JOURNAL OF SYNCHROTRON RADIATION

ISSN 1600-5775

A new method for studying sub-pulse dynamics at synchrotron sources

James Wingert,* Andrej Singer and Oleg G. Shpyrko

Department of Physics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA. *Correspondence e-mail: jcwingert@physics.ucsd.edu

Received 14 May 2015
Accepted 21 July 2015

Edited by G. Grübel, HASYLAB at DESY, Germany

Keywords: coherent X-ray scattering; speckle; statistics; ultrafast dynamics; XPCS.

The possibility of studying dynamics at time scales on the order of the pulse duration at synchrotron X-ray sources with present avalanche photodiode point detection technology is investigated, without adopting pump–probe techniques. It is found that sample dynamics can be characterized by counting single and double photon events and an analytical approach is developed to estimate the time required for a statistically significant measurement to be made. The amount of scattering required to make such a measurement possible presently within a few days is indicated and it is shown that at next-generation synchrotron sources this time will be reduced dramatically, *i.e.* by more than three orders of magnitude. The analytical results are confirmed with simulations in the frame of Gaussian statistics. In the future, this approach could be extended to even shorter time scales with the implementation of ultrafast streak cameras.

1. Introduction

Studying fast dynamics with X-rays is of ultimate importance in many areas of research including chemistry (Bressler *et al.*, 2009), condensed matter physics (Trigo *et al.*, 2013) and atomic physics (Young *et al.*, 2010). The recent development of new generation X-ray sources allows for the combination of atomic resolution due to X-ray wavelengths and fast time scales due to the pulsed nature of these sources. Conventional X-ray photon correlation spectroscopy (XPCS) measurements rely on correlating sequential intensity measurements. For this reason, the time separating the individual pulses of the source (of the order of 100 ns) is often viewed as a lower limit on the time scales that can be measured. Pump–probe experiments at storage ring sources are inherently limited by the pulse duration, about 100 ps, and suffer from the low repetition rate of pump laser systems. To go beyond this limit, slicing techniques are applied with an increase of temporal resolution by two orders of magnitude, however, at a significant expense of photon flux (Möhr-Vorobeva *et al.*, 2011). The development of X-ray free-electron lasers (XFELs) (Ackermann *et al.*, 2007; Emma *et al.*, 2010; Ishikawa *et al.*, 2012; Allaria *et al.*, 2012) with ultrashort pulse durations (10–100 fs) has led to an improvement in time resolution by about three orders of magnitude. However, today XFELs suffer from intrinsic fluctuations and are still rare and highly overbooked. Here we propose an alternative method of utilizing coherent X-ray scattering to enhance the time resolution at synchrotron sources and go below the pulse duration without sacrificing flux, and to study equilibrium dynamics of various systems.

Coherent X-ray radiation is a superb probe for studying dynamics on the atomic scale. In a typical XPCS experiment, speckle is observed due to interference (Sutton *et al.*, 1991)

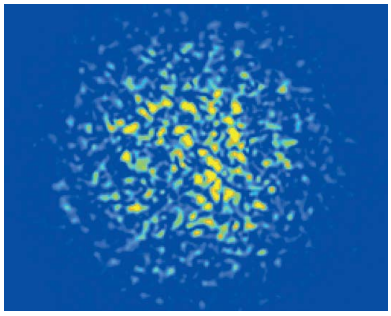

and small changes of the sample can be tracked using intensity correlation techniques. Dynamics for times ranging from fractions of a microsecond to hours (Grübel & Zontone, 2004; Shpyrko *et al.*, 2007; Sikharulidze *et al.*, 2003) have been studied using XPCS. However, similar to intensity interferometry experiments (Brown & Twiss, 1956) at synchrotron sources, the time resolution is limited by the pulse duration of the source (Ikonen, 1992). This is significantly shorter than the time scales that can be observed using conventional XPCS, with which it is not possible to study time scales shorter than the interpulse separation. Synchrotrons are nowadays the principle sources of high-brilliance X-ray beams and we anticipate that our technique can extend the application of these sources to the sub-100 ps time regime, without using pump–probe methods. At present, only pump–probe experiments are able to access time scales shorter than the pulse separation, but these are inherently out-of-equilibrium studies.

