{\text{int}}| \cdot \bar{\mu}_{\text{solids}}.
 \end{aligned} \tag{2.7}$$

By taking the ray through the center of the column, for which $|\ell_{\text{int}}|$ is known, we obtain an estimate of $\bar{\mu}_{\text{solids}}$.

Full column references for fluidized beds Eq. (2.5) requires a reference I_{full} , i.e., a cylinder with a bed that attenuated with $\bar{\mu}_{\text{solids}}$. Traditionally, the reference I_{full} is obtained with a packed bed, without any gas injection. However, this approach can introduce noise due to local sharp variations in the medium caused by the packing structure of the particles (see, e.g., [53]). An alternative is to use a homogeneously fluidized bed, where the motion of the particles reduces the noise, since the signal intensity is averaged over the exposure time. Moreover, using a fluidized bed leads to a more accurate determination of $\bar{\alpha}_{\text{solids}}$, because its value decreases due to the expansion of the bed, as shown in Fig. 2.4. The histogram of 0 cm/s, a packed bed, depends only on the mean attenuation and noise, whereas fluidized beds have a shifted distribution due to the expansion of bed.




Figure 2.3: Raw measurements in a 550×160 pixel detector region, showing (a) two spherical phantoms (diameters equal to 10 mm and 23 mm), (b) a full column reference (not fluidized), and (c) an empty column. The red and dashed rows are used to display noise statistics in Fig. 2.11.




Figure 2.4: Detector count statistics for 1000-1200 measurements of a central pixel, for different superficial gas velocities of a fluidized bed. Vertical lines indicate the modes.

Measuring I_{full} in a fluidized bed is not straightforward, as the measurement during fluidization could contain bubbles. One possible strategy is to measure the column at a gas velocity right before the formation of bubbles. While this expands the bed somewhat, it can only be used for lower gas velocities, and can be challenging to achieve in practice. For our dynamic experiments in Section 2.5.3, we have therefore applied a heuristic approach, which uses the whole time series of projections I_t from an experiment to estimate I_{full} . The heuristic retrieves I_{full} by using a pixelwise mode of the data. To see this, note that a direct average or median value over all projection images does not yield an image of a full column, as bubbles increase the average X-ray intensity. Yet, when considering a single pixel of the detector, most values show a background (i.e., an X-ray that has passed through the bed without bubbles), and perturbations due to bubbles occur only in a relatively small number of frames. The pixelwise mode corresponds to the most commonly observed value, and therefore to the background I_{full} . It can be extracted by building a histogram of observed pixel values, and selecting the bin with the highest count, i.e., the *largest mode* of the distribution of I_t . We show this in Fig. 2.4, where the mode is computed using the maximum of an interpolating polynomial of degree 20 in the central detector pixel. Before reconstruction, we compute I_{full} for all pixels and detectors. Estimating I_{full} in this way is particularly useful when the bed expands significantly during fluidization, for instance with Geldart A type particles. Another advantage is that it does not involve acquisition of data using a secondary experiment, and therefore eliminates a possible error source. The disadvantage of such heuristic is that a reliable statistic requires a large number of time samples.

2.4.2 X-ray tomographic reconstruction

We refer the reader to Chapter 1 for the introduction on reconstruction techniques. Algebraic techniques are often used with sparse-angular or noisy data (see chapter 1 or [4]). SIRT suits our three-angle set-up, therefore, well. Analytic techniques, such as the filtered-backprojection (FBP), are faster, but cannot produce sufficient image quality in such an exceptional sparse-angular case. SIRT, in addition, serves as a baseline for more involved techniques with additional regularization and constraints.

The most straightforward way to obtain time-resolved reconstructions from fluidized beds is to run SIRT independently for each 3-tuple of projections, yielding a sequence of 3D reconstructions. Using the GPU-accelerated *ASTRA Toolbox*, reconstruction takes about a minute per timeframe. Alternatively, methods that reconstruct multiple frames simultaneously can incorporate spatio-temporal constraints, which can improve the image quality considerably. These constraints, however, often put more requirements on memory and computation time [78].

