



Time-resolved XPCS analysis across broad time-scales using multi-tau two-time correlations

Fabio Brugnara, Marco Cammarata, Yuriy Chushkin, Irene Festi, Massimo Granata, David Hofman, Matteo Leonardi, Daniele Marzi, Denis Nabari, Angela Trapananti, Simone Ziglio, Federico Zontone, Marco Zanatta and Giacomo Baldi

J. Synchrotron Rad. (2026). **33**, 698–707



IUCr Journals

CRYSTALLOGRAPHY JOURNALS ONLINE

This open-access article is distributed under the terms of the Creative Commons Attribution Licence <https://creativecommons.org/licenses/by/4.0/legalcode>, which permits unrestricted use, distribution, and reproduction in any medium, provided the original authors and source are cited.





Time-resolved XPCS analysis across broad time-scales using multi-tau two-time correlations

Fabio Brugnara,^{a,*} Marco Cammarata,^b Yuriy Chushkin,^b Irene Festi,^a Massimo Granata,^c David Hofman,^c Matteo Leonardi,^{a,d} Daniele Marzi,^a Denis Nabari,^{a,d} Angela Trapananti,^{e,f} Simone Ziglio,^a Federico Zontone,^b Marco Zanatta^{a,d} and Giacomo Baldi^{a,d,*}

Received 9 December 2025

Accepted 26 February 2026

Edited by A. Madsen, European XFEL, Germany

Keywords: X-ray photon correlation spectroscopy; multi-tau algorithm; synchrotron data analysis.


^aDepartment of Physics, University of Trento, Via Sommarive 14, 38123 Trento, Italy, ^bESRF – The European Synchrotron, 71 avenue des Martyrs, 38000 Grenoble, France, ^cLaboratoire des Matériaux Avancés – IP2I, CNRS, Université Claude Bernard Lyon 1, 69100 Villeurbanne, France, ^dINFN, Trento Institute for Fundamental Physics and Applications, 38123 Povo, Trento, Italy, ^ePhysics Division – School of Science and Technology, University of Camerino, Via Madonna delle Carceri 9, 62032 Camerino (MC), Italy, and ^fINFN, Sezione di Perugia, Via A. Pascoli, 06123 Perugia, Italy. *Correspondence e-mail: fabio.brugnara-1@unitn.it, giacomo.baldi@unitn.it

We present a multi-tau two-time correlation (MT-2TC) analysis for X-ray photon correlation spectroscopy that enables efficient analysis of large-scale, high-frame-rate datasets collected at modern synchrotron beamlines. The method combines the advantages of the multi-tau autocorrelation and two-time correlation approaches, maintaining sensitivity to non-stationary dynamics while reducing computational complexity and memory requirements from quadratic to linear scaling with the number of frames. This formulation not only improves processing efficiency but also extends the accessible correlation time range, allowing analysis over much broader time scales than conventional techniques. The algorithm is implemented in both dense and sparse data representations, allowing direct processing of sparsified photon-counting data from ultra-fast detectors. We demonstrate the performance of the MT-2TC scheme using experiments on a bulk GeO₂ glass and on a v-Ta₂O₅ thin film conducted at the ESRF ID10 beamline with a 1 kHz Eiger 4M detector. The results show perfect agreement with standard correlation analyses while extending the temporal correlation range by nearly two decades and improving signal-to-noise ratios at long lag times. The MT-2TC approach thus provides a scalable and flexible tool for XPCS at fourth-generation synchrotrons, where enhanced coherent flux and detector speed demand correlation strategies capable of spanning microsecond-to-day timescales.

1. Introduction

X-ray photon correlation spectroscopy (XPCS) is a powerful method for probing slow, nanometre-scale structural dynamics in condensed matter systems. The technique exploits the coherent fraction of synchrotron or free-electron laser X-rays: when a coherent beam is focused onto a sample, the scattered radiation produces a speckle pattern, *i.e.* a granular interference fingerprint of the instantaneous spatial arrangement of scatterers. Recording the temporal fluctuations of this pattern with a fast area detector enables direct measurements of dynamical processes such as diffusion, relaxation, and flow (Sutton, 2008; Shpyrko, 2014; Madsen *et al.*, 2016).

At modern synchrotrons, XPCS takes advantage of the exceptional brightness and coherence of the focused beam. To mitigate radiation damage and to maximize weak scattered signals, pixelated detectors are employed to improve the signal-to-noise ratio (SNR) at fixed dose by averaging over multiple coherence areas (multi-speckle XPCS) (Lumma *et al.*,



OPEN ACCESS

Published under a CC BY 4.0 licence

2000; Falus *et al.*, 2006). The correlation function SNR is given by

$$\text{SNR} = \langle I(q) \rangle C (\tau_{\min} T N_{\text{px}})^{1/2}, \quad (1)$$

which increases linearly with the average scattered intensity $\langle I(q) \rangle$ and scales with the square root of the number of pixels N_{px} , the minimum sampling time $\tau_{\min}$, and the measurement time T . Here, C denotes the contrast, *i.e.* the speckle visibility expected for a fully static sample, and is directly related to the coherence of the X-ray beam (Pecora, 1985).

Dedicated XPCS beamlines are now standard installations at leading synchrotron facilities worldwide, such as ID10 at ESRF, 8-ID-E at APS, 11-ID at BNL and P10 at PETRA III. Over the past decade, these facilities have seen major developments:

(i) *Coherent flux at fourth-generation synchrotron radiation sources.* Third-generation synchrotron radiation sources typically delivered a coherent flux of 5×10^9 to 5×10^{10} photons s^{-1} in the 8–11 keV range [*e.g.* ID10 at ESRF (Favre-Nicolin *et al.*, 2017), 8-ID-E at APS (Jiang *et al.*, 2023), P10 at PETRA III (Sprung *et al.*, 2023)]. The advent of fourth-generation facilities, such as ESRF-EBS and APS-U, has produced a substantial increase in brilliance, and consequently in coherent flux, by up to two orders of magnitude, mainly due to the dramatic reduction of horizontal emittance in the electron beam (Favre-Nicolin *et al.*, 2017; Raimondi *et al.*, 2023; Cornet *et al.*, 2024).

(ii) *Ultra-fast photon-counting pixel array detectors (PADs).* Photon-counting PADs have become the detectors of choice for XPCS thanks to their large dynamic range, negligible readout noise and dark current, and large collection areas (Zhang *et al.*, 2018). The Eiger 4M is today's standard detector, offering frame rates up to 4 kHz (DECTRIS AG, 2025). Recent developments in ultra-fast PADs now enable access to even shorter delay times. Continuous measurements have been demonstrated at 50 kHz for about 50 s and at 1.2 MHz for 4 s in burst mode using the UFXC32k chip (Zhang *et al.*, 2016; Zhang *et al.*, 2017) and the related Rigaku XSPA-500k detector (Zhang *et al.*, 2018). Burst-mode operation of the AGIP detector has reached 195 ns delay times over 700 μs (Jo *et al.*, 2021), while the TEMPUS detector has recently achieved ~ 450 ns in continuous operation using its event-driven mode (Chushkin *et al.*, 2025).

