Imperial College London

Ultrafast measurements using x-ray free-electron lasers

Alvaro Sanchez Gonzalez

April 2017

Thesis submitted in partial fulfilment of the
requirements for the degree of
Doctor of Philosophy of Imperial College London
and the
Diploma of Imperial College

Laser Consortium
Quantum Optics and Laser Science Group
Department of Physics
Imperial College London
London, SW7 2AZ
United Kingdom

Abstract

X-ray free-electron lasers (XFELs) are unique tools for ultrafast science, allowing one to generate tunable pairs of few-fs x-ray pulses with photon energies between ∼10 eV and ∼10 keV. XFEL machines are however much larger and much more complex than table-top optical lasers, spanning from several hundreds of metres in length up to a few kilometres. This increased size and complexity makes them very prone to fluctuations and single-shot jitter in the parameters of the generated x-ray pulses such as photon energy, intensity and time delay, of up to ∼1%, ∼100% and ∼15 fs, respectively. Using data from different experiments carried out at the Linac Coherent Light Source (LCLS) and the SPring-8 Angstrom Compact Free Electron Laser (SACLA), the author will show examples of the influence that these fluctuations can have on the experimental signals, proposing solutions to separate this influence from actual meaningful experimental results. The thesis starts with a brief introduction to x-ray science and the working principles of XFELs followed by an extensive description of the common configurations that allow generating pairs of pulses for time-resolved experiments. The first experimental results are then presented for two successful experiments measuring nuclear dynamics in the 10-100 fs timescales in the acetylene and C_{60} molecules. While successful, these experiments already show the effects of the fluctuations, effects that, as shown in the following chapter, are amplified when attempting to use few-fs pulses to measure few-fs dynamics. Although the presence of the fluctuations may be unavoidable, this thesis shows, using data from three additional ultrafast experiments, that the fluctuations can be circumvented by performing a full single-shot characterization of all x-ray pulses and implementing sophisticated data analysis techniques to sort the experimental data. Finally, having exposed the importance of single-shot x-ray characterization, the author proposes and demonstrates a technique based on machine learning to enable such full x-ray characterization at the next generation of high-repetition-rate XFELs.

Declaration

I hereby certify that the entirety of the contents presented in this thesis are the result of my own work unless otherwise referenced or acknowledged in the body of the text or the caption of the figures. A summary of the author contributions is provided in section 1.2.

Declaration of copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

All third party figures obtained from external sources include the corresponding reference in the figure caption. Explicit permission to include each figure has been obtained from the copyright holder as indicated in Appendix A. Additionally, all external figures used fall under the 'Fair Dealing Copyright Exceptions' from the UK Copyright, Designs and Patents Act 1988 §29 and §29A[1]; and 'Fair Use' under the Copyright Law of the United States §107[2]. For any reuse of these figures, researchers must seek permission from the relevant original source as indicated in Appendix A.

[1]UK Intellectual Property Office, `https://www.gov.uk/guidance/exceptions-to-copyright`. Accessed: 03/03/2017.

[2]US Intellectual Property Office, `http://www.copyright.gov/title17/92chap1.html#107`. Accessed: 03/03/2017.

Acknowledgements

First of all I would like to thank my supervisor, Prof. Jon Marangos, for his support and continuous guidance through this Ph.D. Getting here would have not been possible without your invaluable advice. Additionally, I would like to thank Dr. Thomas Barillot, for being my "go-to" postdoc, and my *de facto* second supervisor, even if bureaucracy does not allow me to make it official. I have really enjoyed our relationship both academically and personally, and will never forget the good times driving around California, including the experiments on ultraviolet absorption at the Santa Cruz beach when x-rays were not available at LCLS.

I want to give my most sincere thank you to everyone that has been part of the Laser Consortium during my time at Imperial College. In particular everyone with whom I had the chance to collaborate during the different experiments: (*alphabetically*) Dane Austin, Vitali Averbukh, Bridgette Cooper, Zsolt Diveki, Taran Driever, Leszek Frasinski, Gedeminas Galinis, Simon Hutchinson, Allan Johnson, Paloma Matia, Lukas Miseikis, Marco Ruberti, Thomas Siegel, Emma Simpson, Christian Strüber, Morgane Vacher and Daniel Walke. Also, special thanks to Paul Micaelli and Charles Olivier for making such a key contribution as master's students to the x-ray predictions using machine learning. I wish the best for your future.

I would also like to thank Susan Parker, Andy Gregory and Peter Ruthven for their great technical support. On a non-scientific, but extremely relevant matter, I want to thank Judith Baylis, Marcia Salviato and Loli Sanchez Rey for always going the extra mile to answer my questions and help me to solve my administrative problems.

None of the experiments presented in this thesis would have been possible without the efforts of **many** other researchers involved in the different collaborations with other departments and institutions. While I will not be able to mention everyone, I want to thank all researchers that I had the chance to collaborate with, and in particular the PIs, for making me always feel welcome to their experiments.

Regarding the different experiments that I have been invited to, I would like to thank Vladimir Petrovic and Chelsea Liekhus-Schmaltz for helping me kick-start my Ph.D. by letting me participate in the data analysis of the acetylene experiment. I would like to thank Nina Rohringer for inviting to my first LCLS beamtime without any prior experience, and for the subsequent collaborations with Clemens Weninger, Laurent

Mercadier, Victor Kimberg, et al. in the non-linear experiments at LCLS and FLASH. I would like to thank Nora Berrah, for letting me play an active role as part of her team together with Thomas Wolf, Li Fang, Razib Obaid, James Cryan, Hui Xiong, et al. in the experiment on the Coulomb explosion of C_{60} at LCLS. I would like to thank Jasper van Thor and Christopher Hutchison for giving me the chance to participate together with Clyde Fare, Takanori Nakane, Ann Fitzpatrick, et al. in my first beamtime at SACLA.

I also really appreciate the high quality support that we have always received from LCLS staff, including beamline scientists, accelerator scientists, data acquisition scientists and technicians. In particular I would like to thank Ryan Coffee, Timur Osipov, Christoph Bostedt, Alberto Lutman, Agostino Marinelli, Tim Maxwell and Chris O'Grady for the long discussions and collaboration before, during and after the experiments.

Additionally, I would like to thank Jan-Erik Rubennson, Conny Såthe, Melanie Muck, Marcus Agåker et al. for making their expertise on x-ray imaging and their x-ray grating spectrometer available for many of the experiments. Similarly, I would like to thank Raimund Feifel and Richard Squibb for contributing with their magnetic bottle spectrometer, and for offering their great expertise for its optimization both as an ion and as an electron spectrometer. I would also like to thank Richard on a personal level for the great times that we spent together in between beamtimes and for being so great in general.

I would also like to acknowledge the Science and Technology Facilities Council (STFC) and, in particular, would like to thank Jim Clarke, David Dunning and Neil Thompson, for providing the main source of funding for this Ph.D. and for the interesting XFEL discussions during my visits to the Daresbury Laboratory.

Last but not least, I would like to thank my family, my girlfriend, my friends and everyone who has been close to me during these years. Your support, and your occasional reminders about the important things in life outside science, have allowed me to preserve my sanity during this journey, and I really appreciate that.

List of publications

The lists below contain all published works derived from results obtained within the extent of this Ph.D. For specific author contributions to each of the publications, see section 1.2.

XFEL related publications

Sanchez-Gonzalez, A., Johnson, A. S., Fitzpatrick, A., Hutchison, C. D., Fare, C., Cordon-Preciado, V., Dorlhiac, G., Ferreira, J., Morgan, R. M., *et al.* Coincidence timing of femtosecond optical pulses in an X-ray free electron laser. *In Preparation* (2017).

Hutchison, C. D., Cordon-Preciado, V., Morgan, R. M., Nakane, T., Ferreira, J., Dorlhiac, G., Sanchez-Gonzalez, A., Johnson, A. S., Fitzpatrick, A., *et al.* X-ray free electron laser crystal structures of dark and light states of a reversibly photoswitching fluorescent protein at room temperature. *In Preparation* (2017).

Berrah, N., Sanchez-Gonzalez, A., Jurek, Z., Obaid, R., Xiong, H., Squibb, R., Osipov, T., Lutman, A., Fang, L., *et al.* Femtosecond-resolved observation of the fragmentation of buckminsterfullerene following X-ray multiphoton ionization. *In Preparation* (2017).

Sanchez-Gonzalez, A., Micaelli, P., Olivier, C., Barillot, T., Ilchen, M., Lutman, A., Marinelli, A., Maxwell, T., Achner, A., *et al.* Accurate prediction of X-ray pulse properties from a free-electron laser using machine learning. *Nature Communications* **8,** 15461 (2017).

Sanchez-Gonzalez, A., Barillot, T., Squibb, R., Kolorenč, P., Agaker, M., Averbukh, V., Bearpark, M., Bostedt, C., Bozek, J., *et al.* Auger electron and photoabsorption spectra of glycine in the vicinity of the oxygen K-edge measured with an X-FEL. *Journal of Physics B: Atomic, Molecular and Optical Physics* **48,** 234004 (2015).

Kimberg, V., Sanchez-Gonzalez, A., Mercadier, L., Weninger, C., Lutman, A., Ratner, D., Coffee, R., Bucher, M., Mucke, M., *et al.* Stimulated X-ray Raman scattering–a critical assessment of the building block of nonlinear X-ray spectroscopy. *Faraday Discussions* **194,** 305–324 (2016).

Rohringer, N., Kimberg, V., Weninger, C., Sanchez-Gonzalez, A., Lutman, A., Maxwell, T., Bostedt, C., Monterro, S. C., Lindahl, A., *et al.* in *X-Ray Lasers 2014* 201–207 (Springer, 2016).

Liekhus-Schmaltz, C. E., Tenney, I., Osipov, T., Sanchez-Gonzalez, A., Berrah, N., Boll, R., Bomme, C., Bostedt, C., Bozek, J. D., *et al.* Ultrafast isomerization initiated by X-ray core ionization. *Nature Communications* **6,** 8199 (2015).

Damiani, D., Dubrovin, M., Gaponenko, I., Kroeger, W., Lane, T., Mitra, A., O'Grady, C., Salnikov, A., Sanchez-Gonzalez, A., *et al.* Linac Coherent Light Source data analysis using psana. *Journal of Applied Crystallography* **49,** 672–679 (2016).

Optical laser publications

Simpson, E., Sanchez-Gonzalez, A., Austin, D., Diveki, Z., Hutchinson, S., Siegel, T., Ruberti, M., Averbukh, V., Miseikis, L., *et al.* Polarisation response of delay dependent absorption modulation in strong field dressed helium atoms probed near threshold. *New Journal of Physics* **18,** 083032 (2016).

Simpson, E. R., Sanchez-Gonzalez, A., Siegel, T., Diveki, Z., Hutchinson, S. E., Miseikis, L., Strüber, C., Austin, D. R. & Marangos, J. P. *A beamline for attosecond pump-probe experiments: towards tracking ultrafast electron dynamics in atoms and molecules* in *SPIE Optical Engineering+ Applications* (2015), 958402–958402.

Contents

Abstract **3**

Declaration **4**

Acknowledgements **5**

List of publications **7**

Contents **9**

List of figures **13**

List of tables **17**

1 Introduction **19**
1.1 Thesis structure 20
1.2 Summary of the author contributions 20

2 X-ray science **23**
2.1 X-ray-matter interaction 24
2.1.1 X-ray diffraction 24
2.1.2 X-ray spectroscopy 27
Energy levels in atoms 27
Energy levels in molecules 28
Photoionization 32
Photoexcitation 33
Near-edge x-ray absorption 34
The decay of an excited state 35
Uses for calibration 37
2.2 Time-resolved spectroscopy 38
2.2.1 Molecular and atomic movies: pump and probe approach 38
2.2.2 An example of a time-resolved ultrafast process 39
2.2.3 The minimum required experimental setup 40

2.3 Earlier x-ray sources . . . 41
2.3.1 From x-ray tubes to undulators . . . 42
2.3.2 Generating short x-ray pulses at synchrotrons . . . 45

3 Experimental aspects of XFELs 49
3.1 XFEL principles . . . 49
3.1.1 Electron-bunch acceleration . . . 51
3.1.2 Electron-bunch compression . . . 52
3.1.3 Undulator and SASE . . . 53
3.1.4 Making a short pulse . . . 54
3.1.5 Seeding . . . 55
Self-seeding . . . 56
External seeding . . . 57
3.1.6 High repetition rate . . . 58
3.2 Multipulse techniques . . . 59
3.2.1 X-ray split and delay . . . 59
3.2.2 Multiple undulators . . . 60
3.2.3 Double-slotted foil . . . 60
Single-bunch mode . . . 60
Double-bunch mode . . . 60
3.2.4 Fresh-slice technique . . . 62
3.2.5 Optical-pump x-ray-probe generation . . . 63
3.3 Relevant diagnostics . . . 63
3.3.1 Gas detector and attenuator . . . 63
3.3.2 X-ray grating spectrometer . . . 64
3.3.3 Time-of-flight spectrometer . . . 66
3.3.4 Magnetic-bottle spectrometer . . . 67
3.3.5 Velocity-map imaging . . . 69
3.3.6 Delay-line anode . . . 70
3.3.7 Time tool . . . 72
3.3.8 XTCAV . . . 73
3.3.9 Beam-position monitors . . . 74
3.4 Comparison to optical lasers . . . 75

4 Measuring nuclear dynamics on 10-100 fs timescales 77
4.1 Ultrafast isomerization in acetylene . . . 77
4.1.1 Experimental setup . . . 80
4.1.2 Data analysis . . . 80
4.1.3 Time dynamics of the proton motion . . . 82
4.1.4 Conclusions and outlook . . . 84

4.2 X-ray-induced Coulomb explosion in C_{60} 84
4.2.1 Physical mechanism and prior experiment 85
4.2.2 Experimental setup 86
4.2.3 Argon calibration 87
Ion time-of-flight calibration 88
Intensity calibration 88
4.2.4 Time-resolved results 90
4.2.5 Checking the reliability of the results 92
Total x-ray energy correlation with the delay 92
Spatial overlap correlation with the delay 94
4.2.6 Conclusions and outlook 96

5 Performing a few-fs time-resolved experiment 99
5.1 An example of few-fs time dynamics in isopropanol 100
5.2 Experimental setup 106
5.2.1 X-ray source 106
5.2.2 X-ray diagnostics 106
5.2.3 Interaction region 109
5.3 Calibration 110
5.3.1 Electron-bunch energy calibration 111
5.3.2 Magnetic bottle calibration 111
5.3.3 X-ray grating spectrometer calibration 115
5.3.4 Pulse energy calibration 117
5.4 Extracting time information: XTCAV 119
5.4.1 Calibrating the XTCAV images 120
5.4.2 Obtaining non-lasing references 121
5.4.3 Obtaining the x-ray power 122
5.4.4 Double bunch configuration 125
5.4.5 Limits of the method 126
5.4.6 Availability of the algorithm at LCLS 128
5.5 Looking at static data in glycine 129
5.5.1 Non-resonant processes 129
5.5.2 Resonant processes 130
5.5.3 Possible water contamination 131
5.6 Searching for a time-resolved signal in O_2 135
5.6.1 Data binning 135
5.6.2 Background subtraction 140
5.7 Conclusions and outlook 144

6 Ultrafast experiments on non-linear x-ray optics and timing 147
6.1 Resonant Raman scattering in an XFEL . . . 147
6.1.1 Prior work: self-stimulated x-ray lasing and SRIXS in an atom . . 147
6.1.2 Stimulated resonant x-ray Raman scattering in a molecule 149
6.1.3 Conclusion . . . 153
6.2 Synchronization between an optical laser and an XFEL at SACLA 154
6.2.1 Experimental method . . . 155
The time-tool branch . . . 155
The experiment branch . . . 157
6.2.2 Stability of the timing . . . 160
6.2.3 The position of time zero . . . 162
6.2.4 Conclusions and outlook . . . 165

7 X-ray characterization, fluctuations and machine learning 167
7.1 The problem of the fluctuations . . . 167
7.2 Introduction to machine learning . . . 174
7.2.1 Supervised learning . . . 175
7.2.2 Linear and polynomial regression . . . 176
7.2.3 Support vector regressors . . . 177
7.2.4 Artificial neural networks . . . 179
7.2.5 Iteratively training a model . . . 180
7.3 Proposed machine learning technique . . . 181
7.3.1 Variables involved in the process . . . 182
Slow single-shot variables from complex diagnostics . . . 182
Fast single-shot variables from simple diagnostics . . . 184
Environmental EPICS variables . . . 184
7.3.2 Numerical treatment of the data . . . 184
Preprocessing of the data . . . 185
Training, validation and testing of the models . . . 186
7.4 Results of the predictions . . . 187
7.4.1 Photon energy prediction for a single pulse . . . 187
7.4.2 Spectral-shape prediction for a single pulse . . . 189
7.4.3 Time-delay prediction . . . 191
7.4.4 Photon-energy prediction for a double pulse . . . 193
7.5 Conclusions and outlook . . . 194

8 Conclusion 197

References 201

A List of figure permissions 217

List of figures

2.1 Bragg diffraction 25
2.2 Reconstruction of the photosystem II protein complex 26
2.3 Bohr's model 27
2.4 Electronic levels in sodium 29
2.5 Molecular potential energy curves of H_2 30
2.6 Photoionization in O_2 32
2.7 Photoexcitation in CO 33
2.8 Near-edge x-ray absorption 35
2.9 Auger decay 36
2.10 Spontaneous and stimulated emission 37
2.11 X-ray tube 42
2.12 Synchrotron radiation 44
2.13 Laser slicing in a synchrotron 46

3.1 Peak brilliance of existing XFELs 50
3.2 Typical XFEL layout 51
3.3 Radiofrequency cavity 51
3.4 Electron-bunch compression 52
3.5 Self-amplified spontaneous emission 53
3.6 Single-slotted spoiler 55
3.7 Examples of SASE spectra 56
3.8 Self-seeding technique 57
3.9 External seeding technique 58
3.10 X-ray split and delay apparatus 59
3.11 Emission in two undulator sections 60
3.12 Double-slotted spoiler for one bunch 61
3.13 Double-slotted spoiler for two bunches 61
3.14 Fresh-slice technique 62
3.15 Gas detector chamber 64
3.16 X-ray grating spectrometer layout 65
3.17 Ion time-of-flight spectrometer trace 66

3.18 Magnetic bottle spectrometer . . . 68
3.19 VMI spectrometer . . . 70
3.20 Delay-line anode . . . 71
3.21 Time tool for optical laser synchronization . . . 73
3.22 XTCAV setup at LCLS . . . 74

4.1 Acetylene isomerization scheme . . . 78
4.2 Carbon-carbon-deuteron angle . . . 82
4.3 Time dynamics of the momentum distribution in the acetylene isomerization 83
4.4 Time dependent ratio of acetylene isomers . . . 83
4.5 XTCAV traces for the fresh-slice mode . . . 87
4.6 Ion TOF calibration . . . 89
4.7 Intensity calibration in argon . . . 90
4.8 Time dynamics in C_{60} . . . 91
4.9 Pulse-energy filtering . . . 93
4.10 Time dynamics in C_{60} after filtering . . . 94
4.11 Spatial overlap verification in argon . . . 95
4.12 Spatial overlap correlation with delay using argon . . . 96

5.1 Eigenstates of the Hamiltonian of the isopropanol cation . . . 100
5.2 Inner valence dynamics in isopropanol . . . 103
5.3 Pump-probe scheme for measuring a hole opening process in isopropanol . 103
5.4 Expected signal from a time-resolved experiment in isopropanol . . . 104
5.5 X-ray diagnostics for two pulses . . . 108
5.6 Interaction region in an XFEL experiment . . . 109
5.7 Electron-bunch energy calibration . . . 112
5.8 Calibration points for the magnetic bottle . . . 113
5.9 Auger electron peaks in O_2 used for calibration. . . . 114
5.10 Calibration points at different retardation potentials . . . 115
5.11 Calibration with oxygen photoelectrons . . . 116
5.12 Using neon for calibration . . . 116
5.13 Gas detector calibration . . . 118
5.14 XTCAV images for a single bunch . . . 119
5.15 Temporal streaking and energy dispersion of the electron bunch . . . 120
5.16 XTCAV reconstruction for a single bunch . . . 123
5.17 XTCAV images for a double bunch . . . 126
5.18 XTCAV reconstruction for two bunches . . . 127
5.19 Non-resonant Auger spectra for a glycine sample . . . 130
5.20 Resonant Auger spectra for a glycine sample . . . 131
5.21 Experimental Auger spectra compared to water and glycine . . . 132
5.22 Signal from the residual gas analyser during the experiment . . . 133

5.23 Ion spectra at the interaction region . . . 134
5.24 Experimental photoelectron spectra compared to water and glycine 134
5.25 Expected signal from a time-resolved experiment in O_2 . . . 136
5.26 Full single-shot diagnostics . . . 137
5.27 Distributions of different x-ray properties . . . 138
5.28 Correlations between x-ray properties . . . 139
5.29 Binning time-dependent data . . . 141
5.30 Single-pulse references . . . 143
5.31 Pump and probe contributions to the background in oxygen . . . 144
5.32 Time-dependent data without background . . . 145

6.1 X-ray lasing scattering in neon . . . 148
6.2 Photon energy dependence of the neon laser . . . 150
6.3 Resonant stimulated Raman scattering in CO . . . 151
6.4 Spectra for the double-pulse seeded XFEL emission . . . 152
6.5 Finding the SRIXS signal based on the intensity dependence . . . 153
6.6 Finding the SRIXS signal in different bins . . . 154
6.7 Experimental setup using the time tool . . . 155
6.8 Time-tool retrieval . . . 157
6.9 SACLA timing drift . . . 158
6.10 Effect of the time-tool correction in Ce:YAG . . . 159
6.11 Drift of the absolute calibration . . . 161
6.12 Model of the Ce:YAG absorption . . . 164

7.1 Distributions of different x-ray properties for a single run . . . 168
7.2 Correlations between x-ray properties for a single run . . . 169
7.3 Correlation between x-ray properties and electron-beam parameters . . . 171
7.4 Correlation between x-ray properties and electron-beam parameters (continued 1.) . . . 172
7.5 Correlation between x-ray properties and electron-beam parameters (continued 2.) . . . 173
7.6 Polynomial models . . . 177
7.7 Support vector regressor . . . 178
7.8 Artificial neural network . . . 180
7.9 Gradient descent algorithm . . . 181
7.10 Scheme of the proposed machine learning technique . . . 182
7.11 Experimental setup for machine learning predictions . . . 183
7.12 Training of the predictive models . . . 185
7.13 Prediction of the photon energy for a single pulse . . . 188
7.14 Spectral-shape prediction . . . 190
7.15 Prediction of the time delay . . . 192

7.16 Prediction of the photon energy for a double pulse 193

List of tables

7.1 Optimal values of the hyperparameters 187
7.2 Results of the predictions 188

Chapter 1

Introduction

As developments in science and technology progress, the need for a better understanding of dynamics and fundamental interactions between light and matter that occur on ultrashort timescales becomes more important. Several processes not only in physics, but also in chemistry, biology and material science, could benefit from the study of the dynamics of the fast nuclear and even faster electronic dynamics to, for instance, create better light harvesting devices or optimize chemical synthesis. In particular, the interaction between x-ray light and matter is of special importance due to the shorter wavelengths, which allow diffraction experiments to image molecular structures, and the increased photon energy allowing an exploration of the deeper energy levels in matter.

The dynamics of interest for some of those processes are extremely fast, typically occurring in the few-fs/attosecond regime (10^{-15}-10^{-18} s)[1, 2]. This is at the limit of what can currently be measured using current technology and, in particular for x-ray science [3], requires the use of hard x-ray free-electron lasers (x-ray FELs or XFELs) [4].

This Ph.D. thesis will elaborate on the use of x-ray free-electron lasers for ultrafast x-ray science, identifying the key aspects required to perform successful ultrafast few-fs experiments making use of the existing XFEL facilities. This will be done in the context of numerous XFEL experiments performed during this Ph.D., mainly at the Linac Coherent Light Source (LCLS) the first operating hard x-ray FEL at SLAC, Stanford, USA [5] but also the SPring-8 Angstrom Compact Free Electron Laser (SACLA), Japan [6] and the Free Electron Laser in Hamburg (FLASH), Germany [7, 8].

Many of the results related to those experiments have been already published or are close to publication (see List of publications) as standalone manuscripts, including exhaustive experimental details and scientific discussion. In this thesis the author will focus on the general challenge of performing few-fs time-resolved experiments at XFELs, concentrating in the particular issues and practicalities of the different XFEL configurations and data analysis techniques, and making only brief introductions to the relevant experiments.

1.1 Thesis structure

The contents of this thesis are divided in six different chapters: two introductory chapters followed by four chapters with experimental results.

The first introductory chapter (chapter 2) is an introduction to x-ray science, with an explanation of the basic physical mechanisms of interest, in particular to x-ray diffraction and x-ray spectroscopy, including time-resolved spectroscopy. This chapter also briefly presents the history of previous x-ray sources, and how they led to the creation of the current XFELs.

The second introductory chapter (chapter 3) discuses all important experimental aspects of XFELs that are relevant to the thesis. This includes the principles of operation of XFELs, the typical detectors used for XFEL experiments, and the multiple possibilities available to generate pairs of pulses for time-resolved experiments.

The first experimental results are shown in chapter 4 consisting of two successful experiments measuring time-resolved nuclear dynamics in the 100-fs timescales at LCLS. This leads to chapter 5, which concentrates on few-fs experiments and the challenges of performing a few-fs experiment in an XFEL through the description of a as yet unsuccessful, but rationally ideal regarding experimental setup, few-fs experiment performed at LCLS. Additionally, with the same goal of achieving few-fs experiments at XFELs, chapter 6 presents the results from another two experiments at LCLS and SACLA, this time focusing on generating shorter and more-coherent pulses, and improving the timing synchronization between x-ray and optical pulses.

Following the first three experimental chapters, which emphasize the importance of having access to accurate single-shot diagnostics of the few-fs x-ray pulses, the final chapter (chapter 7) discusses how to achieve this in the next generation of high repetition-rate XFELs, where many of the complex diagnostic devices will not be able to work at the full repetition rate. In this chapter the author will describe and demonstrate a technique developed during this Ph.D. that uses machine learning to predict variables typically measured with complex diagnostics, such as the inter-pulse delay, the photon energy or even the spectral shape. The author will show how this can be used to obtain a full x-ray characterization for every event by using only data from simple diagnostics as input, and a small fraction of fully characterized events.

1.2 Summary of the author contributions

Experiments at XFELs are typically performed in intense campaigns of only three to ten days of duration, requiring large collaborations of researchers so different responsibilities can be distributed between different groups of people. Some of these responsibilities may be, for example, preparing the sample, installing and maintaining a sample delivery device or a detector, controlling the recording of data, keeping the lab book, etc. The contributions of the author of this thesis have varied through the different experiments

according to the author's experience, but they have always been mainly focused on the data analysis and interpretation of the results. This included both real-time work during the experimental campaigns to verify the quality of the data as it was being collected, and off-line work during the few-month or even year periods after the experiments, to refine the data analysis and interpretation of the signals.

Regarding the experiment on acetylene isomerization (LCLS, June 2013, section 4.1), as it was carried out before the beginning of this Ph.D., the author could only contribute to the data analysis, in particular optimizing peak finding algorithms to improve the detection efficiency and accuracy of the delay-line anode (Hexanode) used for ion detection. This work eventually led to a publication [9].

The experiments on x-ray lasing and resonant Raman scattering (section 6.1) were performed in two different experimental campaigns. During the first campaign (LCLS, February 2014) the author contributed to the optical laser setup for molecular alignment (not used in the end), and played a secondary role in the real-time monitoring and analysis of the x-ray grating spectrometer signal. The results from this first experiment are available in reference [10]. In the second campaign (LCLS, July 2015) the author played a main role in the real-time monitoring, analysis and visualization of the results during and after the experiment, performing most of the experimental data analysis relevant to reference [11]. This work included calibrating the x-ray grating spectrometer, working with XFEL parameters and improving x-ray absorption techniques to deal with XFEL fluctuations.

The results on few-fs experiments (chapter 5) and x-ray fluctuations and machine learning (chapter 7) were based on an experiment led by our group with the initial purpose of measuring ultrafast dynamics in glycine (LCLS, February 2015). The author led the real-time monitoring and data analysis before, during, and after the experiment. This included retrieving, calibrating and analysing electron time-of-flight traces (Auger spectroscopy, photoelectron spectroscopy), ion time-of-flight traces, x-ray grating spectrometer data (x-ray characterization, near-edge x-ray absorption) and X-band transverse deflecting cavity (XTCAV) traces (x-ray power and delay retrieval) using all information available to attempt to extract a time-resolved signal. The analysis of the static Auger spectroscopy data in glycine (section 5.5) was published [12]. The software tools developed by the author for the analysis of XTCAV images (section 5.4) were also part of a publication about the LCLS data analysis framework [13]. Finally, the development of a new technique led by the author that uses machine learning to predict x-ray pulse properties, using also additional data from another experimental campaign (April 2015), has been recently accepted for publication [14]. This work involved studying correlations between machine parameters and x-ray properties and optimizing the machine learning process.

Concerning the experiment on the Coulomb explosion of C_{60} (LCLS, February 2016, section 4.2) the author took a main role in the experimental data analysis both during the

experiment and after the experiment. In this case, the data analysis involved optimizing the ion time-of-flight retrieval, studying the fluence and spatial overlap of the pulses at the focus, extracting the ion-yield time-dependent signals, and studying the changes in the signal as a function of the fluctuating x-ray parameters. Experimental results from this experiment are ready and currently waiting for the theoretical modelling performed by a different part of the collaboration before they are published [15].

Similarly, regarding the experiment about optical laser timing (LCLS, July 2016, section 6.2), the author had a major role in the data analysis and retrieval of timing information between the optical and the x-ray pulses, during and after the experiment. This involved performing the retrieval (with the corresponding calibrations) of delay fluctuations from time tool images, and the modelling of the time-resolved absorption in a Ce:YAG crystal used as an additional time diagnostic. While the original purpose of the experiment (time-resolved crystallography) was not successful, new insights about optical x-ray timing obtained from the analysis of the timing data [16] and static crystallography data [17] are being prepared for publication.

Finally, although not discussed in the context of this thesis, during the first year of this Ph.D., the author was also very actively involved in optical-laser-based attosecond transient absorption experiments at Imperial College, with two related publications [18, 19]. The work included building optical setups in vacuum, analysing spectral interference data for precise time-stamping, developing a new software tool to automatically control laboratory devices and record data, performing the experiments, doing Fourier analysis of the time-resolved data, and creating the conceptual design for a new magnetic bottle electron time-of-flight spectrometer that is currently being built, by modelling electron trajectories.

Chapter 2

X-ray science

Since x-ray radiation was discovered by William Röntgen in 1895, it has been widely used for many applications in science related to medicine, physics, chemistry, biology and material science. There are three properties of x-rays that make them unique for these applications: short wavelengths, long penetration depths and high photon energies. The most well-known application in medicine is x-ray imaging, including more sophisticated techniques such as computerized axial tomography (CAT) scans, which takes advantage of the penetrating properties of the hard x-rays to image parts of the human body. Chemistry, biology and material science benefit from x-ray diffraction images, possible thanks to the short wavelengths that allow imaging of crystals and molecules to the atomic level. Physics and chemistry also take advantage of the high photon energies available to access energy levels and trigger processes that otherwise would not be accessible with optical, infrared or ultraviolet radiation.

Formally, x-ray radiation is the name given to the part of the electromagnetic spectrum with wavelengths between 10 nm and 10 pm, or equivalently, photon energies between 124 eV and 124 keV. Together with gamma rays, and extreme ultraviolet (XUV), x-rays fall under the category of ionizing radiation due to their large photon energies making them capable of removing electrons from matter. The x-ray spectrum is typically divided in two different regions: *soft x-rays* and *hard x-rays*. Soft x-rays correspond to the less energetic part to the spectrum, which is strongly absorbed by matter, and in particular by the atmosphere, with typical penetration lengths of only a few micrometres. On the other hand hard x-rays correspond to the high-energy end of the spectra and have the property of very easily penetrating thick sections of otherwise opaque materials, including air. The boundary between the hard x-rays and the soft x-rays is not defined sharply, as the transition between absorption and transmission is smooth, and typically different for different materials. Generally, x-ray radiation with photon energies above 5-10 keV is considered "hard" x-ray radiation.

In this chapter the author will explain some of the main interactions between x-rays and matter and how to make use of them in ultrafast science through the so called *pump*

and probe experiments. The history of the different existing x-ray sources that led to the development of the current x-ray free-electron lasers will also be briefly presented.

2.1 X-ray-matter interaction

There are many types of interaction between x-rays and matter of interest in the x-ray science performed at XFELs that can be divided into two main groups: x-ray diffraction and x-ray spectroscopy.

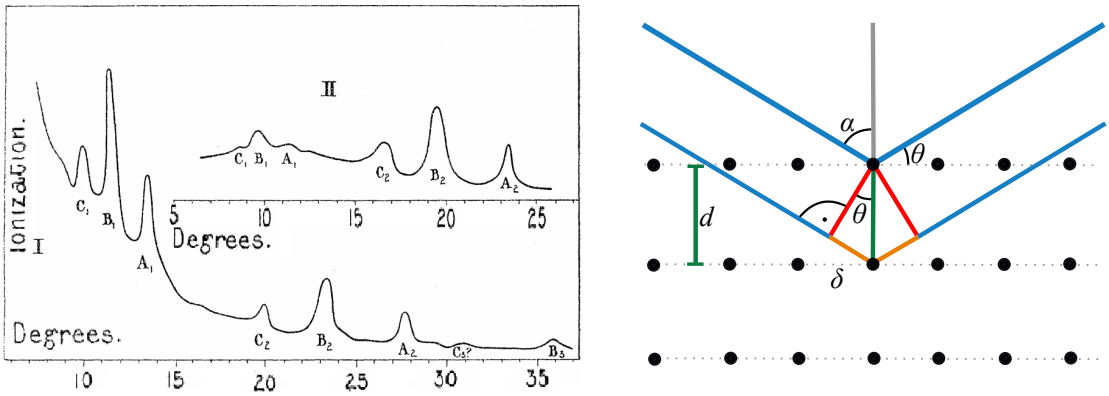
2.1.1 X-ray diffraction

X-ray diffraction [20] consists of using x-ray light to illuminate a sample while collecting the scattered light transmitted or reflected by the sample in an spatially/angularly resolved manner. The collected light resulting from the interaction of the x-ray radiation and the sample is denominated a *diffraction pattern* and can be used to infer information about the geometry of the sample. Similarly to optical diffraction, which requires the wavelength of the light to be similar to the size of the diffracting objects to occur, x-ray diffraction allows one to extract information about structures between 10 nm and 10 pm in size, depending on the wavelength. This range includes the typical inter-atomic distances of a few ångströms (1 Å = 0.1 nm = 100 pm) found in molecules and crystals. As a consequence, hard x-rays with wavelengths in the order of 1 Å can be used to image molecules and crystals to the atomic level.

This diffraction effect was measured and quantified for the first time by W.H. Bragg et al. [21] in 1913 by studying the angle-dependent reflection of x-ray light in different crystals. Peaks in the x-ray reflection were observed at very specific incidence angles (figure 2.1a), later called Bragg peaks, which led to the development of the Bragg diffraction model. This model describes the x-ray interaction with the crystal as multiple reflections occurring at the different planes of atoms inside the crystal (figure 2.1b). These multiple reflections are prone to interference effects when they overlap, with the character of the interference (constructive or destructive) being determined by the difference in the path length of each of the reflections. Figure 2.1b shows that this difference in path length (orange region) can be expressed in terms of the separation between the planes d and the angle θ, complementary to the incidence angle α, as $2d\sin(\theta)$. Additionally, for a constructive interference to occur, the difference in path length has to correspond to a whole number of times n the wavelength λ. This leads to the famous Bragg condition for reflection,

$$2d\sin(\theta) = n\lambda. \tag{2.1}$$

Since the discovery of Bragg diffraction this procedure has been widely used to characterize all kind of simple crystalline periodic structures, obtaining the lattice structure and the lattice parameters [22].



(a) (b)

Figure 2.1: Bragg diffraction. (a) Strength of the reflected x-ray beam from a 100 face (I) and a 111 face (II) surface in rock salt as function of the incidence angle. The figure shows that the reflectivity peaks at certain angles. These peaks labelled as $\{\mathrm{A, B, C}\}_{\{1,2\}}$ correspond to the first and second order reflections from three different planes of the crystal. (b) Bragg model of x-ray diffraction. The difference in optical path between the reflections from different planes of the crystal 2δ is related to the distance d between the planes and the incidence angle α (or its complementary θ). Bragg's law predicts that maximum reflection is achieved when the interference between the reflections from different planes interfere constructively. Figure (a) extracted from [21]. Figure (b) is a modified public domain figure originally made by Matthias M.

X-ray diffraction is however not limited to simple periodic structures, but can also be used in the Fraunhofer approximation of diffraction [23], where the intensity of a diffraction pattern $U(\mathbf{k})$ corresponds to the magnitude of the Fourier transform of the spatial features,

$$U(\mathbf{k}) \propto |\int_{\text{size}} a(\mathbf{r}) e^{-i\mathbf{k}\cdot\mathbf{r}} d\mathbf{r}|, \tag{2.2}$$

where $\mathbf{r}$ are the spatial coordinates, $\mathbf{k}$ are the wave vectors corresponding to different angles, and $a(\mathbf{r})$ is the shape of the sample. This is consistent with Bragg diffraction, as the Fourier transform of a periodic structure is also periodic, corresponding to the different orders n of the Bragg peaks. While no information about the phase is available when measuring the diffraction patterns, different phase retrieval [24] algorithms relying on an oversampled diffraction pattern can be used to extract the spatial information.

The main limitation to this approach is that the probability of diffracting a photon by a single molecule is extremely low. One approach to increase this probability is to crystallize the molecule of interest into a large periodic lattice with multiple repetitions of the molecule. Since in the spatial domain this corresponds to convolving the molecular shape with a Dirac comb, using the convolution product property, the diffraction pattern is the same of that of a single molecule, multiplied by a very finely spaced Dirac comb, which cannot be resolved unless the diffraction pattern is oversampled. This means that all of the molecules in the crystal can contribute constructively to the diffraction pattern, enhancing the signal. While that approach works in many cases, some molecules cannot

be made into large crystals, or will change their shape or properties when crystallized. Assuming that the single-photon diffraction probability cannot be enhanced in these cases, the alternative approach is to increase the number of photons by building brighter x-ray sources. However, we will see in the next section that x-ray photons can interact with matter through many more processes apart from the elastic scattering that occurs during diffraction, and that many of those mechanisms lead to the complete destruction of the sample. By increasing the number of photons, that destruction is guaranteed to happen.

Although the problem of imaging a single molecule or a nano-crystal by diffraction may seem hopeless due to the unavoidable destruction of the molecule, destruction, like any other process in nature, takes some time. If we can provide all the photons needed to produce a diffraction pattern within an extremely short amount of time, we may be able to take an image of the molecule just before it gets destroyed... by the same x-ray pulse used to take the image. This approach to x-ray diffraction in XFELs is known as diffraction-before-destruction [25, 26](figure 2.2), and it is one of the main justifications for the existence of x-ray free-electron lasers, which as we will see later, are the only sources that are bright enough to provide a very large number of x-ray photons ($\sim 10^{12}$) within a few femtoseconds of time.