Studying dynamics at time scales of the order of the pulse duration at synchrotron sources will necessitate working in a low photon regime. At the very low photon limit, the basis for assessing the dynamics within a sample will naturally be cast in terms of the relative frequency of single and double photon events measured at a detector. Single and double photon events can be counted using an avalanche photodiode (APD) point detector (Kishimoto *et al.*, 1998). Such a device can be read out rapidly, so that a separate measurement is possible for each synchrotron pulse. For synchrotron X-ray intensities, the vast majority of the recorded intensities will be dark noise, along with some rare single photon events. Also, very rarely, an intensity corresponding to two photons from the same pulse will be recorded. We will call such an event a 'double', whereas measurements registering a single photon will be called 'singles' and those with a dark noise value 'zeros'. Fig. 1(*a*) depicts these three possibilities.

Recently, the feasibility of making XPCS measurements within the time scale of a single synchrotron pulse has been considered by employing a 'streak camera' (Namikawa *et al.*, 2009; Itatani *et al.*, 2002; Chang *et al.*, 1996). Simplistically, such a device 'sweeps' the signal along a column of pixels over the duration of a single pulse so that the vertical direction represents time [see Fig. 1(*b*)]. Unlike the point-detector approach, this allows the time separating individual photons that arrive in a single pulse to be identified. This can be done with an area detector as well so that multiple measurements are made simultaneously. To be attractive, such a device would need to be read out at the ring repetition rate of present synchrotron sources (often about 6.5 MHz), which may prove the most difficult aspect of its implementation. However, it should be noted that this might be made less challenging by the extremely sparse nature of the data, which may lend itself to reduction by some appropriate hardware.

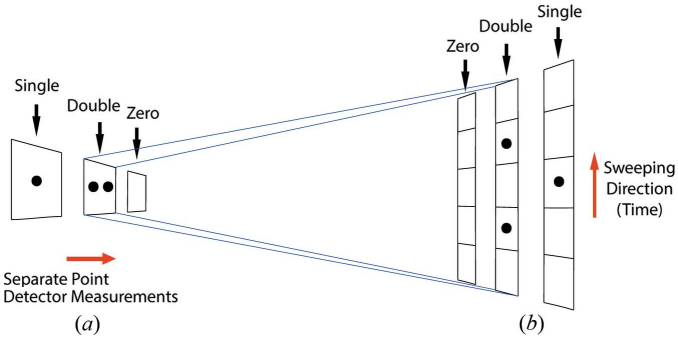

Figure 1
Two approaches to studying intrapulse dynamics: (*a*) depicts separate measurements made with a simple point detector, and three possible outcomes. (*b*) shows the streak-camera approach to the same measurements. A single point detector measurement is equivalent to integrating the counts on the streak camera.

The approach we develop for studying sub-pulse dynamics does not require the implementation of the streak camera concept, although it is a very exciting prospect that we discuss. We show that sample dynamics can still be studied without resolving the temporal separation of the two photons that comprise the doubles, and this is our most important result. Existing APD point detectors are the best candidate for the first measurements to implement our approach. By varying the time scale (τ) of the dynamics of interest, for instance by changing the temperature, one could establish a crossover point for the type of relaxation being studied. The signature of the crossover will be a transition in the relative frequency of the doubles.

Because the APD technique does not depend on correlating individual intensity measurements in time, it becomes possible to go below the 'fundamental' limit of roughly 100 ns imposed by the pulse frequency all the way into the 100 ps regime. This is an impressive three orders of magnitude improvement in the measurable time scale. However, because all of the photons are concentrated into short pulses, no extra flux is required to make a measurement with comparable signal-to-noise as the case where intensities from sequential pulses are being correlated. The proposed technique fully utilizes the capabilities of synchrotron sources.

The feasibility of studying dynamics that occur within a single pulse has already been demonstrated at XFEL sources (Gutt *et al.*, 2009; Hruszkewycz *et al.*, 2012). The much larger coherent flux from such sources results in multiple photons frequently being counted per pixel. Our method is very similar conceptually, but our analysis differs as a result of the lower flux from synchrotrons, which necessitates that we only consider one and two photon events.