(iii) *Raw data sparsification.* To handle the huge data rates produced by ultra-fast PADs, raw data sparsification was introduced at the 8-ID-E beamline (Zhang *et al.*, 2021) and is now standard at ID10 (ESRF), 8-ID-E (APS) and 11-ID (BNL). The approach, motivated by the measured low-count-rate frames of typical XPCS experiments (Chushkin *et al.*, 2012), stores only non-zero values and their positions, drastically reducing data throughput while preserving the entire measurement information.

(iv) *Beamline stability.* Improvements in mechanical and thermal stability have significantly extended the upper dynamical limit of XPCS. Relaxation times of several hours can now be reliably measured, as shown in slow-dynamics experiments (Ruta *et al.*, 2013).

The combination of increased coherent flux and ultra-fast PADs dramatically expands the probed range towards faster times. Concurrently, improvements in instrumental stability and data sparsification pipelines enable measurements over hours or even days at high frame rates. For instance, we recently continuously tracked the relaxation dynamics of vitreous GeO_2 and vitreous Ta_2O_5 for up to 16 h with 1 ms resolution, generating over 40 million frames with the Eiger 4M detector at ID10 (see Section 4). These advances allow XPCS to span an unprecedented dynamical window, from microseconds to days. This opens new scientific opportunities, such as capturing both α and β relaxations within a single correlation function, or simultaneously resolving fast and slow dynamics in aging systems. However, these capabilities introduce significant computational challenges, demanding increased processing power and memory for the calculation of correlation functions.

Three distinct approaches are commonly used for data reduction in multi-speckle XPCS: standard intensity correlation (IA) (Lumma *et al.*, 2000), multi-tau intensity autocorrelation (MT-IA) (Schätzel, 1987; Cipelletti & Weitz, 1999), and two-time correlation (2TC) (Sutton *et al.*, 2003; Bikondoa, 2017; Cipelletti *et al.*, 2003). The multi-tau scheme is widely employed for systems where the dynamics do not change significantly during the measurement. It improves the SNR at long delay times and dramatically reduces computational complexity [from $O(N^2)$ to $O(N \log N)$] by correlating progressively time-averaged frames. In contrast, the two-time correlation approach is necessary for systems with time-dependent dynamics, such as those driven by external stimuli, undergoing aging, or relaxing toward equilibrium. This method calculates the correlation as a function of both measurement time and delay time, yielding a 2D correlation map.

In this paper, we introduce a new data reduction scheme that combines the advantages of both methods: the *multi-tau two-time correlation* function (MT-2TC). This approach mitigates computational complexity and largely reduces the memory limitations that constrain the number of frames in standard two-time calculations, and provides a more direct quasi-logarithmic visualization of the correlation matrix. Furthermore, the measurement-time average of our function reduces to the classical multi-tau correlation function, allowing for the straightforward recovery of the standard autocorrelation, as in the two-time formalism.

2. Multi-speckle correlation functions

2.1. Intensity autocorrelation (IA)

Photon correlation spectroscopy with area detectors was first demonstrated in XPCS by Dierker *et al.* (Dierker *et al.*, 1995) and in dynamic light scattering (DLS) by Kirsch *et al.* (Kirsch *et al.*, 1996), and later further developed by Cipelletti *et al.* (Cipelletti & Weitz, 1999) and Lumma *et al.* (Lumma *et al.*, 2000). In the standard multi-speckle scheme, the intensity autocorrelation (IA) function $g_2(q, \tau)$ is calculated by averaging over detector pixels with the same scattering vector q ,

$$g_q^{(2)}(\tau) = \left\langle \frac{\langle I_p(t) I_p(t + \tau) \rangle_p}{\langle I_p(t) \rangle_p \langle I_p(t + \tau) \rangle_p} \right\rangle_t, \quad (2)$$

Here, $I_p(t)$ is the intensity at the pixel p and time t ; the brackets $\langle \cdot \rangle_p$ and $\langle \cdot \rangle_t$ denote ensemble (pixel) and time average, respectively. The order of averaging and normalization is important: taking the ensemble average first preserves information on ergodicity (Joosten *et al.*, 1990), while the chosen normalization suppresses fluctuations in the incident beam (Duri *et al.*, 2005)¹.

In a typical XPCS experiment, the detector records a series of speckle patterns (frames) at constant spacing Δt . The data can be arranged as an intensity matrix $\mathbf{I}$ of dimensions $N_f \times N_{px}$ where N_f is the number of frames and N_{px} the number of pixels. When stored as a dense matrix, besides the challenges of data storage, a direct calculation of the numerator in the above expression scales as $N_f^2 N_{px}$, which quickly becomes computationally prohibitive.

As depicted in the *Introduction*, when detectors operate at high frame rates and scattering signal is low, most entries of $\mathbf{I}$ are zero. Using sparse storage formats, mainly compressed sparse row (CSR) or column (CSC) representations (Scott & Tuma, 2023), can drastically reduce the computational cost of evaluating equation (2), since only nonzero elements are processed.² However, row or column indexing in sparse matrices is intrinsically slow, limiting the efficiency of a direct element-wise evaluation of equation (2). A more practical solution is provided by two-time correlation functions exploiting the sparse linear algebra.

2.2. Two-time correlation (2TC)

Omitting the time averaging in equation (2), the two-time correlation (2TC) function can be defined as (Sutton *et al.*, 2003)

$$G_S^{(2)}(t_1, t_2) = \frac{\langle I_p(t_1) I_p(t_2) \rangle_p}{\langle I_p(t_1) \rangle_p \langle I_p(t_2) \rangle_p}, \quad (3)$$

where t_1 and t_2 are the times of the correlated frames. Independently, Cipelletti *et al.* (2003) introduced an equivalent form, $G_C^{(2)}(t_w, \tau)$, expressed in terms of the waiting time $t_w = (t_1 + t_2)/2$ and the lag time $\tau = |t_1 - t_2|$. Using the symmetric definition of the waiting time (Bikondoa, 2017), $G_C^{(2)}(t_w, \tau)$ corresponds to a 45° rotation of the $G_S^{(2)}(t_1, t_2)$ matrix. Within both representations, the standard IA can be computed by time averaging (Duri *et al.*, 2005),

$$g^{(2)}(\tau) = \langle G_C^{(2)}(t_w, \tau) \rangle_{t_w} = \langle G_S^{(2)}(t, t + \tau) \rangle_t. \quad (4)$$

Equations (3) and (4) give a practical computational pathway to compute both 2TC and IA taking advantage of the efficiency of matrix multiplication algorithms,

¹ Different definitions of $g_2(q, \tau)$ exist in the literature, varying in both the order of averaging and the choice of normalization (Cipelletti & Weitz, 1999; Lumma *et al.*, 2000; Duri *et al.*, 2005; Chushkin *et al.*, 2012; Bikondoa, 2017).