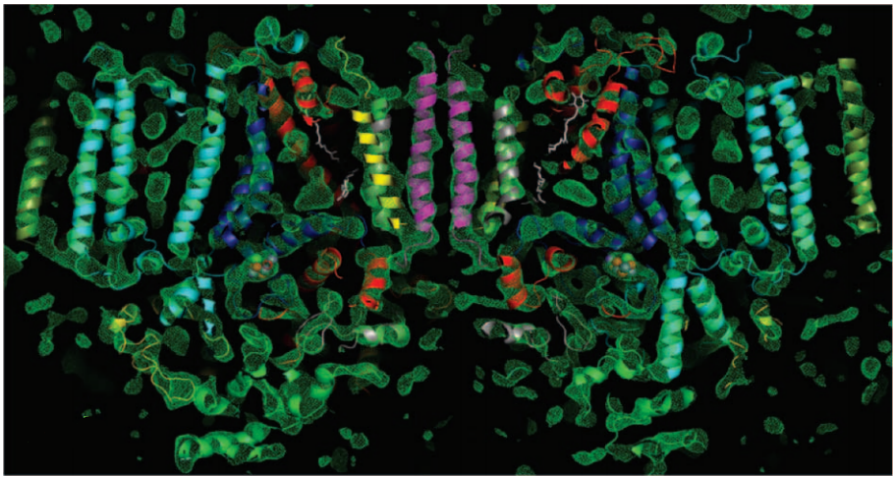


Figure 2.2: X-ray diffraction experiment in photosystem II, a protein complex involved in the photosynthesis. The electron density (green) is retrieved by analysing the diffraction pattern of x-ray pulses (50 fs pulses at LCLS) by nanocrystals of the complex. The short duration of the pulses allows imaging the sample before it is destroyed. Figure extracted from [27].

The timescale in which the destruction of the sample happens (at the nuclear and electronic level) is still a matter of discussion. The author will later present two experiments in the context of this thesis which aimed to gain a further understanding of this (sections 4.1 and 4.2).

2.1.2 X-ray spectroscopy

X-ray spectroscopy [28] consists of using x-ray radiation to study the behaviour of molecules and atoms at the microscopic level by inducing energy transitions. Information about the quantum states of the system can be then extracted by measuring the properties of any emitted electrons (total yield, energy, angular distribution), ions (charge, mass, energy, angular distribution), or photons (total yield, photon energies), as well as any changes in the photon absorption.

Energy levels in atoms

We have known about the quantum character of the electron energies in atoms since Bohr proposed his model of the structure of the atom in 1913 [29, 30]. This model succeeded to explain the discrete emission lines observed in hydrogen, previously modelled using the empirical Rydberg formula, by interpreting the atom as a series of electrons orbiting the nucleus in a discrete number of stable orbits, with the different emission lines corresponding to the transition of the electron between a higher energy orbit into a lower energy orbit (figure 2.3). This model already introduced the *principal quantum number* $n = 1, 2, ...$ to describe the allowed quantum orbits, and showed that the radius r_n and the binding energy E_n associated to the electron orbits were proportional to n^2 and to $1/n^2$ respectively. Orbits with principal quantum numbers 1, 2, 3, etc. are typically also known as the K-shell, L-shell, M-shell, etc.

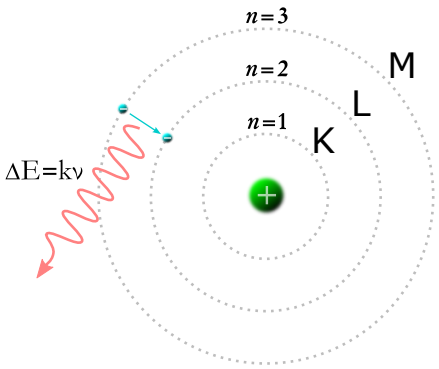


Figure 2.3: Bohr's model of the hydrogen atom. The atom is interpreted as a positively charged nucleus surrounded by a number n of discrete stable electron orbits (K, L, M, etc.). Modified public domain figure of unspecified author.

A few years later, the discovery of the Schrödinger equation [31, 32] allowed physicists to find more realistic analytical forms of the orbits of the hydrogen atom, expressing the electron density probability distribution $\psi(\mathbf{r})$ in spherical coordinates as a spatial wavefunction,

$$\psi_{n,l,m}(\mathbf{r}) = \psi_{n,l,m}(r, \theta, \phi) = R_{n,l}(r) Y_l^m(\theta, \phi), \tag{2.3}$$

where $R_{n,l}$ corresponds to the radial part of the solution (which can be expressed in terms of the Laguerre polynomials) and Y_l^m the angular part of the solution known as the spherical harmonics. These solutions are parametrized in terms of 3 quantum numbers $n = 1, 2, ...$; $l = 0, 1, 2, ..., n-1$ and $m = -l, -l+1, ..., l$, known respectively as the *principal*, *azimuthal* and *magnetic* quantum numbers. Different values of the azimuthal quantum number $l = 0, 1, 2, 3$, indicating states with different spatial angular momentum, are also known as s, p, d, e orbitals. The magnetic quantum number corresponds to the different allowed projections of that angular momentum across an arbitrary axis. Since for the hydrogen-like atoms the energy of the state only depends on the n quantum number, states with the same n value but different l and m values are denominated *degenerate states*. Apart from the spatial wavefunction, electrons possess an intrinsic form of angular momentum called spin with a value of 1/2. This spin can also have different projections $s = -1/2, 1/2$, known as the spin quantum number. Due to its angular nature, the spin can magnetically interact with the orbital angular momentum, in what is called the spin-orbit coupling which can slightly shift the energy levels breaking the degeneracy.

While the Schrödinger equation can be solved exactly for single-electron hydrogen-like atoms, there is no general analytical solution $\psi(\mathbf{r}_1, \mathbf{r}_2, ...)$ for atoms with more than one electron [32]. A common approximation is to solve the equation for a Hamiltonian independent of the electron-electron interaction, leading to a solution described as a product of single electron solutions (more specifically, as a sum of products), and then add a perturbative term corresponding to the electron-electron repulsion when calculating the energy of the states. This approach allows one to keep the definition of the four quantum numbers for each electron independently but breaks the degeneracy when adding the perturbation terms for the calculation of the energy. Furthermore, the Pauli exclusion principle for fermions does not allow two electrons to occupy the same quantum state, which in this case is equivalent to not allowing two electrons with the same 4 quantum numbers. As a consequence, this principle sets the maximum electron capacity for the different shells to 2, 8, 18,... for the K, L, M,... shells respectively. This allows one to define the electronic configuration notation to specify the occupied states in an atom $n\{s, p, d, ...\}^{\{0,1,2,...\}}$ where the superscript indicates the number of electrons present in the n shell, in the $\{s, p, d, ...\}$ orbital. For example the electronic configuration for the ground state of neon is $1s^2 2s^2 2p^6$. The energy levels of the different occupied states are typically shown in diagrams similar to figure 2.4.

Energy levels in molecules

In the case of the molecules, the situation becomes more complicated, as the wavefunction of the system cannot be expressed in terms of the electronic positions $\mathbf{r}_1, \mathbf{r}_2, ...$ alone, but also the position of the nuclei $\mathbf{Z}_1, \mathbf{Z}_2, ...$ Even for the simplest possible

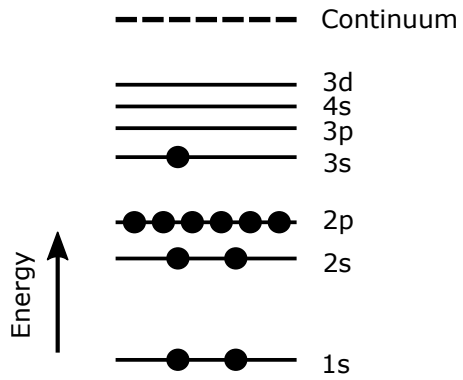


Figure 2.4: Energy diagram of the ground state electronic configuration of sodium: $1s^2 2s^2 2p^6 3s^1$.

molecule H_2^+, with just two nuclei and a single electron, there is no analytical solution to the Schrödinger equation for the molecular wavefunction $\psi(\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{Z}_1, \mathbf{Z}_2, ...)$, and the Born-Oppenheimer approximation is required [33]. This approximation is based on the large difference in mass between the nuclei and the electrons, and assumes that due to the low mass of the electrons, they can instantaneously adapt to any change in the position of the nuclei. Under this approximation the total wave function can be expressed as a product,

$$\psi(\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{Z}_1, \mathbf{Z}_2, ...) = \Phi(\mathbf{r}_1, \mathbf{r}_2, ...; \mathbf{Z}_1, \mathbf{Z}_2, ...)\chi(\mathbf{Z}_1, \mathbf{Z}_2, ...), \tag{2.4}$$

where χ is the nuclear wave function, only depending on the position of the nuclei, and Φ is the electronic wave function, which depends on the coordinates of the electrons as variables, and the coordinates of the nuclei as parameters. As a consequence, the electronic part of the equation can be solved for different fixed values of the nuclear coordinates, finding the allowed energy states, and leading to molecular potential energy curves, like the ones shown in figure 2.5. The shape of these energy curves define the binding or dissociative nature of the electronic states, depending on whether there is a local minima in the energy of the system for a certain geometrical configuration, or the lowest energy is achieved when the internuclear distances tend to infinity. In the case of H_2 (figure 2.5), the ground state, for which the Born-Oppenheimer approximation is the most accurate, shows a stable molecule with an internuclear distance of approximately 1.4 atomic units or 0.74 Å.

In order to gain more understanding of the electronic configuration associated with any of the calculated levels, it is possible to make the *linear combination of atomic orbitals* (LCAO) approximation [33]. The LCAO allows the expression of the molecular orbital for a single electron as a superposition of atomic orbitals, filling the different molecular orbitals with the electrons available. In a diatomic molecule starting from only s and p atomic orbitals we can form different types of molecular orbitals depending on the symmetry and the parity. From the point of view of symmetry, if a given superposition

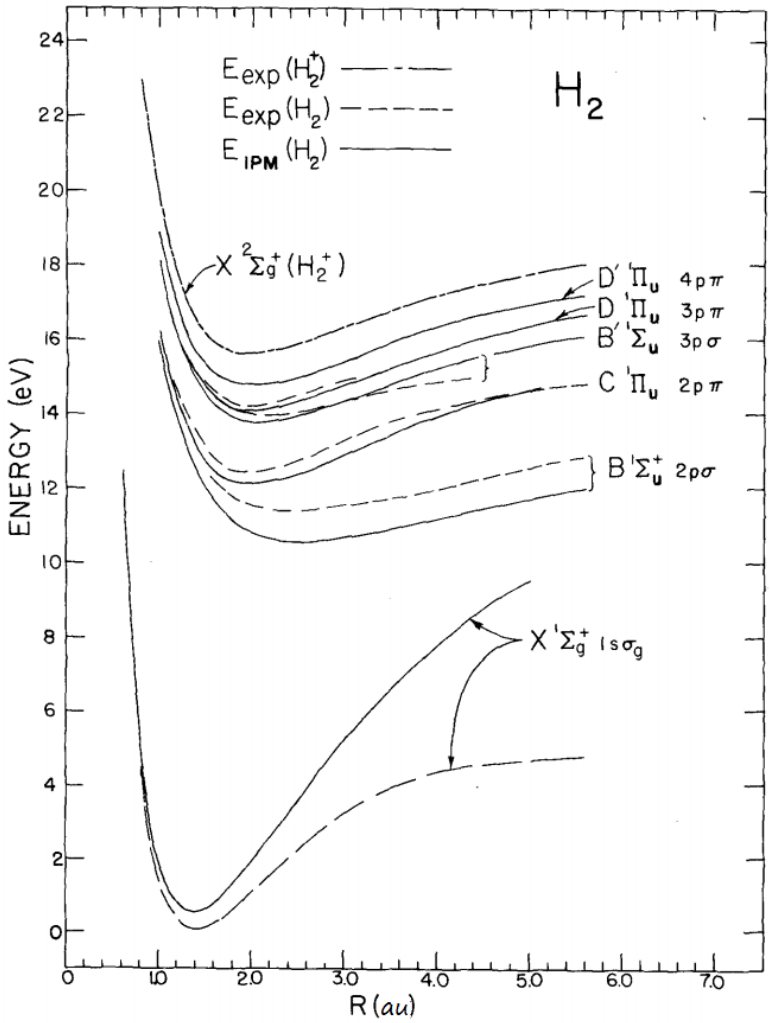


Figure 2.5: Molecular potential energy curves of H_2 represented as function of the separation R between the hydrogen atoms in atomic units. The minimum of the lowest energy curve situated at 1.4 atomic units or 0.74 Å corresponds to the internuclear distance of the hydrogen molecule. Figure extracted from [34].

of the atomic orbitals yields a wavefunction symmetric for rotations around the inter-nuclear axis, then the molecular orbital is of the σ type. This is the case when combining s atomic orbitals or p_z atomic orbitals, where z is parallel to the inter-nuclear axis. On the other hand, combinations involving p_x or p_y atomic orbitals lead to non-symmetric wavefunctions named π orbitals. From the point of view of the parity, if the superposition of orbitals is indistinguishable under inversion of the electron coordinates $\psi_+ = \phi_A + \phi_B$ the orbital is denominated even or *gerade* (g), while it is denominated odd or *ungerade* (u) in the opposite case $\psi_- = \phi_A - \phi_B$. In the case of the combination of s atomic orbitals, g combinations lead to *bonding* orbitals with a higher electron density in the space between the nuclei and u combinations lead to *antibonding* orbitals with a smaller electron density in the between the nuclei.

When calculating the energy of molecular orbitals, it is typically only necessary to take into account the electrons in the *valence* shell, i.e. the last occupied electron shell, as the orbitals associated with the lower energy shells are much smaller and similar to the atomic orbitals, as they are not affected by the other surrounding nuclei. Also, only combinations of atomic orbitals that have similar energies in the first place lead to a significant change in the energy levels [33]. Taking this into account, the ground state of the molecule can be found by accommodating all the valence electrons into the molecular orbitals that would yield, after taking into account electron-electron repulsion, the smallest total energy. For example, in the case of O_2 ($Z=2 \cdot 8$), 2 electrons from each atom correspond to the $1s^2$ atomic orbitals, leaving a total of 12 valence electrons accommodated as $2\sigma_g^2 2\sigma_u^2 3\sigma_g^2 1\pi_u^4 1\pi_g^2 3\sigma_u^0$. This type of electronic molecular configuration allows one to define two important concepts, the highest occupied and lowest unoccupied molecular orbitals (HOMO and LUMO). In the previous example, the $1\pi_g$ is the HOMO while $3\sigma_u$ is the LUMO. Additionally, the molecular electronic configurations can be classified collectively using molecular spectroscopy notation, which in the case of the ground state of O_2 is ${}^3\Sigma_g^-$. In that notation 3 correspond to the spin multiplicity, calculated as $2S+1$, g corresponds to the overall parity, the minus sign $-$ indicates the mirror symmetry under a plane containing the inter-nuclear axis, and Σ denotes the overall rotation symmetry around the internuclear axis, similarly to σ for the individual molecular orbitals. Other possible symmetries yield Π, Δ or Φ states. As an example, adding another electron to the HOMO of O_2, would leave the configuration $2\sigma_g^2 2\sigma_u^2 3\sigma_g^2 1\pi_u^4 1\pi_g^3$ with an odd number of electrons in π-like orbitals, and an overall Π symmetry. This particular valence configuration is of special importance in x-ray spectroscopy as it can be obtained by exciting an electron from the inner shell, in what is known as the $1s$ to π^* resonance. Examples of this resonance will be presented later in this section, as well as in sections 5.3 and 6.1. Additionally, molecular states obtained by exciting an electron into a higher principal quantum number orbital are called Rydberg states. These types of electronic states are typically very diffuse, extending for very long distances compared to the valence shell, making it possible to treat them in a similar way to atomic orbitals. For further information about molecular orbitals, see reference [33].

The existence of the different energy levels allows for a wide range of possibilities of interaction between light and matter, including photoionization, photoexcitation, and different kinds of relaxation paths for the excited states. The models described above allow the creation of diagrams of the energy levels of an atom or a molecule considering the contribution of all the electrons together towards the energy of the state, as in figure 2.5, or taking into account the individual contributions from each individual electron, as shown in figure 2.4. While the first approach is more accurate, as it does not require knowledge of the precise location for each of the electrons (which in some cases is not well defined, specially for electrons in the outer valence and in excited configurations),

the second type of diagram will sometimes be more useful to depict the relevant electron transitions in the following sections.

Photoionization

Photoionization, or single-photon ionization, is the process by which a single photon is absorbed by an atom or a molecule releasing an electron and leaving the molecule/atom in a positively-charged state (figure 2.6a). This process is the underlying mechanism of the photoelectric effect [35], and the electrons released are denominated photoelectrons.

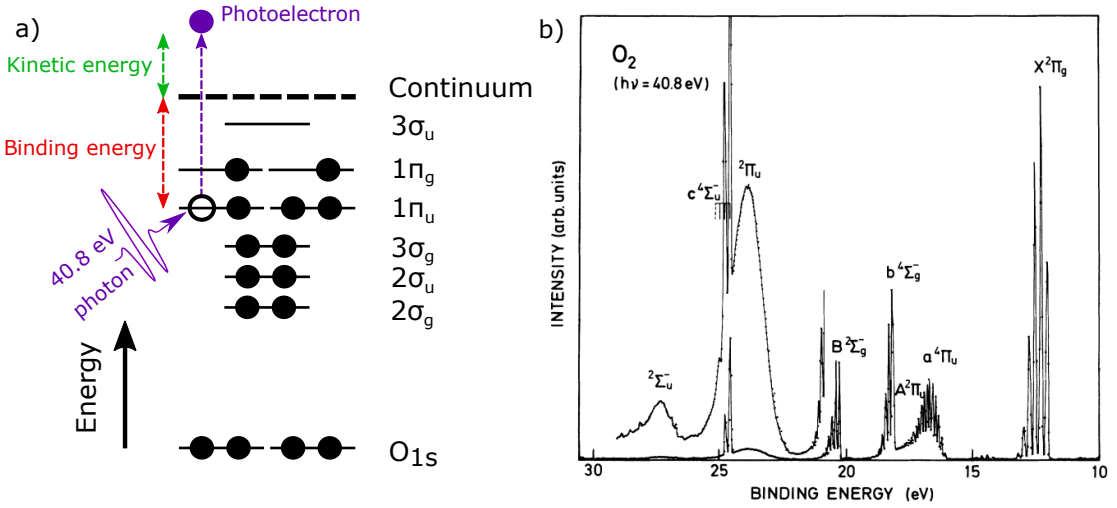


Figure 2.6: (a) Photoionization mechanism in O_2 where a 40.8 eV photon from a helium lamp ionizes an electron from the valence of O_2, producing a photoelectron with a kinetic energy corresponding to the difference between the energy of the photon and the binding energy of the electron. (b) Photoelectron spectra of O_2 obtained with a 40.8 eV helium lamp showing different photoelectron lines corresponding to photoionization from different levels with different binding energies. Figure (b) extracted from [36].

In order to satisfy energy conservation, the energy of the photon must be equal or larger than the binding energy of the electron that is removed. Furthermore, the difference between the energy of the photon and the binding energy of the electron, is transformed into kinetic energy, which due to the large difference of mass between the ion and the photoelectron is mainly carried by the latter (figure 2.6a). This allows one to relate the photon energy E_{photon}, the binding energy of a given electron B_e and the kinetic energy of that electron after being photoionized K_e as

$$E_{\text{photon}} = B_e + K_e. \tag{2.5}$$

Using this relationship, it is possible to select a fixed known photon energy to induce photoionization from all energetically accessible atomic levels, and analyse the energy spectrum of the emitted electrons to map the different molecular/atomic states in what is known as *photoelectron spectroscopy* [37, 38]. An example of a photoelectron signal in O_2

as a function of the binding energy for different $O_2 \rightarrow O_2^+$ transitions, after photoionization with a helium lamp (40.8 eV), is shown in figure 2.6b.

The amplitudes of the photoelectron peaks are proportional to the probability of the process, which in turn is proportional to the *photoionization cross-section* [39, 40]. This cross section is strongly dependent on the particular electronic states involved in the ionization as well as the exact photon energy used for the photoionization. In general for a given electronic state the ionization cross section is zero for photon energies below the binding energy of the state, maximal for the photon energy that exactly matches the binding energy of the state and slowly decays exponentially for photon energies above the binding energy.

A particular type of photoionization of interest to this thesis is core-ionization, which happens when an electron is photoionized from one of the inner shells of the atom/molecule instead of the valence shell, leaving the ion in a highly excited state. Depending of the main quantum number of the shell involved it can be classified as K-edge (n=1), L-edge (n=2),... ionization.

Photoexcitation

Photoexcitation is the process occurring when a molecule/atom absorbs a photon, using the energy of the photon to increase the internal electronic energy by promoting one electron from a low-energy to a higher-energy level. Contrary to photoionization, which leaves the atom/molecule in a charged state, photoexcitation simply leaves the molecule in an excited state. Figure 2.7a shows an example of the excitation process $CO \rightarrow CO^*$ promoting a core electron from the oxygen 1s K-shell to the 2π molecular orbital in the well-known π^* resonance existing at 534.5 eV for CO.

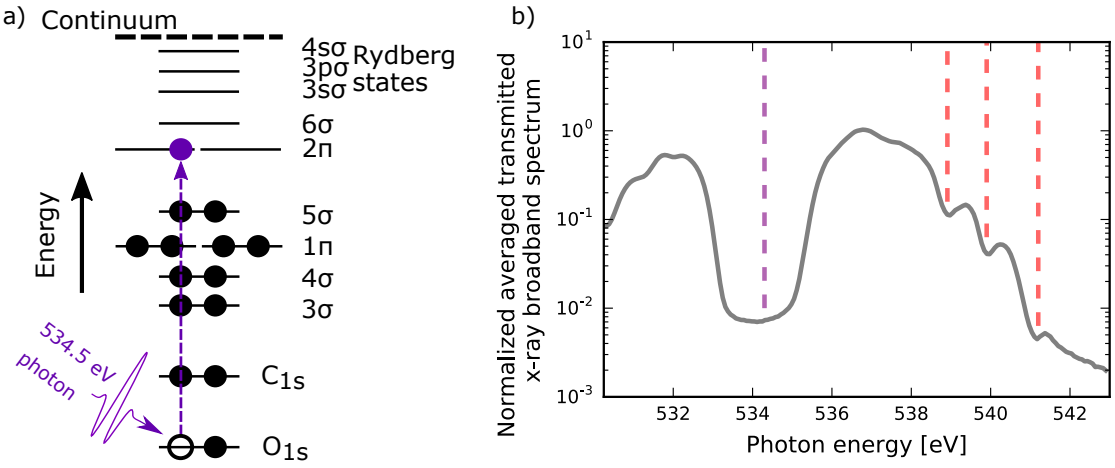


Figure 2.7: (a) Example of a photoexcitation process in CO. (b) Transmitted XFEL spectra through a CO sample showing the absorption corresponding to the 1s$\rightarrow \pi^*$ excitation (purple dashed line) and the 1s$\rightarrow$ Rydberg states (red dashed lines).

As there is no particle released able to carry any remaining energy, in order to have a significant *photoexcitation cross-section* it is required that the photon energy is tuned to the energy of the particular electronic transition. Furthermore, the *photoexcitation cross-section* is directly proportional to the dipole matrix element $|\langle b|\hat{\mathbf{d}}|a\rangle|^2$ between the initial state $|a\rangle$ and the final state $|b\rangle$. These matrix elements are zero for many transitions not satisfying certain *selection rules*, including angular momentum conservation, which requires the molecule/atom to change its total angular momentum to compensate for the absorption of a photon with spin 1. The transitions with null dipole matrix elements are forbidden and the involved excited states are called dark states.

The different allowed resonant photoexcitation transitions can be studied by looking at the transmission of light of different wavelengths through a sample. Figure 2.7b shows an averaged broadband XFEL spectrum after being transmitted through a CO sample obtained from an experiment that will be discussed later in section 6.1. In this figure we can identify different resonant absorption lines corresponding to CO excitations, including the strong 1s$\rightarrow \pi^*$ at 534.3 eV (purple dashed line), and the weaker 1s$\rightarrow$ Rydberg states transitions at 538.91 eV, 539.90 eV and 541.2 eV (red dashed lines). In general photoexcitation is the base of *resonant spectroscopy* or *absorption spectroscopy* [41].

Similarly to core-ionization, core-excitation occurs when the electron is promoted from one of the inner shells of the atom/molecule, as in the previous examples.

Near-edge x-ray absorption

Another way to study the photoionization and photoexcitation in a sample is through the x-ray absorption near-edge structure (XANES) [42]. XANES profiles are obtained by measuring the x-ray absorption of a sample at different photon energies around the different absorption edges. The effect of the photoionization in the XANES profiles shows as a sudden increase in the absorption as the photon energy exceeds the different absorption K-edges that decays exponentially after the edge. Figure 2.8a shows a XANES profile for the glycine molecule (NH_2-CH_2-COOH) obtained by taking into account each of the atomic contributions to the molecule. As glycine contains two carbon atoms, two oxygen atoms and only one nitrogen atom, we observe two strong edges at 284 eV and 543.5 eV, and a weaker edge at 410 eV corresponding to photoionizations from the carbon, oxygen and nitrogen 1s orbitals respectively, which, as described earlier, are not allowed for energies below the location of the respective atomic edges.

To observe the effect of the photoexcitation processes in the XANES profiles, we can look at actual experimental absorption data obtained for glycine zoomed around each of the three edges (figure 2.8b-d). The first difference we observe compared to the previous figure is that the real locations of the K-edges for the molecular 1s-like orbitals, indicated as dashed red lines, are shifted by more than 5 eV when compared to the atomic values. The second difference is that, apart from the increase in the

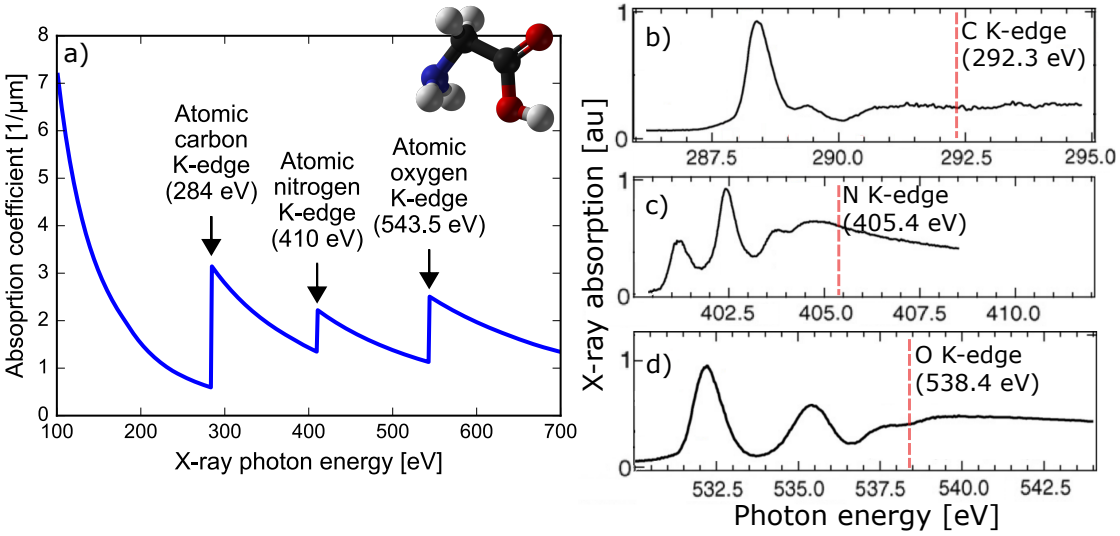


Figure 2.8: (a) Calculated x-ray absorption for the glycine molecule (NH_2-CH_2-COOH) using the atomic cross sections obtained from [43]. As the photon energy surpasses each of the K-edges, the absorption suddenly increases due to photoionization. The K-edges indicated in the figure correspond to individual atomic values for carbon, nitrogen and oxygen. (b, c, d) Experimental XANES profiles obtained in glycine at the carbon, nitrogen and oxygen K-edges respectively. The real location of each of the K-edges in glycine is indicated in the figures. The absorption peaks situated below each of the edges correspond to photoexcitation processes. Experimental data from figures (b, c, d) extracted from [44].

absorption above the edges, the XANES profiles also show different absorption peaks below the edges. These peaks correspond to resonant core-photoexcitation processes, as described in the previous subsection, from the core carbon, nitrogen and oxygen 1s core levels into high energy molecular orbitals close to the continuum. In particular, the photoexcitation processes corresponding to the absorption peaks shown in the figures occur into π^*, σ^* and Rydberg orbitals. For more information about the correspondence between each of the photoexcitation peaks and the molecular states involved, including the exact energies, see reference [44].

The decay of an excited state

Following ionization or excitation, a molecule/atom can be left in an excited state, with an electron hole in the inner-valence shell, or even in the core shell, if a core-ionization or core-excitation has occurred. Such excited states can decay through different channels including the emission of an electron, the emission of a photon or even fragmentation.

The dominant process for light elements (Z<∼30), whenever it is energetically allowed, is the emission of an electron through what is called an *Auger-decay* [45] originating from electron-electron interaction. This process occurs when an electron in the outer shell moves to the existing hole in the excited molecule, releasing enough energy in the process to auto-ionize another electron, denoted as *Auger-electron* (figure 2.9a). There are three electron states involved, the initial state of the hole, the initial state

of the electron moving into the hole, and the initial state of the Auger electron, with corresponding binding energies B_h, B_1 and B_A. Taking into account these values the final kinetic energy of the Auger electron K_A is

$$K_A = B_h - B_1 - B_A. \tag{2.6}$$

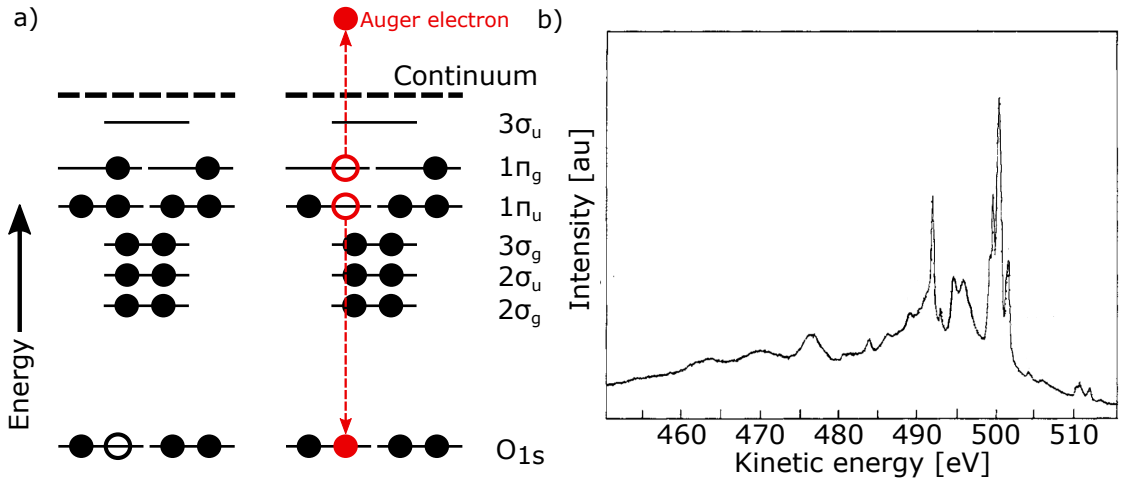


Figure 2.9: (a) Example of an Auger decay process $O_2^+ \rightarrow O_2^{2+}$ where an inner shell hole is occupied by a valence electron. (b) Experimental Auger-electron spectrum from a core ionized state of O_2^+ showing different emission lines corresponding to the different Auger transitions allowed. Figure (b) extracted from [46].

In order for the process to occur, it must be energetically possible for the states involved leading to $K_A \geq 0$, which typically happens when the hole is deep in the valence shell or in the core shell, and only when B_h is larger than the double ionization potential. Excited states meeting this condition are called *Auger-active states*. In an Auger-active state, there are often many electronic states that can participate with different probabilities, leading to very rich characteristic Auger spectra, like the one shown in figure 2.9b for a core ionized state $O_2^+ \rightarrow O_2^{2+}$. The use of Auger processes to study the energy levels of atoms and molecules is known as *Auger spectroscopy* [47, 48, 49, 50].

Typical Auger-decay lifetimes are very short for light elements, in the order of a few-fs [51], overtaking almost any other kind of decay. However, when an excited state is not Auger active, other types of decay can play a role, including the process of emission of a photon through the so-called *spontaneous emission* mechanism. Similar to photoexcitation, spontaneous emission can occur when a low-energy level $|b\rangle$ with an electron vacancy is coupled through a non-zero dipole matrix element $|\langle b|\hat{\mathbf{d}}|a\rangle|^2$ to a higher energy state $|a\rangle$ populated with electrons, as shown in figure 2.10a. Under this condition, an electron can decay from $|a\rangle$ to $|b\rangle$ emitting a photon with the energy of the difference between the initial and the final state (figure 2.10b). The decay probability

$N(t)$ follows an exponential distribution $N(t) = e^{-A_{12}t}$, where A_{12} is the spontaneous emission lifetime of the state, also known as the *Einstein A coefficient*, which is typically on the order of a few hundred of picoseconds, or even nanoseconds. The emission of light through spontaneous emission is also known as *fluorescence* [52, 53]. In the cases where the excited state is created by the absorption of a photon, resulting in an inelastic photon scattering with a fixed energy difference between the absorbed and the emitted photon, the process is termed *Raman scattering* [54].

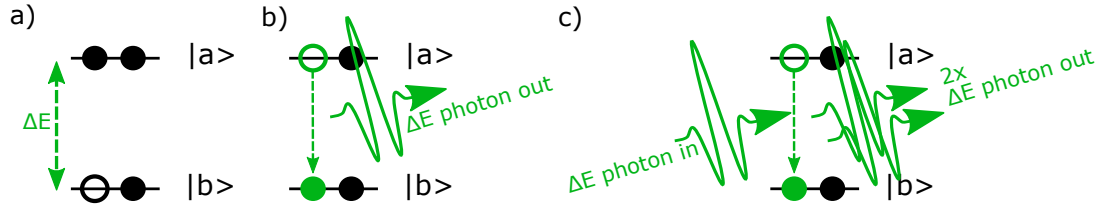


Figure 2.10: (a) Two-level system in an excited state. (b) Spontaneous emission process resulting in the emission of a photon. (c) Stimulated emission also resulting in the emission of a photon. Stimulated emission can only occur in the presence of radiation tuned to the transition level and becomes the dominant process under population inversion condition.

While the process of spontaneous emission through fluorescence occurs in very long timescales ($>$ 100 ps) [53], a similar process, known as *stimulated emission* [55], the principle behind the optical lasers, can occur if photons tuned to the energy of the transition are present and the number of molecules/atoms in the high-energy level is larger than the number of systems in the low-energy level. Under this condition of population inversion, instead of exciting electrons from the low-energy level to the upper energy levels, the photons can dominantly stimulate the emission of identical photons at a rate that is proportional to the intensity of the external electric field, the difference in the population between the upper level and the lower level and the *Einstein B coefficient* B_{12} (figure 2.10c). This allows one to increase the photon stimulated emission rate to any arbitrary value (as long as there is population inversion) by increasing the intensity of the stimulating field. This is of relevance for the results shown in section 6.1, as both stimulated emission and spontaneous emission can be combined to achieve efficient photon decay in otherwise Auger-active states, creating ultrashort x-ray pulses in the process.

Uses for calibration

Apart from the spectroscopic interest of the previously described interactions, they are also very useful for calibration. Well-known Auger spectra, which are independent of the ionization/excitation photon energy can be used for absolute calibration of electron detectors. Photoelectron spectra can be used to calibrate the central photon energy of the source, as the photoelectron peaks shift with the photon energy. Finally resonant absorption as well as fluorescent emission can be used to calibrate x-ray grating spec-

trometers across a wide range of photon energies. Some examples of different types of calibration will be discussed in section 5.3.

2.2 Time-resolved spectroscopy

Many process of interest in molecules involve both nuclear and electron dynamics that often induce changes in spectroscopic properties of the sample. These changes may modify the absorption spectrum, the fluorescence emission spectrum, the photoelectron spectrum or the shape of the Auger-electron spectrum. Time-resolved spectroscopy [56, 57, 51, 58, 59] relies on measuring the time dependence of these properties to infer information about the underlying physical mechanism driving the process, by comparing the changes in the observed spectroscopic signal to the theoretically predicted changes obtained for a model of the mechanism. This is typically achieved by making use of the *pump and probe* approach.

2.2.1 Molecular and atomic movies: pump and probe approach

While the straightforward approach to study the dynamics of a process would be to initiate the process and "record a movie" of the evolution, this is not always possible in ultrafast science, as there are no cameras fast enough, nor electronic systems, to acquire a series of images or signals within a few femtoseconds ($1\,\mathrm{fs} = 10^{-15}\,\mathrm{s}$) with exposure times short enough to resolve the dynamics of interest. Instead, we typically have to resort to a pump and probe approach [60, 56, 61].

Pump and probe is a very general technique that relies on the capability of starting the same physical process of interest multiple times. This is achieved by *pumping* the system: inducing a change in the system well defined in time that triggers the dynamics of interest. The pumping mechanism can range from a mechanical force to an electrical discharge, although throughout this thesis we will concentrate on using optical laser or x-ray laser pulses for this purpose as they allow generating the shortest possible pumping events, of down to a few tens of attoseconds in length [62, 63].

Assuming the clock of the dynamics of interest can be reset at any time via the pumping process, it is possible to follow the time evolution of the system by measuring a property of the system at a fixed time delay after the pump event and repeating the process for many different time delays. This is called *probing* the system. The probing mechanism is chosen to interact with the evolving system in a way that produces a signal dependent of the instantaneous state of the system: the *probe signal*. Similarly to the pump, probe mechanisms can range from a camera taking a picture to the activation of an electric field, typically requiring ultrashort optical or x-ray pulses in the case of x-ray science. The probe signal usually consists of any product resulting from the light-matter interaction, such as measuring any electrons, ions or photons created or absorbed.

Sorting the measured signal as a function of the pump-probe time delay, it is possible to obtain the equivalent of a "movie" of the process frame by frame.

2.2.2 An example of a time-resolved ultrafast process

There are many ultrafast processes in molecules and atoms that can be measured using the pump and probe approach, making use of very different pump and probe mechanisms, each of them justified in terms of quantum mechanical models. While we will see many examples of these processes in the following chapters, due to the experimental character of this thesis we will not go into the quantum-mechanics of light-matter interaction much beyond what has been presented in section 2.1. However, in order to gain some intuition on how microscopic quantum effects can be probed macroscopically, we will show an example based on attosecond transient absorption used to measure the quantum interference of a superposition of quantum states of different energies.