2. Estimating number of measurements and time required

We reduce the correlation analysis applied in XPCS to counting single and double photon events. The statistical distribution that generates the speckle will determine how rare the doubles are relative to the singles and will form the basis of our analysis. Our goal here is to estimate the minimum amount of time required in order to draw statistically significant conclusions about the presence of dynamics in a sample of interest. To do so, we perform some straightforward

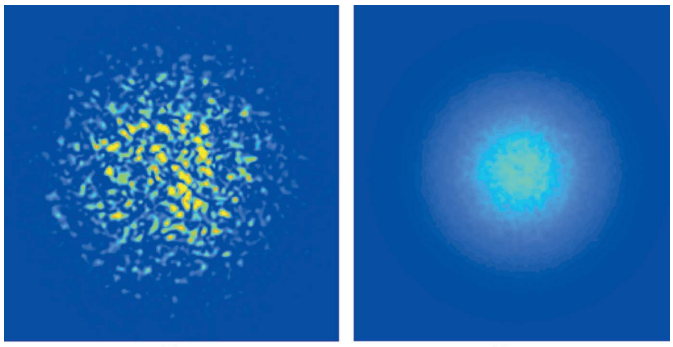

(*a*) (*b*)

Figure 2
(*a*) A simulated speckle pattern from a quasi-static sample. (*b*) The average of 1000 such random speckle patterns, representing a dynamic sample.

statistical calculations comparing the relative frequency of doubles to singles for two extreme cases: speckle from a sample with dynamics significantly slower, and much faster, than the pulse length. In these extreme cases, the spacing of the two photons that make up the double is irrelevant (all spacings are equally probable), so it is best to have in mind point detector measurements for the following.

We begin our analysis by presenting the distributions corresponding to the two extreme cases. A quasi-static sample (static on the time scale of the pulse) will generate a persistent speckle pattern, as in Fig. 2(*a*), whereas a highly dynamic sample will produce a speckle pattern that is 'washed out', as in Fig. 2(*b*). For speckle generated by a quasi-static sample, illuminated by a beam with r transverse coherent modes and scattering an average of m photons per speckle, the probability of k photoevents[1] per speckle is given by the negative binomial distribution (Goodman, 2007)

$$P_{\rm S}(X=k)=\left(\frac{r}{m+r}\right)^{r}\frac{\Gamma(k+r)}{k!\,\Gamma(r)}\left(\frac{m}{m+r}\right)^{k}, \qquad (1)$$

where $\Gamma(k)$ is the gamma function. The negative binomial distribution applies to the intensity distribution of *each point in a single speckle pattern*. However, it will also describe the distribution in time for *one individual speckle in that pattern* if the entire pattern varies on time scales (even many decades) longer than the pulse duration. This will clearly be the case for any system being investigated in the ultrafast regime. It would also be true for a *completely* static sample if either the sample or detector were translated occasionally.

For X-rays scattered by a sample that de-correlates much faster than the duration of a pulse, there will no longer be genuine speckle, and the distribution of intensities will simply be given by the well known Poisson distribution

$$P_{\rm D}(X=k)=\frac{\exp(-m)\,m^{k}}{k!}. \qquad (2)$$

The ratio of doubles to singles forms the basis of our analysis, as it allows one to determine whether some given data are generated by one distribution or the other. The probability of a double, in the quasi-static case, is $P_{\rm S}(X=2)$, given by equation (1). If τ is the time scale of the physical process of interest, then in the $\tau\rightarrow 0$ limit (the 'dynamic' sample), the probability of a double will be $P_{\rm D}(X=2)$, according to equation (2). In both cases, the number of doubles will be binomially distributed (Ross, 2006), with the expected number of doubles given by $\mu_j = NP_j(X=2)$, and a variance of $\sigma_j = NP_j(X=2)[1-P_j(X=2)]$ around this value, where j is D or S. In the low photon limit, both distributions have an expected number of singles equal to Nm.[2]