² A related approach based on event correlators was proposed by Chushkin *et al.* (2012), showing dramatic speed-ups.

$$\mathbf{G}_S^{(2)} = [(\mathbf{I} \cdot \mathbf{I}^T) \odot_r \mathbf{n}] \odot_c \mathbf{n}^T \quad \text{with } [\mathbf{n}]_i = \left(\sum_{p_j} I_{t_i, p_j} \right)^{-1}. \quad (5)$$

The matrix product $\mathbf{I} \cdot \mathbf{I}^T$ computes the pixel-averaged correlations in the numerator of equation (3). The subsequent element-wise multiplications, $\odot_r$ (row-wise) and $\odot_c$ (column-wise) with the normalization vector $\mathbf{n}$ ensures the proper normalization. The entries of the resulting matrix ($N_f \times N_f$ large) correspond directly to the 2TC: $[\mathbf{G}_S^{(2)}]_{i,j} = G_S^{(2)}(i\Delta t, j\Delta t)$.³ The IA vector is then obtained by averaging over the offset diagonals of this matrix, *i.e.* averaging over constant lag times: $[\mathbf{g}^{(2)}]_k = (1/N_f - k) \sum_{i=0}^{N_f-1-k} [\mathbf{G}_S^{(2)}]_{i,i+k}$.

This matrix-multiplication approach leverages optimized linear algebra routines and is especially powerful for sparse data, where direct indexing is slow. However, it comes at the cost of memory: $\mathbf{G}_S^{(2)}$ is a dense matrix of dimension $N_f \times N_f$. For large datasets, this rapidly exceeds available resources (*e.g.* $\mathbf{G}_S^{(2)}$ from 1 million frames requires several terabytes). A general workaround is time-binning, where data are coarse-grained to longer frame-times, however, at the expense of losing access to the fastest dynamics. Alternatively, multi-tau autocorrelation offers a way to overcome time-scale limitations.

2.3. Multi-tau autocorrelation (MT-IA)


Multi-tau calculation of intensity autocorrelation was first proposed by Schätzle (Schätzle, 1983) for single area detectors, and later extended to multi-speckle DLS by Cipelletti (Cipelletti & Weitz, 1999). Instead of following the prescription of equation (2), which requires evaluating all possible frame pairs and thus scales as N_f^2 , multi-tau algorithms maintain a nearly constant ratio between the delay time τ and the sampling time Δt , still maintaining the temporal resolution $\Delta t/\tau \ll 1$. This leads to a quasi-logarithmic sampling of lag times: the correlation function spans several decades of τ with far fewer points, reducing both computational cost and memory requirements, as well as increasing the SNR at large τ as a consequence of the non-constant Δt .

Let us first focus on the calculation of the un-normalized IA for a single pixel. The multi-tau algorithm organizes the computation into a cascade of linear correlators (levels), each with a fixed number of channels (four in the example of Fig. 1).

(i) *First correlator (corr = 0)*. The pixel intensity stream, sampled at Δt , is fed into the channels. All channels except the first (which corresponds to zero lag time, *i.e.* the speckle pattern variance) compute correlations at linearly spaced delays $\tau_{ch} = ch * \Delta t$.

(ii) *Subsequent correlators (corr = 1, 2, ...)*. Before feeding the signal to the next correlator, pairs of consecutive frames are summed, effectively doubling both the sampling time and the accessible delay times. In these correlators, only the second half of the channels are used, since the first half corresponds to lag times already computed at higher temporal resolution.

³ Since $\mathbf{G}_S^{(2)}$ is symmetric, only one half of the matrix needs to be computed.


Figure 1

Multi-tau correlator scheme for a single pixel p with four channels. Data sampled at Δt are correlated in $corr = 0$ by all channels but the first, then binned and passed to subsequent correlators, each doubling the delay times. Only new lag times (second half of channels) are computed at each stage (Cipelletti & Weitz, 1999).

(iii) *Iteration*. This binning–correlation process is iteratively repeated, each stage doubling the frame time and extending the correlation function to longer τ . The multi-tau autocorrelation (MT-IA) function is obtained by merging the outputs of all correlators.

Finally, the MT-IA functions calculated for all pixels with the same scattering wavevector are averaged and normalized.⁴ Multi-speckle multi-tau autocorrelation is the starting point of our novel correlation scheme.

3. Multi-tau two-time correlation (MT-2TC)

The 2TC matrix defined in equation (5) is a standard tool in XPCS data analysis. It provides information on sample ergodicity and stationarity, allows verification of measurement stability, and highlights temporally correlated noise sources (Bikondoa, 2017) (e.g. spurious high-count signals from environmental radioactivity appearing in consecutive frames). Once stationarity is established, the IA is obtained by time-averaging over the selected stationary time ranges. As already noted, the main limitation of this method lies in the size of the $\mathbf{G}_S^{(2)}$ matrix. In addition, when the measurement time far exceeds the sample decorrelation time, visualizing the entire matrix becomes uninformative. In such cases, data are either inspected through sub-matrices or, as in recent studies (Perakis *et al.*, 2017; Chèvrement *et al.*, 2024), reformulated in the $G_C^{(2)}(t_w, \tau)$ representation and plotted with a logarithmic τ axis. To overcome these limitations, we propose a multi-tau

definition of the $G_C^{(2)}(t_w, \tau)$ function, quasi-logarithmically spaced in τ . This approach extends the standard multi-tau scheme of Fig. 1 by avoiding temporal averaging.

We denote frame-shifts as follows: $\mathbf{I}[i:]$ is the sub-matrix of $\mathbf{I}$ starting from row i , and $\mathbf{I}[: -j]$ excludes the last j rows. Frame binning for the $corr$ -th correlator is defined as $[\mathbf{I}^{corr}]_{f,p} = \sum_{k=0}^{2^{corr}} I_{f \cdot 2^{corr} + k, p}$.⁵ The non-normalized multi-tau 2TC is then defined as a set of vectors of different sizes, one for each correlator index $corr = 0, 1, \dots$ and channel index $ch = 0, 1, \dots, N_{ch} - 1$:

$$C_{MTcorr, ch}^{(2)} = \sum_{\text{pixels}} \mathbf{I}^{corr}[: -ch] \odot \mathbf{I}^{corr}[ch :], \quad (6)$$

where $\odot$ is the element-wise (Hadamard) product. For each correlator and channel, the intensity matrix is first binned by 2^{corr} , then element-wise multiplied with its frame-shifted copy (by ch rows), and finally summed over pixels. Normalization follows equation (3),

$$G_{MTcorr, ch}^{(2)} = C_{MTcorr, ch}^{(2)} \odot \mathbf{n}^{corr}[: -ch] \odot \mathbf{n}^{corr}[ch :]. \quad (7)$$