2.4.3 Visualization

In our results, we show measurements and reconstructions in a single format (see, for example, Fig. 2.5). In panel (a), we show a region of interest in the projections y , the result of Eq. (2.5). In panel (b), we display an iterate x of SIRT, Eq. (1.8), using a 3D density plot where the opacity is linearly increased from 0% to 20%. A horizontal and vertical cross-

sectional slice of x are shown in panel (c). Dashed lines are drawn between each source and the center of its corresponding detector.

2.5 Results

Tomographic reconstructions of fluidized beds provide, next to visual inspection, a better insight in bubble morphology and bubble dynamics than raw projection images do. While SIRT is commonly used for reconstruction, the unusually low number of projection angles in our set-up may cause distinct image artifacts [79]. These artifacts can usually be linked to parameter choices, the positions of bubbles, or the presence of noise. In Section 2.5.1, we investigate three commonly encountered artifacts through numerical simulations. In Section 2.5.2, we perform an experimental validation using spherical phantoms to mimic the bubbles in a fluidized bed. Section 2.5.3 presents experimental results of a Geldart B bubbling fluidized bed.

2.5.1 Numerical simulations

Numerical set-up In our simulations, we inspect reconstructions of solitary bubbles, of multiple bubbles, and of a bubble with noise. We use a simulation model that closely resembles our experimental set-up. In particular, we use the geometry of sources and detectors that is obtained after calibration of our set-up (Section 2.3.2). For each simulation, a ground truth is first constructed as a discretized volume. The projections, y , are then generated by the linear forward projection model for the conebeam geometry, i.e., the operator A in Eq. (1.5), via the implementation available in ASTRA Toolbox [80]. For reconstruction, we discretize on a $(300, 300, 1303)$ grid, with an isotropic voxel size of $550/300 \approx 1.83$ mm. SIRT is run with a masked area (cf., Eq. (1.8)) using a discretized exact cylinder of 50 mm diameter. In our experiments with real data, we use a larger mask to compensate for a slight tilt of the column. In all other aspects, the simulation model is equivalent to the reconstruction model in Section 2.5.3, and therefore enables experimentation with numerical data without the effects of noise, preprocessing, or calibration error. In the following, SIRT is run for 2000 iterations, and we do not apply any post-processing routines.

Solitary bubbles To show artifacts in the simplest circumstances, we start by studying *solitary* bubbles, i.e., bubbles for which a projection does not overlap with any of the other bubbles. In our experiments, this is frequently the case when no other bubbles travel at the same height. When modeled as spheres, solitary bubbles may be reconstructed to surprisingly high accuracy. Fig. 2.5 shows that three projections of a full-gas spherical bubble are sufficient for an accurate reconstruction; only minor perturbations are visible at the bubble interface in the directions of projection. Of notable interest is the horizontal cross-sectional plot, where, despite the three-angle view, the bubble boundary is almost perfectly round.

We found that only in the case of noiseless data and $b = 1$, SIRT recovers the bubble to this level of accuracy. In many other situations, bubbles are reconstructed as irregular convex hexagons, or have a cupping-like distortion of the bubble fraction. An example is given in




Figure 2.5: Reconstruction of a simulated sphere with a diameter $D = 20$ mm and a bubble fraction $b = 1.0$. The visualization format is explained in Section 2.4.3.




Figure 2.6: Reconstruction x of a spherical bubble x^* with $b = 0.95$ and diameter $D = 20$ mm. The bubble in panel (b) has a smoothed interface and is less affected by reconstruction artifacts.

2 panel (a) of Fig. 2.6, where $b = .95$ rather than $b = 1$. Fortunately, this does not always pose a problem. The artifact only shows in the horizontal plane, where limited information is present due to the sparse-angular resolution. The severity of this artifact furthermore depends on the bubble fraction, the noise on the projections, and the effect of the SIRT constraints, i.e., $0 \leq x \leq 1$ in Eq. (1.8). In particular, artifacts are less pronounced when there is some spatial variation in the bubble fraction. A bubble with a smoothed interface, shown in panel (b) of Fig. 2.6, is less affected by hexagon and cupping artifacts. This shows that the use of additional regularization and constraints has the potential to significantly improve the quality of the results.