We now want to calculate how many individual measurements need to be made to discriminate between data coming from the quasi-static and dynamic samples for a given set of parameters. To do so, we suppose that our sample is in fact quasi-static (our null hypothesis), and ask how long one must measure in order to rule out the $\tau\rightarrow 0$ case at a statistically significant level. A standard t-test (Ross, 2006) for this null hypothesis would then compute $t=(Y-\mu_{\rm D})/\sigma_{\rm D}$, where Y is the actual number of doubles collected. The expected value of this quantity, if the sample is truly quasi-static, will be

$$E(t)=\frac{\mu_{\rm S}-\mu_{\rm D}}{\sigma_{\rm D}}. \qquad (3)$$

Using equation (2), we calculate the number of measurements N required to discriminate between a quasi-static and dynamic sample. We choose to use a 3σ confidence level for the calculation. Fig. 3 shows a plot of the value of N that solves $E(t)=3$ for different numbers of modes r on a log–log scale. The linear shape of the curve at low m indicates an inverse-square dependence. In fact, it can be shown by a straightforward calculation that for $m\ll 1$, the solution can be approximated by

$$N\approx\frac{18}{m^{2}(r^{2}+r-1)^{2}}. \qquad (4)$$

Notice that there is a minimum in N for each of the curves, which occurs at about $m=0.19$ for the $r=1$ case. The reason for the minimum is that, as m increases beyond this value, the probability of a double for the Poisson distribution approaches that of the negative binomial. The probability of a double becomes *more* likely with Poisson statistics when m reaches about 0.47, for the $r=1$ case (see Fig. 3).

To estimate actual measurement times for a realistic example we consider the Advanced Photon Source (APS), which is a state-of-the-art synchrotron source. In our calculation we assume full coherence, so that $r=1$. We then estimate the coherent flux, m, and compute the time required to discriminate between quasi-static and dynamic systems for four situations, which are summarized in Table 1. The first case corresponds to an approximation for a generic experiment performed at APS with current operating parameters. We consider optimized experimental geometry with the speckle size being matched to one pixel (Falus *et al.*, 2006) and m is calculated by assuming ten photons scattered into our point

[1] Here, the term photoevent means the detection of a single photon.

[2] In principle, one could consider number of triples, *etc.*, but in the low photon limit considered here these will almost never occur.

Table 1
Time required for a 3σ result for selected values of m.

Source	Photons per pixel per pulse, m	Measurements required, N	Time required
APS, typical experiment	1.5×10^{-6}	8×10^{12}	330 h (150 ns pulse spacing)
APS, partial direct beam	0.015	9×10^{4}	7 ms (150 ns pulse spacing)
APS upgrade, typical experiment	7.5×10^{-5}	3.2×10^{9}	4 min (75 ns pulse spacing)
APS upgrade, partial direct beam	0.19†	1400	240 µs (75 ns pulse spacing)

† m should actually be 0.75 in this approximation; the smaller value used corresponds to a minimum in N [see Fig. 3(*b*)].

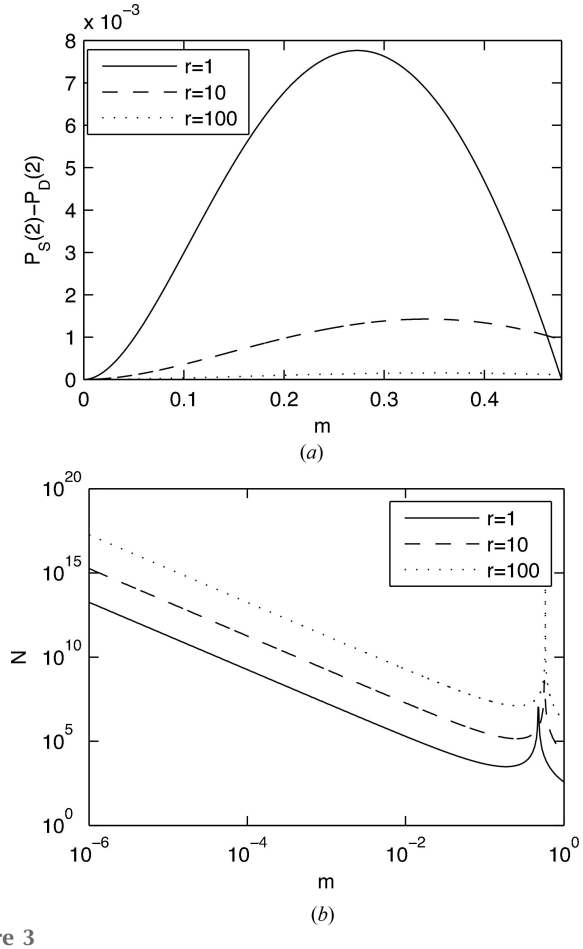

Figure 3
(*a*) The difference in the probability of a double for the two distributions. (*b*) The expected number of measurements, N, which must be collected for a 3σ result.