As in the standard multi-tau scheme, the correlator and channel indices determine the lag time τ , while the i th element of the vector $G_{MTcorr, ch}^{(2)}$ corresponds to a waiting time t_w ,

$$\begin{aligned} \tau(corr, ch) &= 2^{corr} ch \Delta t \\ t_w(i, ch, corr) &= \left(\frac{ch + 1}{2} + i \right) 2^{corr} \Delta t \end{aligned} \quad (8)$$

Computationally, each Hadamard product in the denominator requires roughly $N_{px} N_f / 2^{corr}$ multiplications per correlator

⁵ The last frames are discarded if 2^{corr} is not a divisor of N_f .

and channel. Computing the full object thus involves about $2N_{\text{px}}N_{\text{ch}}N_{\text{f}}$ operations.⁶ Similarly, the number of elements in $G_{\text{MT corr, ch}}^{(2)}$ scales as $2N_{\text{f}}N_{\text{ch}}$. For large datasets, both the computational cost and storage requirements are therefore reduced from quadratic to linear in N_{f} . This enables calculation and storage of 2TC functions irrespective of frame rates and measurement times, thus allowing operation at the full detector speed. Visualization is naturally achieved in a $t_{\text{w}}-\tau$ plot with logarithmic τ , which conveniently displays the full dynamics of the sample—even when relaxation times are much shorter than the total acquisition time—and facilitates verification of stationary ranges.

As an example, Fig. 2 compares the standard 2TC (top) with the MT-2TC (bottom), computed on a v-Ta₂O₅ 128 frames measurement at low frame rate $\Delta t \simeq 4$ s (details given in Section 4.2). In the multi-tau calculation, we used 16 channels, which is sufficient to avoid time-blurring effects as, apart from the first half of the first correlator's channels, they always keep $\tau/\Delta t > 8$. The MT-2TC can be viewed as a 45° rotation of the 2TC matrix, but with the lag time τ quasi-logarithmically sampled: the rectangle rows corresponding to each correlator and channel double in both height and width every eight rows (except for the first correlator, which has 16 channels). In this example, both the 2TC and the MT-2TC reveal the same physical information, namely the sample dynamics evolving with an increasing relaxation time. However, important practical differences arise:

- (i) Although both are computed from the same 128 frames, they contain very different numbers of elements. The 2TC has $N_{\text{f}}^2 = 16384$ entries, while the MT-2TC contains fewer than $2n_{\text{ch}}N_{\text{f}} = 4096$ entries. This gap grows quadratically with N_{f} .
- (ii) To perform this comparison, we simulated 4 s frames by binning the original 1 ms data, so that the standard 2TC achieves a reasonable SNR. At much higher frame rates the 2TC would become unusable, dominated by noise. By contrast, the MT-2TC remains robust: for each lag time it shows the same SNR irrespective of the frame rate, since its SNR naturally improves with increasing τ .

Finally, the IA function, consistent with that of Section 2.1, can be recovered by averaging over t_{w} , and errors can be computed accordingly,⁷

$$g_{\text{MT corr, ch}}^{(2)} = \frac{1}{N_{t_{\text{w}}}} \sum_{t_{\text{w}}} G_{\text{MT corr, ch}}^{(2)}. \quad (9)$$

3.1. MT-2TC from sparse matrices

In principle, equation (7) enables the calculation of both IA and 2TC for arbitrarily large N_{f} . In high-frame-rate and long-duration measurements, however, data are typically stored in sparse formats to circumvent memory limitations. Therefore, an efficient implementation of equation (7) must explicitly

⁶ This estimate follows a geometric series and overestimates the true number of multiplications.

⁷ When averaging over a restricted observation-time range, care must be taken to truncate at the appropriate correlator index, since higher-order correlators may correlate frames lying outside the chosen t_{w} interval under the symmetric definition.




Figure 2

Comparison of the standard 2TC (top) and MT-2TC (bottom) computed for a v-Ta₂O₅ measurement at a low frame rate ($\Delta t \simeq 4$ s), chosen to improve the SNR (see Section 4.2 for details). The labels indicate which correlator and which set of channels correspond to the rows of the MT-2TC. Notably, the first correlator consists of 15 channels (the first channel is omitted, as it corresponds to the self-correlation of frames), while the other two correlators consist of eight filled channels each.

operate on sparse data structures. A key difficulty arises in the binning step: binning sparse intensity matrices is highly inefficient, as it neither reduces memory usage nor simplifies subsequent calculations. In fact, binning a sparse matrix by a factor of 2 usually results only in a frame re-indexing of the events, since two consecutive frames rarely record photons in the same pixel. To address this, we propose a two-step implementation of equation (7) that fully advantage of the sparse formalism.

3.1.1. Step 1: sparse-domain 2TC sub-matrices

We first compute the linear 2TC matrices using equation (5), but restricted to small sub-matrices of $\mathbf{I}$. These are defined by a parameter s as $\mathbf{I}_n = \mathbf{I}[n2^s:(n+1)2^s]$ with $n = 0, 1, \dots, N_{\text{f}}/2^s$. For each n , two consecutive 2TCs are calculated: (i) the autocorrelation of $\mathbf{I}_n$, and (ii) the cross-correlation of $\mathbf{I}_n$ with $\mathbf{I}_{n+1}$.⁸ These correspond to the diagonal blocks of the entire $\mathbf{G}_S^{(2)}$ matrix (Fig. 3, left). Each 2TC sub-matrix is then converted into its multi-tau version, subsequently filling the first correlators of $G_{\text{MT corr, ch}}^{(2)}$ by appending the correlation values at the end of its vectors. For each correlator, this conversion is achieved by binning the 2TC sub-matrices, both by rows and columns, by 2^{corr} , and extracting the channel vectors from the corresponding ch -offset diagonals (Fig. 3, right). The cross-correlation matrices $\mathbf{I}_n\mathbf{I}_{n+1}$ are treated analogously, except that their first correlators and channels

⁸ The extension of equation (5) to this case replaces the transposed $\mathbf{I}_n$ matrix with $\mathbf{I}_{n+1}$. The same happens in the normalization.




Figure 3

Illustration of the two-step sparse implementation of the multi-tau 2TC. Left: decomposition of the full $\mathbf{G}_S^{(2)}$ into diagonal and off-diagonal sub-matrices computed from sparse blocks of $\mathbf{I}$. Right: mapping of a 2TC sub-matrix into the multi-tau representation via successive binning and extraction of ch -offset diagonals. In this example, $s = 4$ and four channels are used. Double-colored entries indicate values used in different multi-tau correlators.

start from the corner rather than from the diagonal. It can be shown that this algorithm fills the first correlators equivalently to equation (7), but moves time-binning after the correlation computation. The advantage of this approach is that it avoids explicit frame binning, still avoiding high memory consumption, since only two small $\mathbf{I}$ sub-matrices are processed at each step, which are then transformed from the linear 2TC into its multi-tau representation.

3.1.2. Step 2: dense-domain computation

The second stage follows equation (7). The intensity matrix $\mathbf{I}$ is time-binned by a factor 2^{s+1} and converted into a standard dense matrix. From this point, subsequent correlators are computed via the Hadamard product and binning scheme of equation (7). In practice, the optimal choice of s is such that $2^s \simeq 1/\text{sparsity}$, *i.e.* binning becomes advantageous once the resulting matrix is effectively dense.