Attosecond transient absorption is a very specific type of pump-probe experiment that aims to resolve dynamics in the few-fs/attosecond regime by pumping the sample with an ultrashort laser pulse, and using as a probe the frequency resolved absorption of an attosecond pulse or train of pulses, as a function of the time delay. Let us assume we have a molecule in the ground state $|\psi_G>$ that is ionized into an intermediate state $|\psi_I>$ through the interaction with a pump laser inducing the photoionization of a valence electron,

$$|\psi_G> \xRightarrow[\text{photoionization}]{\text{pump pulse}} |\psi_I> . \tag{2.7}$$

Under certain circumstances, for example if the pump pulse has a large bandwidth, this intermediate state $|\psi_I>$ can be a linear superposition of eigenstates of the Hamiltonian $|\phi_1>$ and $|\phi_2>$ with energies E_1 and E_2,

$$|\psi_I> \;=\; \frac{1}{\sqrt{2}}|\phi_1> + \frac{1}{\sqrt{2}}|\phi_2> \tag{2.8}$$

that will evolve over time with the corresponding phase terms,

$$|\psi_I(t)> \;=\; \frac{e^{-\frac{\imath E_1 t}{\hbar}}}{\sqrt{2}}|\phi_1> + \frac{e^{-\frac{\imath E_2 t}{\hbar}}}{\sqrt{2}}|\phi_2> . \tag{2.9}$$

As the global phase of a quantum state is unphysical, the state can be rewritten as

$$|\psi_I(t)> \equiv \frac{1}{\sqrt{2}}|\phi_1> + \frac{e^{-\frac{\imath \Delta E t}{\hbar}}}{\sqrt{2}}|\phi_2>, \tag{2.10}$$

where $\Delta E = E_2 - E_1$.

Equation 2.10 shows the behaviour of the dynamics of interest, which in this case consists of changes in the phase of the second term oscillating at the frequency corresponding to the energy difference between $|\phi_1>$ and $|\phi_2>$, however, it does not provide

a way to measure the dynamics of the phase. For this purpose we can make use of a probe pulse, which again, will to induce a transition, for example a resonant photoexcitation from that state $|\psi_I(t)>$ into a different final state $|\psi_F>$. The transition probability P in the perturbation regime (being $\hat{H'}$ the perturbation in the Hamiltonian) is given by Fermi's golden rule [32] and proportional to

$$\mathrm{P} \quad \alpha \quad |<\psi_F|\hat{H'}|\psi_I(t)>|^2, \tag{2.11}$$

where $\hat{H'}$ is proportional to the momentum operator $\hat{p_x}$ for the absorption of a photon in the dipole approximation. This momentum operator can be easily expressed in terms of the position operator by working in the eigenstate base of the non-perturbed Hamiltonian $\hat{H_0}$ (for which $|\phi_1>$ and $|\phi_2>$ are eigenstates) and using the commutator $[\hat{x}, \hat{H_0}] = \frac{\imath\hbar\hat{p_x}}{m}$.

Expanding the transition probability in terms of the superposition of states we obtain

$$\mathrm{P} \quad \alpha \quad |I_1 + e^{-\frac{\imath\Delta Et}{\hbar}} I_2|^2 \tag{2.12}$$

$$\alpha \quad |I_1|^2 + |I_2|^2 + 2|I_2||I_1| \cos\frac{\Delta Et}{\hbar}, \tag{2.13}$$

where $I_j = \frac{<\psi_F|\hat{x}|\phi_j>}{\sqrt{2}}$ are constants, corresponding to the dipole matrix elements.

By looking at equation 2.13 we observe that, as long a both dipole matrix elements are not zero, the transition probability oscillates in time with a period of $T = \frac{2\pi\hbar}{\Delta E}$, which for an energy difference of 2 eV, typical of inner valence levels, yields oscillation periods of around 2 fs. The transition probability involves the absorption of a photon, so by performing the experiment in many molecules simultaneously and simply measuring the delay-dependent absorption of the probe pulse, we can obtain a macroscopic observable that mimics the evolution of the phase of the superposition of the quantum states.

2.2.3 The minimum required experimental setup

There are many key elements to a pump-probe experimental setup. It is worth noting that, for experiments using radiation in the vacuum ultraviolet or the soft x-ray ranges, the experimental setup may need to be partially or fully enclosed in vacuum chambers to avoid absorption of the radiation by air.

The first elements are the ultrashort pump and probe pulses, which can be generated from the same source, such as an optical laser, and then divided in two using a beam splitter, or generated independently from different sources, as may be the case when the wavelength requirements for each of the pulses are very different, including optical-pump x-ray probe experiments. In the cases where the pump and probe pulses do not propagate collinearly, it is necessary to introduce a recombination element, such as a beam splitter for optical pulses, or "holey" mirrors and D-shape pickoff mirrors for x-ray pulses to avoid x-ray reflection or refraction. Additionally it is necessary to provide a way to control and change the time delay between the pump and probe pulses. This is

typically achieved by changing the optical path length for either the pump or the probe pulses, modifying the actual propagation distance or the refractive index of the medium. In sections 3.1 and 3.2, we will study the different existing approaches to generate pairs of ultrashort pulses at free-electron lasers of the required characteristics.

The next relevant element is the sample where the dynamics of interest take place situated at the *interaction region*. The sample can be a solid, a liquid, a gas or even more exotic states such as plasmas. Most of the experiments presented in this thesis make use of gas samples, delivered using effusive jets, pulsed jets synchronized with the pump pulses, high-pressure gas cells or ovens to sublimate a solid sample into a gas. Many engineering solutions are required to successfully use liquid, solid and high-density gas samples in vacuum, including differential pumping, sample rastering, and different kinds of liquid jets that are out of the scope of this thesis.

The final key element is the probe signal, which can be obtained from characterizing the spatial or spectral distribution of any electrons, photons or ions generated by the probing mechanism. Some examples of typical detectors used in XFEL experiments that allow measuring these properties are time-of-flight spectrometers, x-ray grating spectrometers or velocity map imaging detectors, all of which will be discussed in section 3.3.

Apart from those minimum requirements, we will show throughout this thesis that any successful XFEL experiment requires measuring many more signals containing information about the fluctuations in the generation of the x-ray pulses in order to be able to reveal the information contained in the probe signal. Section 3.3 will also include some of the required diagnostics used for this purpose, sections 4.2.5, 5.6 and 6.1.2 will show different strategies to use that information to improve the analysis of the results, and chapter 7 will present a novel approach to extract that information.

2.3 Earlier x-ray sources

One of the challenges of x-ray science is the creation of x-ray sources in the first place, and in particular sources allowing one to produce ultrashort pulses. Other well-known approaches that work well for generating light in the infrared, visible and VUV ranges are not compatible with x-ray radiation. For example, generating 10 keV hard x-ray light using the black body radiation principles, as done in incandescent light bulbs, would require temperatures up to 10^7 K, comparable to the temperature of the core of the Sun. Using light-emitting diodes (LED) would require semiconductors with band gaps in the order of $10^2 - 10^4$ eV, when the actual band gap for most semiconductors is in the order of a few eV. Finally, mimicking optical laser principles requires finding atomic or molecular levels separated by $10^2 - 10^4$ eV. While such electronic levels are available in many atomic species, pumping the transition to achieve population inversion is very challenging, and in some cases require an external x-ray source of similar energies [64].

The typical mechanism used to generate x-ray radiation relies on the principle of *bremsstrahlung* or braking radiation, the radiation emitted by charged particles when they are subject to acceleration. This approach has the advantage of not being dependent on resonant transitions to generate certain wavelengths, but instead presents the flexibility of tuning the generated photon energies by means of the amount of deceleration induced on the charged particles, typically electrons, which can be controlled continuously. Historically different sources have made use of the *bremsstrahlung* radiation in different ways, taking us to the current state of the art sources, the XFELs.

2.3.1 From x-ray tubes to undulators

The simplest way to induce *bremsstrahlung* is by applying a strong linear deceleration to a charged particle. This is the principle of x-ray tubes, similar to the device that was originally used by W. Röntgen for the discovery of the x-rays in 1895. X-ray tubes [65] (figure 2.11a) consist of a cathode C, where electrons are generated when a voltage U_h generating a high current is applied, and an anode A, set a very high voltage U_a with respect to the cathode, accelerating the generated electrons. As these components are enclosed in a vacuum tube, the electrons can travel freely toward the anode where they collide suffering a very abrupt linear deceleration, emitting x-ray radiation in the process.

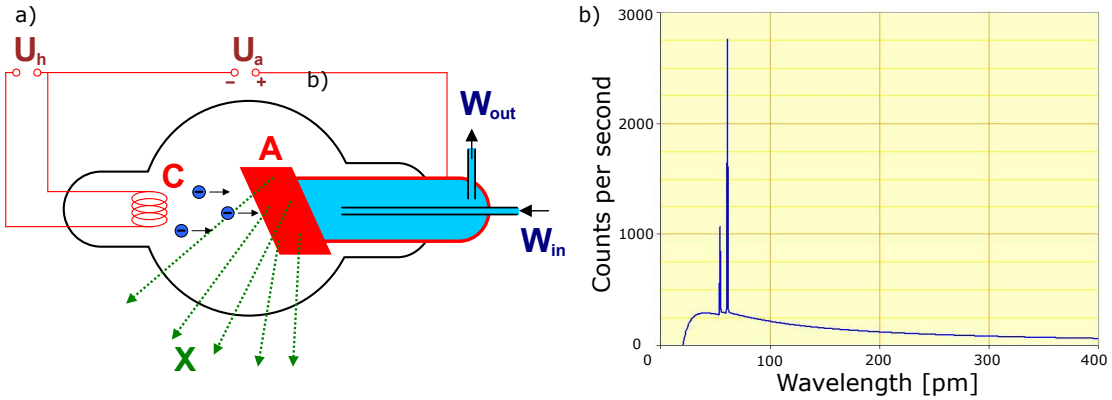


Figure 2.11: (a) Diagram of a water-cooled ($\mathbf{W}_{\text{in}}$, $\mathbf{W}_{\text{out}}$) x-ray tube consisting of a cathode C where electrons are generated by applying a voltage $\mathbf{U}_{\text{h}}$ and an anode A where the generated electrons collide after being accelerated by a voltage $\mathbf{U}_{\text{a}}$. (b) X-ray spectrum corresponding to an x-ray tube operating at 60 kV with a rhodium anode. The spectrum shows a continuum x-ray spectrum with characteristic peaks corresponding to core impact ionization followed spontaneous emission processes. Public domain images.

The characteristics of the x-ray spectrum depend on the acceleration potential (30 to 150 kV) and the material of the anode. The upper photon energy cut-off matches the acceleration potential U_a applied to the tube, corresponding to an electron being accelerated to U_a that completely slows down to 0 kinetic energy in a single interaction, emitting a single photon. In practice, electrons can decelerate in multiple interactions,

emitting several lower-energy photons in the process, and leading to a continuous x-ray spectra (figure 2.11b). Furthermore, electrons may induce impact ionization, removing inner-shell electrons from the anode, leading to emission of characteristic spectral lines as shown in figure 2.11b. Typical peak powers for x-ray tubes can range between 0.1 and 100 kW, however, creating pulsed x-ray radiation using x-ray tubes requires fast switching of the strong electric fields, which typically only allows one to generate pulses in the millisecond regime.

The limitations in terms of pulse durations and the little control over the generated broadband spectra led to the development of more sophisticated x-ray sources, based not on the linear deceleration of electrons, but on the centripetal acceleration experienced by an electron in a circular trajectory (figure 2.12a). As this kind of emission does not completely decelerate the electrons, in order to generate x-ray radiation of the same wavelength, it is necessary to accelerate the electrons to higher energies, up to several GeV. As a consequence, x-ray sources have developed hand in hand with electron accelerators.

Accelerating electrons to a few GeV is much more challenging, and not achievable by simply applying a static accelerating field as in x-ray tubes. The acceleration is typically performed gradually, making the electrons follow cyclic trajectories. The first type of such cyclic accelerator is the cyclotron [66] (figure 2.12b) where a magnetic field of fixed amplitude, and a periodic electric field of fixed frequency is used to induce rotation and acceleration in an electron beam. As the electrons gain speed the radius of the trajectory becomes larger according to the cyclotron radius r_c,

$$r_c = \frac{mv}{qB}, \tag{2.14}$$

where m, q and v are the mass, charge and speed of the electron, and B is the strength of the magnetic field. The spiral trajectory of the electron in cyclotrons makes them however very demanding in terms of surface size, limiting the maximum energy achievable to a few hundreds of MeV, even for much more sophisticated designs with variable magnetic fields. In order to solve this problem, the second generation of cyclic accelerators, known as *synchrotrons*, use electromagnets to adapt the strength of the magnetic field in time to the speed of the electrons to keep the cyclotron radius constant (figure 2.12c). Using this approach, electrons have been accelerated up to a few hundreds of GeV. Furthermore as x-ray generation does not require completely decelerating the electrons, synchrotrons dedicated to x-ray generation are build inside storage rings, which are machines designed to keep accelerated electrons circulating for large periods of time of several hours. Figure 2.12d shows the experimental diagram of SOLEIL, a 354 m circumference 2.75 GeV storage ring near Paris. Before entering the storage ring, electrons are accelerated to 100 MeV using a linear accelerator (section 3.1.1), and then to the

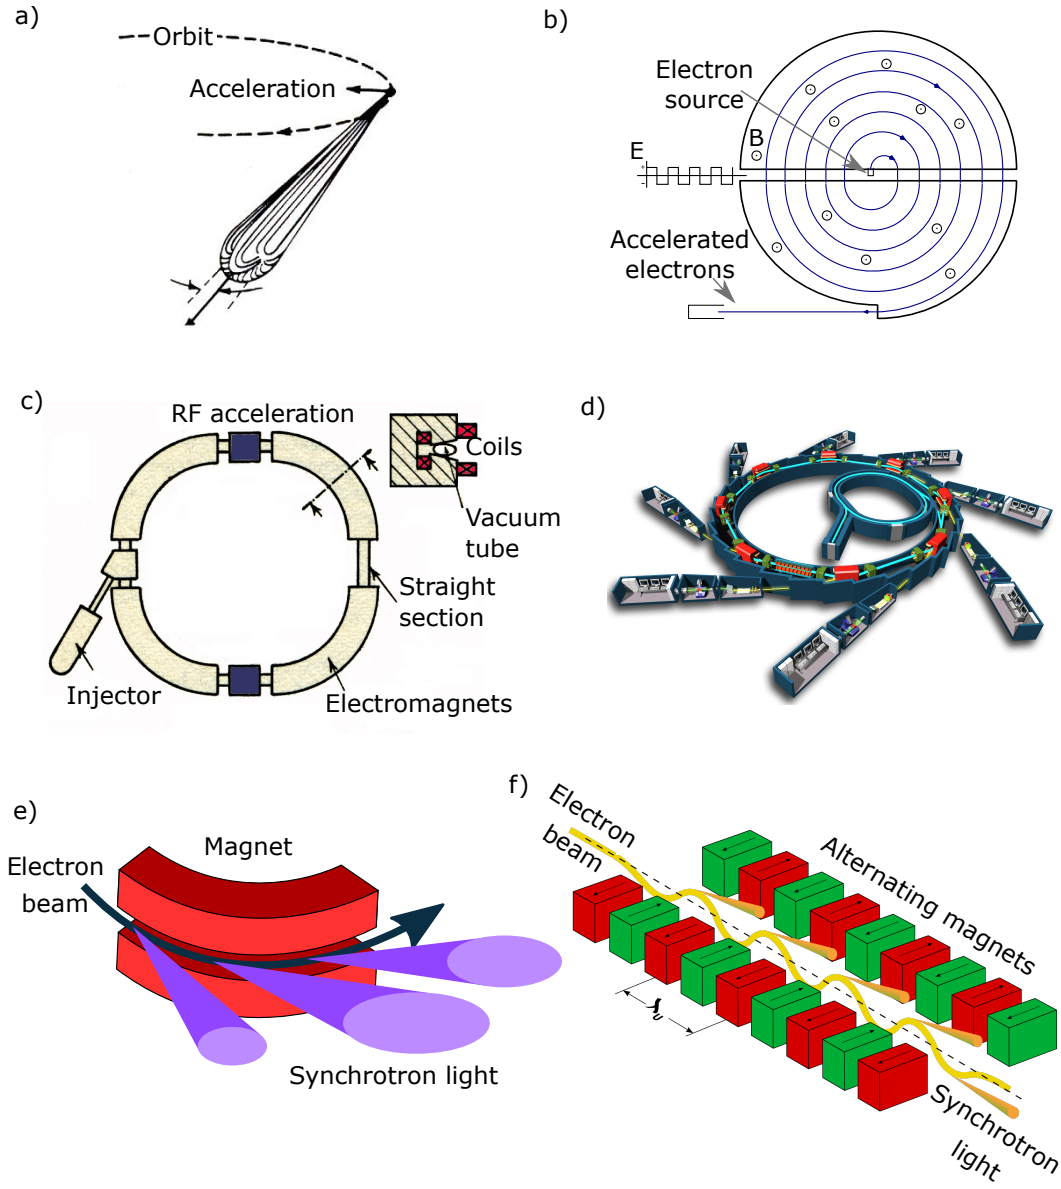


Figure 2.12: (a) Radiation generated by the rotation of a charged particle. (b) Diagram of a cyclotron electron accelerator. The uniform magnetic field B keeps the electrons rotating in a spiral trajectory, while the periodic electric field E accelerates them at each of the rotations. (c) Diagram of a synchrotron electron accelerator. As the electrons are accelerated, the strength of the magnetic fields is increased to keep the radius of the trajectory constant. (d) Drawing of the storage ring SOLEIL. Multiple beamlines around the storage allow the users to perform multiple experiments simultaneously. (e) High divergence radiation emitted by a bending magnet. (f) Directional radiation emitted by a wiggler. Figure (a) is a public domain figure by R. Bartolini. Figures (b), (c) and (f) are extracted from Wikipedia under a creative commons license authored by TNorth, Antivolt, and Bastian Holst respectively. Figure (d) is under copyright EPSIM 3D/JF Santarelli, Synchrotron Soleil and is reproduced with permission.

full 2.75 GeV using a synchrotron-like accelerator (small ring in figure 2.12d). Storage rings for x-ray generation are commonly just referred to as *x-ray synchrotrons* [67, 68].

As the electrons travel around the storage ring at high speeds, they are driven through magnetic devices inducing x-ray generation, which is delivered to different beamlines as shown in figure 2.12d. The simplest device achieving this is the bending magnet [69] (figure 2.12e) where a uniform magnetic field is used to induce a single circular trajectory. The spectrum of the generated x-ray light is also broadband and continuous, as the change in curvature is gradual, and the pulses are of the same duration as the electron-bunch length (typically a few tens of ps). The main limitation of bending magnets is that the radiation is produced in many different directions, as shown in figure 2.12e. To solve this problem, a different kind of device known as *wiggler* [70] consisting of a series of alternating magnets as shown in figure 2.12f is used to induce wiggles in the trajectory of the electrons, producing x-ray radiation at each of the turns. The radiation produced at the different points of the trajectory adds up linearly in the same direction, achieving a much higher directionality in the emission, and a much higher intensity compared to single bending magnets, while keeping the other properties constant. In this sense a wiggler can be seen as a series of alternating bending magnets. Current existing x-ray synchrotrons use an evolved version of the wiggler called *undulator* of between 1 and 5 meters of length that induces smaller oscillation amplitudes in the electron bunches. These smaller amplitudes allow interference patterns across the different turns in the undulator leading to narrower energy bands that can be narrowed further by using monochromator devices to control the spectral properties. Using electron-beam energies of up to 8 GeV at a few hundred MHz of repetition rate, x-rays with photon energies up to 300 keV and pulse durations from a few to a few hundreds of ps can be generated, providing a total flux of up to 10^{15} photons/s.

We will see in section 3.1.3 that similar but longer undulators of up to tens or even hundreds of meters, used in conjunction with low emittance electron bunches lasing in an exponential amplification regime, are one of the basis of the current XFELs.

2.3.2 Generating short x-ray pulses at synchrotrons

X-ray pulses from synchrotrons as described in the previous sections are limited in terms of pulse duration by the duration of the electron bunch to around a few tens of ps. Contrary to powerful femtosecond optical lasers, which have been available since the mid 1980s, x-ray spectroscopy was limited to picosecond timescales until the beginning of the 2000s with the introduction of electron-bunch modulation techniques. Different approaches were proposed to reduce the bunch length by for example, applying a side kick to the electron bunch, and using a spatial filter to achieve sub-picosecond bunch lengths [71], with the most successful approach using femtosecond optical pulses to modulate the electron bunch with the *laser slicing* method [72].

In the first demonstration of laser slicing, 0.4 µJ 100-fs 800-nm optical pulses co-propagated at a small angle with standard 30-ps electron bunches through the wiggler (figure 2.13a). The interaction between the optical pulse and the electron bunch induced a temporal modulation in the energy of the electron bunch that after the first bending magnet became also a spatial modulation as shown in figure 2.13b. The spatial modulation consists of a broadening of the beam width due to the additional lateral momentum that the electrons gain/lose from the electric field of the optical pulse. The temporal modulation limits the width of the broadening to the pulse duration of the optical pulses. Furthermore the dispersion properties of the storage ring causes the spatial energy modulation responsible for the diagonal slice in figure 2.13b. This spatial modulation allows to pick different spatial parts of the beam using a knife edge to obtain pulses of different durations. Figure 2.13c shows the central part of a cross-correlation trace obtained for the inner part of the beam, evidencing the existence of a temporal hole due to the slicing. Cross-correlations measurements for the light from the sides of the beam (figures 2.13d-e) show however evidence of much shorter 300-fs pulses.

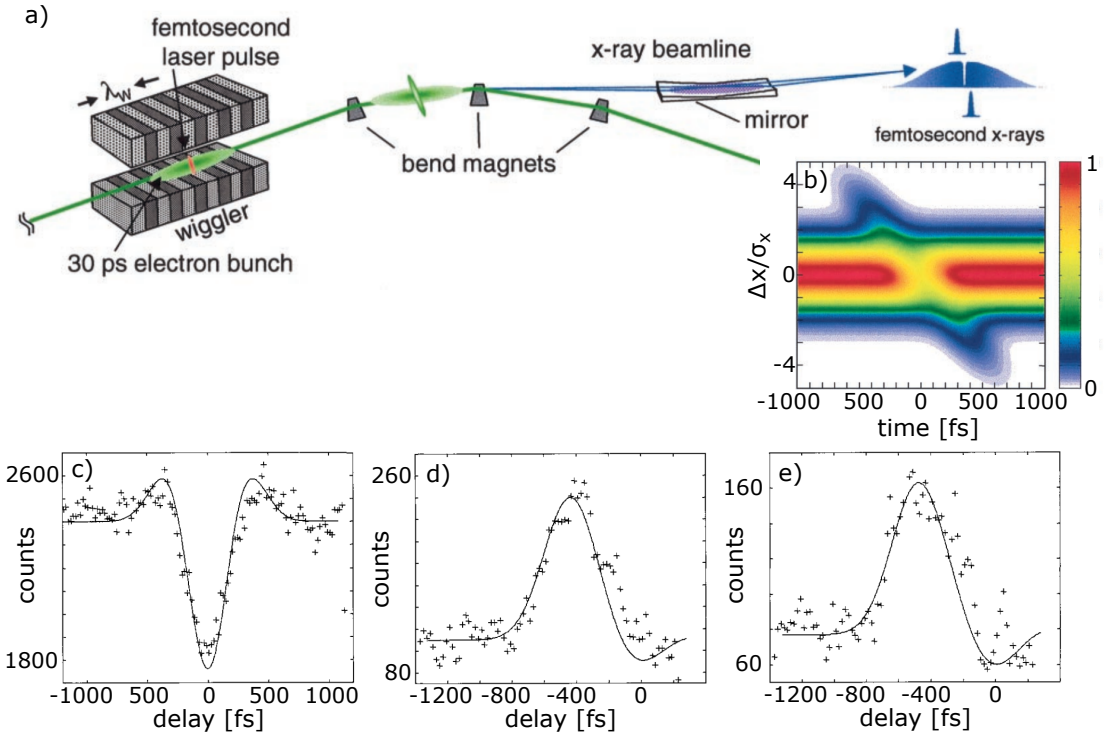


Figure 2.13: (a) Experimental setup of a laser slicing beamline at the Advanced Light Source synchrotron. (b) Model calculation of the spatial electron-bunch distribution due to the laser slicing. (c,d,e) Cross-correlation scans of the synchrotron light for different spatial regions: $-3\sigma_x$ to $3\sigma_x$, $3\sigma_x$ to $8\sigma_x$ and $4\sigma_x$ to $8\sigma_x$ respectively. The solid lines correspond to the model calculation from (b). Figures extracted from [72].

This approach led to the creation of laser slicing synchrotron beamlines [73, 74] with typical pulse durations of 100 fs and 150 fs resolution. This type of beamline allowed for the first time measurements of the dynamics of the electron rearrangement in the

photo-induced insulator-to-metal phase transition [75], the light induced spin crossover dynamics [76], and in general were the building block to the development of femtosecond time-resolved x-ray absorption spectroscopy [77].

The main drawback of the laser slicing technique is the low photon flux. As only a small portion of the electron bunch is used, the generated photon flux may be as low as 10^2 to 10^4 photons per pulse, and may be reduced by a further two orders of magnitude due to the typical low transmission of the beamlines [73, 74]. These low photon fluxes require very long acquisition times, and make diffraction experiments practically impossible, as well as any kind of non-linear x-ray spectroscopy. Furthermore, many processes of interest, particularly regarding pure electronic dynamics happen at a much shorter, few-fs or even attosecond timescales. However, attempting to reduce the pulse duration following the same approach would just decrease the photon flux even further.

Solving these issues was the main motivation behind the creation of the next generation of x-ray sources, the x-ray free-electron lasers. Current XFELs, which will be discussed in detail in the next chapter, can provide fluxes of up to 10^{12} photons per pulse for pulse durations between a few and a few tens of femtoseconds.

Chapter 3

Experimental aspects of XFELs

As introduced in section 2.3, the development of XFEL technology was driven by the need to generate very large numbers of photons contained within very short periods of time. This requirement can be quantified using the *peak brilliance* parameter of a source,

$$\text{brilliance}\left[\frac{\#\,\text{photons}}{\text{mm}^2\cdot\text{mrad}^2\cdot s\cdot 0.1\%\text{BW}}\right]=\frac{\#\,\text{photons}}{\text{Area}\cdot\text{Divergence}\cdot\text{Time}\cdot\text{Bandwith}}. \quad (3.1)$$

This parameter takes into account not only the energy, but also the spatial, temporal and spectral properties of the generated light. A high brilliance is of extreme importance for x-ray science, and in particular for non-linear spectroscopy, which requires a very large number of photons packed into a short period of time, with a small size-divergence product to obtain large intensities when focused and, in some cases, with a well defined narrow bandwidth.

XFELs are specifically designed to provide very high brilliance. Looking at figure 3.1, we see that current XFELs improve the peak brilliance by 10 orders of magnitude compared to synchrotron radiation, while covering the same spectral range. However, this improvement comes at a cost: while synchrotrons typically operate at repetition rates of few hundreds of MHz, most of the currently operating XFELs only generate up to a hundred pulses per second.

In this chapter the author will summarize the principles and important devices key to the generation of XFEL pulses. It will concentrate in particular on the different possibilities to generate pairs of pulses for time-resolved experiments, and on all the relevant devices used to characterize the x-ray pulses and measure experimental signals.

3.1 XFEL principles

The main basic components of an XFEL consist of an electron accelerator to create an accelerated electron bunch, an electron-bunch compressor to reduce the time-duration of the bunches, and an undulator to generate the x-ray radiation. As an example, figure

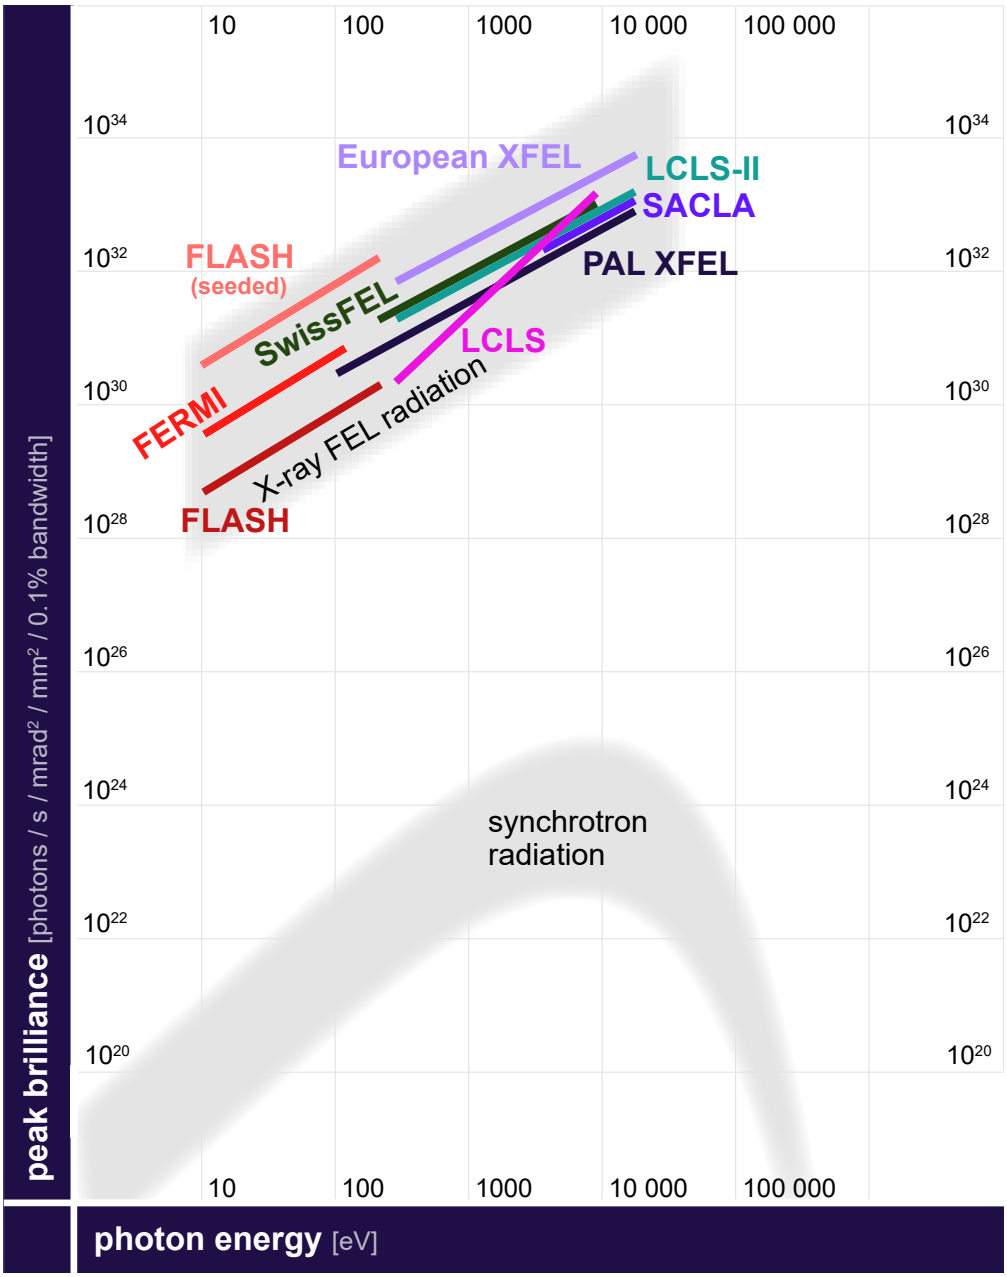


Figure 3.1: Peak brilliance and photon energies for some available XFELs. Figure extracted from [78] (Copyright European XFEL) updated with current data by the author.

3.2 shows the layout of FLASH, a soft x-ray FEL located in Hamburg, spanning about 300 m in length.

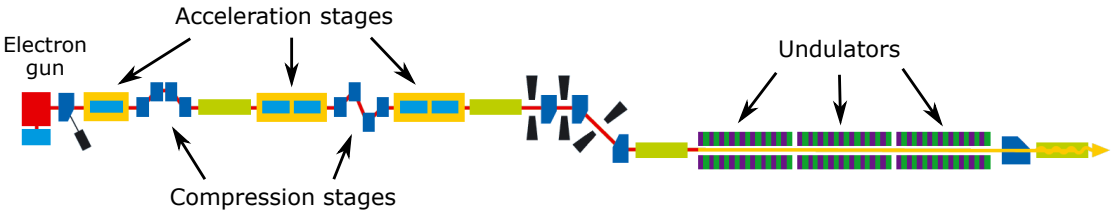


Figure 3.2: Layout of the XFEL FLASH spanning 300 m in length. After the electron bunches are generated at the electron gun, they are driven through different acceleration and compression stages before they generate x-ray pulses at the undulators. These different main parts in the layout will be described in the next sections. Figure extracted from [79] and modified.

3.1.1 Electron-bunch acceleration

The electron-bunch acceleration stage typically consists of a series of linear accelerators (LINACs), each of them formed by one or more radiofrequency (RF) cavities. Due to its charge, an electron can easily be accelerated to a certain energy, for example 100 eV, simply by making it travel through a potential difference of 100 V. This approach however becomes impractical when the target energy is as high as several GeV, since a voltage of 1 GV would be needed. Since it is not actually the voltage, but the electric field, that accelerates the electron, it would be sufficient to create a device that allows the electron to always be in the presence of a negative electric field and consequently be accelerated forward.

Following this idea, the RF cavities (figure 3.3) are formed of different sections that couple to an electromagnetic field of the right frequency, which happens to be in the RF domain, creating a standing wave that allows the electron to travel the length of each section in exactly the same time that takes for the field to change sign. By including several of these RF cavities tuned to different energies it is possible to accelerate electrons bunches from relatively low energies to relativistic speeds.

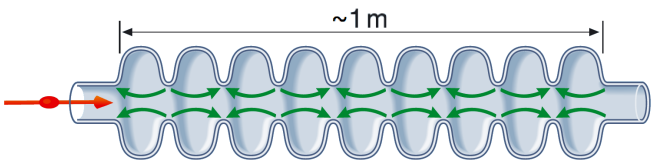


Figure 3.3: Radiofrequency (RF) acceleration cavity. The RF field couples to the cavity, creating alternating positive and negative electric fields at each alternating subsections that flip sign over time at the period of the RF field. The RF cavity is tuned so the time that electron bunch takes to travel through each of the subsections is equal to that period. As a result the electron always experiences an electric field of the same sign and gets accelerated in the same direction for the full length of the cavity. Figure extracted from [79].

3.1.2 Electron-bunch compression

Due to the finite length of each electron bunch, the acceleration process may accelerate different electrons in the bunch differently, inducing a certain electron energy spread that after propagation becomes dispersion in time. Electron-bunch compressors are used to compensate for this dispersion.

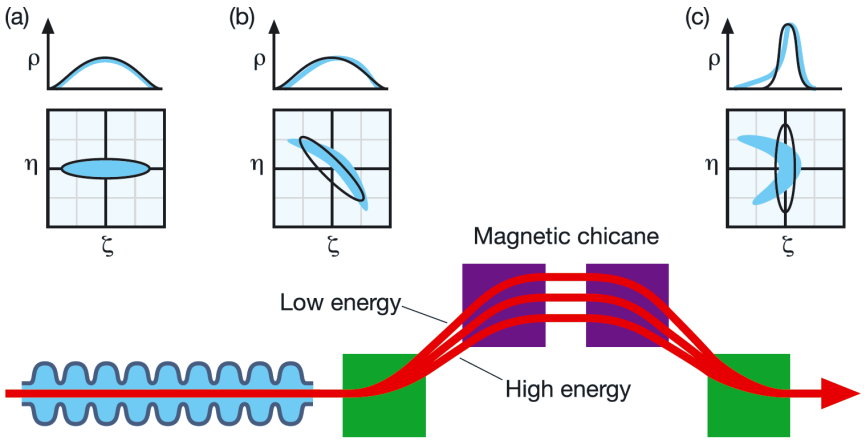


Figure 3.4: Electron-bunch compression using a magnetic chicane formed of 4 magnetic dipoles (green and purple squares in the figure). The insert plots show the instantaneous electron current ρ and spectrogram of the electron bunch (time ζ vs electron energy η). The initial electron bunch with a flat energy as function of time (a) gets chirped as it propagates through the RF cavity (b). The magnetic chicane approximately compensates for this chirp by reordering the electrons in time (c) resulting in a shorter electron bunch for which the instantaneous electron current increases. Figure extracted from [79].

The compression can be achieved via RF cavities by fine tuning of the phase between the electric field and the electron bunch. When used for acceleration this phase is set so the electrons experience the lowest possible negative electric field, however, for compression this phase is set so the electron always cross each section when the value of the electric field is near zero but presents a positive slope in time. While this would not affect a perfectly monochromatic and compressed electron beam, it affects a dispersed bunch, as the electrons at the front of the pulse are slowed down by the small, but positive value for the electric field, and the electrons at the tail of the pulse are accelerated due to the small but negative value of the field, flipping the sign of the initial dispersion. (figures 3.4a and 3.4b).

For non-relativistic bunches, a bunch where the fastest electrons are in the tail will naturally compress itself after propagation, however, in the relativistic regime, where the speed of the electrons is very close to the speed of light, it is necessary to further compress the bunches by inducing spatial dispersion in a magnetic chicane. By using four magnetic dipoles (bottom-right in figure 3.4), the chicane device deflects differently electrons with different energies in a way that the trajectories for lower energy electrons (now situated at the front) are longer than that for higher energy electrons. In this way,

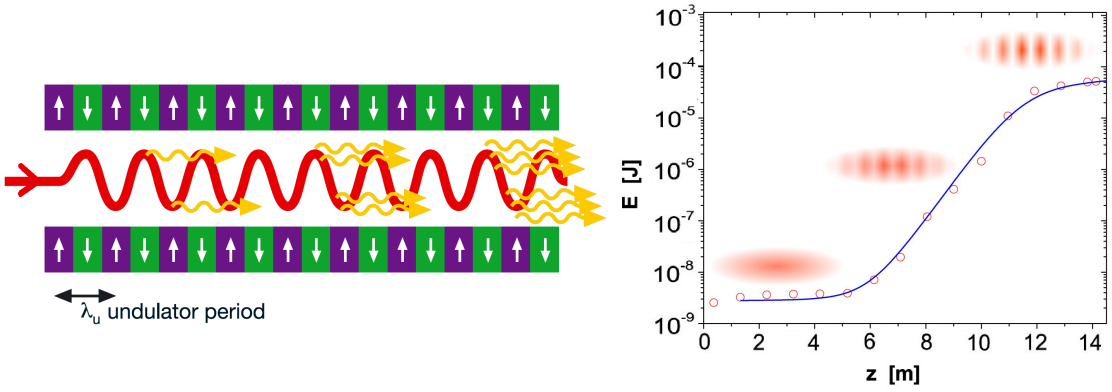
it takes longer for the electrons at the front to leave the chicane, effectively compressing the bunch (figures 3.4b and 3.4c).