Bubble-bubble reconstruction interference Reconstructions become more complicated when projections contain overlapping bubbles. A reconstruction of multiple bubbles is different from the addition of reconstructions of solitary bubbles: while the X-ray model is linear (Eq. (1.5)), the constraints in Eq. (1.7) lead to a non-linear relationship between y and x^* . Fig. 2.7 illustrates this with a simulation in which, on Detector 3, a smaller bubble overlaps with a larger bubble. Intuitively speaking, SIRT relies on Detector 1 and Detector 2 to separate the bubbles spatially, however, cannot do that uniquely. In the figure, intensities from the sides of the bigger bubble are therefore erroneously displaced into small patches, and both bubbles show hexagon artifacts. The problem is more pronounced with multiple overlapping bubbles, but the magnitude of the artifacts depend on the shapes, distances, and bubble fractions.

Geometry of noise artifacts The last simulation clarifies the origin of a star-shaped streaking artifact, a phenomenon that was previously observed for our double line detector set-up [53]. We confirm that one cause of this artifact is noise, through a simulation of additive Gaussian noise on the projections. In experimental data, other error sources, i.e., mismatches between Ax and y in Eq. (1.5), could also lead to similar artifacts. Fig. 2.8 shows the reconstruction of a spherical bubble, with a diameter $D = 10$ mm, a bubble fraction $b = 1.0$, in which a zero-mean Gaussian noise with variance equal to $2.6 \cdot 10^{-2}$ was added to the projections. Panel (a) shows the average reconstruction, panel (b) the bias, and panel (c) the variance, using 100 reconstruction samples. The figure shows increased variance along projection lines and around the bubble, which is again an effect of non-linear reconstruction: in a linear method, the spatial distribution of the variance would not depend on x , and hence a variance plot would not show a bubble. This does, however, not suggest using one: a linear method would have higher variance. At the end of the following section, we quantify the variance of the noise in our set-up.

2.5.2 Phantom experiments

We proceed with an experimental validation of our set-up. First we repeat the bubble-bubble reconstruction of last section, using two spherical phantoms. We then validate the accuracy of velocity estimations, and compute noise statistics from our measurements.




Figure 2.7: Reconstruction of two simulated bubbles with diameters $D_1 \approx 1.23$ cm and $D_2 \approx 0.83$ cm (74 and 50 voxels, respectively), for which the projections overlap on Detector 3. The visualization format is explained in Section 2.4.3.




Figure 2.8: Reconstruction (a), bias (b) and variance (c), from a simulated spherical bubble, with a diameter $D = 10$ mm and a bubble fraction $b = 1.0$. Zero-mean Gaussian noise with a variance of $\sigma^2 = 2.6 \cdot 10^{-2}$ has been added to the projections. (b) and (c) are computed using 100 samples.




Figure 2.9: Reconstruction of two polystyrene phantoms, with diameters $D_1 = 23$ mm and $D_2 = 10$ mm, that overlap on the second detector. The reconstruction, panel (b), has been rotated to match Fig. 2.7. The detector row marked in blue is used to compute noise statistics in Fig. 2.11. The visualization format is explained in Section 2.4.3.

Static bubble phantoms Fig. 2.9 displays a two-ball phantom experiment that repeats the numerical simulation on bubble artifacts (Fig. 2.7). Here, two expanded-polystyrene foam balls are imaged in a 50 mm diameter column, filled with Geldart B particles (specifications in Table 2.1), without the inflow of gas. An example of a phantom is displayed in Fig. 2.2. The phantoms are spherical, feature sharp boundaries, and have negligible X-ray attenuation. They therefore model idealized bubbles in a static snapshot of a fluidized bed. The column is rotated, using a rotation table, so that the smaller, 10 mm phantom, overlaps with the larger, 23 mm phantom on Detector 2. Without fluidization, the static packing structure of the solids adds an additional noise-like pattern to I and I_{full} . We will further discuss this in Section 2.5.3.