3.2. Implementation and memory considerations

As described in Section 3.1, the MT-2TC calculation for sparse data is divided into two main stages. In the first stage, standard two-time correlation (2TC) sub-matrices are computed from successive batches of 2^s frames. At this stage, only a limited number of intermediate quantities are required. Two arrays, I_0 and I_1 (each of dimension $2^s \times N_{\text{px}}$), store the frame batches used to compute diagonal and off-diagonal 2TC sub-matrices. The corresponding correlation results are stored in an intermediate matrix $G_{\text{sub}}^{(2)}$ of dimension $2^s \times 2^s$. The output of this stage consists of (i) a time-binned version of the original frame sequence, denoted I_{dense} (of dimension $N_f/2^s \times N_{\text{px}}$), and (ii) a partially filled MT-2TC object, in which only the channels associated with the first correlators are populated.

The first-stage computation is performed by iterating over the dataset in steps of 2^s frames. At each iteration, two consecutive frame batches, I_0 and I_1 , are loaded into memory. These batches are used to compute both the diagonal and off-diagonal two-time correlation (2TC) sub-matrices, which are

temporarily stored in the intermediate array $G_{\text{sub}}^{(2)}$. The resulting sub-matrices are then transformed into their multi-tau representation and accumulated in the corresponding channels of the MT-2TC object. Finally, the frames I_0 are time-binned by summation to produce the corresponding dense frame, which is stored in the vector I_{dense} for subsequent processing in the dense stage of the algorithm.

In the second stage, the MT-2TC is completed by operating on the dense (time-binned) frames. This step involves matrix slicing and dense matrix elementwise multiplications, followed by successive binning of I_{dense} by a factor of two at each iteration, as described by equations (6) and (7).

In the present implementation, the computation of sparse 2TC sub-matrices relies on Intel MKL sparse matrix-multiplication routines (Intel Corporation, 2023), while the frame-binning step and all dense operations are implemented using the *NumPy* library (Harris *et al.*, 2020).

3.2.1. Computational time

The computational time of the MT-2TC algorithm scales linearly with the number of frames N_f . In the first stage, the number of loop iterations scales as $N_f/2^s$. In the second stage, the initial number of dense frames to be processed is also $N_f/2^s$, and the first iteration dominates the computational cost, since subsequent iterations operate on progressively halved frame sequences. As a result, the total computational time increases approximately linearly with N_f .

This behavior contrasts with the standard 2TC approach, whose computational cost and memory usage scale quadratically with the number of frames. However, the standard matrix-multiplication-based 2TC benefits from highly optimized linear algebra routines and remains faster until memory limitations come into play. As N_f increases and the full 2TC matrix can no longer be stored in memory, the standard approach becomes impractical. In this regime, the MT-2TC retains linear scaling with N_f , providing a decisive advantage in terms of feasibility for large datasets.

3.2.2. Memory considerations

The main advantage of the MT-2TC method is the strong reduction of memory usage. In standard 2TC calculations, the dominant memory limitation arises from storing the full 2TC matrix, whose size scales as $N_f \times N_f$. For example, on typical beamline computing infrastructure such as that available at the ESRF, storing the full 2TC matrix for a dataset of 200 000 frames using 32-bit floating-point numbers would require approximately 160 GB of memory, already at the limits of workstations equipped with a few hundred gigabytes of RAM.

In contrast, the MT-2TC correlation object scales linearly as $2 \times N_f \times N_{\text{ch}}$. For the same dataset and using 16 channels, the MT-2TC object occupies only about 25.6 MB. This directly addresses the dominant memory bottleneck associated with the size of the two-time correlation matrix in large-frame XPCS datasets.

Additional memory is required for intermediate quantities. The sparse frame batches I_0 and I_1 each contain of the order of




Figure 4 MT-2TC (top) and IA (bottom) computed on X-ray-damaged v-GeO₂ at 400 K, using 16 channels to preserve temporal resolution. Top: two-time correlation function on 1 000 000 frames at 1 ms, plotted with lag time on a logarithmic axis, allowing visualization of the relaxation process across the entire measurement window. Bottom: comparison of the multi-tau IA (orange error bars) with the standard linear definition of equation (2) (blue dots), the latter computed from 20 000 frames at 50 ms to avoid memory limitations.

N_{px} non-zero values and therefore require limited memory.⁹ The largest intermediate objects are the correlation sub-matrix $G_{\text{sub}}^{(2)}$ and the dense, time-binned frame sequence I_{dense} , whose sizes depend on the choice of the parameter s . For the example shown in Fig. 4 (top), the optimal choice is $s = 14$. In this case, computing the MT-2TC for 1 million frames requires approximately 1 GB of memory for each of $G_{\text{sub}}^{(2)}$ and I_{dense} . Given the number of frames involved, a total memory requirement of a few gigabytes represents a substantial reduction compared with standard two-time correlation calculations.

Importantly, the algorithm does not require the full dataset to be loaded into memory at once and therefore supports chunked processing, in which data are handled in successive time blocks. At any given time, only the frame batches I_0 and I_1 need to reside in memory. As a result, the method supports datasets that approach or exceed available RAM, with practical limitations progressively determined by the speed of data loading from disk and computation, rather than by memory capacity.

4. Experimental demonstrations and perspectives

From a practical standpoint, the MT-2TC approach is particularly well suited to high-frame-rate XPCS experiments

⁹ Here we assume the optimal choice $s \simeq \log_2(1/\text{sparsity})$, as discussed previously.

involving large numbers of frames. By preserving access to the full two-time information while drastically reducing memory requirements, it enables efficient storage and visualization, as well as flexible *a posteriori* selection of time windows without recomputing correlations, simplifying both online inspection and post-processing of large datasets.

In this section, we illustrate the versatility and effectiveness of the proposed tool through two recent experiments we carried out at the ID10 beamline of ESRF using the Eiger 4M detector working at a frame rate of 1 kHz. The first case concerns the relaxation dynamics of v-GeO₂, where our aim is to study the temperature dependence of the contrast and its connection to the non-ergodicity of the glass. The second example addresses X-ray-induced dynamics in a thin film of v-Ta₂O₅, obtained by ion beam sputtering, comparing as-deposited and annealed samples. Here, the focus was on relaxation processes occurring on the millisecond timescale, which are particularly relevant because they dominate the noise response of ground-based gravitational-wave detectors, such as Advanced Virgo (Acernese *et al.*, 2014), Advanced LIGO (Aasi *et al.*, 2015) and KAGRA (Aso *et al.*, 2013), at frequencies in the range between 10 and 1000 Hz (Steinlechner, 2018).

For both oxide samples, the dynamics observed at temperatures below the glass transition are dominated by X-ray-induced processes, *i.e.* relaxations activated by the interaction of the intense coherent X-ray beam with the sample, primarily through bond breaking due to the radiolysis process, which superimpose on the intrinsic material dynamics (Ruta *et al.*, 2017; Martinelli *et al.*, 2023).