detector (or each speckle) per second. Whether this amount of scattering can be achieved in a geometry that is suitable for measuring length scales that correspond to the time scales in question can be determined by the experimentalist. However, this level of scattering represents an important cutoff, as it results in the measurement being achievable on the order of several days, which is the length of a typical beam time. The second situation would correspond to 10^9 photons s^{-1} of coherent flux and a fraction of 10^{-4} of the direct beam being scattered into the detector.[3] The third and fourth scenarios are the same as the first two, but are estimated for the proposed APS upgrade, where coherent flux is expected to be about two orders of magnitude higher. In all cases, if the sample is indeed quasi-static, and one collects the number of pixel values listed in Table 1, the expected total number of doubles is about 20, compared with about ten in the Poisson case.

It is worth discussing how the measurement times computed here change if we do not assume full coherence. This is a practically important question, because synchrotrons are not currently fully coherent sources. It is possible, however, to apply spatial filtering to reduce the number of coherent modes at the expense of flux. Our calculations show that if one is in a linear regime where modes and flux can be scaled by the same factor,[4] it is always preferable to have fewer coherent modes. However, for low values of m, such as those listed as typical experimental values here, there is almost no change (<0.3%) in measurement time. For the values of m corresponding to partial direct beam (see Table 1), measurement times could be reduced by more than a factor of two for spatially filtered beams. Importantly, spatial filtering becomes less effective in the high-coherence regime (Singer & Vartanyants, 2011) and an optimization of the experimental parameters will be required in this case.

3. Numerical simulations

To confirm our analytical findings we perform simulations in the frame of Gaussian statistics (Goodman, 2007). In particular we simulate an ensemble of scatterers, whose dynamics can be represented by a predefined time scale. If we illuminate such a system with coherent X-rays and the typical geometrical conditions for XPCS are met (Falus *et al.*, 2006), the time dynamics can be studied by intensity evolution of the speckle. For simplicity, we consider only a single momentum transfer vector and assume the evolution of the speckle behaves according to Gaussian statistics. This can be simulated using a model presented by Pfeifer *et al.* (2010). The starting point of the model is a signal in time with constant amplitudes and completely randomized phases. To induce correlations in time, one can apply filtering in the Fourier space, *i.e.* the signal is Fourier transformed, multiplied by a filter function, which is a Fourier transform of the desired autocorrelation function in time, and Fourier transformed back. To obtain a negative exponential autocorrelation function of the form $\exp(-\alpha t)$, which is typical for diffusive dynamics with a time constant $\tau = 1/\alpha$ and is found in many systems (Berne & Pecora, 2000), we use a filter function of the form $1/(\alpha + i\omega)$, where ω is the frequency. By modifying the width (τ) of the autocorrelation function, we are able to define the time scale of the system. Fig. 4 shows the simulated time evolution of a speckle for the

[3] This second situation is not intended to represent an achievable scattering cross section for a realistic experiment, but rather to set a sort of lower bound on the measurement time for comparison.

[4] So that one could choose between 10 photons s^{-1} and one coherent mode or 100 photons s^{-1} and ten modes.

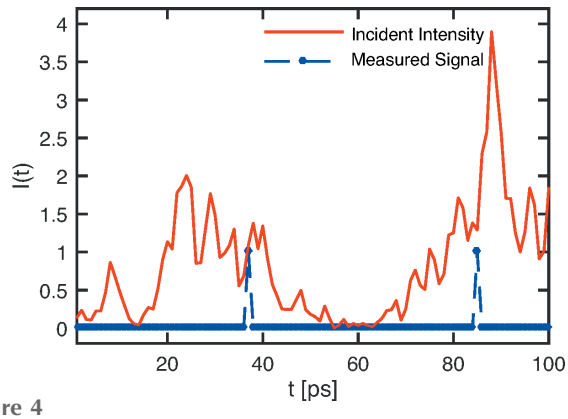

Figure 4
A time series from the random Gaussian signal that generates the data and the location of the discrete photoevents from the Poisson filter.

case of $\tau = 10$ ps time dynamics, if we interpret each data point as corresponding to 1 ps. The width of the 'spikes' corresponds to the aforementioned time scale. Note that for the limiting case of fast dynamics, this intensity profile is simply replaced with a flat line.