3.1.3 Undulator and SASE

The last key element is the undulator where the x-rays are created. Similar to wigglers (section 2.3.1), undulators are formed by a series of alternating magnets that create an alternating magnetic field (figure 3.5a). By choosing the right spacing and magnetic field intensity for a given electron energy, one electron that enters an undulator can wiggle, describing a periodic trajectory with a turning point at the centre of each section. As a consequence of the strong acceleration that the electron experiences at the turning points of this trajectory, it emits synchrotron radiation in the overall forward direction, which is tangential to the trajectory at those points. The emitted wavelength can be expressed as [79]

$$\lambda_{phot} = \frac{\lambda_u}{2\gamma^2}(1 + \frac{K^2}{2}), \tag{3.2}$$

where λ_u is the undulator period, γ is the Lorentz factor and K is the dimensionless undulator strength parameter. This undulator can be expressed as $K = \frac{eB\lambda_u}{2\pi m_e c}$, where B is the amplitude of the magnetic field, m_e and e the electron mass and charge, and c is the speed of light. Wigglers typically present values of K larger than one, leading to bigger oscillations that lase independently, while undulators present values of K much smaller than one.



(a) (b)

Figure 3.5: (a) Diagram of an undulator. The series of alternating magnets in the undulator (green-purple) with a periodicity λ_u induce periodic oscillations in the electron bunch trajectory (red). X-rays are produced at each of the turns, and amplified in the subsequent turns through the self-amplified spontaneous emission (SASE) process. (b) X-ray energy E generated as a function of the propagation distance z. The energy grows exponentially until it saturates when the electron bunch (shown in red) is microscopically ordered. This process is known as microbunching. Figures extracted from [79].

The situation becomes more complex when more than one electron is involved in the process as, ideally, the emission of radiation from all electrons should be collectively constructive. For an initial random microscopic distribution of electrons in the bunch, this is not possible as each electron lases at a different time according to their individual trajectories, however, the self-amplified spontaneous emission (SASE) [80] naturally solves this problem.

SASE can be explained in a picture where the electron bunch is travelling in the undulator along with some x-ray wiggler-like radiation produced spontaneously in the first sections of the undulator. Due to the presence of the propagating electromagnetic wave (with, for example, 2.4 nm of wavelength for 500 eV x-ray pulses), and similarly to what happens in an RF cavity, the electrons are accelerated or slowed down by the field depending on where they sit with respect to the wave. This process has two main consequences. First, as the electron bunch is now longer than the radiation wavelength itself, the process tends to group the electrons together in smaller bunches separated by the radiation wavelength. This is known as micro-bunching. Second, since the electrons are grouped by the existing field, the photons emitted at each turn are in phase with the existing field, creating a constructive interference.

As it can be seen in figure 3.5b this collective process leads to an exponential gain as a function of the length of the undulator in a way that would be expected for stimulated amplification. This kind of collective effect is what differentiates undulator emission in a synchrotron from undulator emission in an XFEL, enhancing the brightness by many orders of magnitude. In order for the process to occur, all electrons have to be very close in space and energy. This is typically quantified using the transverse *emittance* ϵ of the bunch, calculated as the position-momentum phase-space area occupied by the electrons. For example, generating 8 keV SASE pulses at LCLS requires a normalized emittance $\gamma\epsilon$ (where γ is the electron-beam energy) smaller than 1 µm, showing no gain for values larger than 3 µm [81].

After saturation of the microbunching process is achieved (the electron bunch is completely ordered) the x-ray emission stops, however, sophisticated techniques as *tapered undulators* can be used to avoid early saturation of the bunch [82]. This is achieved by increasing the magnetic field towards the end of the undulator to follow the energy loss of the electron bunch and maintain the resonance condition at the wavelength of interest. Also, undulator emission as described here presents linear polarization parallel to the plane where the electrons wiggle, however, a different type of undulator, called *Delta undulator* can be used to achieve a full control of the polarization [83, 84].

3.1.4 Making a short pulse

Typical electron bunches used for x-ray generation in the undulator normally span a few hundreds of femtoseconds in time, yielding radiation pulses of similar length. A common

approach to reduce the pulse length is by shortening the emission length of the electron bunch using a single-slotted spoiler [81]. This device consists of a ∼10-µm-thick foil made typically of aluminium, beryllium or carbon, with a small slit (∼ 250 µm) that is introduced into a magnetic chicane (figure 3.6). This induces Coulomb scattering of the electrons that travel through the foil, *spoiling* the emittance of the tail and the front of the electron bunch and limiting the SASE amplification to the fraction of the electrons that travel through the slit, subsequently reducing the pulse duration to a few-fs/as [81].

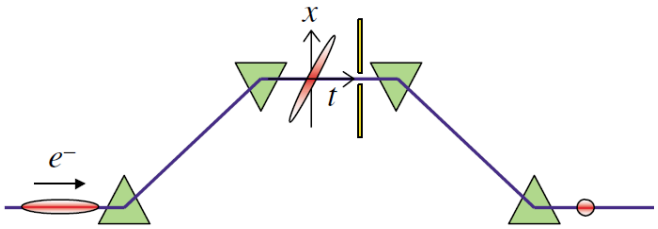


Figure 3.6: Single-slotted spoiler. When a chirped electron bunch travels through a chicane, the different energy components, also mapped into different times due to the chirp, become spatially separated. This allows the slotted spoiler (yellow) to filter in space a narrow energy band, effectively reducing the bunch length in time after the chicane. Figure extracted from [81].

Nevertheless, this technique does not always guarantee pulses as short as the final bunch length. Especially in soft x-ray generation, where the electron-bunch speeds are slower, combined with the non-straight trajectories of the electrons, a significant time-slippage between the generated x-rays and the electron bunch (from a few to a few tens of femtoseconds per 100 meters of undulator) can happen, leading to longer pulses. The effect of the slippage can be minimized by using a smaller number of undulator sections, however, this typically reduces the total x-ray energy, not saturating the SASE process, and leading to a more unstable emission. In practice it is necessary to find a compromise between pulse energy, pulse length and stability.

3.1.5 Seeding

One of the problems of XFELs is that the amplification is performed in a single pass through the undulator, without any kind of resonant cavity. As a consequence, a very large bandwidth can be amplified efficiently. While a large bandwidth is required for generating short pulses, what happens in practice is that the SASE process tends to prioritize one or a series of individual random frequencies within the SASE bandwidth for different parts of the bunch. This limits the longitudinal coherence and generates an x-ray spectrum with spectral spikes. The particular position and amplitude of the spikes depend on the microscopic structure of each electron bunch, which is not controllable or even measurable, meaning that for most purposes they have to be considered stochastic. Figure 3.7 shows typical SASE spectra for four consecutive single shots. Each of the spectra shows individual spikes apparently distributed in a random manner. When

averaged over a 1000 shots, it can be observed that the distribution of spikes follows a smooth Gaussian-like shape.

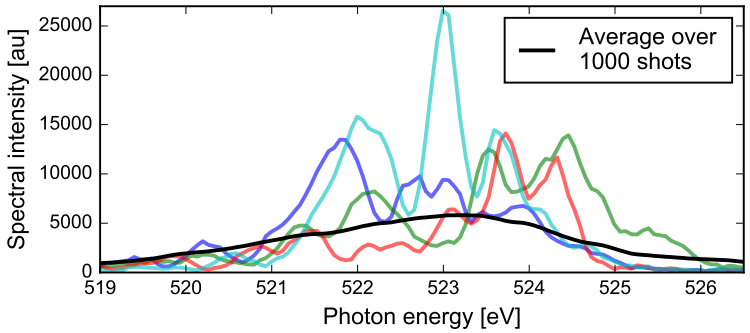


Figure 3.7: Examples of x-ray spectra in arbitrary units[1] for 4 consecutive single shots (red, green, cyan, blue) under SASE operation measured at LCLS. An average of 1000 shots is shown in black. The individual single shot spectra are very rich in structure due to the SASE fluctuations. This structure is however stochastic and disappears when averaging multiple events.

The most straight forward method to overcome this problem relies either on averaging the data over thousands of shots, or on monitoring the spectra on a shot-to-shot basis. However, experiments that require a fixed photon energy, e.g. tuned to a resonant transition may have to discard the majority of the shots not presenting the right spectral properties. In order to provide additional spectral control, seeding techniques are used where the generation is not based on the amplification of spontaneous emission, but based on a controlled seed of light that starts the amplification process.

Self-seeding

In this case the emission process is seeded by the XFEL itself [85], but under controlled conditions in a three step process (figure 3.8). A first undulator section generates a few µJ of x-ray radiation through SASE. This radiation is then driven through a monochromator with a slit that filters everything but a photon energy of choice within the SASE bandwidth. Simultaneously the electron bunch is driven through a magnetic chicane with a double purpose: to remove the micro-bunching induced in the first undulator section, and to compensate for the difference in path length of the light inside the monochromator. Finally, the small amount of filtered light is injected along the electron bunch through a second undulator section, where this time a carefully selected narrow spectral region seeds the amplification process.

This methods guarantees that most of the energy of the final pulse lies within a narrow spectral window, removing most of the photon energy fluctuations. However, the fluctuations are transferred into pulse energy fluctuations: since the generation of the seed is still stochastic, the shots with a spectral shape overlapping well with the monochromator after the first undulator are amplified much more efficiently. If the

[1] [au] stands for "arbitrary units" throughout the length of the thesis unless otherwise indicated.

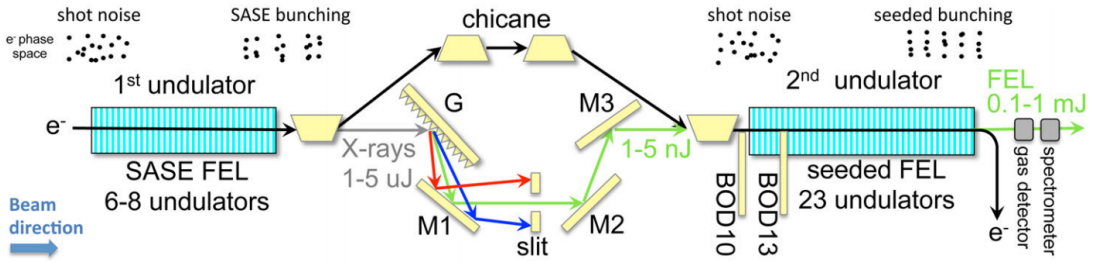


Figure 3.8: Scheme of the self-seeding technique for spectral control in XFEL radiation at LCLS. After the electron bunch generates SASE radiation at the first undulator, a monochromator is used to filter a narrow band of this radiation. The output of the monochromator is then recombined again with the electron bunch in a second undulator, seeding the generation process, and enabling a fine control of the photon energy and a narrow spectral shape. Figure extracted from [85].

second undulator is not driven under saturation conditions this may produce fluctuations in the pulse energy up to 90% of the total maximum energy. Nevertheless monitoring fluctuations in the total x-ray energy can be achieved easily using gas detectors (section 3.3.1), making this approach very useful for resonant experiments, where the measured intensity fluctuations can even be used to automatically perform intensity scans.

Additionally this approach still does not solve the problem of the longitudinal coherence as two different SASE spikes may have the same photon energy, but be generated by distant parts of the bunch. More information about this mode will be shown in section 6.1.

External seeding

In order to provide XFELs that can compete with optical lasers in terms of stability and coherence, the straight forward approach is to use an optical-laser-based solution. However, there are two fundamental differences between XFELs and optical lasers. First, while optical lasers work efficiently in the near infrared range, in particular around 800 nm (1.55 eV) for the popular Titanium-sapphire lasers, XFELs can cover the full UV to x-ray range from a few eV to a few tens of thousands of eV. On the other hand approaches such as high harmonic generation (HHG) [86] allow one to generate light up to a few hundreds of eV. However, in most of these schemes the number of photons per pulse is not larger than 10^8 for the XUV, and drops to less than $10^4 - 10^6$ at 300 eV, which takes us to the second main difference: XFELs can generate on the order of $10^{11} - 10^{12}$ photons per pulse.

The intermediate solution is to use an optical laser as a seed for the XFEL using different techniques to up-convert the frequency. This approach, known as high-gain harmonic generation (HGHG), was originally proposed more than 25 years ago [87] but not proven experimentally until recently [88]. In this example, an electron bunch tuned at the right energy (∼1.2 GeV) (figure 3.9) is driven along with the optical seed (260 nm,

third harmonic of Ti-sapphire in this case) through a short undulator section (modulator) tuned at the wavelength of the seed. This induces a coherent energy modulation in the electron bunch which is transformed into a temporal modulation (microbunching) using a magnetic chicane. Although this modulation is dominated by the wavelength of the seed, it creates strong harmonic components, and when driving the bunch through a second undulator section tuned at the nth harmonic component, the electron bunch emits coherently at the corresponding harmonic frequency.

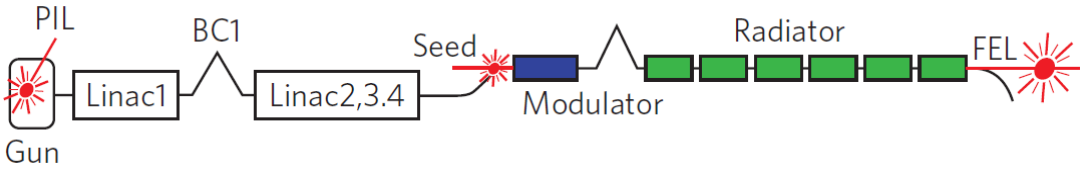


Figure 3.9: Scheme of the external seeding technique for creating longitudinally coherent, spectrally stable x-ray pulses at FERMI. An optical seed pulse is used to induce an energy modulation in the electron bunch at the modulator, which transforms into temporal modulation after propagation through a chicane. This allows the modulated electron bunch to lase very efficiently at any one of the harmonics frequencies of the seed pulse when driven through an undulator tuned at the relevant frequency. Figure extracted from [88].

This approach provides pulses with a very high longitudinal and spatial coherence, similar to that of the seed pulse. It also removes most of the photon energy and pulse energy fluctuations, as the stable seed laser microscopically modulates the electron bunch. The main limitation of the technique is however the photon energy. Currently, the maximum achievable photon energy is on the order of 10^2 eV, and it is planned to reach 10^3 eV in the next few years [88], which is still far from the $2 \cdot 10^4$ eV that can be achieved with regular hard x-ray SASE FELs. Nevertheless there are many important energy edges in the 0-1000 eV energy range, including carbon, nitrogen, and oxygen K-edges making HGHG a very attractive option for future spectroscopy experiments.

3.1.6 High repetition rate

As mentioned in the introduction of this chapter, most of the current XFELs are limited to a repetition rate of around 100 Hz, including SACLA and LCLS, the two main hard x-ray FELs used in this thesis. This limitation is characteristic of the first generation of XFELs and it is related to the large amount of heat generated by the standard LINAC and RF cavities. The next generation of XFELs, now on its onset, makes use of superconductive LINACs which minimize the heat loss and allow much higher repetition rates. Second generation XFELs include the European XFEL currently under commissioning, with up to $27 \cdot 10^3$ pulses per second, and LCLS-II currently under construction, with up to 10^6 pulses per second.

High-repetition rate XFELs, however, will present many other challenges, specially related to the large amount of data generated at MHz repetition rate, and the use of vital

detectors that will not work at the increased repetition rates. Chapter 7 will present a technique developed in the context of this Ph.D. to overcome that problem, allowing to extrapolate some devices and strategies used for first generation XFELs into the next generation of MHz XFELs.

3.2 Multipulse techniques

There are many different approaches allowing to create two few-fs x-ray pulses delayed by a controllable time delay.

3.2.1 X-ray split and delay

The simplest technique from a conceptual point of view, known as x-ray split and delay (XRSD) [89, 90], makes use of a beam splitter that consists of a grazing incidence two part mirror where the focusing x-ray beam is split in two, as seen in figure 3.10, inducing a difference in path length and arrival time. This method allows the generation of a pair of one-colour pump-probe pulses in the range of 0 to 100 fs with a precision of 0.1 fs. On the downside, there are large power losses in the gap between the two parts of the mirror for low delays ($<\sim$ 10 fs). This limits the capabilities of XRSD for few-fs experiments. Also due to the spatial arrangement of XRSD, the overlap between the two beams needs to be re-optimized at each new delay, leading to potential systematic errors. Due to the fact that the splitting does not involve the electron bunch, this is the configuration that minimizes the jitter in the time delay.

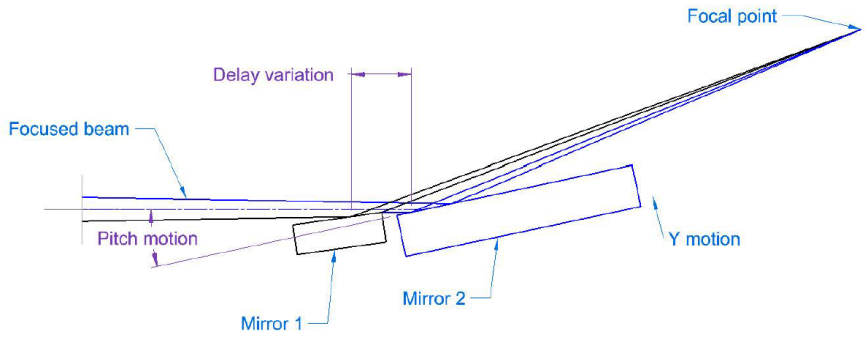


Figure 3.10: X-ray split and delay apparatus. As the x-ray beam is focusing, it is split in two by making it reflect off two different mirrors. This induces a relative difference in the optical paths between each part of the beam resulting in a time delay at the focus. Figure extracted from [89].

Some of these problems can be solved, while keeping a sub-fs time accuracy, by using more sophisticated designs that use up to 8 different mirrors [91]. However, the additional reflections also decrease the overall transmission.

3.2.2 Multiple undulators

This technique [92] uses a single electron bunch and two undulator sections potentially tuned at different wavelengths, with a magnetic chicane in between (figure 3.11). After the electron bunch generates an x-ray pulse in the first undulator, it passes through a chicane with a double purpose: removal of the micro-bunching, and introduction of a time delay between the electron bunch and the first pulse. Then, the same electron bunch lases in the second undulator, creating a second pulse at a certain time delay. While this method presents the advantage of having two-colour capabilities, it presents a higher jitter (few-fs) in the delay, and higher fluctuations (∼50%) shot-to-shot in the relative intensities due to the fact that the lasing process does not always reach saturation. It also presents only partial temporal overlap, with a minimum delay of 3 fs, due to the time slippage between the x-ray pulses and the electron bunch between the undulators.

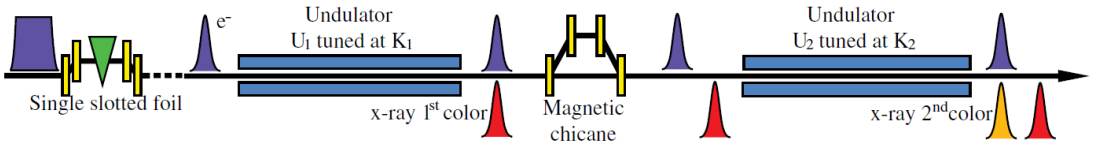


Figure 3.11: Emission in two undulator sections. An short electron bunch spoiled with a single slotted spoiler generates an x-ray pulse of one colour at the first undulator with undulator parameter K_1. The same bunch is delayed with respect to the generated pulse in a magnetic chicane before it generates another pulse of different colour in a second undulator with undulator parameter K_2. Figure extracted from [92].

3.2.3 Double-slotted foil

Single-bunch mode

This mode [92] makes use of a double-slotted spoiler to split the initial long electron bunch into two bunches delayed by a time that depends on the characteristics of the double-slit. These two bunches can lase separately and, combined with a multiple-undulator approach (figure 3.12), can achieve saturated two-colour pump-probe emission. The advantages are similar to that of the previous technique, with the difference that since the lasing process saturates, the fluctuations in the pulse energy are smaller (∼10%). On the other hand, temporal overlap is not achievable in the two-colour mode due to a minimum separation (∼10 fs) in the electron bunches.

Double-bunch mode

In this approach two independent electron bunches are generated at the electron gun by illuminating the cathode with two time-delayed ultraviolet laser pulses [93, 94]. Using different linear acceleration stages combined with magnetic chicanes, it is possible to accelerate the electron bunches to different final energies to generate two different colours,

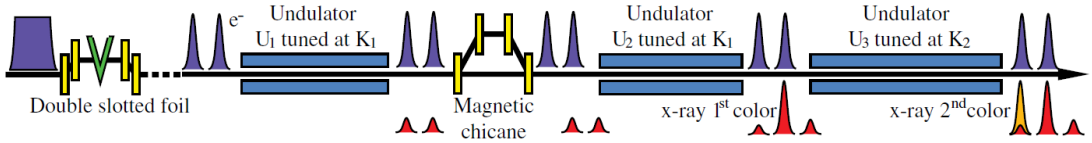


Figure 3.12: Double-slotted spoiler for one bunch. A double-slotted spoiler splits the long initial electron bunch into two short electron bunches. The two electron bunches generate a pair of weak pulses at the first undulator with parameter K_1, however, a magnetic chicane is used to shift the electron bunches with respect to the pulses in time by the pulse separation. This enables the amplification and saturation of the second pulse alone in a second undulator tuned at the same parameter K_1. Finally, the second bunch generates a different colour pulse in a third undulator with a different parameter K_2. Figure extracted from [92].

while at the same time having control of the energy dispersion (chirp) and compression of the electron bunched at different stages (figure 3.13). This allows one, in particular, to drive the two chirped bunches through a magnetic chicane, and filter them to obtain shorter pulses. As a dispersive element, the magnetic chicane spatially separates electrons with different energies. This has a double effect. First, each of the electron bunches follows a different trajectory allowing for independent filtering. Second, due to the linear chirp within each electron bunch, different parts with slightly different energies, also follow slightly different trajectories. As the linear chirp maps time into energy, the effect of the magnetic chicane allows mapping time into space through spatial dispersion of different energies. A double-slotted version of the foil used by P. Emma et al. [81] can be used to filter a narrow energy band within each of the bunches, spoiling the rest of the bunch and effectively shortening the length of emission for both electron bunches down to a few fs.

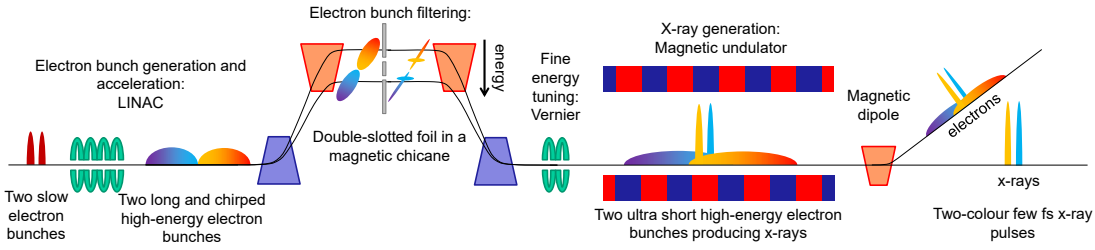


Figure 3.13: Double-slotted spoiler for two bunches. Each of the bunches is "spoiled" independently by a different slit. This is a simplified diagram, in reality there is an additional magnetic chicane and multiple acceleration stages interleaved with the chicanes to allow more tuning flexibility.

Once the electron bunches have been prepared for emission, they are driven through a final linear acceleration stage, the *Vernier*, which adjusts the exact final energy allowing for fast photon energy tuning, and then through a single undulator section, where the unspoiled part of each of the electron bunches emits an ultrashort x-ray pulse with a few tens of microjoules per pulse. While using more undulator sections would give more energy per pulse, it would also yield longer pulses due to the time-slippage between

the x-ray pulses and the electron bunches after propagation. Once the light has been generated, the electrons, are separated using a magnetic dipole.

3.2.4 Fresh-slice technique

This technique developed by Lutman et al. [95] relies on accurate control of a 50 fs electron bunch to make different parts of the same electron bunch lase in different undulators (figure 3.14). An orbit corrector (magnetic dipole) is used before the first undulator, to align the tail of the bunch and make it travel straight, while the front of the bunch oscillates through the undulator. This causes the spatial overlap with the generated light to happen only for the tail of the bunch, suppressing all the emission from the rest of the bunch. After driving the bunch through a chicane, a similar process is performed for the head of the bunch in a second undulator. Since each of the undulators can be tuned to different undulator parameters, the colours of each of the pulses can be tuned independently in a simple way.

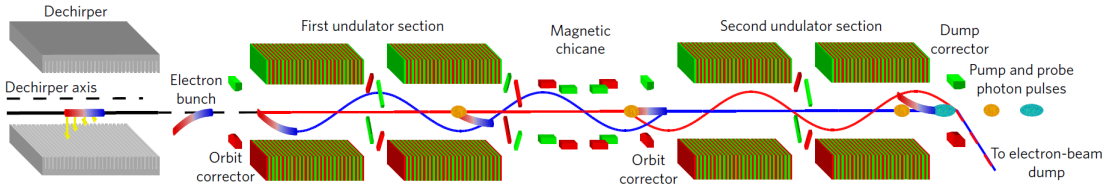


Figure 3.14: Scheme of the fresh-slice technique for double-pulse generation at LCLS. The trajectory of a single electron bunch is controlled to induce lasing at the tail and at the front of the bunch independently in two different undulators. A magnetic chicane is used in between the two to control the time delay between the pulses. Figure extracted from [95].

Although this approach is similar to the multiple undulator approach presented in section 3.2.2, it has several advantages. First, the process can be driven to saturation in both undulators, providing a much higher stability of the final pulse energy. Secondly, as each of the pulses is emitted by a fresh portion of the bunch, the emission of the second pulse is not affected by the generation efficiency of the first one, making both pulses more independent. Thirdly, since the first pulse is generated further back in the bunch, it is possible to achieve temporal overlap by making the separation between the head and the tail of the bunch match the delay introduced by the chicane minus the time slippage between the first generated pulse and the bunch between the two undulators.

Alternatively, the mode can be set for the head of the bunch to lase in the first undulator, and the tail in the second undulator. This way, once the first pulse is generated, it never overlaps with the bunch again, removing any potential interference in the second undulator. Avoiding this is specially important when both undulators are tuned to the same energy, as it could introduce a delay-dependent emission efficiency due to the first pulse partially seeding the process in the second undulator for zero delay. The downside of this configuration is, however, that the minimum time delay that can be achieved

corresponds to the distance between the head and the tail of the bunch plus the time slippage, corresponding to approximately 25 fs in total. Experimental data taken in this mode will be shown in section 4.2.

3.2.5 Optical-pump x-ray-probe generation

In many cases it is desirable to use a pump pulse that excites energy levels separated by a few eV, and then probe the system using core ionization or diffraction, with x-ray wavelengths. However, the methods explained in the previous sections only allow for a detuning of up to 2-3% the photon energy between the different colours. In order to provide such pump-probe capabilities, the only option is to use an optical laser (or its harmonics) to provide one of the pulses, and the XFEL to provide the other one.

This is achieved by using optical lasers triggered on the same timing signals as the XFEL [96], however, keeping a few-km machine and an optical laser free of drift and jitter to fs precision (1 µm corresponds to 3 fs of optical path), is practically impossible, and parallel experiments have to be designed to monitor and characterize this jitter shot-to-shot [97, 98, 99]. This will be further discussed in sections 3.3.7 and 6.2.

3.3 Relevant diagnostics

Due to the nature of the experiments, involving light and matter interaction, the resulting signal channels available usually measure the outcoming light, the released charged particles (electrons or ions) or both at the same time. This is the case of, for example, attosecond transient absorption by core electrons [100] where, apart from measuring directly the absorption signal in the laser radiation (which depending on the target may not be possible), it is possible to measure the Auger electronic signature of the decay of core-hole levels as a function of the time delay, and potentially use this to retrieve the absorption information. Furthermore many other detectors are necessary to perform a full characterization of each the x-ray shots, necessary to monitor the fluctuations.

3.3.1 Gas detector and attenuator

One of the most fundamental properties to know about any pulsed light source is the amount of energy contained in each pulse. The usual approach to measure this in XFELs is to use a gas detector [101]. Gas detectors consist of a cylindrical differentially-pumped chamber containing typically molecular nitrogen (figure 3.15). When an x-ray pulse travels through this chamber, a fraction of the photons photoionize electrons from the nitrogen molecules leaving them in different excited states, which in turn may emit secondary electrons through Auger relaxation. Both, photoelectrons and secondary electrons starting with high kinetic energies interact with the rest of the gas until they are thermalized. In this process, the excited nitrogen gas relaxes through the emission of UV light in the 300-400 nm range, which is partly measured with a photomultiplier

tube. This fluorescent radiation correlates strongly with the total x-ray energy, however, a calibration to obtain an exact absolute value in millijoules is necessary. More details about this calibration will be shown in section 5.3.4.

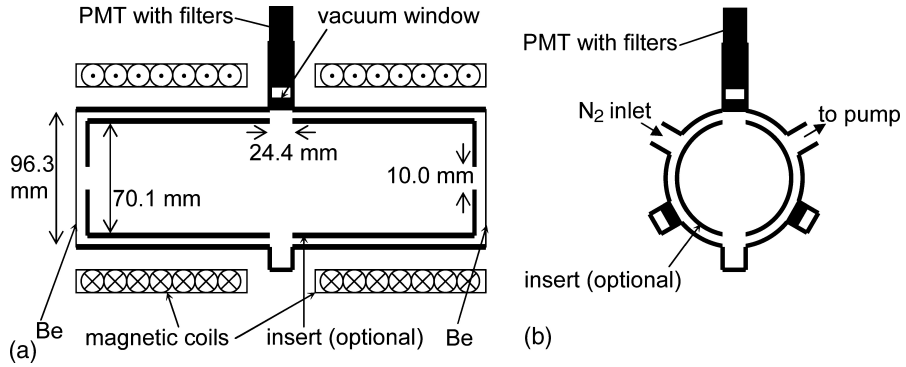


Figure 3.15: Diagram of a gas detector chamber filled with N_2 used to characterize XFEL pulse energies. The x-ray-induced fluorescence caused by the x-ray pulses is measured with a photomultiplier tube (PMT) and calibrated to obtain the total x-ray energy. Figure extracted from [101].

This approach to detecting the pulse energy has the advantage that, when used at low pressures (few mbar), gas detectors are nearly transparent to the generated radiation, and can be used inline with the XFEL as non-intrusive monitors of the pulse energy on a shot-to-shot basis. Additionally, using a similar approach with a longer chamber (a few meters) and slightly larger pressures, gas detectors transform into gas attenuators designed for a fine control of the pulse energy by means of tuning the pressure. Gas attenuators are usually interleaved with gas detectors to monitor the attenuation.

3.3.2 X-ray grating spectrometer

Photon spectrometers for light near the visible domain have been used and improved since the discovery of the capability of prisms, and later gratings, to disperse light. Grating spectrometers in the regions close to the visible domain are typically arranged in a 4-f configuration, imagining an entrance pinhole (1-d spectrometers) or slit (2-d imaging spectrometers) onto a film or charge-coupled device (CCD) [102] camera. A diffraction grating is placed at the Fourier plane of the system, in a fashion that it is the light coming from the first order of the grating that is used to form the image. This allows mapping wavelength into space and, after a proper calibration, can be used to retrieve the spectrum of the light.

This scheme presents, however, some issues when measuring XUV and soft x-ray pulses. First of all, those wavelengths cannot easily propagate through optical elements, so lenses cannot be used. Secondly the reflectivity of common dielectric and metallic mirrors at 45° of incidence is almost negligible at those wavelengths, so it is necessary to work at grazing incidence ($\sim 1 - 3°$ with the surface) and minimize the number of

mirrors. Also, due to the small intensity of the relevant signals and other wavelength requirements, a regular CCD camera may not be adequate to measure the signal.

As a solution to the mentioned issues, flat-field grating [103] x-ray spectrometers (figure 3.16) can be used. This type of spectrometer [104, 105, 106] makes use of a single reflective concave grating with a variable line spacing at grazing incidence to act both as an imaging element and as a dispersive element. By placing a slit, pin-hole or a point light source in the object plane, the system can form an aberration-free image distributing different wavelengths in a plane along the dispersive axis using the first, or a higher order reflection, of the grating. Additionally, the perpendicular axis to the dispersive axis can be left unmodified to preserve spatial information.

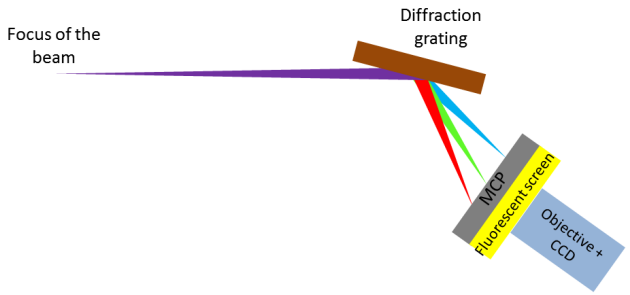


Figure 3.16: X-ray grating imaging spectrometer layout. The x-ray beam is diffracted off a flat-field x-ray grating (curved and with variable line space). This creates a dispersed image of the beam focus (or a slit at the entrance of the spectrometer) in a plane where a detector (MCP or x-ray camera) is placed. The recorded image contains spectral information along one of the axis and spatial information along the other axis.

Regarding the detection, the CCD camera can be replaced by a micro-channel plate (MCP) [107], a device able to amplify weak photon or particle signals, followed by a fluorescent screen, which transforms these signals into visible light that can be imagined by a regular optical camera. Via secondary emission of electrons in the MCP, the incident photon radiation can be amplified and, depending on the applied voltages, the MCP may be used to detect up to single photons with resolutions in the order of 100 µm. Alternatively specialized x-ray CCD cameras [108] can be used.

This type of measurement is considered intrusive, as the light that enters the spectrometer cannot be delivered to an experiment. As a consequence, these spectrometers are typically installed at XFELs as a separate diagnostic towards which light is directed during the tuning process or fed with a small fraction of the light split off the main beam [109]. They can also be installed downstream from the interaction region of a experiment to monitor the x-ray energy when working with low density samples [12], or to monitor the absorption of high density samples [11].

3.3.3 Time-of-flight spectrometer

A time-of-flight (TOF) spectrometer simply consists of a drift tube with a certain length situated between the interaction region where particles are generated, and a detector where particles are detected. *Time-of-flight* refers to the time taken by the particles to fly through the length tube. The drift tube needs to be at high vacuum (in the order of 10^{-6} mbar) for the travelling particles not be scattered, and be long enough for the particles to take a measurable time to reach the detector, as the energy resolution for a given particle grows linearly with the length. The detection mechanism usually consists of an MCP that creates a signal spike for each single electron or ion arriving at the detector, connected to a time-to-digital converter [110], with resolutions as good as a fraction of a nanosecond (10^{-10} s), producing a TOF trace (figure 3.17).

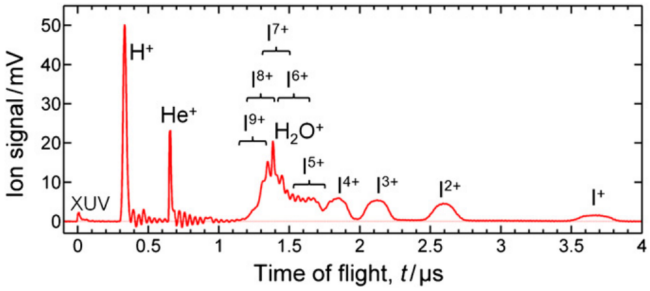


Figure 3.17: Ion time-of-flight spectrometer trace. Different ion species reach the detector at different times. Figure extracted from [111].

The working principle of TOF spectrometers is based on the classical equation of the kinetic energy,

$$E = \frac{1}{2}mv^2, \tag{3.3}$$

where E is the energy, v is the speed and m is the mass of a particle. This equation can be rewritten in terms of the time-of-flight t that a particles takes to travel a distance L,

$$E = \frac{1}{2}m\frac{L^2}{t^2}. \tag{3.4}$$

Using a calibrated equation of this form, this spectrometer allows one to perform energy-resolved electron spectroscopy [112, 113, 114], calculating the electron kinetic energy E as function of the time-of-flight t. Additionally for the detection of fast high-energy electrons, a retardation voltage may be applied at the beginning of the drift tube to reduce the speed of the particles and make the temporal detection possible, typically allowing energy resolutions better than 1% of the electron energy. More exhaustive details about the calibration of electron TOF spectrometers will be presented in section 5.3.2.

TOF spectrometers can also be used to perform ion spectroscopy [115, 116, 111]. The main typical purpose in this case is to identify the different ion species present

in an experimental sample after an ionization or a fragmentation event. In this case, the charged ions are pushed towards the detector with a fixed accelerating voltage V, allowing one to rewrite the time-of-flight equation in terms of the charge q of the ion,

$$qV = \frac{1}{2}m\frac{L^2}{t^2}. \tag{3.5}$$

After calibration, this equation can be solved to obtain the charge to mass ratio q/m as a function of the time of flight t. Furthermore, after identifying the mass and charge associated with each time-of-flight peak (figure 3.17), the kinetic energy of the ions emitted in the direction of the detector can also be extracted from the equation by adding an additional energy term. More details about ion time-of-flight spectroscopy and its calibration will presented in section 4.2.3.

The collection efficiency of TOF spectrometers depends on the type of spectroscopy. In the case of ion spectroscopy, the accelerating voltage pushing the ions towards the detector allows a collection efficiency of up to 100% depending on the geometry, the maximum kinetic energy of the ions and the magnitude of the voltage. In the case of electron TOFs the situation is much less favourable as the collection efficiency is determined by the solid angle subtended by the surface area of the detector at the interaction region and is in consequence extremely low. For a typical detector with 6 cm of diameter, situated at a distance of 75 cm, the collection efficiency would be as low as 0.04%. This can be improved by using electrostatic or magnetic lenses to steer a higher fraction of the electrons towards the detector at the cost of greatly degrading the final energy resolution as the electric field of the lenses may alter the energy of the electrons. A better option is to use a magnetic bottle spectrometer.

3.3.4 Magnetic-bottle spectrometer

The magnetic bottle time-of-flight spectrometer (MTOF) [117] is a variation of the TOF spectrometer where a magnetic field is used to steer the charged particles, typically electrons, towards the detector.

Classical electrodynamics describe the interaction of a charged particle with a magnetic field using the magnetic part of Lorentz force $\mathbf{F_B}$,

$$\mathbf{F_B} = q\mathbf{v} \times \mathbf{B}, \tag{3.6}$$

where $\mathbf{B}$ is the magnetic field, q is the charge of the particle, and $\mathbf{v}$ is the velocity of the particle. As the force is always perpendicular to the speed of the particle, the magnetic field cannot modify the velocity of the particle just its direction. In a uniform magnetic field, this results in the well-known circular trajectories when the magnetic field is perpendicular to the initial velocity ($\theta = \pi/2$), straight trajectories when they are parallel($\theta = 0$), and helical trajectories in all other cases. For an electron of charge

e and mass m the helical trajectory presents angular frequency $\omega = eB/m$, and radius $r = v\sin(\theta)/\omega$.