4.1. XPCS at ms time resolution on v-GeO₂

The experiment was performed in the small-angle configuration of the ID10 beamline: the detector was centered at 1.75° using a 9.7 keV X-ray beam, giving access to the 1.2–2.2 nm^{−1} Q -range. To slow down the X-ray-induced dynamics such that it fell within our measurement time window, the 1.16×10^{12} photons s^{−1} beam was defocused to 30 μm × 40 μm on the ~100 μm-thick v-GeO₂ slab sample. XPCS scans were collected at a 1 kHz frame rate for 15 temperatures in the 300–1100 K range. The samples were prepared starting from a bulk cylinder of v-GeO₂, obtained with the standard melt-quenching method. The cylinder was cut into small slices and thinned with polishing paper, reaching 100 ± 10 μm of thickness. The X-ray-induced relaxation time never exceeds 6 s; however, because of the very low scattered signal (~0.2 photons pixel^{−1} s^{−1}), we took ~1 h scans exploiting the stationary dynamics of the sample to improve the SNR (Cipelletti & Weitz, 1999). Regardless of temperature (below T_g), the sample initially displayed non-stationary dynamics: during the first ~200 s of beam exposure the relaxation slows down, from millisecond to second time scales. After this transient regime, the dynamics became stationary. In the experiment, we investigated the dynamics of this ‘damaged’ (or ‘rejuvenated’) glass (Baglioni *et al.*, 2024; Alfinelli *et al.*,

2023), thus measuring on the same irradiated spot, within the regime of stationary dynamics.

The stationarity of the dynamics is quantified using the novel MT-2TC. An example, computed for the 400 K scan at $Q = 1.7 \pm 0.1 \text{ nm}^{-1}$ (1000000 frames at 1 ms), is shown in Fig. 4. The two-time correlation function, naturally represented with the lag time on a logarithmic scale, allows the relaxation process to be monitored across the entire measurement time, even when the relaxation itself is much shorter than the measurement time. As expected, correlations at shorter lag times are noisier due to the limited number of correlated photons, while at longer lag times the statistics improve and averaging is no longer needed to obtain stable values. Once stationarity is established, the $g^{(2)}$ function is obtained by averaging over the rows of the $G_{MTcorr, ch}^{(2)}$ object, as shown in the lower panel of Fig. 4. Here, the agreement with the Kohlrausch–Williams–Watts (KWW) relaxation function over six decades in time is significant.

The multi-tau IA is compared with the standard linear definition of equation (2), computed from 20000 frames obtained by binning the original 1 million frames by a factor of 50 in order to avoid memory limitations. Otherwise, when using 32-bit floating-point numbers, the full 2TC matrix would require approximately 4 TB of memory. The two methods give consistent results, but the multi-tau calculation provides a significant advantage in terms of SNR, especially at long lag times, and, more importantly, extends the accessible time window by almost two decades. Given our experimental objectives, reliable access to the low-lag-time regime, *i.e.* lag times much shorter than the X-ray-induced relaxation time, is essential to resolve the plateau before the initial decay of the correlation function and thus to accurately determine the contrast, which would otherwise be compromised by excessive frame binning. A detailed analysis of the studied temperature dependence of the non-ergodicity factor will be presented in a future publication.

4.2. X-ray induced relaxation dynamics in v-Ta₂O₅

The experiment was carried out in the wide-angle configuration of the ID10 beamline, with the detector angle varied between 1.5° and 10° to probe the Q -dependence of the induced relaxation within the range $3\text{--}20 \text{ nm}^{-1}$, using a working energy of 21.7 keV. The samples consisted of two $2.5 \mu\text{m}$ -thick v-Ta₂O₅ films deposited on Si substrates by ion-beam sputtering at the Laboratoire des Matériaux Avancés (LMA): one in the as-deposited state and the other annealed at 500°C for 10 h (additional details on sample preparation will be provided in a forthcoming publication, as well as a discussion on the XPCS results). The $0.8 \times 10^{12} \text{ photons s}^{-1}$ X-ray beam was tightly focused to a spot size of $10 \mu\text{m} \times 7 \mu\text{m}$ in order to increase the dose rate and accelerate the beam-induced dynamics. The experiment is performed in transmission geometry through the 0.5 mm Si substrate, which is almost transparent at this working energy.

Subjected to such X-ray dose, the relaxation dynamics evolves on the scale of a day, requiring measurements to be

monitored over several hours. Fig. 5 shows this evolution for the as-deposited sample at $Q = 20 \text{ nm}^{-1}$. The MT-2TC plot (top), computed from 60000000 frames at 1 ms, captures the full evolution of the dynamics across the measurement window. The characteristic relaxation time starts at a few tens of seconds and rapidly increases to thousands of seconds, and the relaxation function progressively stretches, evidencing the heterogeneous nature of the sample dynamics. The MT-2TC representation also guides the choice of stationary time-windows for computing autocorrelation functions and extracting KWW relaxation parameters. In practice, it suggests exponentially increasing window widths during the first hours, followed by equally large windows at later times (vertical red dashed lines). The $g^{(2)}$ functions are then obtained by t_w -averaging over the appropriate segments of the MT-2TC [according to equation (8)], ensuring that the averaged correlations are calculated from frames belonging to the selected stationary windows.

From a practical perspective, the MT-2TC for this dataset, comprising 60 million frames, was computed only once and required approximately 30 min on a standard analysis workstation. Subsequent calculations of the $g^{(2)}$ functions for different waiting times and window selections involve only cuts and averages of the precomputed MT-2TC, and therefore incur negligible additional computational cost. This illustrates




Figure 5

X-ray-induced relaxation dynamics of the as-deposited v-Ta₂O₅ sample at $Q = 20 \text{ nm}^{-1}$. Top: MT-2TC function for lag times $\tau > 1 \text{ s}$, computed from 60000000 frames at 1 ms, capturing the evolution of the relaxation dynamics during beam exposure. Vertical dashed red lines indicate the stationary time windows selected for IA analysis. Bottom: corresponding $g^{(2)}$ functions obtained by averaging over the selected windows. Access to faster dynamics is limited to the wider time windows, which exhibit improved SNR; therefore, only $g^{(2)}$ from these windows are shown for short lag times ($\tau < 1 \text{ s}$). Inset: extracted relaxation times and stretching exponents, which progressively increase and decrease, respectively, with dose, both following an approximately exponential trend.

a key practical advantage of the method for long-duration experiments, where repeated recalculation of two-time correlations would otherwise be prohibitive.

Fig. 5 (bottom) shows the resulting $g^{(2)}$ functions calculated for the selected time windows. As expected from equation (1), the SNR improves with increasing window width. Consequently, access to faster dynamics is restricted to the wider time windows, since the correlations computed over shorter intervals are dominated by statistical noise. The fitted t_w dependence of the relaxation time τ and the stretching exponent β is presented in the inset. For the as-deposited sample, the relaxation time progressively increases with exposure, following an approximately exponential trend.