In this experiment, Fig. 2.9, the larger bubble in the horizontal cross-sectional plot of x displays the hexagon artifact, as well as a slight decrease in intensity. The spherical shape is still recovered well, thanks to the high resolution of the detectors in the vertical dimension. In the 3D plot, a few low-intensity patches are observed between the two bubbles, although not as pronounced as in the numerical example. The cross-section cuts in Fig. 2.9 show the appearance of streaks, which are associated with the noise. In essence, comparing figures 2.9 and 2.7, we see that the reconstruction of the phantom-bubble is a more fuzzy version of the numerical bubble, and Fig. 2.9 can be interpreted as a representative signature of an actual (spherical) bubble.

Bubble diameter estimation To explore our set-up's ability to reconstruct spheres, we attempt to recover the known diameters of the 23 mm and 10 mm phantoms, as well as two simulated counterparts, using a template-matching procedure. Letting $B \in \mathbb{N}_d^{N_d^3}$ denote a binary template of a 3D sphere with diameter d , we compute a matching score between the template and a reconstruction volume x with

$$\text{score} = \frac{1}{N_d^3} \sum_{i=1}^{N_d^3} \left(x_i B_i + (1 - x_i)(1 - B_i) \right). \quad (2.8)$$

The first term in the summation matches the sphere, consisting of 1.0-valued voxels, whereas the second term matches the background, with 0.0 values – this avoids matching a template of a small sphere with a large bubble. Eq. (2.8) simplifies to the Dice similarity coefficient when x is binarized, and can take values between 0.0 and 1.0. The score is 1.0 in the case of a perfect match with the template. As $0 \leq x \leq 1$, we can omit binarization, which avoids the subjective choice of a segmentation threshold parameter, and lets gray values contribute to the score with a value that is proportional to their distance to the template.

Fig. 2.13 shows template-matching results for both the phantom and simulation experiment. A spherical template with diameters ranging between 1 and 35 mm is translated through the volume. The best score encountered during this translation is plotted in panel (a) as a function of diameter. In the noiseless simulation, bubble diameters are recovered to a high accuracy, which suggests that the impact of hexagon artifacts is limited in this experiment. In the phantom experiment, the 10 mm phantom is underestimated as 9.1 mm

and the 23 mm phantom is overestimated as 23.5 mm. Their respective scores of 0.70 and 0.86 reflect a substantial impact of noise. Panel (b) shows a volume rendering of the best matching templates together with the reconstructed images of the phantoms, and panel (c) shows an intersection in the horizontal plane with the simulations in the lower half. Note that in panel (c) diameters are plotted smaller than found, as the intersection plane lies in-between the two phantom centers.

The experiment is an indication that the technique has the potential to quantify bubble volumes within a reasonable error margin, likely due to the information that is present in the bubbles' 3D neighborhoods. To explore how reliably the shapes of real bubbles can be reconstructed, phantom experiments with more complex shapes (e.g., concave or ellipsoid forms) will be designed in a future study.

Moving phantom Fluidized beds are reconstructed by computing x_t for a 3-tuple of projections at each time t . However, to extract dynamic characteristics of bubbles, such as the bubble velocity, measurements need to be taken at prescribed (equidistant) intervals and with synchronized detectors. This aspect of the set-up is tested end-to-end, by performing a time-resolved reconstruction of a moving phantom. First, a 25 mm glass phantom is pulled upwards through an empty column, using a motor-driven traverse. Then, in each timeframe t , a position $p_t \in \mathbb{R}^3$ of the phantom is inferred by fitting a 3-dimensional sphere to the reconstruction x_t . The experiment was repeated for traverse speeds of 62 and 125 mm/s and for three different horizontal starting positions of the phantom.

Fig. 2.10 shows the 125 mm/s experiment, with the phantom starting in the center of the column. The plot confirms an accurate upward speed of the phantom. Using a least-squares interpolation of p_t , the phantom was estimated to have a speed of 128 mm/s, an overestimation of 2.4%. This error was consistent for all other starting positions of the phantom. The results of the 62 mm/s experiment were similar, now with a overestimation of 1.6% in all starting positions. Since the error correlates with the traverse speed, but not with the phantom starting position or travel distance, we expect the error to be due to an inaccuracy in the experimental equipment.