Early magnetic bottle spectrometers made use of a varying magnetic field from about 1 T in the interaction region to around 1 mT in the drift tube as shown in figure 3.18. Electrons generated at the interaction start following a trajectory of radius r_i and frequency ω_i according to the dynamics dictated by the initial emission angle θ_i, the initial magnetic field B_i and the initial speed v. These initial conditions set the initial angular momentum of the trajectory l_i,

$$l_i = \frac{m^2 v^2 \sin^2(\theta_i)}{eB_i}. \tag{3.7}$$

This angular momentum is conserved as the electron adiabatically moves towards low magnetic field regions, in particular, for a final position with field B_f, the final angle of the particle θ_f satisfies

$$\frac{\sin(\theta_f)}{\sin(\theta_i)} = \sqrt{\frac{B_f}{B_i}}. \tag{3.8}$$

As shown in figure 3.18, at a final position where the magnetic field is less intense, the angle of the particle becomes smaller, and the trajectory much more parallel to the magnetic field. As the initial velocity of the particle cannot change due to the interaction with the magnetic field, this implies a net transfer of momentum from the perpendicular direction to the direction parallel to the axis of the magnetic bottle.

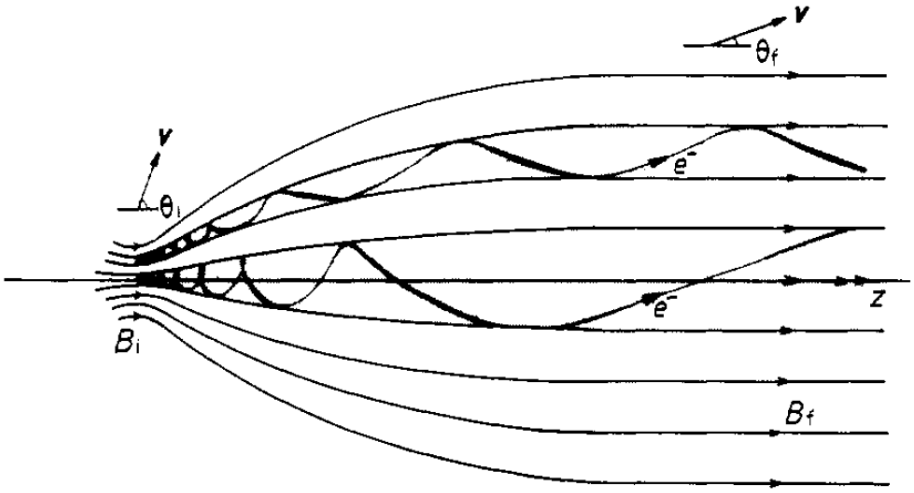


Figure 3.18: Magnetic field lines in a magnetic bottle spectrometer. Diagram showing the helical motion of electrons emitted initially at different angles as they are steered towards the same direction by the magnetic fields. Figure extracted from [117]. © IOP Publishing. Reproduced with permission. All rights reserved.

Following angular momentum conservation, it is also possible to show that the ratio between the initial and final magnetic fields determine the magnification factor $\frac{r_f}{r_i}$ of the

magnetic bottle,

$$\frac{r_f}{r_i} = \sqrt{\frac{B_i}{B_f}}, \tag{3.9}$$

with fields of 1 T and 1 mT yielding a magnification factor of ∼30. This is consistent with the adiabatic approximation as it shows that the magnetic flux enclosed in the orbit $B\pi r^2$ stays constant for the length of the trajectory.

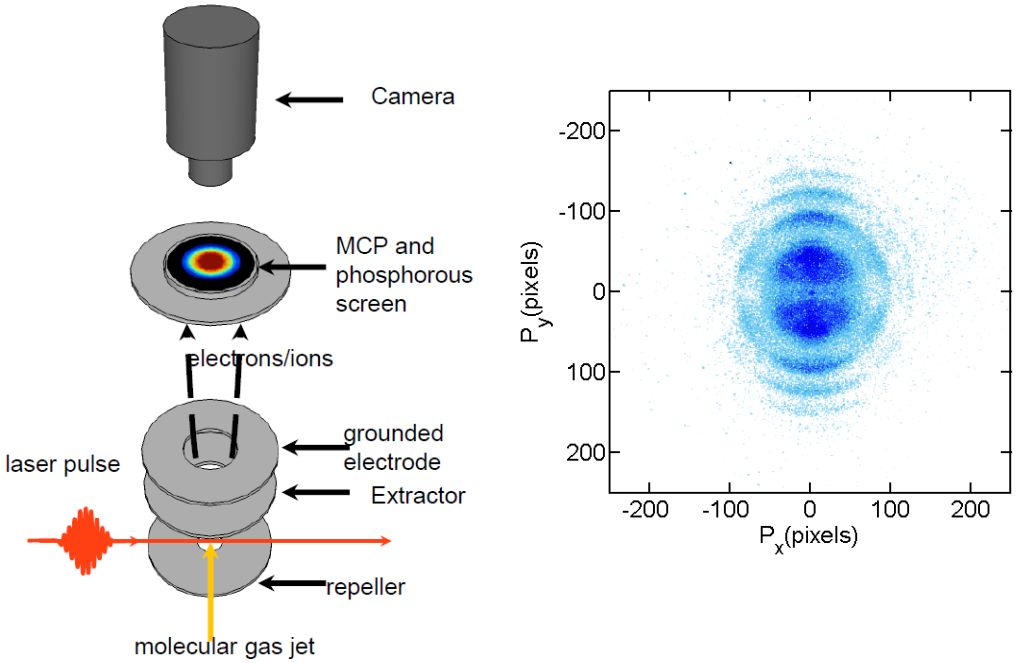
Magnetic bottles designed as in figure 3.18 make use of a sophisticated electromagnet surrounding symmetrically the interaction region. As a result the magnetic field lines are parallel at the interaction region, and electrons generated with angles larger than $\pi/2$ are steered towards the opposite direction, allowing for collection efficiencies around 2π steradians. Modern magnetic bottles, making use of large permanent magnets situated at only one side of the interaction region, allow electrons emitted with angles larger than $\pi/2$ to reflect off the magnetic field, steering all of them towards a common direction and achieving up to full 4π steradians collection efficiency [118]. These magnetic bottles are commonly used in XFEL experiments for both ion and electron spectroscopy [114, 119, 120, 121, 12].

This increased electron collection compared to regular TOF detectors allows for a reduction in the total amount of charged particles needed for a measurement, potentially decreasing the space-charge effects around the interaction region. A detector of this kind was conceptually designed by modelling electron trajectories for different emission angles using a commercial ion and electron optics simulator (SIMION [122]) as part of this Ph.D. An engineering design of the detector was later carried out by Z. Diveki and G. Galinis, and it is currently being assembled for use in photoelectron spectroscopy experiments.

3.3.5 Velocity-map imaging

A velocity-map imaging (VMI) device [123] is a system that images the projection of the momentum distribution at the interaction region onto an MCP with a fluorescent screen (figure 3.19a). Making use of a number of electrodes acting as extractor, repeller and electrostatic lenses, the particles emitted with the same momentum are focused to the same position on the detector, regardless of the particular initial location within the interaction volume. Assuming a symmetric emission around an axis perpendicular to the projection axis, the image obtained (figure 3.19b) can be inverted and the full three-dimensional momentum distribution of the emitted particles (of the same kind) can be obtained.

While this system allows one to measure the momentum distribution for large groups of particles, it does not unequivocally retrieve the momentum direction and/or kinetic energy for each individual particle, as in this case one of them is needed to obtain the other. Pure VMI devices using MCP detectors have not been used for the results presented in this thesis, instead delay-line anodes in VMI configuration have been used.



(a) (b)

Figure 3.19: VMI spectrometer. (a) VMI experimental setup. A repeller is used to push the particles towards the detector. Different sets of lenses modify the trajectory of the particles to image the projection of the momentum in the interaction region at the detector plane. (b) Examples of above-threshold ionization (ATI) photoelectrons from xenon. Different photoelectron peaks, corresponding to different momenta/kinetic energies, show as rings in a VMI trace. Figures provided by Thomas Barillot.

3.3.6 Delay-line anode

The delay-line anode [124] is a device that can be used in conjunction with an MCP to measure both time of flight and position for each particle reaching the detector. A two-layer delay-line anode consists of two different delay lines placed after an MCP detector as shown in figure 3.20a.

When a particle hits the MCP with the two-layer delay-line anode behind, five different timing signals are generated: two signals at the ends of the horizontal delay line t_l and t_r, two signals at the ends of the vertical delay line t_u and t_d, and a signal at the MCP t_{MCP}. By comparing the time difference between the signals generated at the end of each delay-line anode, information about the horizontal x and vertical y position of the particles can be obtained, in particular,

$$x = a_x(t_l - t_r) + b_x \tag{3.10}$$

$$y = a_y(t_d - t_u) + b_y, \tag{3.11}$$

where a_x, b_x, a_y and b_y are scaling and offset constants obtained from the manufacturer and refined by calibration. Additionally the corresponding time-of-flight of the particle is measured by t_{MCP}.

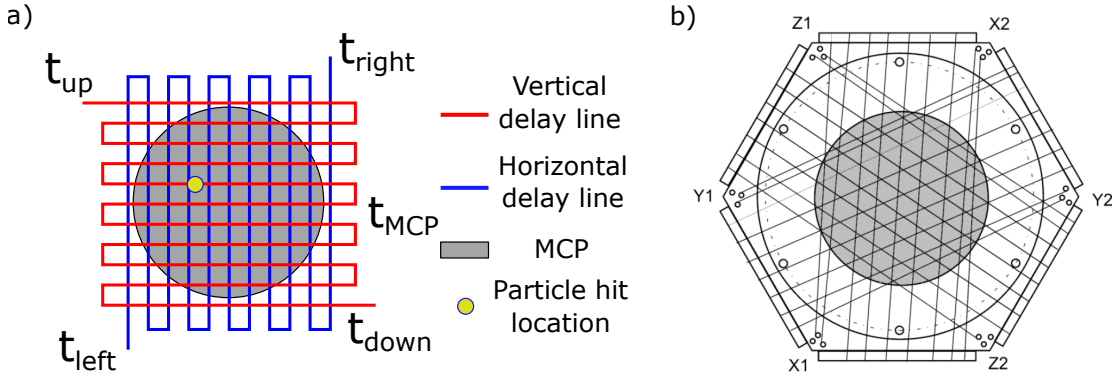


Figure 3.20: (a) Diagram of a two-layer delay-line anode placed behind an MCP. When a particle hits the MCP, five different signals are generated: one at the back of the MCP (t_{MCP}) and four more at the ends wires forming the layers of the delay-line anode (t_{up}, t_{down}, t_{left}, t_{right}). (b) Sketch of a three-layer delay-line anode (hexanode) extracted from [125] © 2011 IEEE.

Before applying these equations, it is important however, to guarantee that all five signals have been generated by the same particle. In order to verify this, the time constant of the individual delay lines, which is the total time that a signal takes to travel the full length of the delay line, can be used. Regardless of the particular location at which a particle hits the detector, the sum of the arrival times at both ends of the delay with respect to the arrival time of the particle must be equal to that time constant. While the exact time when the particle hits the detector cannot be measured, the time measured at the MCP t_{MCP} is independent of the position, and also serves as an absolute reference. This allows one to write the following equalities that must be satisfied when all five signals are generated by the same particle,

$$(t_l - t_{\mathrm{MCP}}) + (t_r - t_{\mathrm{MCP}}) = t_l + t_r - 2t_{\mathrm{MCP}} = c_x \tag{3.12}$$

$$(t_d - t_{\mathrm{MCP}}) + (t_u - t_{\mathrm{MCP}}) = t_d + t_u - 2t_{\mathrm{MCP}} = c_y, \tag{3.13}$$

where c_x and c_y are time constants obtained from the manufacturer and refined by calibration.

One of the limitations of the two-layer delay-line anode is that it cannot distinguish between the x and y coordinates of two particles arriving at two different locations at exactly the same time. In order to solve this problem, an additional delay line is added in the *hexanode* device, presenting a total of three delay lines at 60 degrees from each other (figure 3.20b). In this case seven signals are generated allowing one to obtain the locations across each delay line, typically denoted as u, v and w, which can then be used

in pairs [126] to provide redundant information about the x and y coordinates,

$$x_{uv} = u + O_x \tag{3.14}$$
$$y_{uv} = \frac{1}{\sqrt{3}}(u - 2v) + O_y \tag{3.15}$$
$$x_{uw} = x_{uv} \tag{3.16}$$
$$y_{uw} = \frac{1}{\sqrt{3}}(2w - u) + O_y \tag{3.17}$$
$$x_{vw} = (v + w) + O_x \tag{3.18}$$
$$y_{vw} = \frac{1}{\sqrt{3}}(w - v) + O_y, \tag{3.19}$$

where O_x and O_y are arbitrary constants. This redundant information can be used when some of the signals are missing, or to resolve the situation when two particles arrive at exactly the same time. While the hexanode presents in principle more advantages than the two-layer delay-line anode and the previous detectors, the retrieval of the position when many particles arrive at the detector is in practice so complex that is only unambiguous as long as not more than two particles arrive simultaneously [125], and not typically used to detect more than 10-20 particles per measurement.

Making use of a hexanode in a VMI configuration it is possible to independently obtain the mass-to-charge ratio, kinetic energy and momentum for each of the arriving particles. This allows one to, for example, obtain the full three dimensional momentum of the multiple charged fragments after the Coulomb explosion of a small molecule (section 4.1).

3.3.7 Time tool

Some experiments require the use of synchronised XFEL radiation and optical laser radiation. However, due to the inherent time-jitter in the XFEL pulse generation the timing between these two sources has not been successfully stabilized yet to better than a few hundreds of femtoseconds.

On the other hand it is already possible to measure this jitter on a shot-to-shot basis and sort the data. This is achieved using time tools that gate the transmission of an optical pulse (correlated to the pulse used for the experiment) with the x-ray pulse using a thin optically-transparent Si_3N_4 membrane that has its optical properties strongly modified by the free-charges generated by the x-ray pulse.

In the spatial domain [98], the x-ray/optical pulse delay can be obtained by crossing the x-ray beam and the optical beam in the membrane at an angle and measuring the transmitted profile of the optical pulse (figure 3.21b). On the other hand, it can also be obtained in the spectral domain [97, 98] by creating a broadband chirped pulse much longer than the x-ray pulse and propagating both of them collinearly through the membrane. Since for the chirped pulse wavelength maps into time, by measuring the

spectrum of the outgoing optical pulse and finding the wavelength range affected by the x-rays, it is possible to retrieve the changes in the relative timing delay (figure 3.21a). These methods allow measurement of the jitter with a precision down to approximately 6 fs, which may be enough to sort the data for some experiments. Additionally, the spatial and spectral approaches can even be combined to provide a sub-fs accuracy [99], however, practical issues of the implementation along with a real experiment typically prevent the main experiment from benefiting from such sub-fs accuracy. For more details about these issues, the retrieval of the timing and the capabilities of the time tool at SACLA, see section 6.2.

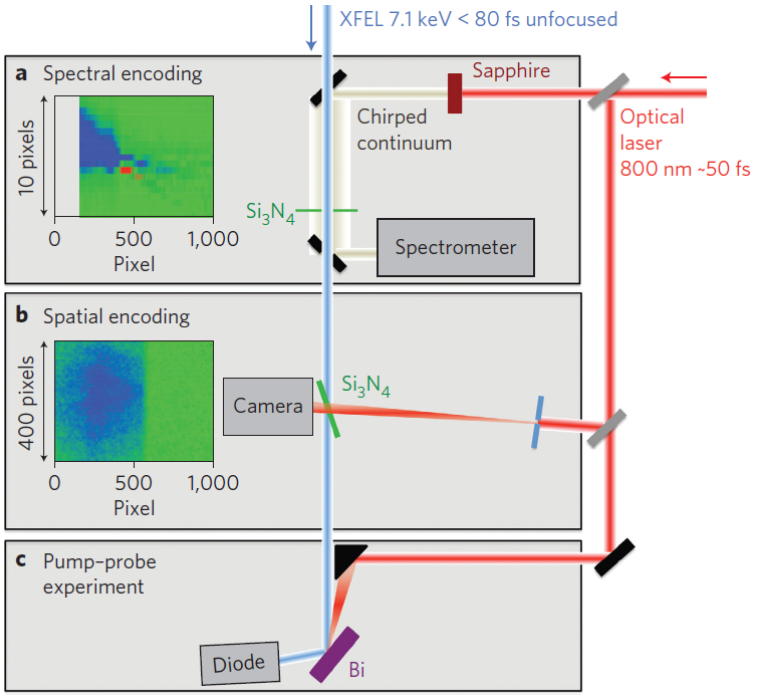


Figure 3.21: Spectral and spatial encoding of x-ray-optical-pulse jitter using a time tool. The x-ray beam and the optical beam are combined in three different experiments running simultaneously. (a) A chirped continuum pulse obtained from the optical beam propagates collinearly with the x-ray pulses through a Si_3N_4 screen. The x-ray pulses modify the transmission differently for different frequencies of the chirped pulse depending on the relative arrival times. This timing information can be extracted by recording the spectrum of the chirped continuum after the screen. (b) In this case the x-ray beam and the optical beam are combined in the Si_3N_4 screen at an angle. This arrangement maps different parts of the screen into different relative arrival times, and allows the extraction of the timing information by recording the spatially-resolved optical transmission through the screen. (c) The x-ray beam and optical pulses interact at the interaction region of the main experiment being carried out. Figure extracted from [98].

3.3.8 XTCAV

An alternative approach to measuring properties of the generated x-ray pulses directly, is to infer some of these properties from the electron bunch/es that generate the pulse/s

in the first place. This presents the advantage that, contrary to the x-ray pulses, the electron bunches are no longer needed after the undulators, and can be probed intrusively. This can be achieved using a transverse deflector cavity called XTCAV at LCLS [127, 128]. XTCAV streaks the electron bunches with a radiofrequency field along one direction perpendicular to the propagation, followed by a magnetic dipole in the other perpendicular direction (figure 3.22). This process produces a two-dimensional trace showing the time-dependent electron energy distribution. By comparing this trace to the one generated in a situation when no lasing is taking place, information about the temporal properties of the x-ray pulses can be retrieved.

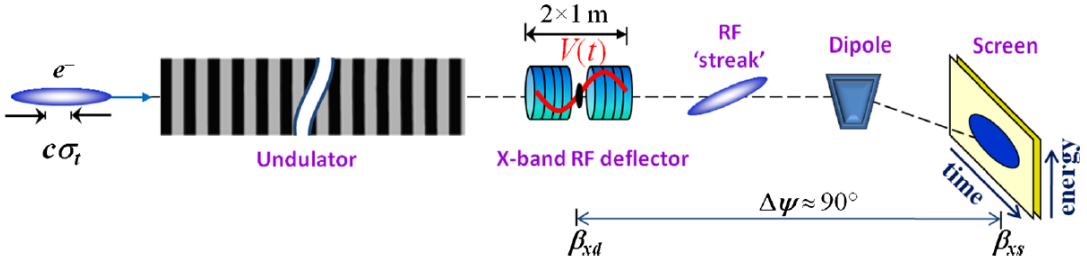


Figure 3.22: XTCAV setup at LCLS. After the undulator, the electron bunch is streaked with an x-band RF deflector along one axis and dispersed using a magnetic dipole along the perpendicular axis. As a result an image of the time-resolved energy distribution of the bunch can be recorded by placing a screen after the dipole. Figure extracted from [127].

XTCAV has been a key tool to obtain the results presented in this thesis. Working with XTCAV traces involve a series of non-trivial calibrations that have been developed and improved in the context of this Ph.D. Much more detailed information about these procedures, including the retrieval of the temporal properties of the x-ray pulses will be provided in section 5.4.

3.3.9 Beam-position monitors

Another very useful electron-beam diagnostic, yet simpler than XTCAV, is the beam-position monitor (BPM) [129]. BPMs are electromagnetic devices consisting of two resonant cavities, a monopole cavity and a double dipole cavity, which probe the electric field created by the passing electron bunches. The monopole cavity is used to measured the total charge present in the electron beam, while the double dipole cavity is used, after calibration, to measure the horizontal and vertical position of the electron bunch. These devices are of critical importance for XFEL alignment and tuning as they are non-intrusive and offer a spatial resolution of up to 200 nm for charges as low as 200 pC [129]. Furthermore, when used in a dispersive region, such as a magnetic chicane, they can be used to measure exact fluctuations in the electron-beam energy mapped into spatial fluctuations. A typical XFEL like LCLS utilizes hundreds of this type of detector, although their information is not always recorded shot-to-shot.

3.4 Comparison to optical lasers

As optical lasers and XFELs are systems of completely different nature, they present many significant differences in the modes of operation and characterization. Some of these differences have already been introduced in the previous sections and are summarized here:

- While the size of an optical laser, including an HHG setup, can fit in a $30\,\mathrm{m}^2$ room, the biggest XFELs available needs dedicated facilities of a few kilometres in length. This also increases the cost of the systems by many of orders of magnitude.

- The wavelengths available from optical lasers are limited by the emission band of the levels used for lasing. This strongly limits the tunability and the maximum achievable photon energy. Other approaches such as parametric amplifiers allow continuous tuning of the wavelength, however, they require powerful driving lasers and typically only allow generating longer wavelengths than that of the driving laser. On the other hand the emitted wavelength in an XFEL depends on the electron-bunch energy, the magnetic field and the undulator spacing. By designing the undulator with the right spacing, and then varying the magnetic field strength and electron energy it is possible to emit radiation from infrared ($\sim 10^{-1}\,\mathrm{eV}$) to hard x-rays ($\sim 10^{4}\,\mathrm{eV}$) while keeping a high degree of tunability for a given design. It is worth noting that even though it is possible to produce soft x-ray pulses through high harmonic generation with optical lasers, due to the low efficiency, the intensities can be several of orders of magnitude lower due to the small number of photons per pulse ($\sim 10^{6}$) compared to that of an XFEL ($\sim 10^{12}$).

- Due to the SASE process, which is initiated by the stochastic emission of spontaneous radiation, a typical XFEL pulse is formed by several temporal spikes within the width of the pulse length that are amplified in the SASE process separately and do not present coherence with respect to each other. Although this can be minimized by working in low charge modes, or using XFEL seeding, the temporal profile of an optical laser is in general better defined and more stable.

- For an experiment only involving optical lasers, it is common to trigger all devices and processes using the laser pulses themselves, allowing for a simple optical synchronization. In case of XFELs, due to the great distance from the electron source to the interaction region, the fact that there are travelling electrons involved in the process and the SASE process itself, there is an inherent shot-to-shot jitter between the main trigger and the arrival time of the x-rays. This fact makes the experiments involving both optical lasers and XFELs at the same time extremely challenging, requiring special techniques to correct or at least measure this jitter.

Being fully aware of these differences is critical when writing research proposals for XFEL beamtimes and planning XFEL experiments, specially for researchers with mainly optical laser experience.

Chapter 4

Measuring nuclear dynamics on 10-100 fs timescales

While this thesis concentrates on performing ultrafast few-fs experiments, many interesting x-ray-triggered nuclear dynamic processes in molecules occur on longer timescales, between a few tens of femtoseconds and a few picoseconds. These relatively long timescales allow more flexibility when choosing the XFEL multipulse operation mode, as they do not normally require more sophisticated techniques such as single-shot time-stamping (sections 5.4 and 6.2), using nominal time-delay values from the machine configuration instead. As a consequence users typically favour modes providing a higher intensity over modes providing very short pulses or with time-stamping capabilities.

In this chapter the author will discuss two different experiments aimed at measuring nuclear time dynamics in acetylene (0-100 fs) and in C_{60} (25-925 fs).

4.1 Ultrafast isomerization in acetylene

An isomerization process occurs when the nuclei forming a molecule rearrange themselves spatially from an initial configuration to a final configuration. The purpose of this experiment was to measure the time dynamics of the isomerization of acetylene molecules after sudden x-ray-induced core-ionization, in particular, the dynamics of the proton motion required for the transition from the initial linear configuration, characteristic of neutral acetylene molecules (figure 4.1a), and the vinylidene configuration (figure 4.1b). Rapid proton migration is a fundamental process for hydrocarbon photochemistry, and its investigation in acetylene, the smallest hydrocarbon that can isomerize, is a key to understanding the non-adiabatic electron-nuclear interactions involved in the isomerization. Furthermore, learning about the dynamics of processes following core-ionization, like the nuclear rearrangement happening during the isomerization, is critical in x-ray science as the majority of the experiments make use of photon energies large enough to ionize from the core, for example in diffraction-before-destruction experiments [25, 26].

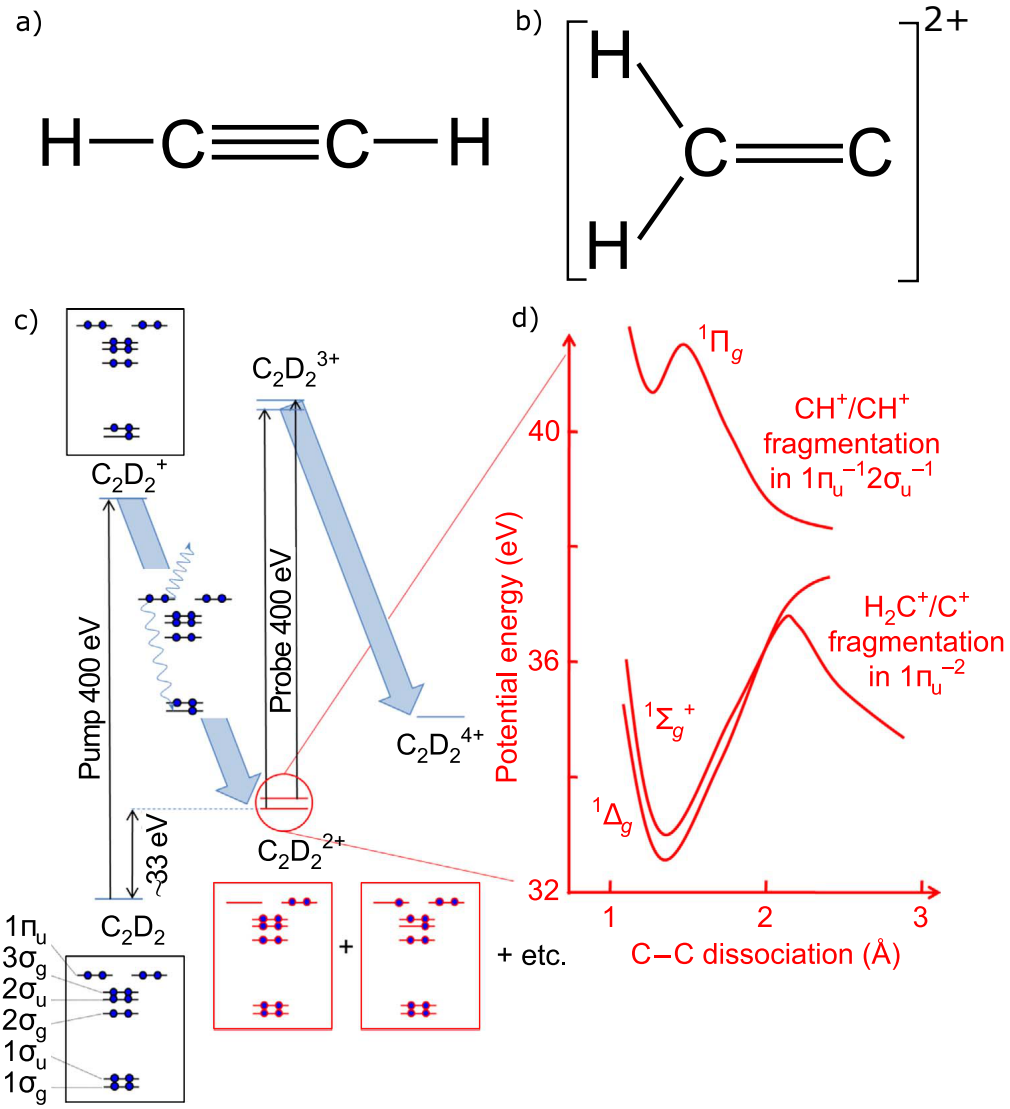


Figure 4.1: (a) Structure of the linear configuration of acetylene. (b) Structure of the vinylidene isomer. (c) Energy level diagram for the pump-probe process. The pump pulse induces core photoionization in the neutral molecule, followed by an Auger decay into a doubly charged molecule. This process can populate mainly three different cationic states with different potential energy surfaces shown in (d) [130]. The probe pulse induces another core-ionization, followed by an Auger decay, leaving the molecule in a 4+ cationic state typically exploding due to Coulomb explosion in 4 different fragments. Figures (c) and (d) extracted from [9].

While this experiment was carried out in June 2013, a few months before the beginning of this Ph.D., the author played a significant role in the analysis of the data led by C.E. Liekhus-Schmaltz. In this section we will briefly introduce the experiment to later discuss the data analysis. Extensive experimental and theoretical information can however be found in the related publications [9, 131].

In order to decrease the speed of the isomerization process to facilitate the measurement we used deuterated acetylene C_2D_2: acetylene molecules where the hydrogen atoms are replaced by deuteron isotopes with double the mass. This also helps distinguishing acetylene protons from hydrogen atoms originated from other background species such as H_2O. The target isomerization process is described in figure 4.1c. A high-energy photon from the pump pulse removes an electron from the core of one of the carbons, leaving the molecule in an excited charged state $C_2D_2^{+*}$ which populates dicationic states $C_2D_2^{2+}$ through an Auger emission process. The potential energy surfaces of some of these states are shown in 4.1d. The most excited state of the three $^1\Pi_g$ corresponds to the fragmentation in the acetylene configuration CD^+/CD^+ (figure 4.1a). Alternatively isomerization into vinylidene may occur if the final state is the $^1\Sigma_g^+$ [130], which can also fragment into C_2D^+/D^+ given enough kinetic energy.

Some time after the isomerization process has been initiated in the $C_2D_2^{2+}$ dication, a photon from the probe pulse removes another core electron from one of the carbons, followed again by an Auger decay and leaving a quadruply charged ion $C_2D_2^{4+}$. This ion, which is highly unstable due to the strong charge repulsion, fragments in one of many possible channels through Coulomb explosion. In particular it can fragment into four different atomic ions $C^+/C^+/D^+/D^+$. By analysing the momentum of the individual ions after the fragmentation, information can be extracted about the geometry of the molecule and the location of the proton[1] at the arrival time of the probe pulse, in what is called Coulomb explosion momentum mapping (CEMM).

The proton motion in the acetylene molecule following photoionization has been previously studied in singly-ionized and doubly-ionized acetylene molecules, making use of optical (800 nm) pulses [132], optical laser based UV pulses (266 nm) [133], XUV pulses (38 eV) generated with a free-electron laser (FLASH) [134], and x-ray pulses (310 eV) from a synchrotron [135, 130]. All of these experiments show significant signs of proton migration within the first 100 fs after the ionization and, in particular, within the first 60 fs for the x-ray initiated case. The presence of proton motion, however, does not necessarily indicate isomerization. The excitation of the molecule following core-ionization and Auger decay may also lead to a more random motion of the C-H bond angle. This undirected proton migration may occur in addition to the directed proton motion that is associated with the isomerization. Molecular vibrations can also cause, for example, protons to move within femtosecond timescales [136]. Furthermore, significant proton motion may occur just before the fragmentation of a molecule [137,

[1]For simplicity, when discussing experimental results we will use "proton" instead of "deuteron".

136] without leading to a stable isomer. Distinguishing between these processes using the CEMM technique is not easy and generally requires supporting theoretical modelling.

Considering the upper limit to the proton motion of 60 fs provided by previous experiments [135, 130], we concentrated our work [9] in the 100 fs after the ionization.

4.1.1 Experimental setup

The experiment was performed at LCLS, in the Atomic and Molecular Optics (AMO) beamline. Single pulses of 10 fs of nominal duration and 100 µJ of energy tuned at 400 eV, well above the carbon K-edge, were generated using the single-slotted foil (reference [81] and section 3.1.4). We used the x-ray split and delay (XRSD) apparatus (references [89, 87] and section 3.2.1) to split the single pulse into the pump and probe pulses. We used the XRSD in four different configurations to introduce nominal time delay values of 12 fs, 25 fs, 50 fs and 100 fs. For the nominal 0 fs time delay we simply reflected the beam off one of the mirrors. The pulses were focused into a 50 μm^2 region, where the sample C_2D_2 (in 1:3 mixture with He) was delivered through a nozzle using a pulsed valve with a backing pressure of 2.4 bar. Since we were interested in measuring the ions in coincidence, i.e. simultaneously measuring 4 different ions from the same molecular fragmentation, the sample pressure and position of the nozzle was adjusted to obtain a single molecular ionization per pulse on average. While the XRSD apparatus must be correctly calibrated to preserve spatial overlap at different delays, similar double core-ionization rates were observed for all times delays, indicating consistency in the spatial overlap as the double core-ionization is a non-linear process in the fluence.

We used a delay-line anode (section 3.3.6) in a VMI (section 3.3.5) configuration as an ion detector. This allowed us to simultaneously measure the time-of-flight and the position in the detector which was then used to obtain the mass, charge, kinetic energy and three-dimensional momentum of each of the ions using the advanced reconstruction technique [125] for the analysis of the delay-line anode signals.

Under these conditions we recorded approximately $2 \cdot 10^6$ events at each time delay, corresponding to approximately 4.5 hours per delay at 120 Hz. As the physical core-ionization mechanism used was not a resonant process, the experiment was insensitive to shot-to-shot jitter in the photon energy. Furthermore, performing the experiment in coincidence allowed us to remove much of the effect of the intensity fluctuations by selecting only events for which the molecule is exposed to two-photon interactions.

4.1.2 Data analysis

The most critical aspect of the data analysis of this experiment was distinguishing the true coincidences from the false coincidences: identify events containing $C^+/C^+/D^+/D^+$ ions coming from the same molecule. This was done in different stages. We first filtered out the events containing more than 6 particles and events not containing at least two C^+

ions and two D^+ ions. After identifying the different possible 4-tuple $C_1^+/C_2^+/D_1^+/D_2^+$ candidates for each event, we applied a constraint to keep only the events with a total momentum p_j close to zero ($< \delta p_j$) in each of the directions,

$$\mid p_{D_{1,j}} + p_{D_{2,j}} + p_{C_{1,j}} + p_{C_{2,j}} \mid = p_j < \delta p_j, \tag{4.1}$$

as this is what we would expect in cases where the four particles truly came from the fragmentation of a single molecule.

Additionally, in order to obtain consistent locations of the protons within the molecular frame, we needed the carbon-carbon (CC) axis to be well defined. This was constrained by requiring that the momentum vectors of the carbons involved in the same fragmentation were approximately parallel and pointing to opposite directions. In particular we checked that the relative angle between the momentum vectors of the carbons was close to 180 degrees,

$$\cos^{-1}\left(\frac{\mathbf{p}_{C_1} \cdot \mathbf{p}_{C_2}}{|\mathbf{p}_{C_1}||\mathbf{p}_{C_2}|}\right) = \theta_{C_1-C_2} > \pi - \delta\theta_{C_1-C_2}. \tag{4.2}$$

This filtering process reduced the number of events per delay value from $2 \cdot 10^6$ to $\sim$1200. Such a large decrease (down by a factor of $\sim$1500) in the number of events is typical in coincidence experiments. The main reason behind this has to do with the detection probability P_i of the individual particles, which is typically between 10% and 20%. Because of this, as the number of particles n to be detected increases, the probability P_{all} of detecting all of them decreases exponentially $P_{\text{all}} = (P_i)^n$.

For the events satisfying all the constraints we calculated the carbon-carbon-deuteron (CCD) angle θ defined as,

$$\theta = \cos^{-1}\left(\frac{\text{sign}((\mathbf{p}_{C_1} - \mathbf{p}_{C_2}) \cdot \mathbf{p}_{D_2})(\mathbf{p}_{C_1} - \mathbf{p}_{C_2}) \cdot \mathbf{p}_{D_1}}{|\mathbf{p}_{C_1} - \mathbf{p}_{C_2}||\mathbf{p}_{D_1}|}\right). \tag{4.3}$$

This CCD angle measures the angle between the vector defined by the CC axis (pointing to the direction of the momentum of the second deuteron D_2) and the momentum of the first deuteron D_1, as shown in figure 4.2a. As the choice of D_1 and D_2 is arbitrary, we calculated two angles for each coincidence, one for each deuteron playing the role of D_1. The totality of the events obtained from all time delays sorted as function of this angle and the total kinetic energy are shown in figure 4.2b. Based on the definition of the CCD angle, values of the angle smaller than 90 degrees should correspond to the vinylidene-like configuration and values above 90 degrees to the linear acetylene configuration. Looking at the kinetic energy distributions for each of these groups independently 4.2c, we observed that the acetylene-like fragmentation has an overall higher total kinetic energy, consistently with the potential energy curves shown in figure 4.1d.

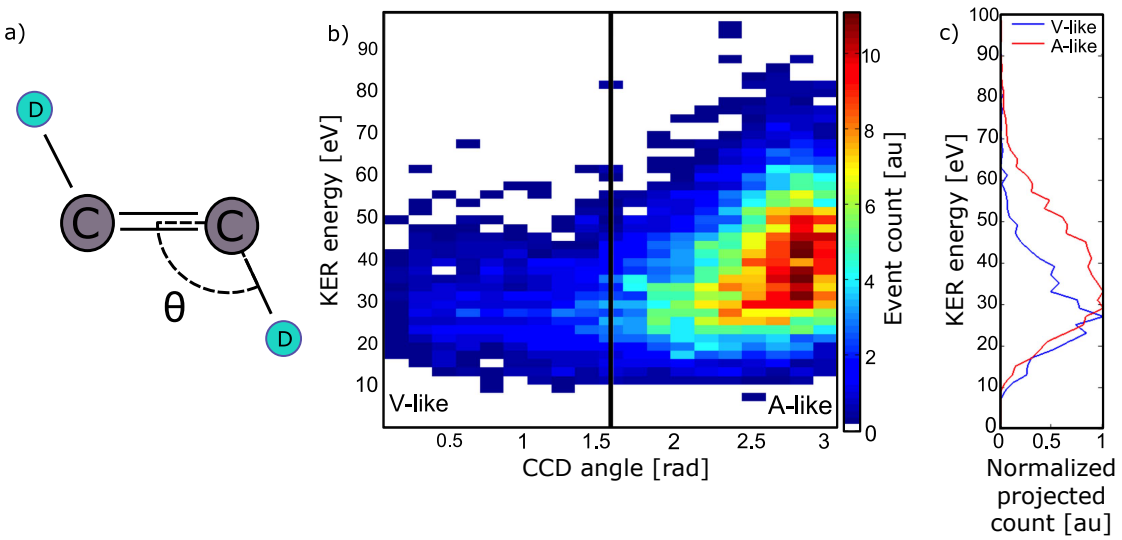


Figure 4.2: (a) Definition of the carbon-carbon-deuteron (CCD) angle θ. (b) Distribution of $C^+/C^+/D^+/D^+$ coincidences as function of the total kinetic energy released and the CCD angle including data for all time delays. Coincidences to the left of the vertical black solid line correspond to vinylidene-like configurations, while data to the right corresponds to the linear acetylene configuration. (c) Coincidences from the acetylene region show an overall higher total kinetic energy than coincidence from the vinylidene region. Figure partially extracted from [9].