In the point detector approach to studying dynamics, it is not possible to measure directly the spacing between photoevents within the sample; it is only possible to say how many photons were incident on the detector per pulse. But by varying a physical parameter related to τ, such as the temperature, the time scale of the dynamics can still be studied. This can be demonstrated by varying the 'spike' width in our simulations. We interpret each point in the time series as a separate pulse/measurement, the absolute duration of which is arbitrary. Fig. 5(*a*) shows the relative probability of a double *versus* τ for different numbers of measurements.[5] For this figure, and the other simulation results shown here, we have chosen $m = 0.01$. As the number of measurements increases, it becomes possible to see the transition from the dynamic to the quasi-static regime. This occurs at about $N = 10^6$ for the simulations shown, which is consistent with our analytical calculations. In future experiments of this nature, given a sufficient quantity of data, it will be possible to estimate relaxation times using this analysis.

If the data instead come from a streak camera device, with time resolution better than the pulse duration, it will be possible to measure directly the spacing of photoevents that make up our doubles, and traditional XPCS methods become possible. This is equivalent to interpreting the data shown in Fig. 4 as a single pulse (which would be a double in this case), rather than a sequence of separate APD measurements (in which case it would count as two singles). In XPCS, the quantity of interest is

$$g_2(\tau) = \frac{\langle I(t)I(t+\tau)\rangle_t}{\langle I(t)\rangle_t^2}, \qquad (5)$$

where $I(t)$ is the intensity measured at time t and τ is the time delay. Typically, g_2 has its maximum value for zero time delay

[5] The time signal is stationary and the simulation in Fig. 4 is statistically equivalent to a series of individual single pulse measurements stacked together in a time series.

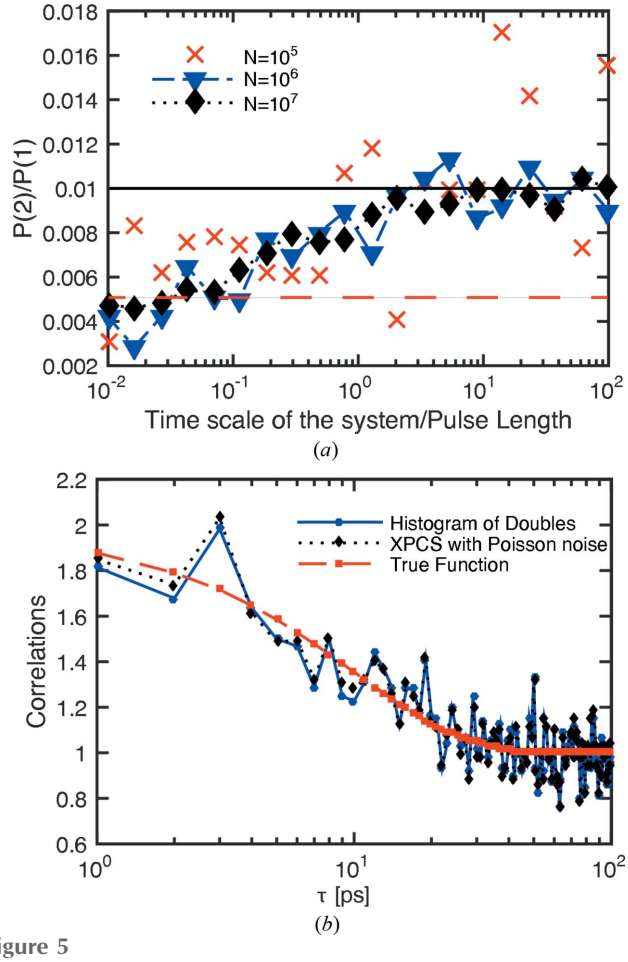

Figure 5
(*a*) The relative frequency of a double *versus* the time scale of the system, for different sample sizes. The horizontal lines represent the limiting quasi-static (solid black) and dynamic (dashed red) cases. (*b*) A normalized histogram of the number of doubles *versus* the time separating them, plotted alongside g_2.