Measurement noise The local sharp variations in the medium, associated with the distribution of the particles, leads to high noise. In a fluidized bed, due to the motion of the particles, the averaging of the signal over the exposure time reduces the noise. However, to capture the fast dynamics of fluidization, detectors need to operate at short exposure times. As a consequence, the projection data still contains high levels of measurement noise, and our reconstructions degrade accordingly (as shown in Fig. 2.8). Measurement noise can be attributed to stochastic fluctuations due to photon scattering and absorption, as well Poisson-distributed detector noise [3]. To quantify how well bubbles can be observed, we compute noise statistics over 1,000 measured frames of the two static bubble phantoms. We take the pixels in a single detector line, displayed in red in panel (a) of Fig. 2.3, and in blue in panel (a) of Fig. 2.9.




Figure 2.10: Least-square interpolation of positions p_t , obtained after presetting the traverse to 125 mm/s, and the detectors to 65 Hz.




Figure 2.11: Panel (a) shows the line profile of raw measurement data of the marked detector row in panel (a) and (b) of Fig. 2.3. The computed noise statistics are shown in red (std.dev. in strong red, min/max values in light red). Panel (b) shows the same detector row, now after preprocessing, see the marked row in panel (a) of Fig. 2.9. Noise statistics of y are in blue.

Fig. 2.11 plots I_{full} , I , and y , and their means and standard deviations. For I , i.e., the raw data, detector counts range from 6,500 to 9,000 in the solids. The corresponding standard deviations vary from about 130 counts, near the inner wall of the PMMA cylinder, to about 110 in the column center. The standard deviation increases in bubble voids, and depends on the size of the void. This is expected, as the variance of counting noise depends on its mean value, which is proportional to the photon flux, and thus inversely proportional to attenuation.

Panel (b) of Fig. 2.11 shows the impact of the random fluctuations on I and I_{full} after computing y from Eq. (2.5). The signal-to-noise ratio for the signal y_i in a pixel i is computed as

$$\text{SNR}_i = 10 \log_{10} \left(\frac{\mu_i^2}{\sigma_i^2} \right), \quad (2.9)$$

where μ_i is the mean and σ_i the standard deviation. This yields an SNR value of 24 dB for the 23 mm phantom, and an SNR of 12 dB for the 10 mm phantom, both measured at the center of the phantom. This confirms the intuition that we should be able to image larger bubbles with a higher confidence. It must be noted that the SNR values are a set-up dependent indication of image quality. Noise varies with the detector framerate, source current, column diameter, and set-up geometry. Furthermore, since bubbles are observed from multiple angles, we expect better signal-to-noise ratios in reconstructed images. An explicit variance analysis of noisy reconstructions, e.g., as conducted for a parallel-beam geometry using the filtered-backprojection [3], cannot easily be applied to our custom geometry and particular image reconstruction method. Unlike the circular parallel beam case, with three source-detector pairs, the voxelwise variance in our reconstructed volume may not be radially symmetric. Moreover, with box constraints in the iterative method, noise will be nonlinearly propagated over the 3D reconstruction, leading to unclear data-dependent transformations to the per-voxel noise distributions. The reader is referred to Chapter 9.8 of Buzug [3].

In our multi-source set-up, one contribution of noise is expected due to cross-scattering of photons, i.e., photons that originate from the non-facing sources on the left and right sides of each detector, and have scattered under an angle of approximately 120° , primarily via Compton scattering (for a 150 kVp X-ray spectrum). To quantify this error, we measure detector counts with the facing source off, and with one or two of the non-facing sources on, and subsequently average over all pixels in 630 frames and all symmetrical source-detector combinations. The detector counts, listed as S and D in Table 2.2, are the contributions due to scattering and darkfield current (Section 2.3.1). Their combined value can reach up to 14% of the measured signal (i.e., $I \approx 6500$, see Fig. 2.11). The scattering of the particulate material has a relatively small contribution compared to the PMMA column, and the dynamic variation of photon scattering during fluidization is therefore expected to be small. In a reconstruction computed with scattering and darkfield correction, we noted only subtle changes in voxel values, between -0.01 and 0.01 . This did not lead to a structural improvement in the background, bubbles, or artifact shapes. Removing a uniform identical bias from both I and I_{full} in Eq. (1.4), with $I_{\text{full}} < I$, leads only to a slightly stronger