4.1.3 Time dynamics of the proton motion

With access to the three-dimensional momentum for each of the four particles involved in true coincidence events, it was possible to separate the components of the momentum of the deuterons in the directions perpendicular and parallel to the CC axis, and plot this for each of the time-delay values (figure 4.3b-f). We can observe a clear difference between the 0 fs delay case (figure 4.3b), where most of the deuteron momentum is in the direction of the CC axis, and the subsequent values of the delay (figure 4.3c-f), where part of this momentum is transferred to the direction perpendicular to the CC axis. This change in the momentum indicates a fast proton motion from being aligned with the CC axis to locations out of the CC axis, and would be compatible with an isomerization between the acetylene and vinylidene configurations (figure 4.3g).

In order to quantify the proton motion, we calculated the ratio between the signal in the perpendicular direction (red region) and the signal in the parallel direction (blue region) of the CC axis as defined in figure 4.4a at each time delay, resulting in Figure 4.4b. This shows that the proton motion already starts before the first experimental point at 12 fs. Additionally we observe a decrease in the angle at 25 fs. While this decrease is small, it is discernible within the experimental error bars, and could be a sign of vibrational coherence, originated by the asymmetric bending mode of the $^1\Sigma_g^+$ state, with a half period of precisely 25.5 fs (654 cm^{-1}) [138].

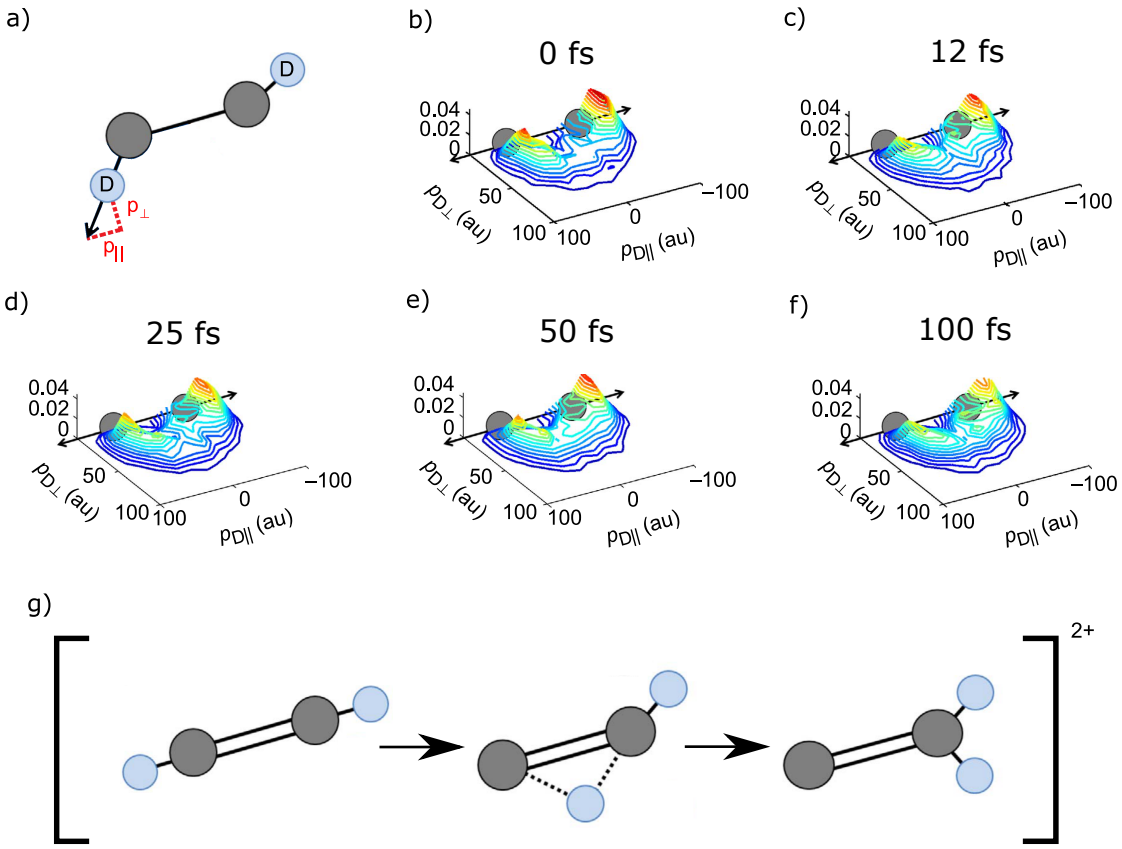


Figure 4.3: (a) Diagram of the momentum direction for each of the deuterons with respect to the CC axis. (b-f) Distribution of deuteron momenta in atomic units with respect to the CC axis at different time delays of 0 fs, 12 fs, 25 fs, 50 fs and 100 fs. As the time delay increases, the momentum very strongly parallel to the CC axis initially, redistributes to directions perpendicular to the CC axis. This would correspond to the transition between the linear configuration and the vinylidene configuration as shown in (g). Figure extracted from [9] and modified.

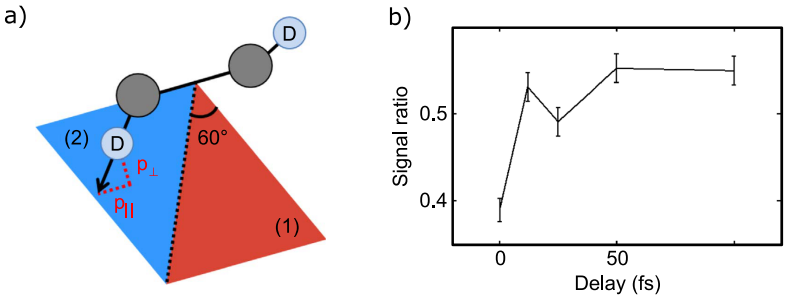


Figure 4.4: (a) Regions defined in the momentum space to differentiate the vinylidene configuration (red region) from the linear configuration (blue region). (h) Ratio of the signal in region 1 divided by the signal in region 2 as defined in (a) for the different time delays. There is a significant increase of this ratio happening as early as 12 fs after the ionization, indicating very fast proton dynamics. Figure extracted from [9] and modified.

4.1.4 Conclusions and outlook

The results show very fast dynamics in the proton motion, occurring already within the first 12 femtoseconds after the core-ionization. These results, consistent with the previously reported upper limit of 60 fs [135, 130], were published as a proof of x-ray core-initiated sub-100-fs isomerization [9].

Further investigation, however, found inconsistencies between the measured sub-100-fs isomerization, and independent theoretical studies estimating that the isomerization of the acetylene dication should occur in much longer timescales [139]. In particular these models indicate that true isomerization occurs in the low-lying states of the dication, and requires overcoming a potential barrier of 2 eV, leading to picosecond timescales and being unfeasible at sub-100-fs timescales.

While this topic is still open for debate, some members of the original collaboration, but not including the author of this thesis, have attempted to solve these inconsistencies by performing additional modelling. The related work [131] concludes that in fact the main process contributing to the sub-100-fs dynamics is not isomerization, but significant proton motion before fragmentation of the molecule. This is in line with what the author already mentioned in the introduction: that proton migration does not necessarily imply isomerization; and calls for careful interpretation of the results obtained with the Coulomb explosion momentum mapping method.

In the context of this thesis, this experiment is the first example of a time-resolved measurement. In this case the experimental procedure allowed the results to be very robust against XFEL fluctuations, examples of which will be shown in the following sections. The above-edge ionization process is insensitive to photon energy fluctuations. The measurement in coincidence reduces the effects of the intensity fluctuations. Finally the generation of a single pulse and the use of the XRSD apparatus to split it into two pulses allows for additional stability compared to two-pulse modes, and in particular, sub-femtosecond stability in the pump-probe delay. As we will learn throughout the rest of this thesis, this is not the typical case, especially for few-fs experiments not using the XRSD apparatus, and it is often required to understand the fluctuations, and perform complex data analysis to explicitly avoid their effects.

4.2 X-ray-induced Coulomb explosion in C_{60}

This experiment is the continuation of the results published by B. F. Murphy et al. [121] in 2014 aiming to understand the Coulomb explosion occurring in the C_{60} molecule after x-ray ionization, by explicitly measuring time-resolved dynamics. Most of the numerical modelling related to this experiment is still under development by the same group that performed the numerical modelling for the previous work [121]. However, the analysis of the experimental data has already been completed and shows clear evidence of time

dynamics that are being prepared for publication [15]. We will use this as another example of the typical data analysis performed for a basic time-resolved experiment.

The C_{60} molecule or buckminsterfullerene consists of 60 carbon atoms arranged in the configuration of a soccer ball, with twenty hexagons and twelve pentagons, justifying its common name: *bucky-ball*. In this work C_{60} is used as a showcase molecule for understanding what happens to the structure of a large molecule right after sudden core-ionization. Being able to model this kind of behaviour is very relevant to assess the reliability of methods like diffraction-before-destruction [25, 26], to determine how short an x-ray pulse needs to be to image a molecule without seeing any structural changes caused by the ionization of the pulse itself. However, modelling very large systems using full quantum mechanical methods is numerically unfeasible, and there is interest to see to what extent a classical mechanical model can be used to predict the behaviour of a Coulomb explosion when compared to experimental results. C_{60} has the advantage of containing only carbon atoms, while at the same time having a similar size to other small/medium size biological molecules regularly used in XFEL experiments.

4.2.1 Physical mechanism and prior experiment

The Coulomb explosion process starts with the photoionization of a core-electron from one of the carbon atoms in the molecule triggered by the absorption of an x-ray photon with photon energy above the carbon K-edge (280 eV). Following the core-ionization, the molecule is in a very excited charged state C_{60}^{+*}, which decays through the emission of an Auger electron into a double charge state C_{60}^{2+}. Under normal circumstances, the C_{60}^{2+} molecule may eventually split into two individual single ionized fragments of different sizes, however due to the large photon flux available at XFELs, many core-photoionization processes followed by Auger emission can occur in the same molecule before it breaks, charging the C_{60} molecule and accumulating potential energy which eventually leads to a Coulomb explosion. This explosion causes the molecule to break into much smaller fragments, which can keep absorbing photons through core-ionization and continue increasing their charge. This eventually leads to an ion distribution consisting of many different fragments sizes (from C_{60} to C_1) each of them in many different charge states.

The distribution of generated ions depends directly on the x-ray pulse properties. In particular the generation of highly-charged single carbon ions, which requires multiple photoionization processes to occur in the exact same atom of the molecule, is favoured only under very specific circumstances. In the experiment by B. F. Murphy et al. [121] they used a single x-ray pulse at 485 eV, with different pulse durations of 4, 30 and 90 fs and similar pulse energies ($\sim$300 µJ) to study the differences in the ion yield. They observed that longer pulse durations generated highly charged ions between C^{2+} to C^{6+} more efficiently. The effect is explained in terms of electron recombination:

for a given number of photons, a similar amount of ionization is expected in a given molecule, however, if all the ionization happens within a short period of time before the molecule explodes, there is a higher probability that the electrons and the ions would recombine together, making it more difficult to produce very highly charged states. In this sense, the different length of the pulses indirectly yields information about the time dynamics of the Coulomb explosion. In line with this, the purpose of the follow-up experiment was to use a pump-probe approach to explicitly measure the time-resolved dynamics of the Coulomb explosion using the same mechanism. More information about the experimental details and the prior results is available in reference [121].

4.2.2 Experimental setup

This experiment was performed using the AMO instrument at LCLS. The fresh-slice mode (section 3.2.4) was used to generate the pulses. This method allowed the generation of pairs of independent ∼5-10 fs pulses with summed x-ray pulse energies of up to 0.9 mJ from a single ∼30 fs long electron bunch. As shown in figure 4.5a the trajectory of the bunch was controlled to induce lasing in the head of the bunch at the first undulator section, and the tail of the bunch at the second undulator. This effect can be readily observed looking at XTCAV images corresponding to the bunch at different stages of the process. Figure 4.5b was obtained by disabling the lasing in both undulators, it shows a very flat profile of the electron bunch in the energy-time phase space, corresponding to the state of the electron bunch before entering the first undulator. Figure 4.5c was obtained by disabling the lasing only in the second undulator, corresponding to the state of the bunch after the first undulator. It shows changes centred at around -5 fs (the head of the bunch), which are a consequence of the lasing process in the first undulator. Finally, figure 4.5d shows the electron bunch after lasing in both undulators where more changes appear at around 10 fs (the tail of the bunch) when compared to figure 4.5c, and corresponding to the generation of the second pulse in the second undulator. More details about how to disable the lasing in the undulator, and XTCAV images in general will be provided in section 5.4.

The photon energy of the pulses was tuned well above the carbon K-edge, at 640 eV. This value was chosen to match previous conditions for the fresh-slice configuration (which was still being commissioned).

The time delay between the pulses was changed by means of a magnetic chicane situated between the two undulator sections delaying the electron bunch with respect to the pulse generated in the first undulator. This limited the minimum delay achievable to approximately 25 fs corresponding to the separation between the tail and the head lasing parts of the bunch (∼15 fs), and the time slippage due to different propagation speeds between the x-ray pulse and the electron bunch, which had a nominal value of approximately 10 fs. This limitation could have been avoided by having the electron

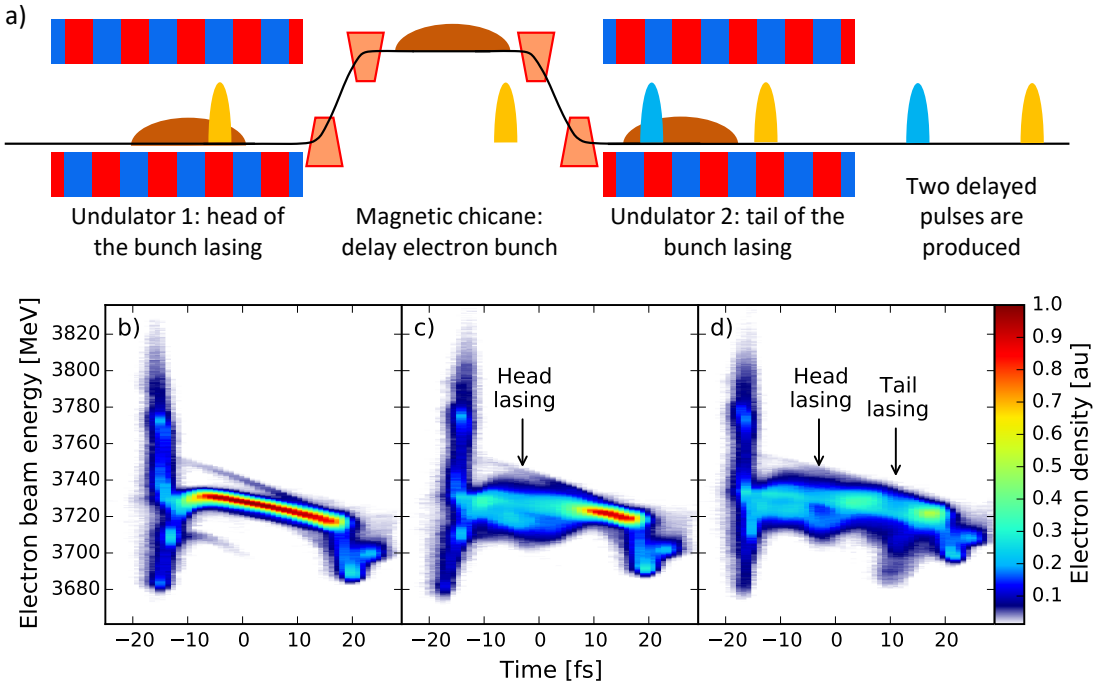


Figure 4.5: (a) Scheme explaining the fresh-slice mode used for this experiment. (b) XTCAV image of the electron bunch without lasing. (c) XTCAV image of the electron bunch after lasing at the head of the bunch in the first undulator. (d) XTCAV image of the electron bunch after also lasing at the tail of the bunch in the second undulator.

bunch lase at the tail in the first undulator and at the head in the second undulator. However, this would cause the first x-ray pulse to overlap with the electron bunch in the second undulator, entangling the generation of both pulses, and preventing us from disabling the generation of the second pulse without affecting the generation of the first one, which was part of our calibration process.

The pulses were focused using a pair of planar-elliptically bent Kirkpatrick-Baez (KB) mirrors [140] into a focal spot with an expected size between 50 and 100 μm^2. Additionally the focusing beam was bounced off a single mirror of the XRSD apparatus (section 3.2.1), which was installed in the beamline as a back-up option for the generation of the delayed pulses. While the XRSD apparatus was not used for introducing any delay, it contributed to a decrease in the overall energy transmission of the beamline.

An oven to deliver the C_{60} sample was installed at the interaction region, and the ions resulting from the ionization were collected and analysed using a time-of-flight magnetic bottle spectrometer (section 3.3.4), under the same conditions as during the prior experiment [121].

4.2.3 Argon calibration

There were two critical calibrations for this experiment, the calibration of the magnetic bottle time-of-flight (MTOF) ion spectrometer and the intensity calibration, both of

them carried out making use of argon delivered to the interaction region using an effusive gas jet. More exhaustive details about other types of calibrations will be discussed in section 5.3.

Ion time-of-flight calibration

The MTOF detector measures the time that it takes for different particles with different masses and different charges to travel between the interaction region and MCP detector situated at a fixed distance after accelerating the particles with a fixed acceleration voltage. For an MTOF with a length L (2 m in this case), and an acceleration voltage V (2000 V in this case), assuming zero initial kinetic energy for the particles, the time of flight t_{TOF} is related to the mass m and the charge of the particle C as,

$$V \cdot C = \frac{1}{2} m \frac{L^2}{t_{\mathrm{TOF}}^2}. \tag{4.4}$$

While the values of L and V are nominally known, in practice it is necessary to treat them as free parameters of the model and fit them to the experimental data. Furthermore the MTOF does not exactly measure times-of-flight, but arrival times for each particle t_{MTOF} with an arbitrary offset. As a consequence, the values measured need to be reference to the ionization time t_0, which is unknown and treated as another free parameter of the model. Taking all this into account and expressing the charge in terms of the electron charge e, $Q = C/e$, and the mass in atomic mass units as a multiple of the mass of the proton m_p $A = m/m_p$, the equation transforms into

$$Q \cdot e \cdot V_0 = \frac{1}{2} A \cdot m_p \frac{L_0^2}{(t_{\mathrm{MTOF}} - t_0)^2}. \tag{4.5}$$

This allows us to obtain the mass-to-charge ratio,

$$A/Q = 2eV_0 \frac{(t_{\mathrm{MTOF}} - t_0)^2}{m_p L_0^2}, \tag{4.6}$$

where t_0, L_0 and V_0 are free parameters of the model which are adjusted to match known ion peaks in argon, and other species in the background as shown in figure 4.6.

Intensity calibration

The purpose of the intensity calibration is to determine what the intensity distribution at the interaction region is. While this value could be in theory obtained taking into account the geometry of the experiment, the reflectivity of the different elements and the pulse duration, in practice there are many unknowns in the exact focal size and reflectivity of the different elements. Alternatively is possible to use a non-linear signal to estimate the intensity.

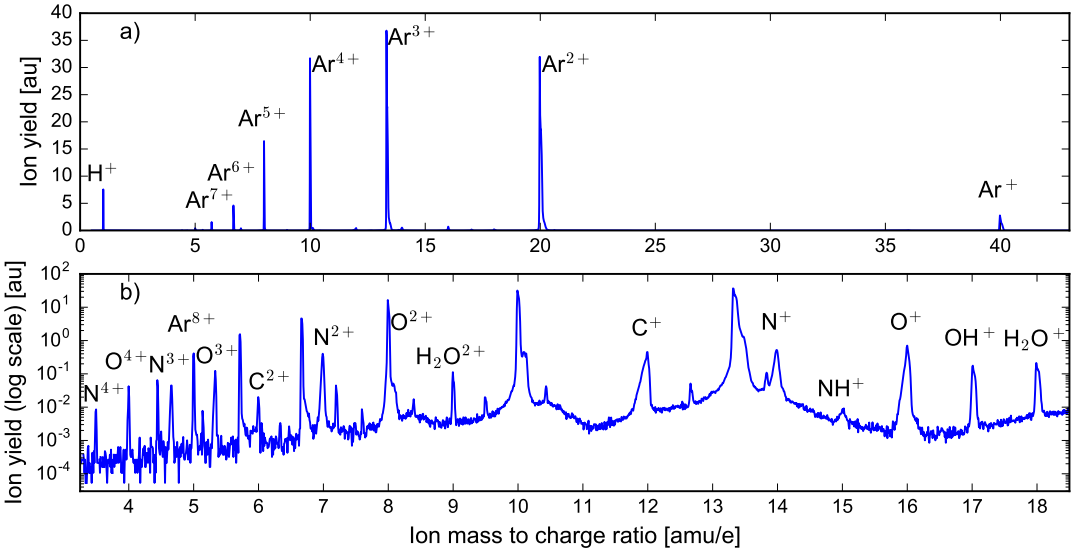


Figure 4.6: Ion time-of-flight calibration using Argon. (a) The main ion peaks correspond (mass-to-charge ratio in brackets, from right to left) to $Ar^{+}(40)$, $Ar^{2+}(20)$, $Ar^{3+}(13.3)$, $Ar^{4+}(10)$, $Ar^{5+}(8)$, $Ar^{6+}(6.66)$, $Ar^{7+}(5.71)$ and $H^{+}(1)$. (b) Using a logarithmic scale we can observe many other species mainly from the background useful for calibration including $H_2O^{+}(18)$, $OH^{+}(17)$, $O^{+}(16)$, $NH^{+}(15)$, $N^{+}(14)$, $C^{+}(12)$, $H_2O^{2+}(9)$, $O^{2+}(8)$, $N^{2+}(7)$, $C^{2+}(6)$, $O^{3+}(5.33)$, $Ar^{8+}(5)$, $N^{3+}(4.66)$, $O^{4+}(4)$ and $N^{4+}(3.5)$.

In this particular case we used multi-photon ionization in argon, yielding highly charged argon ions as our non-linear process. To illustrate that this process is non-linear, let us call Γ_{core} the core-ionization cross section for argon and F the fluence of the pulse in terms of photons per unit of area. Ignoring all other channels, the probability of an ion being core-ionized to produce an Ar^{2+} ion (the photoionization is followed by an Auger emission process) $P(Ar^{2+})$ can be calculated as

$$P(Ar^{2+}) = \Gamma_{\text{core}} F, \tag{4.7}$$

which is linear in the fluence. Starting from there, and approximating the core-ionization cross section from the Ar^{2+} ion to be equal to Γ_{core}, the probability of generating an Ar^{4+} ion can be expressed in terms of the conditional probability of having a second photo ionization from the Ar^{2+} ion and the probability of having an Ar^{2+} ion as

$$P(Ar^{4+}) = P(Ar^{4+}|Ar^{2+})P(Ar^{2+}) = \Gamma_{\text{core}} F \cdot P(Ar^{2+}) = \Gamma_{\text{core}}^2 F^2, \tag{4.8}$$

which is quadratic in the fluence, and so on for more highly charged argon ions.

It is worth noting that due to the argon ions being long-lived states this process is non-linear in the fluence and not in the intensity, so technically this is a fluence calibration and not an intensity calibration. The intensity can however be related to the fluence by taking into account the temporal properties of the pulse.

The ion yield for the different argon charged states created through ionization from a single x-ray pulse at different nominal pulse energies are shown in figure 4.7. In this case we only detected argon ions up to the 10th charged state, as opposed to the 14th charge state which was detected for the previous experiment [121]. We believe this difference to be due to an overall decrease of the transmission of the beamline due to degradation of the focusing mirrors plus the additional reflection by the XRSD mirror, combined with the shorter pulse durations used for the second experiment.

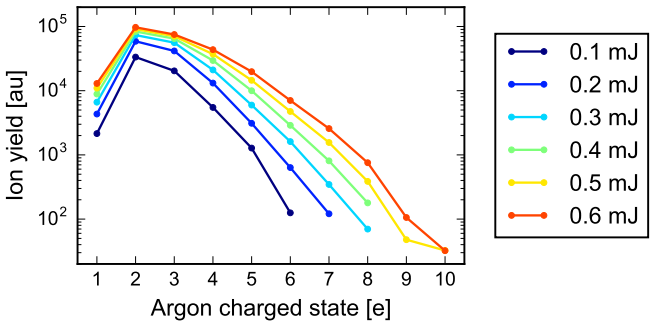


Figure 4.7: Ion yield for different charged states of argon obtained from single-pulse ionization at different nominal x-ray pulse energies measured after the undulator. As the generation of highly charged argon ions is highly non-linear, by increasing the pulse energy by a factor of six the ion yield increases by a few orders of magnitude. This effect is not seen in the low charge ions ($Ar^+(40)$, $Ar^{2+}(20)$).

While the information from figure 4.7 has an arbitrary vertical axis, a Monte-Carlo model simulating the ionization process can be fitted using that data to match the relative ion yield between the different ions. This allowed the determination of the spatial properties and absolute fluence calibration in the previous experiment [121] and it is in development for this experiment.

4.2.4 Time-resolved results

Similar to the increase of the yield of highly-charged carbon ions for longer pulse durations obtained in the previous experiment [121], in this case we were expecting an increased yield for longer delays between the pump pulse and the probe pulse. With the calibration available, measuring this was as simple as collecting ion data at different nominal pump-probe delays set with the magnetic chicane. In particular for this experiment we worked at different values of the delay between 25 and 925 fs, recording 10 minutes of data (∼70000 events per run at 120 Hz) at each time delay. In this case we only detected carbon ions up to the 4th charged state, as opposed to the previous experiment where the 6th charged state was detected, as consequence of the overall lower transmission of the beamline. The results are shown in figure 4.8.

Consistently with the previous results, the ion yield increases very strongly with time delay for the C^{2+}, C^{3+} and C^{4+} ions, and decreases for the C^+ ion. This is consistent

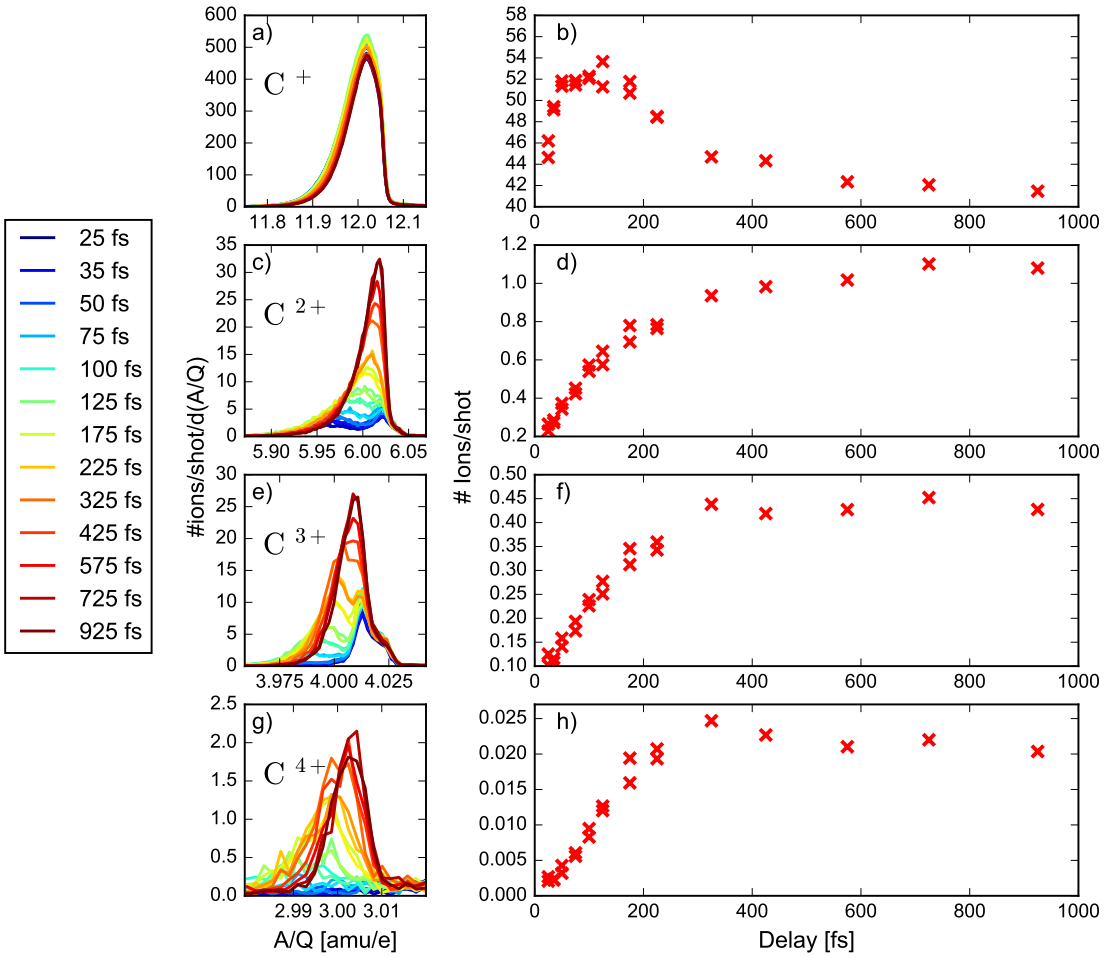


Figure 4.8: Time dynamics appearing in the ion yield of different charged states of carbon obtained after sudden x-ray ionization. (a,c,e,g) Shape of each of the ion peaks for different delays. The shape of the different ion peaks extend around the exact value of the mass to charge ratio due to differences in momentum for the different events, which in the magnetic bottle map into time of flight in a non-trivial way. The small fixed peak at 4.012 in (e) corresponds to O^{4+} contamination. (b,d,f,h) Delay-dependent ion yield for each of the charged states obtained from integrating the figures on the left. The figures show a smooth increase in the ion yield with the time delay for ions C^{2+}, C^{3+} and C^{4+}, with a sudden increase followed by a smooth decrease for ion C^{1+}. All events in the experimental runs are used. The error bars obtained from the statistical analysis of the number of events are negligible, however significant systematic errors exist, as the points taken at redundant values of the delay show small differences.

with the expected behaviour as explained in section 4.2.1. This is also consistent with the results from the previous experiment [121], with a small change in the pivoting point dividing the high-charge ions and the low-charge ions range (For the previous experiment ions between C^{3+} and C^{6+} increased, C^{1+} decreased and C^{2+} stayed constant) due to the decrease in fluence.

It is worth noting that the C^{4+} ion and the C^{3+} (ignoring the oxygen background) change from a regime where they are virtually non-existent, to a regime where they are produced very strongly, most of this happening within the first 200 fs (figures 4.8e-h). There also seems to be an unexpected very fast initial increase of the C^{1+} ion yield (figure 4.8b) which we will discuss in the next section.

4.2.5 Checking the reliability of the results

If there is a lesson that anyone should learn from this thesis, it is that when doing data analysis from XFEL experiments one must be very self-critical with the results found or the dynamics measured. XFELs are very complicated machines which can present very complex correlations between different parameters (see section 5.6.1 and section 7.1) and in particular with the delay. As mentioned before, there is an unexpected initial increase in the C^{1+} ion yield (figure 4.8b). This may be a real effect or an artifact due to some of this correlations. In this section we will introduce some validations that can be performed to discriminate between real effects and fake effects.

Total x-ray energy correlation with the delay

Any LCLS experiment typically has access to the total x-ray energy information from the gas detector (section 3.3.1) for every single shot. Taking this information for all events recorded at each of the delays, and plotting the mean value as function of the delay we obtain figure 4.9a. This figure resembles very strongly the shape of the dynamics measured in figure 4.8b, and in particular the fast rising time. In fact a change in average x-ray energy would precisely yield an overall larger ion yield, independently of any real dynamics.

In order to remove this effect from the data, we can apply a filter to the events included in the time-dependent analysis. Looking at the distribution of the total x-ray energy values at each time delay (figure 4.9b) we can observe that even though the mean value changes significantly, the distributions have a great overlap. Choosing an fixed energy region, in this case between 0.76 and 0.78 mJ, we can filter out all events outside the region, and plot the mean x-ray energy value to verify that the correlation between the x-ray energy and the time-delay disappears (figure 4.9c).

Applying the same filter to the time-dependent data, we obtain the results from figure 4.10, which essentially show the same behaviour as before except for the initial increase of the C^{1+} ion yield, which has completely disappeared. It is worth noting that

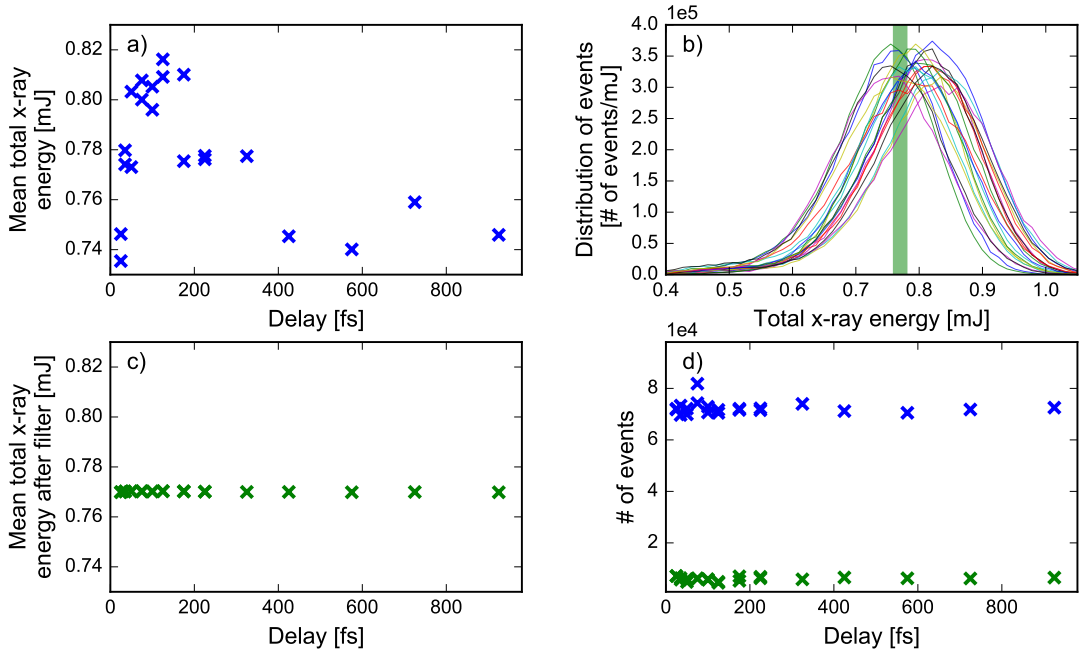


Figure 4.9: (a) Mean value of the total x-ray energy for all events contributing to each of the time delay points used for figure 4.8. (b) Distribution of the total x-ray energy values for the different runs at different time delays. (c) Same as (a) after applying a filter to keep only the events with pulse energies between 0.76 and 0.78 mJ, highlighted in green in (b), resulting in the dynamics shown in figure 4.10. (d) Event statistics for each time delay before (blue) and after (green) applying the filter.

the procedure followed takes a toll on the event statistics, decreasing them by a factor of 10 in this case (figure 4.9d). This approach is a simple example of the binning technique, which will be discussed in more detail in section 5.6.1.

It is worth noting that by filtering in total x-ray energy, we are not removing any effect due to a potential change in the energy balance between both pulses. While this was not straight forward to achieve due to the lack of individualized pulse energy diagnostics, we believe this is not a problem mainly because of two reasons. First, the remaining dynamics are qualitatively consistent with the results obtained in the previous experiment [121]. Second, the mode used for generating the pulses makes the generation of the second pulse completely independent of the first pulse. We often monitored the energy of the first pulse alone by suppressing the second pulse, to verify that its emission was very stable, suggesting that it was the second pulse causing most of the jitter and drift in the total x-ray energy. Furthermore we developed a machine learning model to predict independent pump and pulse energies, using single-pulse data for training, that allowed filtering on pump and probe pulse energies independently to some extent, still obtaining the same dynamics as a result. More information about this type of processing using machine learning will be discussed in chapter 7.

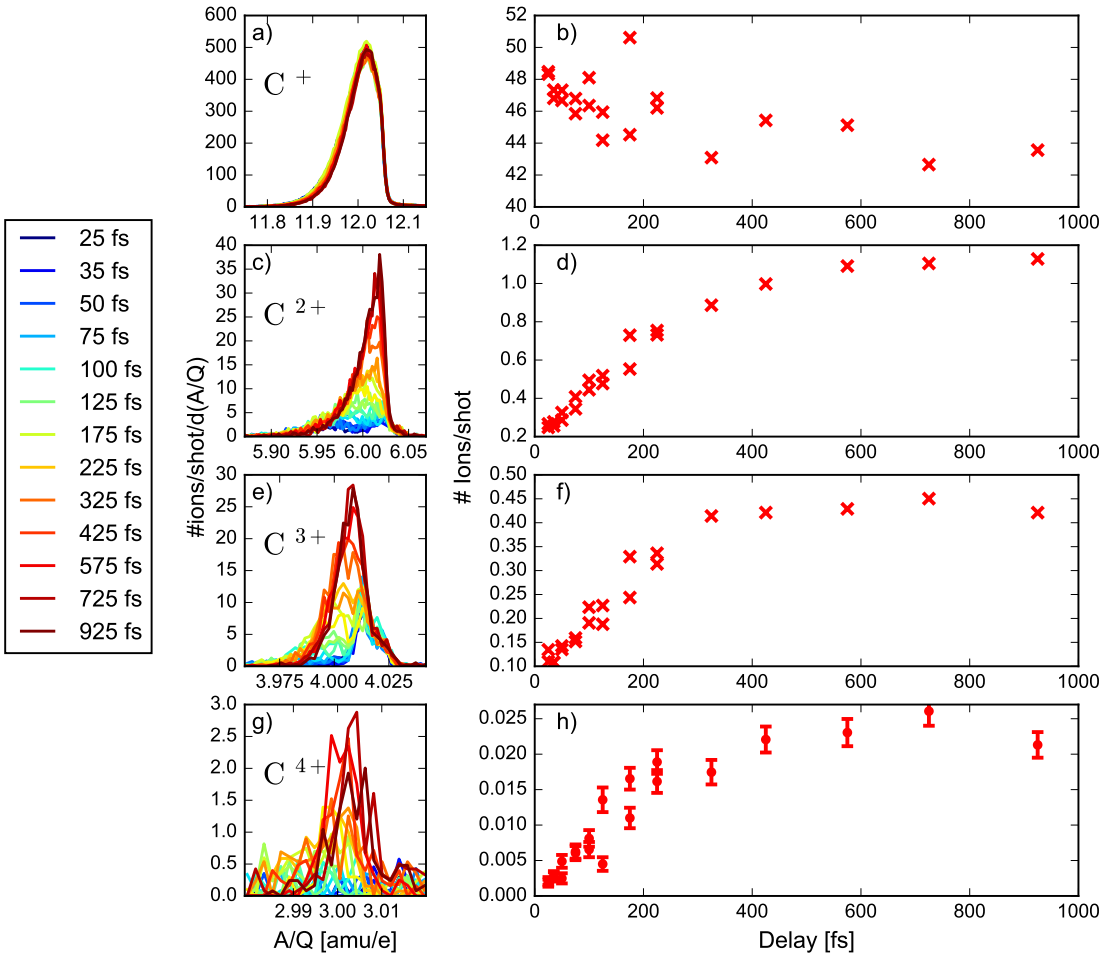


Figure 4.10: Analogous results to figure 4.8 after applying a filter to keep only the events with total pulse energy values between 0.76 mJ and 0.78 mJ. The filter removes the effects of the fluctuations, showing that in reality ion C^{1+} only presents a small decrease with time delay. The filtering process, however, reduce the number of events available, making the statistical error bars become significant in the time dynamics of C^{4+}.