and the width of the function represents the time scale of the dynamics of the system. In the extremely low photon regime we consider here, the numerator will take on values of 0 or 1 almost exclusively, because nearly all pixel values are either 0 or 1. In this case, g_2 simply becomes the probability of a double at fixed τ plus one and should, therefore, reduce to a simple (normalized) histogram of the number of doubles *versus* separation time, τ. Any small difference between the histogram values and the g_2 function arise from the presence of a few pixels with a value greater than 1. Fig. 5(*b*) shows a comparison of equation (5) with the histogram approach for the simulations for $\tau = 10$ ps and a time resolution of 1 ps. That the histogram approach is essentially equivalent to traditional XPCS methods in the low photon regime may prove important, as it potentially allows for the reduction of very large sets of data.

4. Conclusion

In conclusion, our analytical findings and numerical simulations show that the analysis of sample dynamics shorter than

a synchrotron pulse duration is possible using current technology. The analysis of sample dynamics in such a regime consists of the study of the relative frequency of single and double photon events. We forsee that with the future implementation of streak cameras, the determination of the complete shape of the autocorrelation function will be achievable. Also, the results typically obtained through XPCS analysis become possible using a simpler, less data-intensive histogram-based approach. Extending the time scales accessible at synchrotron sources into the picosecond range would be a major advancement; many important physical and chemical processes (Vaterlaus *et al.*, 1991; Kamat *et al.*, 1998; Harris *et al.*, 1988; Beaurepaire *et al.*, 1996; Scholl *et al.*, 1997; Raksi *et al.*, 1996; Shank *et al.*, 1982; McCammon *et al.*, 1979; Cavalleri *et al.*, 2001; Serpone *et al.*, 1995) take place in an intermediate regime that is much slower than FEL pulse durations (femtoseconds), but much faster than the synchrotron pulse spacing (nanoseconds).

Acknowledgements

This work was supported by US Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract DE-SC0001805. We acknowledge many helpful discussions with Stephen Kevan (University of Oregon/LBNL).