	One source			Two sources		
	<i>D</i>	<i>S</i>	<i>Rel.</i>	<i>D</i>	<i>S</i>	<i>Rel.</i>
No column	324	135	7%	324	266	9%
Empty column	324	252	9%	324	501	13%
Full column (not fluidized)	324	302	10%	324	603	14%

Table 2.2: Photon scattering originating from the X-ray sources not facing a detector, e.g., photons originating from source 2 and 3 that are scattered into detector 1, cf. Fig 2.1. Listed are detector counts measured when either a single or both of the non-facing sources are switched on. The values *S* (scattering) and *D* (dark current) are averages over all detector pixels and all possible source-detector combinations. *Rel.* denotes the percentage of $D + S$ with respect to the minimum signal, taken as $I := 6500$ detector counts.

contrast between the intensities. Hence, the result suggests that the stochastic fluctuation of noise has a larger impact on the quality of reconstruction than the scattering bias. A future study to mitigate the impact of scattering, using advanced noise correction techniques [58], is warranted.

2.5.3 Geldart B bubbling fluidized beds

In the following experiment, the Geldart B bed is brought to the fluidization regime using a superficial gas velocity of 17 cm/s. This combination of particle size and gas velocity leads to distinctive bubbles, and is therefore suitable as a demonstration of our technique. At lower superficial gas velocities, bubbles are smaller, whereas at higher velocities, the flow transitions into the slugging regime. X-ray projections are recorded above the gas inlet (see Fig. 2.14 for a single frame). SIRT was run with 200 iterations as stopping criterion to prevent fitting to noise. Detector framerates were set to 65 Hz, i.e., an exposure time of 15 ms. This is the maximum framerate, which depends on the detector ROI. It can be increased by changing the source-detector distances or by imaging with a smaller column – although the latter would increase bubble interactions with the column wall. The supplementary video (Fig. 2.16) in the publication associated with this chapter animates frames 450 to 650, corresponding to three seconds of the experiment. The interval contains differently shaped and sized bubbles, as well as a variety of dynamical behaviours.

Fig. 2.12 shows a solitary bubble extracted from timeframe 475. Together with Fig. 2.9, this confirms that (i) the bubbles are reliably localized and reconstructed, and that (ii) the artifacts (streaks and hexagon-shapes) are an intrinsic limitation of the sparse-angular reconstruction. In comparison to phantom experiments, hexagon artifacts are less pronounced in fluidized beds, due to a smoother variation of b at the bubble interface (cf. Fig. 2.6). (Smoothly varying and irregular interfaces are known physical phenomena [81].) A second observation is that projections in Fig. 2.12 panel (a) contain less noise than in Fig. 2.9 panel (a). This is a combined effect of the creation of a full column reference (Section 2.4.1) and an averaging effect due to the motion of the solids during the exposure time of a single detector timeframe.




Figure 2.12: A solitary bubble that is extracted from frame 475 of the Geldart B fluidized bed reconstruction with a superficial gas velocity of 17 cm/s. The visualization format is explained in Section 2.4.3.

The video, and Fig. 2.15 with timeframes 482-488, display the dynamic behaviour of the field $b(x, y, z, t)$. This enables the quantification of bubble dynamics [81]. We observe that:

- Small bubbles form at the gas inlet and merge with other bubbles on their paths upwards through the column. The cores of the larger bubbles are usually reconstructed with full gas ($b = 1$, yellow), but, as before, bubble interfaces are not sharp (blue-green).
- The flux of interstitial gas through the solids plays an important role, leading to: (i) the existence of regions where b takes “intermediate values”; (ii) the growth and/or shrinkage of bubbles; (iii) the coalescence and/or breakup of bubbles and; (iv) the non-local interactions between bubbles.
- While bubbles often appear to be convected in quasi-steady ways upward, the dynamics are in fact 3-dimensional and intrinsically unsteady, as they are driven by the flow of gas through the column. Two examples are visible in Fig. 2.15. In the middle, a large bubble emerges, while in the top a smaller bubble disappears, i.e., disperses in the solids. Both occur in a very short time, and with little upwards displacement.