We note that, when the induced decorrelation becomes very long, in contrast to the v-GeO₂ measurements, the relaxation over the full dynamical range cannot be adequately described by a single KWW function. In particular, in the fast lag-time range between 10 ms and 10 s, the correlation function exhibits a slow decay that precedes the main relaxation and cannot be captured by a stretched exponential form. We are currently investigating whether this behavior is indicative of the presence of a secondary relaxation process, possibly related to the thermal Brownian noise in gravitational waves detectors. Access to this extended time window is therefore essential to correctly characterize the dynamics and to avoid oversimplified fits based on a restricted lag-time range.

4.3. Further potential applications of the MT-2TC

Beyond the experimental demonstrations presented above, the MT-2TC approach is expected to be broadly applicable to XPCS measurements involving large datasets, high frame rates, and time-evolving dynamics. By providing access to the full two-time information while drastically reducing data size, the method enables flexible *a posteriori* analysis and efficient handling of long-duration measurements.

More generally, one of the most promising applications of the MT-2TC is the detection of two-step relaxation processes in glassy systems, such as the coexistence of α and β relaxations. While such phenomena are well known in light scattering and XPCS studies (Rössler & Becher, 2025), they have not yet been clearly resolved with X-rays in the regimes targeted here. Access to correlation functions over lag-time ranges significantly wider than the characteristic decorrelation region is essential to identify these behaviors and to avoid masking slow or intermediate relaxation components.

In addition, the MT-2TC representation is particularly valuable for experiments exhibiting non-stationary or heterogeneous dynamics, where the identification of stationary time windows is nontrivial. In these cases, it provides an efficient and intuitive way to monitor the temporal evolution of the dynamics and to guide the selection of appropriate averaging intervals, especially in dose-dependent or aging experiments.

In practice, the advantages of the MT-2TC become most apparent when large numbers of frames must be handled simultaneously. While for small datasets the standard matrix-

multiplication-based two-time correlation remain faster due to highly optimized linear algebra libraries, this situation changes as soon as memory limitations prevent computing and storing the full $N_f \times N_f$ correlation matrix. The MT-2TC is specifically designed to operate in this regime, which is increasingly common in high-frame-rate XPCS experiments involving long acquisition times and evolving dynamics.

5. Conclusions

The multi-tau two-time correlation (MT-2TC) algorithm presented here constitutes a significant advance in the computational analysis of XPCS data. By integrating the multi-tau scheme into the two-time correlation formalism, the method preserves sensitivity to non-stationary and time-evolving dynamics while drastically reducing computational and memory demands. Its quasi-logarithmic sampling in lag time enables the computation of correlation functions over very wide temporal ranges, providing natural visualization of dynamical processes spanning many orders of magnitude in time. Moreover, the MT-2TC matrix can be computed once and stored, allowing subsequent analyses, such as stationary time window selection, averaging and fitting, to be performed flexibly without recomputation. The sparse-domain implementation further ensures compatibility with modern fast photon-counting detectors and event-based data streams.

Applied to experimental data from vitreous GeO₂ and Ta₂O₅, the MT-2TC approach yields correlation functions consistent with conventional methods but with improved SNRs and extended dynamical coverage. The technique enables reliable identification of ergodic intervals and efficient extraction of relaxation parameters, even for measurements exceeding millions of frames.

The MT-2TC framework is therefore particularly well suited for upcoming fourth-generation synchrotron sources, where increased coherent flux and detector bandwidth can produce data volumes beyond the limits of traditional correlation algorithms. By combining computational scalability and full temporal sensitivity, the method provides a robust and flexible tool for exploring both stationary and time-dependent phenomena across broad dynamical ranges.

6. Code availability

The implementation of the multi-tau two-time correlation method presented in this work is publicly available under an open-source license at https://github.com/FabioBrugnara/XPCS_library.

Acknowledgements

During the preparation of this work, the authors, not being native speakers, used CHATGPT to rephrase some sentences and improve the readability of the text. After using this tool, the authors reviewed and edited the content as needed, and they take full responsibility for the content of the publication.