Spatial overlap correlation with the delay

Another usual problem of multipulse XFEL experiments is that typically there are not any diagnostics installed to quantify the spatial overlap between the pulses. In this case we can again make use of the non-linear multiphoton ionization in argon to gain some insights of the spatial overlap, and even model it. The idea behind this is simple: two pulses with perfect spatial overlap should behave in the same way as a single pulse with the same total x-ray energy when ionizing argon.

Using the experimental setup at the minimum delay (25 fs) we recorded some argon ionization data separately for a single pulse and for two pulses. Under these conditions, the XFEL pulse energy fluctuations generated some single-pulse events more energetic in terms of total x-ray energy than some of the double-pulse events. This allowed us to

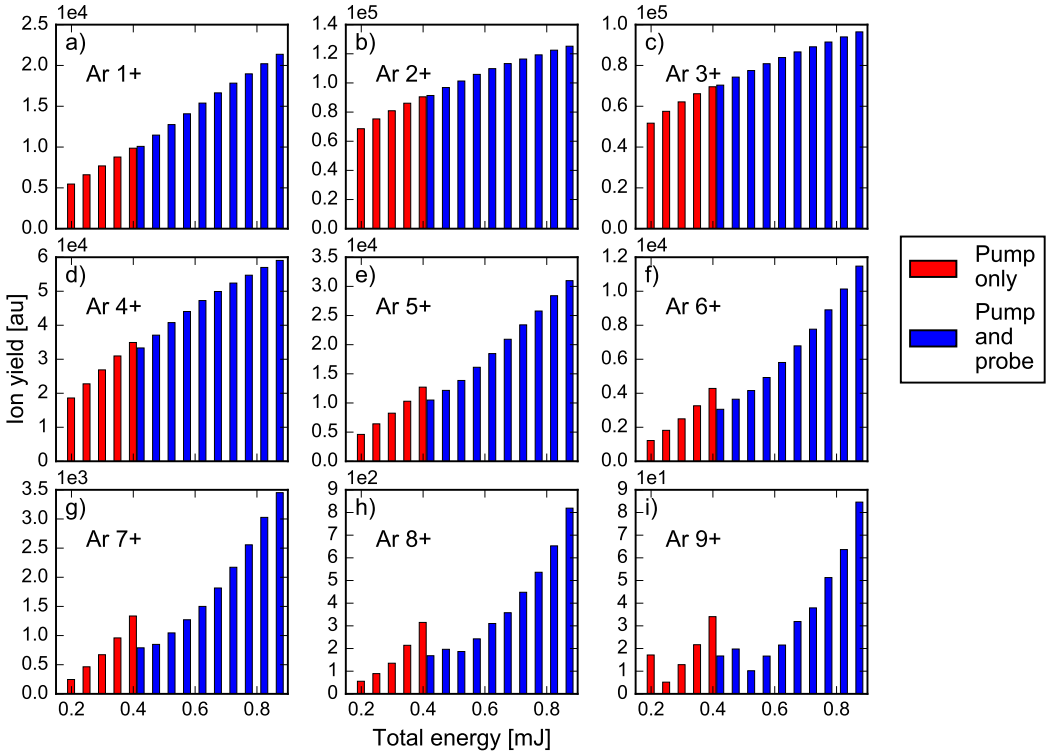


Figure 4.11: Ion yield for different charged states of argon for different nominal total x-ray energies. Experimental data obtained in a single-pulse mode (red) and in a double-pulse mode (blue) are shown together to estimate the degree of spatial overlap. For a perfect spatial overlap, the ion yield obtained for two pulses with a given total x-ray pulse energy should be exactly the same as the ion yield obtained for a single pulse with the same total pulse energy. As the charge state increases, so does the non-linearity of the generation process, amplifying the differences between the double-pulse generation and the single-pulse generation. This evidenced a reasonable, but imperfect, spatial overlap that need to be quantified using theoretical modelling.

sort the events as a function of the total x-ray energy, plotting the yield of each of the argon ions as function of the total x-ray energy for the single-pulse and the double-pulse cases alongside (figure 4.11).

In that figure, we can observe a perfect smooth transition between the single-pulse (red) and the double-pulse (blue) data for the low charge ions (figures 4.11a-b). While this does not give any information about the spatial overlap, since as discussed in section 4.2.3 these ions are mainly generated from single photon processes, it shows consistency in the method. On the other hand as the ion charge increases the transition becomes less smooth. In the case of Ar^{4+} (figure 4.11d), which is already generated through multiple ionization processes, we still obtain a relatively smooth transition, as a sign of having spatial overlap. However, in the case of the ion Ar^{8+} (figure 4.11h), a single pulse of 0.4 mJ can yield as many ions as a double pulse with a total energy of 0.65 mJ, manifesting an imperfect spatial overlap. Making use of this data qualitatively, it would be possible to model the degree of spatial overlap between the pulses.

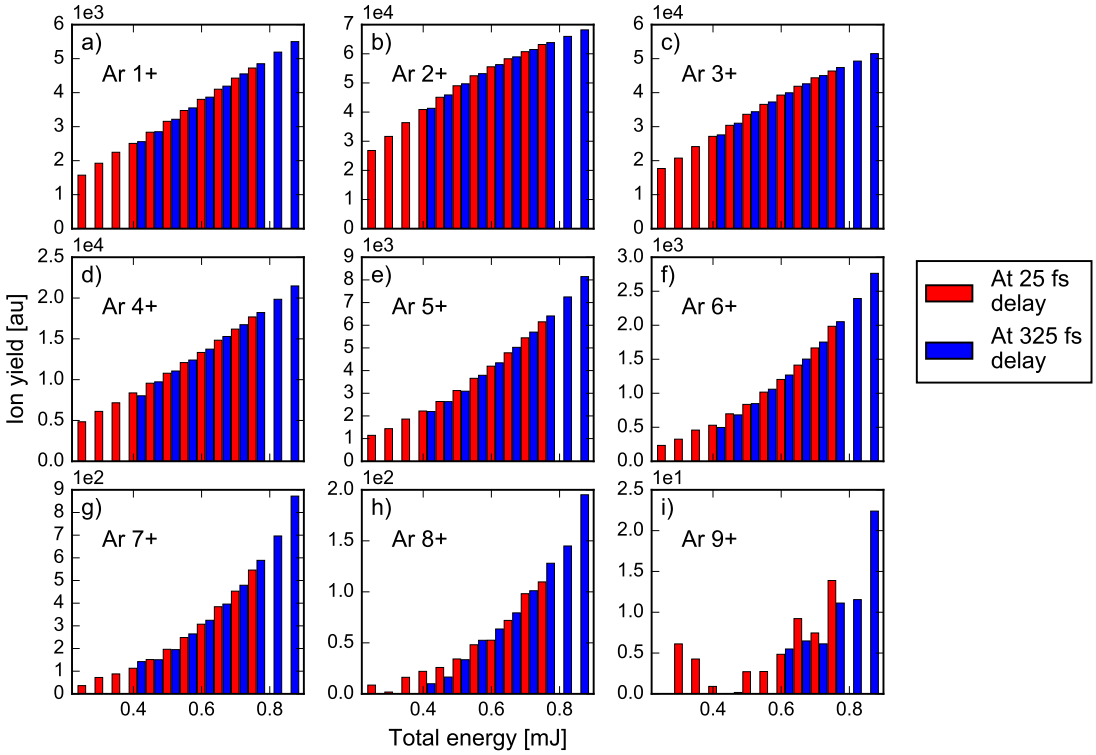


Figure 4.12: Ion yield for different charged states of argon for different nominal total x-ray energies with two pulses. Experimental data obtained at 25 fs (red) and at 325 fs (blue) of delay are shown together to eliminate the possibility of a variation of the spatial overlap with the time delay.

Additionally, since the pulses are generated at different undulators and the exact trajectory of the electron bunch is changed for different delays by the magnetic chicane, some of the dynamics we observe in figure 4.10 could be a result of a systematic change of the spatial overlap as function of the delay. To discard this potential problem we followed a similar strategy comparing the argon ion yield at two different delays, in particular at 25 fs (red) and 325 fs (blue) (figure 4.12). In this case the data from the different delays appears perfectly interleaved, as a sign that the spatial overlap stays stable as function of the delay.

4.2.6 Conclusions and outlook

After unmasking the time-resolved results shown in figure 4.8 to separate the influence of the x-ray fluctuations from the true dynamics extracted in figure 4.10, we still observe statistically significant time-dependent dynamics. In particular we observe a significant increase in the ion yield of the C^{2+}, C^{3+} and C^{4+} ions, and a slight decrease of the C^{+} ion within the first 400 fs. Based on preliminary theoretical modelling performed before during the proposal stages of the experiment, this would correspond to the timescale of the Coulomb explosion, as explained in section 4.2.1.

However, our experimental results agree only qualitatively with that preliminary theoretical modelling. First, the theoretical model predicted the dynamics to happen within the first 100 fs. And second, an increase in ion yield of the C^{3+}, C^{4+} and C^{5+} ions was expected, with a decrease in the yield for the C^{+} and C^{2+} ions. We believe that these inconsistencies are due to differences between the x-ray parameters used for the preliminary model and the actual x-ray parameters available during the experiment. We know, for example, that the preliminary model made use of 4 fs pulses at 485 eV, while the experiment was performed with 5-10 fs pulses tuned at 640 eV. Furthermore, the exact pulse energy at the interaction region of the experiment is not well known, but the fact that the generation of C^{5+} ions was expected and not observed, and similarly for higher-charge argon ions, indicates that the intensity assumed in the theoretical modelling was significantly higher than the intensity used in the experiment. While we are currently waiting for updated theoretical results before we can draw more conclusions and publish the experimental results [15], qualitatively we expect the effect of the lower pulse energies to shift the expected dynamics of the ion yield towards lower charge states as observed in the experiment. This should also cause an overall decrease in the total charge accumulated in the C_{60} molecules, slowing down the dynamics of the explosion and making it more consistent with the experimentally observed results.

Apart from the importance of the experimental results in terms of the physical mechanism, this experiment is also the first example in this thesis where XFEL fluctuations influenced an experimental time-resolved signal. In this case this influence was only partial, allowing us to still observe some of the real dynamics. We will see in later chapters that, as we attempt to approach few-fs timescales, the effect of these fluctuations can completely dominate the signal, requiring much more sophisticated x-ray characterization and data analysis techniques to separate the real signals from the large background of fluctuations.

Chapter 5

Performing a few-fs time-resolved experiment

As we have seen in chapter 4, performing a successful time-resolved experiment on the 10-100 fs timescales can be relatively straightforward in terms of data collection and data analysis. Taking time-resolved XFEL experiments to the few-fs timescales however presents different challenges.

The first challenge is that current XFEL technology is stable to the few 100-fs regimes in terms of timing between optical and x-ray pulses, and to 10-20 fs in terms of timing between pairs of x-ray pulses. Furthermore, generating few-fs pulses usually takes a toll in terms of the stability of the pulse energy and the photon energy, as often the lasing process does not saturate, and the x-ray pulses consist of a single SASE spike, making it much more sensitive to the stochastic SASE fluctuations.

Another challenge is that, while 100-fs dynamics can be studied using ion spectroscopy, as there is enough time for nuclear rearrangement, this is not typically the case for few-fs dynamics, where the electrons play a more important role. Electron spectroscopy is however more challenging than ion spectroscopy as electrons are more difficult to detect, fly much faster and, unlike ions, present a single fixed value of the mass and charge, typically leading to scenarios where the electronic signal in the background is much larger than the time-resolved electron signal.

In this chapter the author will describe how to perform a typical few-fs experiment using data from a real experiment aiming to measure electron-correlation-driven ultrafast dynamics. This will include learning how to deal with the problem of the fluctuations and sometimes even turn them to our favour by performing a full single-shot characterization of the x-ray pulses, and discussing some strategies to find small electronic signals on top of strong fluctuating backgrounds.

5.1 An example of few-fs time dynamics in isopropanol

In this section we will discuss the possibility of measuring inner-shell few-fs hole dynamics after photoionization in isopropanol (C_3H_8O, figure 5.1a) using an XFEL. Similar dynamics have been shown to also appear in glycine and are discussed in detail in reference [100]. The theoretical modelling of isopropanol was performed by Dr. Bridgette Cooper.

As we mentioned in section 2.1.2, the full molecular wavefunctions can be expressed to first approximation as a combination of multiple molecular orbitals $|\phi_i\rangle$ where the different electrons are located. Similarly, it should be possible to express the molecular wavefunction of a cation in terms of *one-hole* (1h) states describing an electron missing from any one of those orbitals $|1h_i\rangle$. However, these one-hole states are not necessarily eigenstates of the Hamiltonian of the cation, and more than one may need to be combined to create an eigenstate. Furthermore, the result of removing one electron from a molecule does not always have to lead to a single hole, but it can also lead to *two-hole-one-particle* (2h1p) configurations $|2h_{ij}1p_a\rangle$, where one of the remaining electrons is excited, creating a second hole. Using these two kinds of contributions, in what is called second order algebraic-diagrammatic construction (ADC(2)) [141], the eigenstates of the cation can be expressed as

$$|\Phi\rangle = \sum_i c_i \,|1h_i\rangle + \sum_{i,j,a} b_{ij}^a \,|2h_{ij}1p_a\rangle \,, \tag{5.1}$$

where c_i and b_{ij}^a are the calculated coefficients, satisfying $\sum_i |c_i|^2 + \sum_{i,j,a} |b_{ij}^a|^2 = 1$.

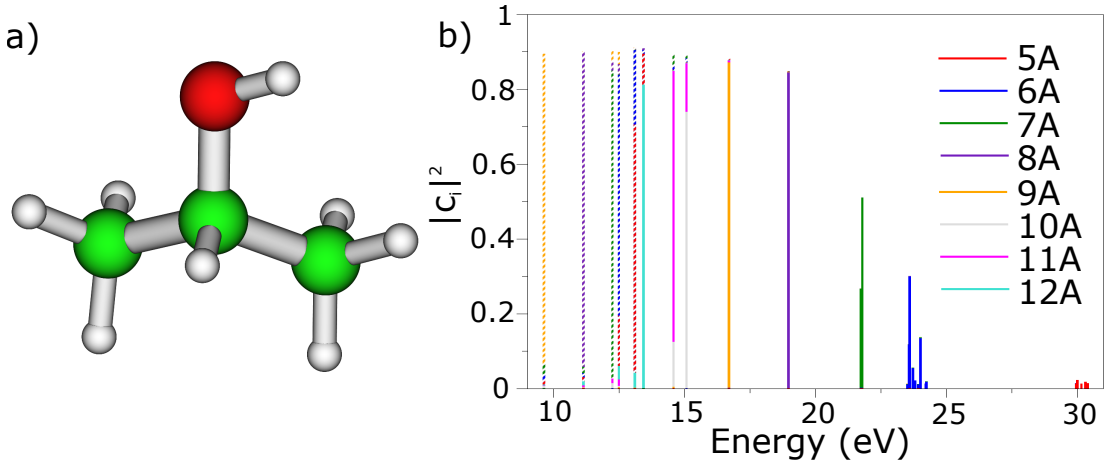


Figure 5.1: Eigenstates of the Hamiltonian of the isopropanol cation. (a) *Gauche* conformer (lowest energy) of isopropanol. (a) Decomposition of the eigenstates of the Hamiltonian expressed in terms of one-hole states (5A, 6A, 7A,...) at different ionization energies. Original figure by Dr. Bridgette Cooper, and modified by the author with permission.

In the particular case of the *gauche* conformer of neutral isopropanol, the 34 electrons are distributed in the following molecular orbitals:

$$(1s^2)_{\mathrm{O}}(1s^2)_{\mathrm{C_1}}(1s^2)_{\mathrm{C_2}}(1s^2)_{\mathrm{C_3}}5a^2 6a^2 7a^2 8a^2 9a^2 10a^2 11a^2 12a^2 13a^2 14a^2 15a^2 16a^2 17a^2$$

where the s orbitals correspond to the cores of the single oxygen and the three carbon atoms and the a orbitals are the valence molecular orbitals. Similarly, the one-hole states of the cation resulting from removing one electron from the $5a$, $6a$, $7a$, ... $16a$, $17a$ molecular orbitals are denoted as 5A, 6A, 7A, ... 16A, 17A states, respectively.

Figure 5.1b shows a diagram with the eigenstates of the isopropanol cation in this formalism. The vertical bars situated at different ionization energies correspond to the different eigenstates. Different colours correspond to different 1h states accessible at different ionization energies, with dashed and solid lines used for states resulting from removing an electron from the outer (17A, 16A, 15A, etc.) and inner (5A, 6A, 7A, etc.) valence respectively. Each of the vertical bars corresponding to the eigenstates is made from 1h state contributions, coloured proportionally to the $|c_i|^2$ coefficients, and spans from 0 to 1, with the remaining uncoloured region corresponding to the contribution from 2h1p states with coefficients $|b^a_{ij}|^2$. Typically, states with lower ionization energies (outer valence ionization) are dominated by single hole contributions, while states with higher ionization energies (inner valence ionization) have stronger 2h1p contributions.

Each of those eigenstates of the Hamiltonian of the cation are stable in time, however, an isopropanol molecular cation is not necessarily generated in an eigenstate. For example, x-ray photoionization with an XFEL, removes an electron from a particular molecular orbital, implying that at time 0, the molecular wavefunction $|\Psi\rangle$ must be similar to a 1h wavefunction:

$$|\Psi(t=0)\rangle \sim |1\mathrm{h}_i\rangle \tag{5.2}$$

However, the time evolution of this function has to be expressed in terms of the eigenstates of the Hamiltonian. For example, let us assume that there are two eigenstates with energies E_1 and E_2 with contributions from the same one-hole $|1\mathrm{h}\rangle$ state and other two-hole-one-particle state $|2\mathrm{h1p}\rangle$,

$$|\Phi_1(E=E_1)\rangle = \frac{1}{\sqrt{2}}|1\mathrm{h}\rangle + \frac{1}{\sqrt{2}}|2\mathrm{h1p}\rangle \tag{5.3}$$

$$|\Phi_2(E=E_2)\rangle = \frac{1}{\sqrt{2}}|1\mathrm{h}\rangle - \frac{1}{\sqrt{2}}|2\mathrm{h1p}\rangle\,, \tag{5.4}$$

and that the total wavefunction at time 0 corresponds exactly to $|1\mathrm{h}\rangle$,

$$|\Psi(t=0)\rangle = \frac{1}{\sqrt{2}}|\Phi_1\rangle + \frac{1}{\sqrt{2}}|\Phi_2\rangle = |1\mathrm{h}\rangle\,. \tag{5.5}$$

In this case the time evolution of the wavefunction would be

$$|\Psi(t)\rangle = \frac{e^{-\frac{\imath E_1 t}{\hbar}}}{\sqrt{2}}|\Phi_1\rangle + \frac{e^{-\frac{\imath E_2 t}{\hbar}}}{\sqrt{2}}|\Phi_2\rangle \equiv \frac{1}{\sqrt{2}}|\Phi_1\rangle + \frac{e^{-\frac{\imath (E_2-E_1) t}{\hbar}}}{\sqrt{2}}|\Phi_2\rangle \neq |1\mathrm{h}\rangle\,, \tag{5.6}$$

which does not preserve the initial one-hole state $|1\mathrm{h}\rangle$ in time.

In order to monitor the presence of the hole, we can define the hole survival probability as $S(t) = |\langle 1\mathrm{h}|\Psi(t)\rangle|^2$, which for this example is

$$S(t) = \left|\frac{1 + e^{-\frac{\imath (E_2-E_1) t}{\hbar}}}{2}\right|^2 = \frac{1}{2} + \frac{1}{2}\cos\left(\frac{(E_2 - E_1)t}{\hbar}\right), \tag{5.7}$$

starting at exactly 1, and presenting periodic survival revivals caused by electron-correlation dynamics.

Taking into account the actual eigenstates of the isopropanol cation, different one-hole states in the inner valence present different types of dynamics (figure 5.2). As there are more states involved in this realistic case, the dynamics are more complex, however, they can still be interpreted intuitively in terms of the eigenstates shown in figure 5.1b. The 8A orbital (purple), corresponding to an almost pure 1h eigenstate, shows a very stable behaviour. The 7A orbital (green), which contributes to two eigenstates very close in energy, shows the beginning of what we believe would evolve as a low-frequency oscillation. Similarly, the 6A orbital (blue) shows faster oscillations, as it mainly contributes to two eigenstates with a larger separation in energy. Finally, the 5A orbital (red) contributes in equal proportion, but very weakly, to many eigenstates, having as a consequence a very quick dephasing of the different contributions from the different eigenstates, resulting in quick drop of the survival probability to 0. Additionally, the 5A state is accessed for energies above the double-ionization threshold (∼25-30 eV), and would typically lead to the fast emission of an Auger electron. It is worth noting that these calculations are performed without taking into account any nuclear motion or ensemble effects that may be present in a real experiment. It is expected theoretically that these contributions would add additional de-phasing effects starting within the first 10-fs of the dynamics, as early as 3 fs after the ionization in some cases [142, 143].

Measuring this type of electron correlation dynamics requires first inducing an inner valence photoionization and then probing the survival probability of the hole. Our proposed experimental scheme making use of two x-ray pulses to measure the dynamics of the 6A hole is shown in figure 5.3. First, an x-ray pump pulse at with central photon energy at 498.8 eV photoionizes an electron from the inner-valence shell of the isopropanol molecule creating a hole in the 6A orbital. Then a resonant probe pulse tuned to 513.8 eV excites an electron from the oxygen 1s shell into the valence hole, leaving the molecule in a excited Auger-active state. As shown in section 2.2.2 the probability of this excitation depends on the time evolution of the state (the one-hole state

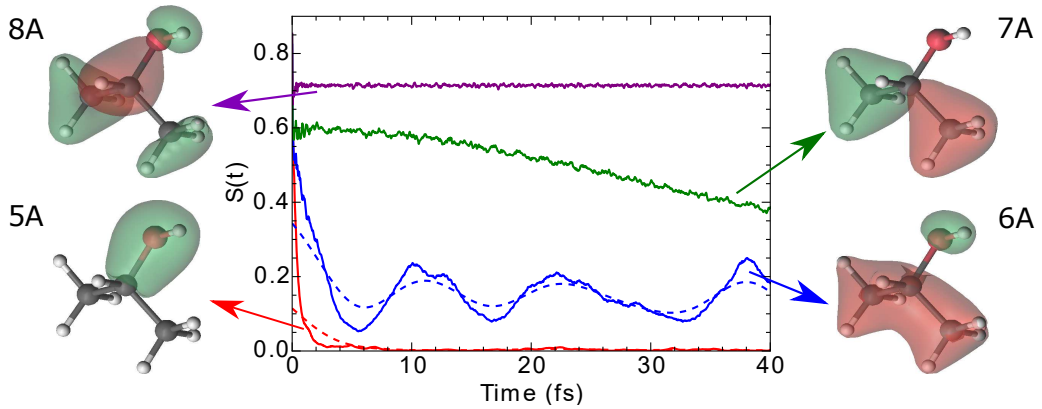


Figure 5.2: Hole survival probability after photoionization of the one-hole states corresponding to 5A, 6A, 7A and 8A molecular orbitals in isopropanol (hole density distribution also shown in figure). The dynamics of the two lower energy isopropanol conformers have been averaged together. The dashed line shows the dynamics convolved with a 5 fs full-width half-maximum (FWHM) Gaussian, accounting for pulse duration and timing uncertainty in a possible XFEL experiment. Theoretical data and molecular orbitals provided by Dr. Bridgette Cooper.

in this case) and is proportional to the survival probability of the hole. Finally this state decays through Auger emission, filling the core hole with a valence electron and emitting a characteristic Auger electron in the process (figure 5.4c). By monitoring the Auger-electron yield as a function of the time delay between the pulses, it is possible to measure the absorption of the probe pulse and, consequently, the existence of the hole.

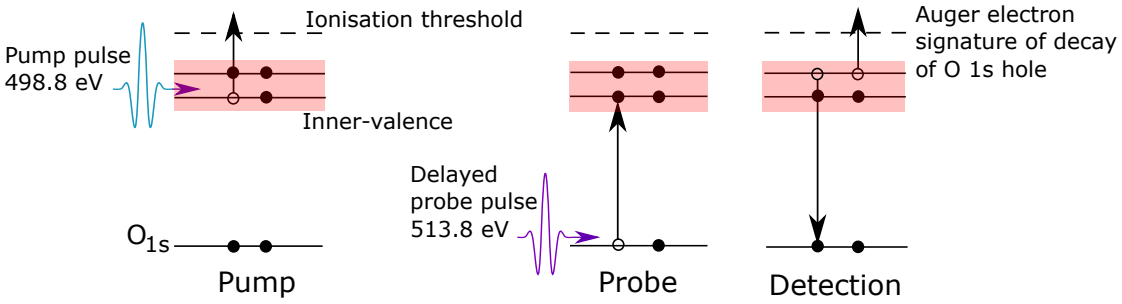


Figure 5.3: Proposed pump-probe scheme for measuring the dynamics of a hole opening process in isopropanol. After the pump pulse ionizes an electron from the inner valence, the probe pulse can excite an electron from the oxygen core level into the inner valence hole. This leaves the system in a highly excited state leading to the emission of an Auger electron. The probe mechanism, however, can only take place when the inner-valence hole exists, correlating the emission of Auger electrons to the survival probability of the hole.

One of the main limitations of this mechanism is that the pump pulse is not selective in the ionization, and it may ionize any electron from the multiple valence states. In fact, taking into account the cross sections for the different valence states, only a fraction of 0.15 of the total valence ionizations occur from the state of interest. This could be solved by using an XUV pulse tuned at 24.2 eV to produce "just-above-threshold" ionization, which would present a much larger cross section for the state of interest. However,

current XFELs do not provide the required tools to pump the experiments with XUV pulses, and are in general much more limited in terms of time resolution (a few ten of fs) when using optical-based pump sources. Considering that only 0.15 of the ionized molecules would be in the right state, the intensity must be carefully controlled to work in a regime where only a small fraction of the molecules are ionized, avoiding double ionization, while at the same time the total number of molecules prepared in the right state is large enough for signal-to-noise purposes.

As a result of the pump ionization from multiple states, a full photoelectron background signal is produced (figure 5.4a) partially overlapping with the Auger electron signal (figure 5.4c). This overlap could be avoided by using much lower photon energies for the pump pulse, however, in practice most of the two-colour XFEL schemes cannot produce pulses separated by more than 3% the central photon energy, 15 eV at 500 eV in this case. As the resonant pulse has to be centered at 513.8 eV, the pump pulse cannot be tuned lower than 498.8 eV.

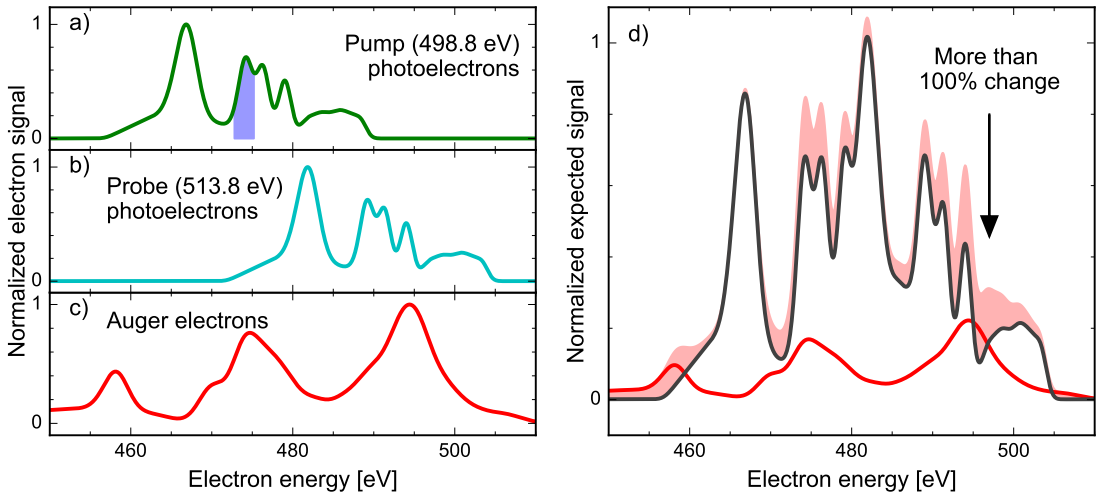


Figure 5.4: Three different electronic signals are measured in the energy range of interest. (a) Photoelectrons produced by valence ionization from the pump, with the photoelectrons ionized from the orbital of interest highlighted in blue. (b) Photoelectrons produced by valence ionization from the probe. Figures in brackets correspond to the required central photon energy for the spectra of the pump and probe pulses respectively. (c) Theoretical Auger spectra produced by the two-photon process presenting the time-dependent dynamics. (d) Expected signal at the detector. The black solid line corresponds to the photoelectron background signal, while the red solid profile and the red highlighted region corresponds to the expected maximum Auger signal assuming a perfect resonant absorption of the probe. The expected signal has been calculated considering a 10% solid angle collection in a direction perpendicular to the polarization. Photoelectron electron data has been extracted from the literature [144]. A description of the Auger calculation methods can be found in reference [12].

In a similar way, the probe pulse may not only induce the desired resonant transition, but it may also ionize from the valence shell. While the cross section for this resonant transition may be several orders of magnitude larger than for the valence ionization,

assuming similar intensities for the pump and the probe pulses, the probe pulse would also produce a similar photoelectron spectra shifted up by 15 eV (figure 5.4b). This background signal presents an almost total overlap with the expected Auger signal, which in this case cannot be avoided by shifting the photon energy due to the resonance condition.

Finally, the amount of emitted Auger electrons cannot be larger than the number of holes generated in the state of interest, which in turn cannot be larger than the number of photoelectrons ionized by the pump from that state. This corresponds to the area of the photoelectron peak highlighted in blue in figure 5.4a. Assuming a perfect efficiency for the probe inducing the resonant transition, the total area of the Auger electron signal cannot be larger than the area of the region highlighted in blue. Luckily, it is possible to perform the electron detection in a direction perpendicular to the x-ray polarisation to enhance the relative amplitude of the Auger electrons (emitted uniformly) with respect to the photoelectrons (emitted parallel to the polarization). Considering a 10% solid angle of collection in that direction, the expected signal is shown in figure 5.4d, with the photoelectron background in black and the scaled Auger signal contribution in red.

Additionally both the pump and probe pulses can produce photoionization from the carbon K-edge, leading to K-edge Auger electron emission. While these photoelectrons and Auger electrons would sit at a completely different energy range, the Auger process may contribute as a different channel producing holes in the 6A orbital, for the double ionized molecule in this case, with different dynamics which could interfere with the main channel. This problem could be solved by using a probe pulse tuned for the carbon K-edge transition instead of the oxygen K-edge transition, however, none of the existing XFELs with few-fs capabilities works well for photon energies as low as the carbon K-edge ($\sim$ 300 eV).

The expected signal from figure 5.4d shows that it should be realistically possible to perform this type of experiment with an XFEL. Considering the pulse durations of XFEL pulses ($\sim$3 fs) and the timing uncertainty in the pump-probe delays ($\sim$1 fs), the expected time-dependent signal should be similar to the dashed lines presented in figure 5.2.

During the course of this Ph.D. a proposal to perform an experiment to measure similar dynamics in glycine at LCLS was accepted and the experiment was actually attempted. This experiment was originally targeted at the carbon K-edge, however, as mentioned earlier, LCLS capabilities did not achieve lasing at 300 eV, forcing us to switch to the oxygen K-edge. Furthermore, due to several reasons (section 5.6) it was not possible to record time-resolved data in glycine, and an alternative sample (O_2) had to be used. We will see at the end of this chapter (section 5.7) that, all things considered, isopropanol is a much better sample than glycine to perform this type of experiment at current XFELs.

5.2 Experimental setup

In order to perform a successful experiment similar to the one proposed in the previous section, it is critical to design a versatile setup and to consider thoroughly all possible scenarios and backup plans achievable within a short amount of time, typically 60 hours. The aim in this case was to create a setup with the capabilities to perform a time-resolved experiment scanning the delay from about -10 fs to 30 fs with two few-fs x-ray pulses near the oxygen K-edge (∼540 eV), using glycine and gas targets as samples, and using electron spectroscopy as the main detection channel.

In order to achieve that, a collaboration was formed with different members focusing on different tasks, from advising on different XFEL modes and supervising the x-ray generation, to contributing with the different equipment, or optimizing the different spectrometers and sample delivery.

5.2.1 X-ray source

From the point of view of the generation of the x-ray pulses, the different compatible approaches listed in section 3.2 were considered. In the end, the double-slotted foil with two independent electron bunches (section 3.2.3) approach was used, as it presented the following advantages over the other alternatives. First, the spatial overlap between pulses is guaranteed, as both pulses are generated at the same undulator section. Second, the use of two chicanes allows for independent control of the time delay between the bunches/x-ray pulses and the spacing between the two colours. Third, it is possible to scan short delay intervals and even to achieve complete temporal overlap. Finally, the mode is fully compatible with fast switching to single-bunch operation and with unspoiled operation, yielding energies up to 500 µJ in this case.

There were however a few disadvantages to using the method as well. This first significant disadvantage is that this XFEL mode was going to be used for the first time in an LCLS user experiment, adding a degree of difficulty, especially considering that the essential double-slotted spoiler was only installed and used for the first time two days before the beginning of the experiment. Also, we learned during the experiment that, from all the available methods, this is the one with the highest jitter. We observed the time delay between the pulses fluctuate by more than 10 fs full-width half-maximum (FWHM) shot-to-shot, the "colour" of each of the pulses can jitter by up to 5 eV FWHM, and the pulse energy ranges from 0 to 30 µJ shot-to-shot, with some correlations between all these instabilities.

5.2.2 X-ray diagnostics

The way to approach the problem of the strong jitter inherent to the chosen x-ray generation scheme was to monitor every jittering variable shot-to-shot. Storing that

information shot-to-shot makes it possible to select data only from the shots where the variables have the desired values, discarding the rest. Furthermore, jitter can even work in our favour by scanning some variables like the time delay almost automatically, and by providing data at different energy ranges that can be used for calibration or background subtraction.

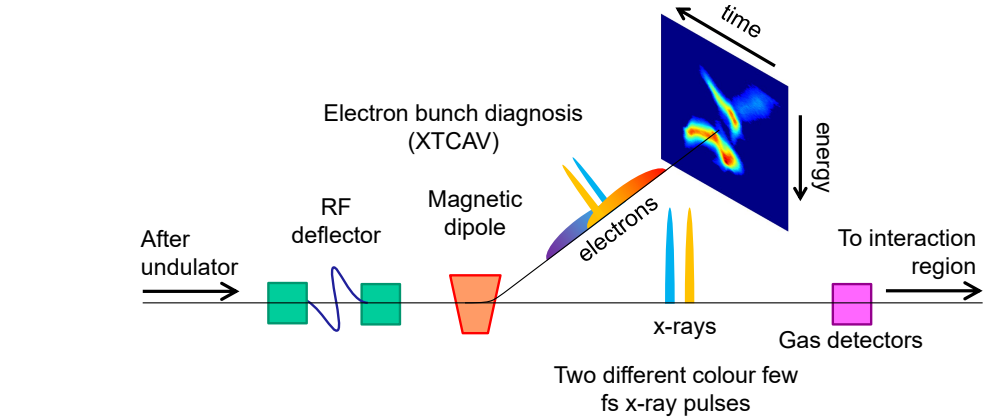
On the downside, the jitter makes it extremely important (and more difficult) to make sure that the recording time is long enough to accumulate the required statistics. To illustrate this, let us assume a process that yields one electron per shot resulting from the two photon process of interest, and that finding this signal on top of other background signals involves measuring a particular feature of the electron energy distribution. As a consequence, within a 100 eV electron spectral range we may be interested in a particular 2 eV region (50 different bins). Assuming we want to measure dynamics within 15 fs, with a resolution of 2.5 fs (6 different bins), and that the photon energy is going to fluctuate by up to 10 eV for each pulse, requiring a resolution of at least 1 eV (10x10 different bins), we are left with a variable space of 4 dimensions with 50x6x10x10=30000 different combinations. Considering a collection efficiency of 10%, that LCLS could only operate up to 60 Hz in this mode at the time, that only 50% of the shots will have enough energy on both pulses to produce the electron and assuming we need at least 10 electrons per bin to overcome the noise levels, the required acquisition time would be 27.7 hours. This shows that while this approach can be used, the fluctuations have to be well characterized to be able to estimate the acquisition time.

Four different types of detectors were used to characterize the x-ray photon energy, pulse energy, and time delay for every shot (figure 5.5).

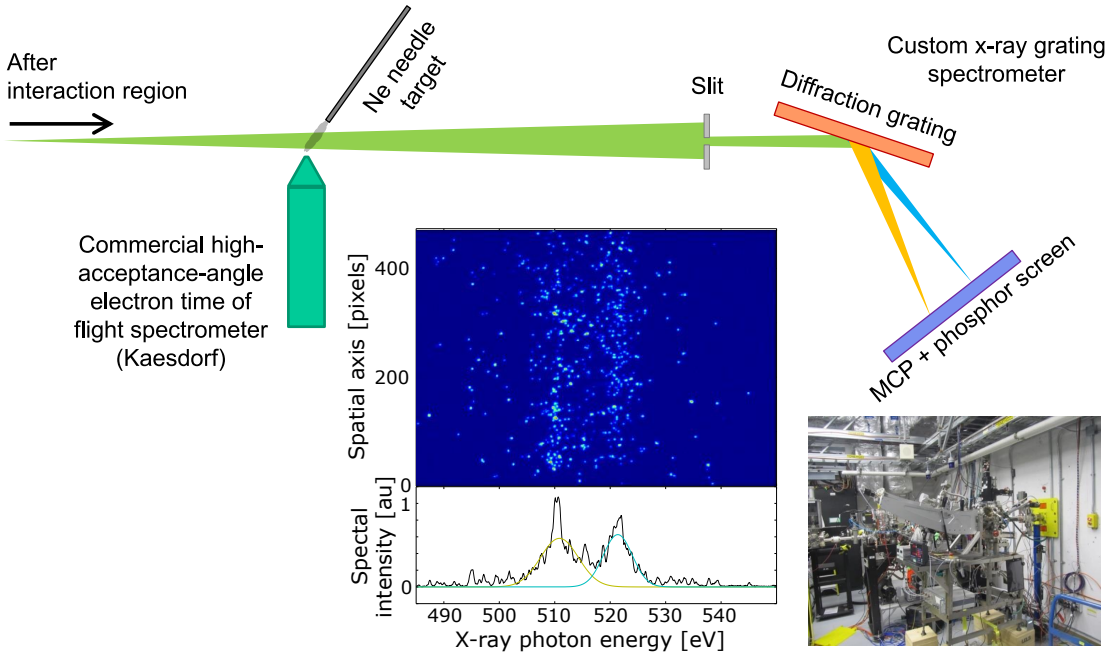
Several gas detectors (section 3.3.1) intersecting the x-ray trajectory before the interaction region (figure 5.5a) were used to measure the total x-ray energy.