References

Ackermann, W. *et al.* (2007). *Nat. Photon.* **1**, 336–342.
Allaria, E. *et al.* (2012). *Nat. Photon.* **6**, 699–704.
Beaurepaire, E., Merle, J.-C., Daunois, A. & Bigot, J.-Y. (1996). *Phys. Rev. Lett.* **76**, 4250–4253.
Berne, B. J. & Pecora, R. (2000). *Dynamic Light Scattering: with Applications to Chemistry, Biology and Physics.* New York: Courier Dover Publications.
Bressler, C., Milne, C., Pham, V.-T., ElNahhas, A., van der Veen, R. M., Gawelda, W., Johnson, S., Beaud, P., Grolimund, D., Kaiser, M., Borca, C. N., Ingold, G., Abela, R. & Chergui, M. (2009). *Science*, **323**, 489–492.
Brown, R. H. & Twiss, R. (1956). *Nature (London)*, **177**, 27–29.
Cavalleri, A., Tóth, C., Siders, C. W., Squier, J., Ráksi, F., Forget, P. & Kieffer, J. (2001). *Phys. Rev. Lett.* **87**, 237401.
Chang, Z., Rundquist, A., Zhou, J., Murnane, M. M., Kapteyn, H. C., Liu, X., Shan, B., Liu, J., Niu, L., Gong, M. & Zhang, X. (1996). *Appl. Phys. Lett.* **69**, 133.
Emma, P. *et al.* (2010). *Nat. Photon.* **4**, 641–647.
Falus, P., Lurio, L. B. & Mochrie, S. G. J. (2006). *J. Synchrotron Rad.* **13**, 253–259.
Goodman, J. W. (2007). *Speckle Phenomena in Optics: Theory and Applications.* Englewood: Roberts and Co.
Grübel, G. & Zontone, F. (2004). *J. Alloys Compd.* **362**, 3–11.
Gutt, C., Stadler, L.-M., Duri, A., Autenrieth, T., Leupold, O., Chushkin, Y. & Grübel, G. (2009). *Opt. Express*, **17**, 55–61.
Harris, A., Brown, J. & Harris, C. (1988). *Annu. Rev. Phys. Chem.* **39**, 341–366.
Hruszkewycz, S., Sutton, M., Fuoss, P., Adams, B., Rosenkranz, S., Ludwig, K. Jr, Roseker, W., Fritz, D., Cammarata, M., Zhu, D., Lee, S., Lemke, H., Gutt, C., Robert, A., Grübel, G. & Stephenson, G. B. (2012). *Phys. Rev. Lett.* **109**, 185502.
Ikonen, E. (1992). *Phys. Rev. Lett.* **68**, 2759–2761.
Ishikawa, T. *et al.* (2012). *Nat. Photon.* **6**, 540–544.
Itatani, J., Quéré, F., Yudin, G., Ivanov, M., Krausz, F. & Corkum, P. (2002). *Phys. Rev. Lett.* **88**, 173903.
Kamat, P. V., Flumiani, M. & Hartland, G. V. (1998). *J. Phys. Chem. B*, **102**, 3123–3128.
Kishimoto, S., Ishizawa, N. & Vaalsta, T. (1998). *Rev. Sci. Instrum.* **69**, 384–391.
McCammon, J. A., Wolynes, P. G. & Karplus, M. (1979). *Biochemistry*, **18**, 927–942.
Möhr-Vorobeva, E., Johnson, S. L., Beaud, P., Staub, U., De Souza, R., Milne, C., Ingold, G., Demsar, J., Schaefer, H. & Titov, A. (2011). *Phys. Rev. Lett.* **107**, 036403.
Namikawa, K., Kishimoto, M., Nasu, K., Matsushita, E., Tai, R., Sukegawa, K., Yamatani, H., Hasegawa, H., Nishikino, M., Tanaka, M. & Nagashima, K. (2009). *Phys. Rev. Lett.* **103**, 197401.
Pfeifer, T., Jiang, Y., Düsterer, S., Moshammer, R. & Ullrich, J. (2010). *Opt. Lett.* **35**, 3441–3443.
Ráksi, F., Wilson, K. R., Jiang, Z., Ikhlef, A., Côté, C. Y. & Kieffer, J.-C. (1996). *J. Chem. Phys.* **104**, 6066–6069.
Ross, S. M. (2006). *Introduction to Probability Models.* New York: Academic Press.
Scholl, A., Baumgarten, L., Jacquemin, R. & Eberhardt, W. (1997). *Phys. Rev. Lett.* **79**, 5146–5149.
Serpone, N., Lawless, D., Khairutdinov, R. & Pelizzetti, E. (1995). *J. Phys. Chem.* **99**, 16655–16661.
Shank, C., Yen, R., Fork, R., Orenstein, J. & Baker, G. (1982). *Phys. Rev. Lett.* **49**, 1660–1663.
Shpyrko, O., Isaacs, E., Logan, J., Feng, Y., Aeppli, G., Jaramillo, R., Kim, H., Rosenbaum, T., Zschack, P., Sprung, M., Narayanan, S. & Sandy, A. R. (2007). *Nature (London)*, **447**, 68–71.
Sikharulidze, I., Farago, B., Dolbnya, I. P., Madsen, A. & de Jeu, W. H. (2003). *Phys. Rev. Lett.* **91**, 165504.
Singer, A. & Vartanyants, I. A. (2011). *Proc. SPIE*, **8141**, 814106.
Sutton, M., Mochrie, S. G. J., Greytak, T., Nagler, S. E., Berman, L. E., Held, G. A. & Stephenson, G. B. (1991). *Nature (London)*, **352**, 608–610.
Trigo, M. *et al.* (2013). *Nat. Phys.* **9**, 790–794.
Vaterlaus, A., Beutler, T. & Meier, F. (1991). *Phys. Rev. Lett.* **67**, 3314–3317.
Young, L. *et al.* (2010). *Nature (London)*, **466**, 56–61.