References

- Aasi, J., Abbott, B. P., Abbott, R., Abbott, T., Abernathy, M. R., Ackley, K., Adams, C., Adams, T., Addesso, P., Adhikari, R. X., Adya, V., Affeldt, C., Aggarwal, N., Aguiar, O. D., Ain, A. *et al.* (2015). *Cl. Quantum Grav.* **32**, 074001.
- Acernese, F., Agathos, M., Agatsuma, K., Aisa, D., Allemandou, N., Allocca, A., Amarni, J., Astone, P., Balestri, G., Ballardin, G., Barone, F., Baronick, J., Barsuglia, M., Basti, A., Basti, F. *et al.* (2014). *Cl. Quantum Grav.* **32**, 024001.
- Alfinelli, E., Caporaletti, F., Dallari, F., Martinelli, A., Monaco, G., Ruta, B., Sprung, M., Zanatta, M. & Baldi, G. (2023). *Phys. Rev. B* **107**, 054202.
- Aso, Y., Michimura, Y., Somiya, K., Ando, M., Miyakawa, O., Sekiguchi, T., Tatsumi, D. & Yamamoto, H. (2013). *Phys. Rev. D* **88**, 043007.
- Baglioni, J., Martinelli, A., Sun, P., Dallari, F., Westermeier, F., Sprung, M., Grübel, G. & Monaco, G. (2024). *Rep. Prog. Phys.* **87**, 120503.
- Bikondoa, O. (2017). *J. Appl. Cryst.* **50**, 357–368.
- Chèvremont, W., Zinn, T. & Narayanan, T. (2024). *J. Synchrotron Rad.* **31**, 65–76.
- Chushkin, Y., Caronna, C. & Madsen, A. (2012). *J. Appl. Cryst.* **45**, 807–813.
- Chushkin, Y., Correa, J., Ignatenko, A., Pennicard, D., Lange, S., Fridman, S., Karl, S., Senftleben, B., Lehmkuhler, F., Westermeier, F., Graafsma, H. & Cammarata, M. (2025). *J. Synchrotron Rad.* **32**, 1220–1227.
- Cipelletti, L., Bissig, H., Trappe, V., Ballesta, P. & Mazoyer, S. (2003). *J. Phys. Condens. Matter* **15**, S257–S262.
- Cipelletti, L. & Weitz, D. A. (1999). *Rev. Sci. Instrum.* **70**, 3214–3221.
- Cornet, A., Ronca, A., Shen, J., Zontone, F., Chushkin, Y., Cammarata, M., Garbarino, G., Sprung, M., Westermeier, F., Deschamps, T. & Ruta, B. (2024). *J. Synchrotron Rad.* **31**, 527–539.
- DECTRIS (2025). *EIGER X 4M Documentation*, <https://www.dectris.com/support/manuals-docs/eiger-x-for-synchrotron/eiger-x-4m/>.
- Dierker, S. B., Pindak, R., Fleming, R. M., Robinson, I. K. & Berman, L. (1995). *Phys. Rev. Lett.* **75**, 449–452.
- Duri, A., Bissig, H., Trappe, V. & Cipelletti, L. (2005). *Phys. Rev. E* **72**, 051401.
- Falus, P., Lurio, L. B. & Mochrie, S. G. J. (2006). *J. Synchrotron Rad.* **13**, 253–259.
- Favre-Nicolin, V., Chushkin, Y., Cloetens, P., da Silva, J. C., Leake, S., Ruta, B. & Zontone, F. (2017). *Synchrotron Radiat. News* **30**(5), 13–18.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C. & Oliphant, T. E. (2020). *Nature* **585**, 357–362.
- Intel Corporation (2023). *Intel Math Kernel Library*, <https://www.intel.com/content/www/us/en/developer/tools/oneapi/onemkl.html>.
- Jiang, Z., Strzalka, J. W., Wang, J., Zhang, Q. & Dufresne, E. M. (2023). *Beamline 8-ID-E, I: Wide-Angle/Small-Angle XPCS*, https://www.aps.anl.gov/Beamlines/Directory/Details?beamline_id=243. Accessed: 2025–09–02.
- Jo, W., Westermeier, F., Rysov, R., Leupold, O., Schulz, F., Tober, S., Markmann, V., Sprung, M., Ricci, A., Laurus, T., Aschkan, A., Klyuev, A., Trunk, U., Graafsma, H., Grübel, G. & Roseker, W. (2021). *IUCr* **8**, 124–130.
- Joosten, J. G. H., Geladé, E. T. F. & Pusey, P. N. (1990). *Phys. Rev. A* **42**, 2161–2175.
- Kirsch, S., Frenz, V., Schärfl, W., Bartsch, E. & Sillescu, H. (1996). *J. Chem. Phys.* **104**, 1758–1761.
- Lumma, D., Lurio, L. B., Mochrie, S. G. J. & Sutton, M. (2000). *Rev. Sci. Instrum.* **71**, 3274–3289.
- Madsen, A., Flueraşu, A. & Ruta, B. (2016). *Synchrotron Light Sources and Free-Electron Lasers: Accelerator Physics, Instrumentation and Science Applications*, edited by E. J. Jaeschke, S. Khan, J. R. Schneider & J. B. Hastings, pp. 1617–1641. Cham: Springer International Publishing.
- Martinelli, A., Caporaletti, F., Dallari, F., Sprung, M., Westermeier, F., Baldi, G. & Monaco, G. (2023). *Phys. Rev. X* **13**, 041031.
- Pecora, R. (1985). Editor. *Dynamic Light Scattering*. New York: Plenum Press.
- Perakis, F., Amann-Winkel, K., Lehmkuhler, F., Sprung, M., Marie-dahl, D., Sellberg, J. A., Pathak, H., späh, A., Cavalca, F., Schlesinger, D., Ricci, A., Jain, A., Massani, B., Aubree, F., Benmore, C. J., Loerting, T., Grübel, G., Pettersson, L. G. M. & Nilsson, A. (2017). *Proc. Natl Acad. Sci. USA* **114**, 8193–8198.
- Raimondi, P., Benabderrahmane, C., Berkvens, P., Biasci, J. C., Borowiec, P., Bouteille, J. F., Brochard, T., Brookes, N. B., Carmignani, N., Carver, L. R., Chaize, J. M., Chavanne, J., Checchia, S., Chushkin, Y., Cianciosi, F., Di Michiel, M., Dimper, R., D'Elia, A., Einfeld, D., Ewald, F., Farvacque, L., Goirand, L., Hardy, L., Jacob, J., Jolly, L., Krisch, M., Le Bec, G., Leconte, I., Liuzzo, S. M., Maccarrone, C., Marchial, T., Martin, D., Mezouar, M., Nevo, C., Perron, T., Plouviez, E., Reichert, H., Renaud, P., Revol, J. L., Roche, B., Scheidt, K. B., Serriere, V., Sette, F., Susini, J., Torino, L., Versteegen, R., White, S. & Zontone, F. (2023). *Commun. Phys.* **6**, 82.
- Rössler, E. A. & Becher, M. (2025). *J. Chem. Phys.* **163**, 054510.
- Ruta, B., Baldi, G., Monaco, G. & Chushkin, Y. (2013). *J. Chem. Phys.* **138**, 054508.
- Ruta, B., Zontone, F., Chushkin, Y., Baldi, G., Pintori, G., Monaco, G., Rufflé, B. & Kob, W. (2017). *Sci. Rep.* **7**, 3962.
- Schätzel, K. (1983). *Opt. Acta: Int. J. Opt.* **30**, 155–166.
- Schätzel, K. (1987). *Appl. Phys. B* **42**, 193–213.
- Scott, J. & Tüma, M. (2023). *An Introduction to Sparse Matrices*, pp. 1–18. Cham: Springer International Publishing.
- Shpyrko, O. G. (2014). *J. Synchrotron Rad.* **21**, 1057–1064.
- Sprung, M., Westermeier, F., Weschke, D., Rysov, R., Roseker, W. & Anthuparambil, N. D. (2023). *P10 – coherence application beamline*, https://photon-science.desy.de/facilities/petra_iii/beamlines/p10_coherence_applications/unified_data_sheet_p10/index_eng.html. Accessed: 2025–09–02.
- Steinlechner, J. (2018). *Philos. Trans. R. Soc. A* **376**, 20170282.
- Sutton, M. (2008). *C. R. Phys.* **9**, 657–667.
- Sutton, M., Laaziri, K., Livet, F. & Bley, F. (2003). *Opt. Express* **11**, 2268–2277.
- Zhang, Q., Bahadur, D., Dufresne, E. M., Grybos, P., Kmon, P., Leheny, R. L., Maj, P., Narayanan, S., Szczygiel, R., Ramakrishnan, S. & Sandy, A. (2017). *Phys. Rev. Lett.* **119**, 178006.
- Zhang, Q., Dufresne, E. M., Grybos, P., Kmon, P., Maj, P., Narayanan, S., Deptuch, G. W., Szczygiel, R. & Sandy, A. (2016). *J. Synchrotron Rad.* **23**, 679–684.
- Zhang, Q., Dufresne, E. M., Nakaye, Y., Jemian, P. R., Sakumura, T., Sakuma, Y., Ferrara, J. D., Maj, P., Hassan, A., Bahadur, D., Ramakrishnan, S., Khan, F., Veseli, S., Sandy, A. R., Schwarz, N. & Narayanan, S. (2021). *J. Synchrotron Rad.* **28**, 259–265.
- Zhang, Q., Dufresne, E. M., Narayanan, S., Maj, P., Koziol, A., Szczygiel, R., Grybos, P., Sutton, M. & Sandy, A. R. (2018). *J. Synchrotron Rad.* **25**, 1408–1416.