Overall, the Geldart B fluidized bed experiment shows the potential of our technique for imaging complex 3D dynamics. A pseudo-3D technique, i.e., using a vertical stacking of 2D slices over an interval of time [82, 83], would not have been adequate because of the unsteady 3D dynamics. To see this, note that in the pseudo-3D technique, subsequent 2D reconstructions of the bubble (while it travels through the slice) are stacked vertically. This only enables an accurate representation of the bubble when the shape of the bubble can be assumed constant during its path upward.




Figure 2.13: Template-matching of spheres with varying location and diameters to 10 mm (blue) and 23 mm (red) diameter bubbles, in a phantom experiment (dotted line) and a noiseless simulation (solid line). Panel (a) shows the best matching-score, Eq. (2.8), obtained for each diameter, panel (b) a 3D plot of best-matching templates overlaid onto the reconstruction of the phantoms (see Fig. 2.9), and panel (c) a horizontal cross-section of the reconstructions of noiseless simulations (bottom part) and phantoms (upper part) in grayscale coloring, onto which the boundaries of the best-matching spheres have been drawn.




Figure 2.14: Projections y of timeframe 485 of the bed fluidized with a superficial gas velocity of 17 cm/s. The projections have been cropped to a region of interest, showing rows 400-1200. The column walls are indicated with dashed lines.




Figure 2.15: Timeframes 482-488 of the bubbling fluidized bed using an superficial gas velocity of 17 cm/s. For this visualization we enhanced the images by postprocessing with a low-weight ($\lambda = 1/0.065$) total variation filter.




Figure 2.16: Still of the 17 cm/s experiment. In the movie, we enhanced the reconstruction by postprocessing with a low-weight ($\lambda = 1/0.065$) total variation filter.

2.6 Conclusion

To measure the gas-solids mixture of fluidized beds, various techniques are used in the literature, each with its own advantages and limitations. In this article, we have introduced a new technique that is fully-3D and simultaneously achieves a high temporal resolution. It consists of a set-up with three stationary X-ray source and flat panel detector pairs, a data-driven geometry calibration procedure, and a tailored image reconstruction method. The technique enables dynamic imaging of large 3D regions, is conceptually simple, and can easily be modified – for instance to image at a high framerate in a region of interest. It therefore enables straightforward experimentation with varying column sizes and different bed materials. Future considerations are an upgrade of the hardware, e.g., an expansion to five source-detector pairs (although this would be a costly and mechanically complex task) and the implementation of an additional beam hardening correction for imaging larger diameter columns [84].

In comparison to an X-ray source and detector on a gantry, such as is common in medical CT, our set-up has a high framerate, but requires a sparse-angle reconstruction for each timeframe. To study this aspect, we have used SIRT as a baseline in numerical simulations and phantom experiments, and analyzed the noise-induced artifacts, hexagon-shaped bubbles, and the degradation of resolving power when multiple bubbles occur at the same detector height. The artifacts are 3D analogues of observations in our previous line-detector set-up [53, 85], and provide directions for tailored post-processing approaches and more sophisticated reconstruction algorithms. Since noise was found to have a substantial effect in the reconstruction, we expect noise-suppressing priors to improve the image quality considerably. We will see that, in Chapter 5, indeed denoising as a preprocess improves the SIRT reconstructions. Depending on the bubbling regime, priors could also incorporate information on the sparsity of the bubbles, their solid contents, or the nearby timeframes. This has the potential to reduce artifacts and achieve a more precise determination of bubble interfaces.

In our results of a Geldart B bubbling fluidized bed, we find clear correspondences between the time-resolved experiment and conducted phantom experiments. Compared to our previous line-detector set-up, we found that the set-up is able to capture the complex 3D evolution of bubbles in the entire volume. This shows the potential of the technique for better quantitative analysis, such as average gas hold-up, as well as a better qualitative analysis, such as coalescence- and break-up processes, of bubbling fluidized beds.