An x-ray grating spectrometer (provided by the Molecular and Condensed Matter Physics division, Uppsala University) was installed downstream from the interaction region (figure 5.5b) for single-shot spectral characterization. It mainly consisted of a grazing incidence grating and a grazing incidence MCP with a phosphor-screen detector. As the x-ray absorption in the sample was negligible due to the low sample pressures, this spectrometer allowed the measurement of the shot-to-shot spectrum of the source x-rays. In particular the exact photon energy for the pulses, separated by approximately 20 eV, and the relative amount of energy in each of the two pulses could be obtained, that combined with the gas detector energy values, was used to calculate the absolute pulse energy for each pulse. A picture of the device, as well as a plot with a sample spectrum are shown in figure 5.5b.

Due to the critical importance of knowing the spectral properties of the x-ray pulses, an additional Kaesdorf electron time-of-flight spectrometer was purchased and also installed downstream along with a neon gas sample. This detector was planned to be used as an x-ray spectrometer by measuring the kinetic energy of neon photoelectrons ionized



(a)



(b)

Figure 5.5: Diagram of the x-ray diagnostics used for shot-to-shot characterization of the double pulses. (a) Diagnostics upstream from the interaction region. The electron bunches are separated from the x-ray pulses and analysed using XTCAV, consisting of an RF deflector, a magnetic dipole and a screen. A sample of an XTCAV trace shows the time-dependent energy distribution of the bunches. The total x-ray energy contained in the pulses is measured using gas detectors. (b) Diagnostics downstream from the interaction region. The spectral properties of the x-rays are measured using two different types of spectrometers. First, an electron time-of-flight spectrometer analyses the kinetic energy of photoelectrons ionized from neon. Second, an x-ray grating spectrometer with a slit directly measures the spectral properties. A sample image from the grating spectrometer and a picture of the spectrometer itself are shown.

by the x-ray pulses. However, in the end the detector did not provide enough signal to work on a shot-to-shot basis and was not used.

The shot-to-shot time delay between the pulses was obtained by characterizing the electron bunches using a transverse cavity (section 3.3.8 and reference [127]) after the undulator. Once the electron bunches are separated from the x-rays, they are dispersed and streaked in perpendicular axes, generating a pattern like the one shown in figure 5.5a. As we will show in section 5.4, these patterns containing the time-dependent electron energy can be used to extract the time-resolved x-ray power profile for each bunch, and in particular the time delay. XTCAV is a device commissioned and maintained by LCLS which could operate at up to 60 Hz before June 2016 and at the full 120 Hz after that.

5.2.3 Interaction region

The interaction region was set to accommodate sample delivery and a magnetic bottle electron time-of-flight spectrometer (MTOF). As shown in figure 5.6, the x-ray beam was focused using a pair of KB mirrors [140] into a focal spot of 5 μm^2 situated at the centre of the interaction chamber (LAMP endstation). This focal size corresponds to a nominal value from a previous knife edge measurement performed under similar conditions, as we could not perform the actual measurement within the allocated beamtime.

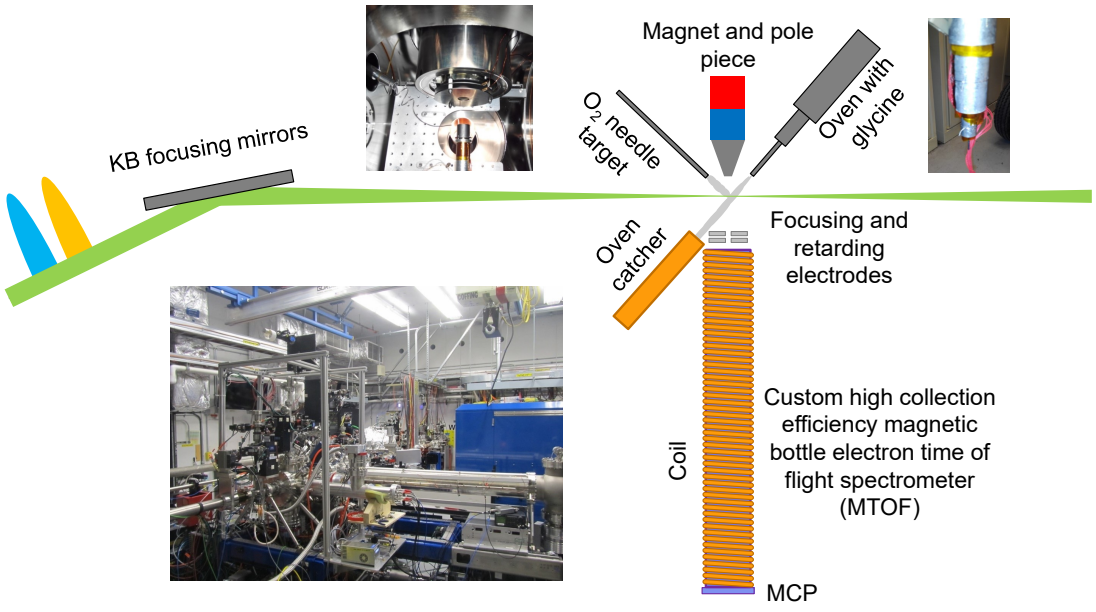


Figure 5.6: Experimental setup at the interaction region. Bottom picture: overall setup with the 2-m-long drift tube of the MTOF. Top left picture: top view of the inside of the interaction where the magnet, the electrodes, the oven catcher and the needle gas-jet can be seen. Top right picture: picture of the outside of the glycine oven.

This interaction chamber consisted of a 6-way cross with several extra ports at different angles to image the interaction region or add additional jets or probes. The whole LAMP endstation was mounted on a 3d motorized translation stage to allow alignment with respect to the x-ray beam.

The MTOF consisted of two different parts. A magnet attached to a conical pole piece was positioned near the interaction region to create a magnetic field of up to 1.3 T near the x-ray focus. This magnet was held by a long rod from one of the side ports also mounted on a motorized 3d translation stage with a bellow, allowing the optimization of the magnet position. A cylindrical electrode and the tip of the magnet were set at a small potential of -0.5 V to repel the electrons in the opposite direction.

On the port opposite to the magnet, a 2m-long drift tube was installed, with an MCP at the end connected to a digitizer to collect the electron signal. Three electrodes were used to focus the electron beam and to apply a retardation potential to the electrons, optimizing the resolution of the device for electrons near 500 eV. The drift tube was surrounded across its whole length by a coil, creating a magnetic field near 1 mT. The gradual change of magnetic field between the magnet and the coil induces a redirection of the trajectories of the electrons towards the detector for a large range of emission angles, drastically increasing the collection efficiency as explained in reference [117].

The gas sample delivery consisted of a needle gas-jet with a leak valve connected through one of the extra ports using also a 3d motorized positioning system with a bellow. During the experiment Ne (calibration) and O_2 (calibration and experiment) were delivered using this gas-jet, with the expected density in the order of about 10^{13} particles/cm^3.

Due to the solid character of the glycine sample, an oven was used to sublimate part of the sample and deliver it to the interaction region from the top port, using again a 3d motorized positioning system with a bellows. The oven consisted of a small cylindrical container loaded with glycine powder connected to a small tube ($\sim$1 mm inner diameter) pointing towards the interaction region. The whole assembly was heated up to 125 degrees Celsius, creating densities of glycine gas in the order of $10^{10} - 10^{11}$ molecules/cm^3. The oven was isolated electrically from the chamber and set to an intermediate potential between the tip of the magnet and the first electrode of the drift tube to avoid perturbations in the field. In order to prevent deposition of glycine in the chamber, an oven catcher consisting of a differentially-heated copper tube was installed in the bottom port for the glycine to be deposited at the bottom of the catcher.

5.3 Calibration

One of the challenges related to the very limited amount of time available to perform an XFEL experiment is that enough data to calibrate the different devices has to also be collected within that time. Finding rough calibration models is often critical during

the course of a beamtime in order to tune the free-electron laser to the right energy or to do online monitoring of the signals.

After the experiment, however, the calibrations have to be refined, often through iterative processes, to guarantee the consistency of all the signals measured across all the detectors. In this section we will discuss the final calibration for all the relevant detectors to this experiment.

5.3.1 Electron-bunch energy calibration

As we discussed in section 3.1.3, the electron-bunch energy is related to the final x-ray wavelength through

$$\lambda_{phot} = \frac{\lambda_u}{2\gamma^2}(1 + \frac{K^2}{2}), \tag{5.8}$$

where λ_u is the undulator period, γ is the Lorentz factor for the electron bunch, and K is the dimensionless undulator strength parameter. It is often simpler to measure the single-shot electron-beam energy using a non-intrusive BPM in a dispersive region (section 3.3.9), than the x-ray wavelength. In practice, however, calculating the wavelength (or photon energy) using that formula directly is difficult due to the inaccuracy with which the actual absolute electron-bunch energy or the K parameter can be measured.

To avoid a problematic *ab-initio* model as described, an empirical linear model can be used instead to obtain the relationship between the photon energy and the electron-beam energy. The linearity of the model can be justified as the variation range of the electron-beam energy is typically below 1% of the central value.

To adjust the parameters of the model, at least two absolute reference values are needed. As in this case we were working near the oxygen K-edge, we used a thin film (0.5 µm) of biaxially oriented polyethylene terephthalate (Mylar), a polymer containing oxygen. With this film placed at the entrance of the x-ray grating spectrometer, we measured the total amount of light arriving at the detector while scanning the photon energy between 527 and 547 eV. After normalizing the single-shot signal on the pulse energy measured by the gas detectors, and averaging the amount of light transmitted through the film at each photon energy, we obtained the experimental XANES profile for Mylar.

We adjusted the parameters of the model until the experimental absorption data (blue line in figure 5.7) matched the theoretical data (green and red lines) at the π^* peak near 531.5 eV and at the oxygen K-edge at around 540 eV for Mylar.

5.3.2 Magnetic bottle calibration

As regular time-of-flight devices, the MTOF relies on measuring the time that a particle takes to travel a fixed distance to determine the particle energy. The corresponding

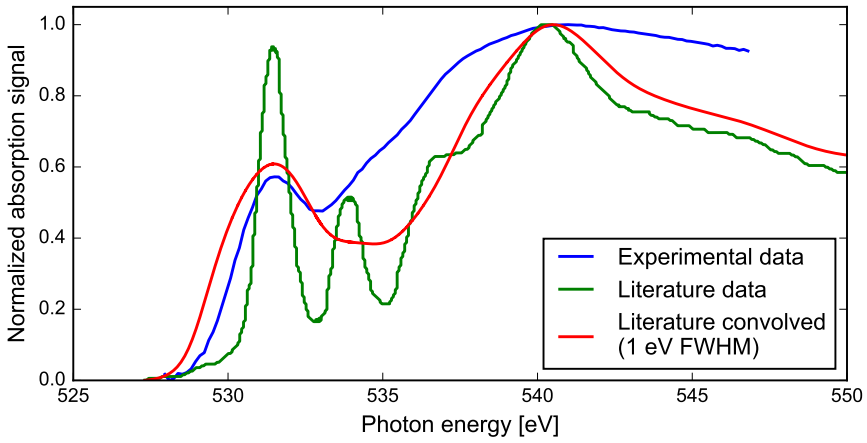


Figure 5.7: Experimental K-edge absorption data for Mylar compared to data in the literature [145] used to calibrate the electron-beam energy. The experimental data fits well the expected curve for Mylar, after convolving the data from the literature using a Gaussian of 1 eV FWHM to account for the bandwidth of the XFEL.

equation would be

$$E(\Delta t) = \frac{1}{2} m \frac{L^2}{\Delta t^2}, \tag{5.9}$$

where E is the energy of the particle, m the mass of the particle, L the length of the distance, and Δt the travelling time.

While in principle this equation does not have any free parameters, there are several reasons why it cannot be used in that form. First, similarly to the ion time-of-flight already discussed in section 4.2.3, the electron trajectories in a magnetic bottle are not straight lines, but elongated helices, and the actual travelling length cannot be easily measured. Additionally, a retardation potential is normally applied to take the electrons to a lower energy regime where their time-of-flight is longer and can be measured more accurately. However, since the actual value of the potential at the generation point of the electron is not easy to measure, the set value of the retardation potential is only *nominal* and may need to be corrected. Finally, to measure Δt it would be necessary to keep a record of the generation time and the arrival time for each electron. Since this cannot be done, instead, the only assumption that can be made is that all electrons are emitted at exactly the same time.

Taking into account all of the above, a more appropriate equation would be

$$E(t) = V + V_c + \frac{1}{2} m \frac{L_e^2}{(t - t_0)^2}, \tag{5.10}$$

where V is the nominal value for the applied retardation potential, V_c a correction to this potential, L_e the effective travelling length, and t_0 the emission time of the electrons. V_c, L_e and t_0 are treated as free parameters in the model and V_c, L_e must be close to

0 and L respectively. These parameters account for other second order effects such as space charge, work function changes and plasma potentials.

Furthermore, even though equation 5.10 works well for a particular retardation potential, it does not work very well for a large range of retardation potentials. The reason behind this is that the effective length L_e of the bottle changes with the retardation potential. This is due to, among other effects, the small distance that the electrons travel at their initial speed before they are slowed down by the applied retardation potential. A more complex model was used in this case to take this into account. This model assumes that two different distances contribute to the time of flight of an electron: L_1 where the electron travels at its original speed, and L_2 where the electron travels at the retarded speed. The equation for the model is

$$t - t_0 = \sqrt{m/2}\left(\frac{L_1}{\sqrt{E - V_c}} + \frac{L_2}{\sqrt{E - V - V_c}}\right). \tag{5.11}$$

This equation presents the difficulty that it cannot be easily inverted to obtain $E(t)$, so it had to be treated numerically.

Different experimental points were used to numerically find the optimal parameters for the models. These points along with the calibration curves for models 5.10 and 5.11 (labelled as 1 and 2 in the figure) are shown in figure 5.8. Different points from multiple physical mechanisms are included:

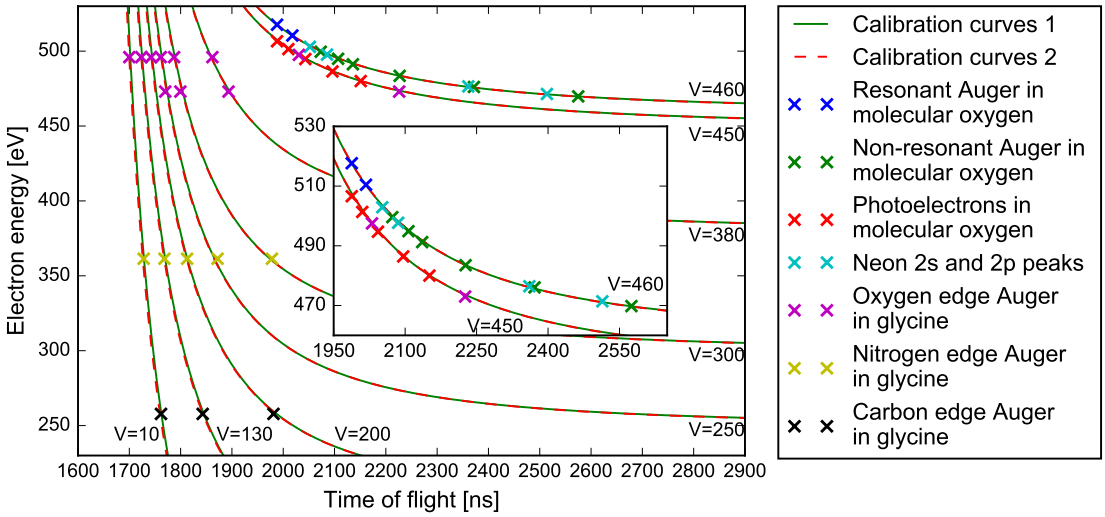
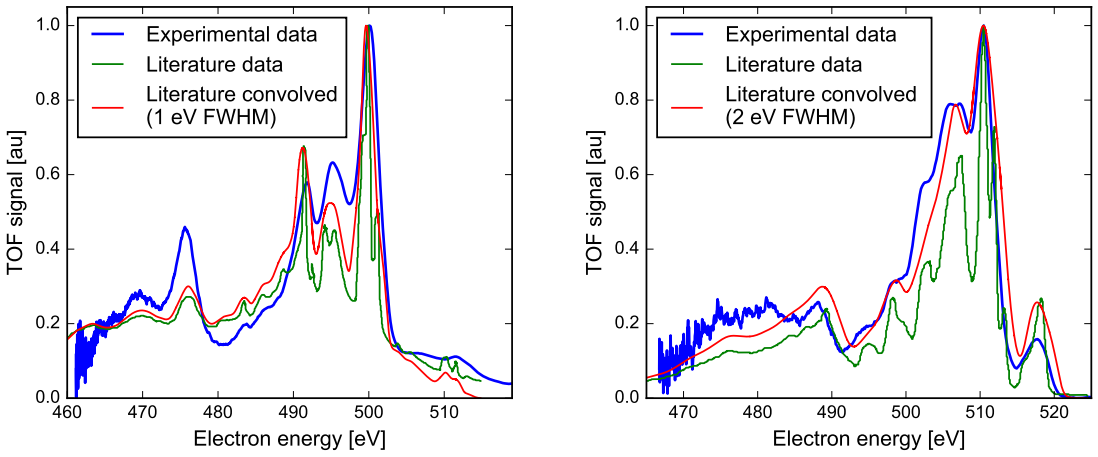


Figure 5.8: Summary of all the points used for the calibration of the magnetic bottle along with the calibration curves 1 and 2 for different retardation potentials. Calibration curve #1 corresponds to model 5.10 and #2 to model 5.11. The two curves only differ slightly at the shortest time-of-flight values (<1800 ns), where curve #2 fits the experimental points more accurately.

- Resonant and non-resonant Auger electron spectra obtained in O_2 by tuning the XFEL near the π^* resonance and above the K-edge respectively and compared

to data in the literature [46, 146]. These are the most reliable points for the calibration, as their position is independent of the x-ray photon energy. Excellent agreement between the calibrated experimental data using model 5.11 and the data from the literature can be observed in figure 5.9.



(a) (b)

Figure 5.9: Auger electron peaks in O_2 used for calibration. (a) Non-resonant Auger peaks obtained at 550 eV and compared to data from the literature [46]. (b) π^* resonant Auger peaks obtained at 529 eV and compared to data from the literature [146]. The shifted peaks around 505 eV correspond to photoelectron peaks with an approximate 2 eV mismatch caused by the fact that the experimental data was taken 2 eV away from the actual π^* resonance near 531 eV. A Gaussian convolution is applied to the data from the literature to account for the resolution of the detector at the different energy regions.

- Photoelectron spectra in O_2 obtained by tuning the XFEL well below the π^* resonance. The absolute energy of these electrons depends on the photon energy, however the spacing between the peaks does not depend on that and can be used for this calibration. In the end once the x-ray grating spectrometer calibration was reliable, the binding energy of the original states of the photoelectrons could be plotted and compared to data in the literature [36, 147, 148] (figure 5.11). While the amplitude of the peaks is different, probably due to the different experimental procedures used to obtain the data, there is quantitative agreement on the position of the peaks.

- Photoelectron spectra in neon obtained by tuning the XFEL at around 520 eV. Again, since it is a photoelectron signal it presents the same caveats as the O_2 photoelectrons. Once the x-ray grating spectrometer calibration was reliable, the binding energies of the neon photoelectron peaks were plotted (figure 5.12c) and compared to literature data for the neon 2s and 2p orbitals (-48.5 eV and -21.6 eV, respectively) [149].

- Auger electrons from the oxygen, nitrogen, and carbon K-edges of a glycine obtained at different retardation potentials. This was especially challenging as there was no experimental data in the literature to compare the spectra. It was however necessary to use this data to calibrate the magnetic bottle, as it was the only data that was taken for a wide range of different retardation potentials. The solution to this was to optimize the calibration in order not to match any particular energy for the peaks, but to ensure that the peak positions remained at the same exact locations across different retardation potentials. In figure 5.8 we can observe that it is for lower retardation potentials that the two models (dashed-red for equation 5.11 and green for equation 5.10) disagree the most, getting a higher accuracy for the more complex model 5.11. Using this model, the spectra for different retardation potentials match the position of the peaks quite well (figure 5.10), even though the relative amplitude of the peaks changes due to energy-dependent collection efficiency of the MTOF.

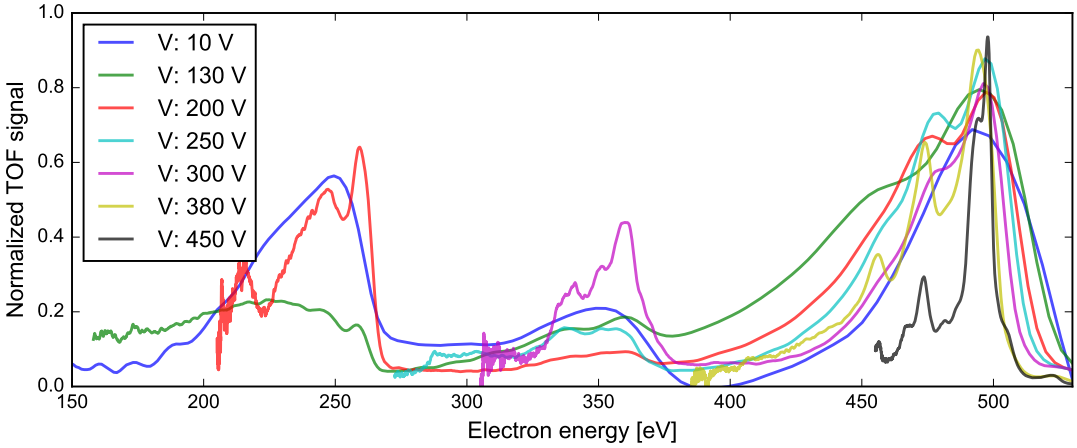


Figure 5.10: Auger electrons in glycine from the carbon edge (near 280 eV), the nitrogen edge (near 400 eV) and the oxygen edge (near 540 eV), for different retardation potentials of the MTOF. The data was obtained with 560 eV single-pulse x-rays. The figure shows that the same calibration is consistent across a wide range of retardation potentials.

5.3.3 X-ray grating spectrometer calibration

The calibration of the x-ray grating spectrometer was carried out using a linear model to convert between pixel number and photon energy.

In order to adjust the two parameters of the linear model, photoelectrons from neon were recorded with the MTOF at the same time as spectral profiles were obtained with the x-ray grating spectrometer using single x-ray pulses. The x-ray grating spectrometer profiles were fitted to Gaussian curves, obtaining a value for the centre of the Gaussian in pixels for every shot. The photon energy of the XFEL was scanned over a range from 519 to 526 eV while recording MTOF spectra and x-ray grating spectrometer traces. The parameters from the model were adjusted until the binding energy obtained (electron

energy measured at the MTOF minus the x-ray photon energy obtained from the x-ray grating spectrometer) for the 2s and 2p photoelectron peaks was consistent with the values in the literature (-48.5 eV and -21.6 eV, respectively) [149] and constant across the whole range of photon energies (figure 5.12).

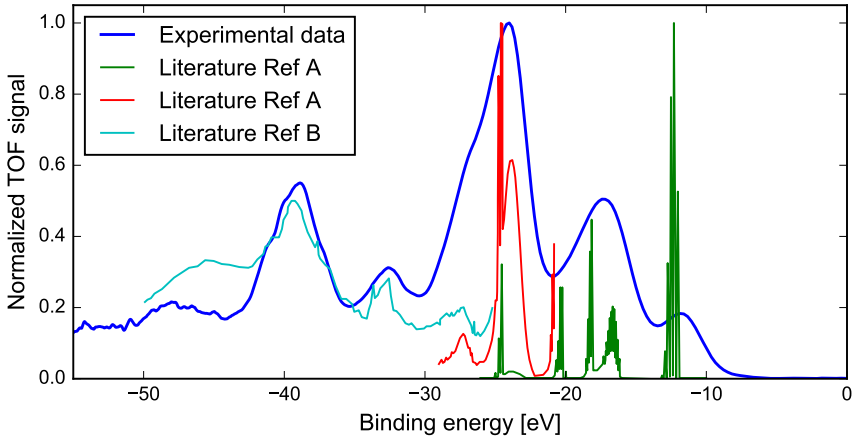


Figure 5.11: O_2 photoelectron spectrum obtained at 520 eV as function of the binding energy compared to data in the literature (Ref. A [36], Ref. B [147, 148]), showing a quantitative agreement in the location of the different peaks.

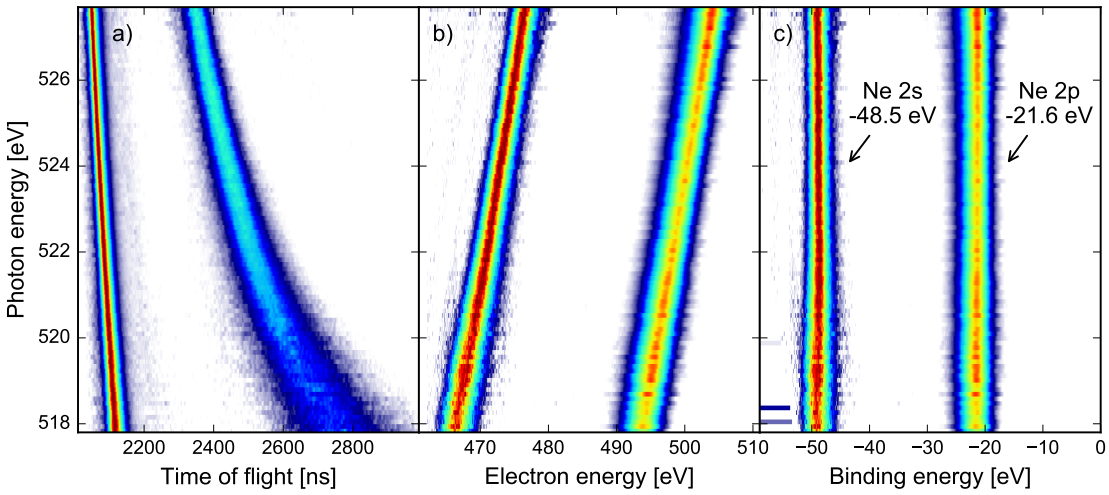


Figure 5.12: Ne photoelectron peaks (2p and 2s) for different values of the incident photon energy. (a) As a function of the electron time-of-flight. (b) As a function of the calibrated electron energy. (c) As a function of the calibrated binding energy. As expected, the experimental values of the binding energy are constant across the whole photon energy range and close to the annotated values found in the literature [149].

Additionally the calibration of the photon energy was also checked against the binding energy of O_2 photoelectrons (figure 5.11).

5.3.4 Pulse energy calibration

The total pulse energy can be measured making use of gas detectors, however, as mentioned in section 3.3.1 these values need to be calibrated. The calibration for these detectors is typically performed as a standard LCLS procedure measuring the electron-beam energy loss. It relies on the fact that total energy of the photons generated at the undulator comes out of the kinetic energy of the electron beam, so the total generated x-ray energy must correspond to kinetic energy lost by the electrons. By using a magnetic dipole with a beam-position monitor, and then inducing oscillations in the electron beam across the undulator, it is possible to change the x-ray generation efficiency and correlate the final electron-bunch energy to the values measured by the gas detector to provide an absolute calibration. It is worth noting that this calibration is dependent on the photon energy because while the electron-beam-energy loss measures absolute energy the gas detectors are effectively counting photons.

A limitation of this procedure is, however, that it needs to be performed under high-energy lasing conditions, typically in the range 0 to 2 mJ, so the electron energy loss is measurable. This means that in practice it has to be done without spoiling the electron bunch. Once the bunch is spoiled, the total pulse energy may be as low as a few µJ, within the uncertainty in accuracy of the gas detector calibration.

In order to provide a more accurate calibration, we compared the values measured by the gas detector to the total amount of light measured in our x-ray grating spectrometer for events with a particular photon energy using single-pulse events generated with the spoiling technique (figure 5.13a). Considering data above 14 µJ as measured by the gas detector, there is a clear linear dependence, however, by extrapolating this linear behaviour, we observe that it does not cross the x axis at (0, 0) but at (8.35, 0). This indicates an offset error in the absolute calibration of 8.35 µJ, which is very significant at these low pulse energies.

Additionally, for energies below 14 µJ, we observe a very clear non-linear behaviour, which can be explained in terms of two effects. First, the standard deviation of the coloured points in figure 5.13a is relatively large, meaning that the shot noise of the measurement is significant. This is due to the low collection efficiency of the x-ray grating spectrometer. Since the data is filtered based on photon energy, and this is obtained from fitting the signal in the x-ray grating spectrometer, only the shots for which the spectrometer collects enough signal to produce a successful fit contribute to the data. This effectively biases the distribution of total signal at the x-ray grating spectrometer for low-energy points, inflating the values. The second effect is however more fundamental: while the default calibration of the gas detector is performed linearly with the energy loss, it may actually not be linear for very low energies, underestimating the actual pulse energy.

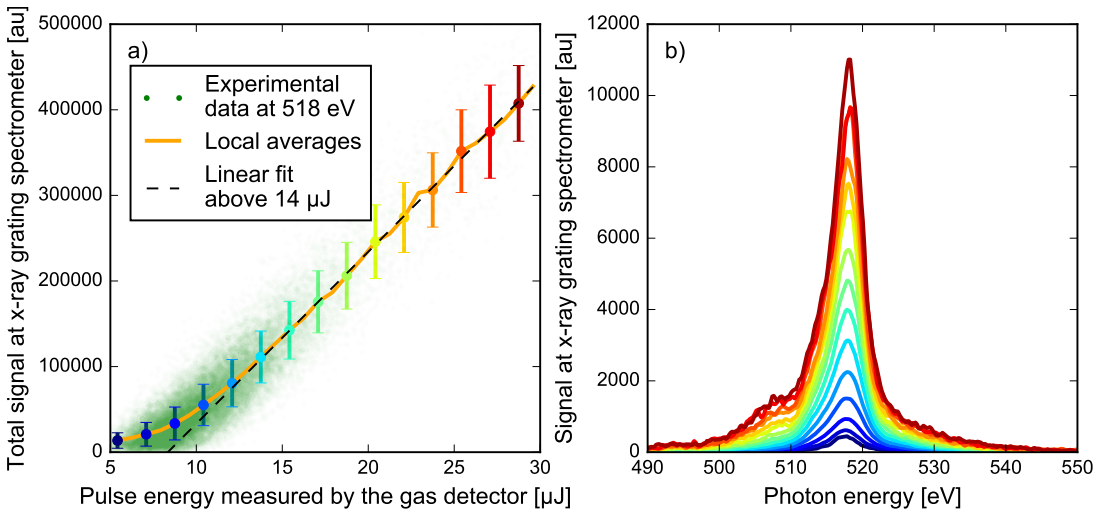


Figure 5.13: Gas detector calibration. (a) Values measured with the default calibration of the gas detector compared to values measured with an x-ray grating spectrometer at 518 eV. The error bars correspond to the standard deviation for each local average. (b) Averaged data from the spectral spectrometers for some of the local averages. The colour of the profiles correspond to the different colours of the points in (a).

While the first effect is significant, it does not explain how the gas detector can ever report values as low as 5 µJ when we have shown that the values are asymptotically off by 8.35 µJ as this would imply negative x-ray pulse energies. After systematically observing this effect, it was determined that indeed the second reason contributes to this, and as a result different improvements were implemented at LCLS. First, different more complex calibration formulas were made available for the low-energy regimes. And second, the actual time-dependent raw signals from the gas detectors photodiodes started to be recorded and made available to future users. In our case, we directly used information from curves similar to the one in figure 5.13a to provide a gas-detector to corrected-gas-detector calibration for the analysis of our data.

Finally, it is worth noting that while we have discussed the measurement of the total pulse energy after the undulator, we normally are interested in the energy at the focus of the interaction region of our experiment. Different reflective elements typically made of boron carbide (B_4C) such as the beam transport mirrors, used to steer the beams towards the different beamlines, or the KB mirrors [140] used to focus the beam, reduce the amount of usable x-ray pulse energy. In the case of soft x-rays (< 2000 eV) at the AMO beamline at LCLS, the nominal total transmission is 66% at 825 eV [150] and gradually decreases as the photon energy approaches the carbon K-edge near 300 eV. Furthermore, due to gradual contamination of the mirrors the reflectance can become greatly reduced. This means that in practice for photon energies near the oxygen edge (~ 540 eV) the transmission may be as low as 10-20% of the initially generated x-ray energy. As a consequence, it is typically required to set-up an additional experiment

using some non-linear mechanism (such as multi-photon processes), as we described in section 4.2.3, to estimate the energy or intensity delivered to the interaction region.

5.4 Extracting time information: XTCAV

A critical device to obtain time-dependent x-ray information is XTCAV (see section 3.3.8 and references [127, 128]). However, the XTCAV images are technically a diagnostic of the electron bunches, so any x-ray information has to be inferred by understanding the effect of the lasing process in the electron bunch and processing the XTCAV images accordingly.

The images provided by XTCAV (after a background subtraction) are proportional to the electron density distribution of the bunches in the time-energy phase space. Figure 5.14 shows examples of such XTCAV images in a single bunch mode using the single-slotted foil (section 3.1.4). To understand the effect of lasing in the bunch we can compare an image where the bunch is lasing to an image where the bunch is not lasing. We can see that for the non-lasing image (figure 5.14a) there is a region where the electron density peaks: this region corresponds to the part of the bunch that goes through the slit of the foil. Looking at the lasing image (figure 5.14b), we can now observe that it looks very similar everywhere except in the region which had the highest electron density. First, there is an overall decrease of the electron energy in that region corresponding to the energy transferred to the x-ray pulses. Second there is a larger energy spread caused by the lasing process. It is possible to study this change in mean energy and energy spread to recover the time-dependent x-ray power.

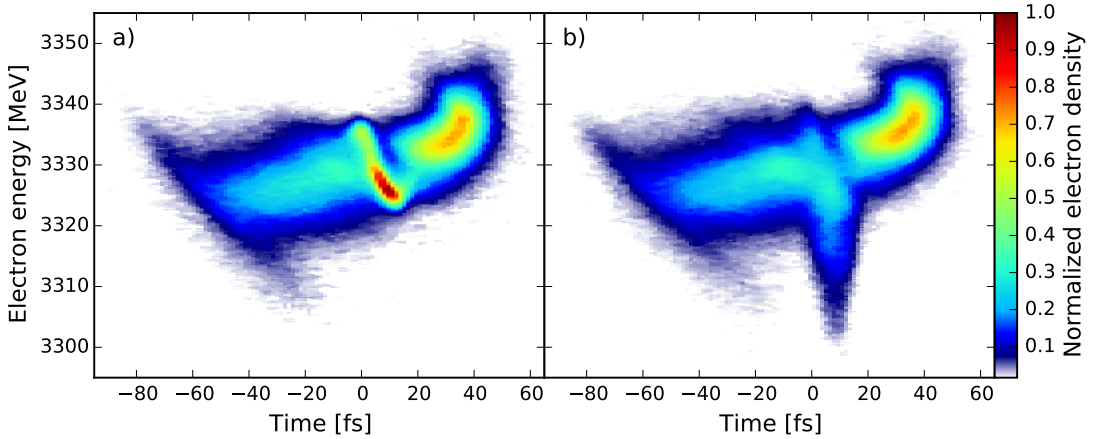


Figure 5.14: XTCAV images for a single bunch spoiled with the single-slotted foil (a) when the bunch is not lasing and (b) when the bunch is lasing. A decrease in the electron bunch energy can be observed in the part of the electron bunch that undergoes lasing, which is the region with the highest electron density before lasing. The images were obtained for two different but similar events. The normalization factor applied to the coloured axis showing the electron density in the time-energy phase space is kept constant across both images.

5.4.1 Calibrating the XTCAV images

The first step consists of converting the pixel information of the axes of the measured image into energy and time to transform the distribution $\mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_t, \mathrm{pix}_E)$ into $\mathrm{Im}_{\mathrm{XTCAV}}(t, E)$ (figure 5.15).

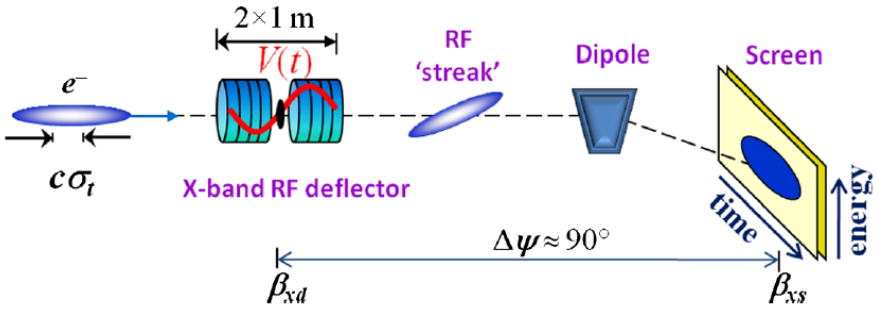


Figure 5.15: Temporal streaking and energy dispersion of the electron bunch using XTCAV at LCLS. The electron bunch is dispersed in energy using a magnetic dipole and streaked in time using an RF field and projected onto a screen to create an XTCAV image. The pixels of the raw image $\mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_t, \mathrm{pix}_E)$ have to be calibrated into time and energy $\mathrm{Im}_{\mathrm{XTCAV}}(t, E)$ using single-shot calibration constants. Figure extracted from [127].

For the energy axis, the scaling constant between pixels and electron energy is obtained as

$$C_{\mathrm{energy}}[\mathrm{MeV/pix}] = C_{\mathrm{dispersion}}[\mathrm{MeV/\mu m}] \cdot C_{\mathrm{camera}}[\mathrm{\mu m/pix}], \tag{5.12}$$

where C_{camera} is the pixel size and $C_{\mathrm{dispersion}}$ is the dispersion induced by the magnetic dipole situated before the XTCAV screen. While this only allows one to calibrate pixel increments into energy increments, an additional calibrated beam-position monitor measures the overall electron-bunch energy $E_{\mathrm{mean\ energy}}$, which can then be assigned to the pixel of the image corresponding to the position of the centre of mass across the energy axis. This pixel can be calculated as

$$\mathrm{pix}_{\mathrm{ECOM}} = \frac{\int_0^{\max(\mathrm{pix}_E)} \mathrm{pix}_E \cdot \mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_E) d\mathrm{pix}_E}{\int_0^{\max(\mathrm{pix}_E)} \mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_E) d\mathrm{pix}_E}, \tag{5.13}$$

where $\mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_E)$ is the marginal distribution,

$$\mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_E) = \int_0^{\max(\mathrm{pix}_t)} \mathrm{Im}_{\mathrm{XTCAV}}(\mathrm{pix}_t, \mathrm{pix}_E) d\mathrm{pix}_t, \tag{5.14}$$

allowing us to finally obtain a calibrated energy axis,

$$E = C_{\mathrm{energy}} \cdot (\mathrm{pix}_E - \mathrm{pix}_{\mathrm{ECOM}}) + E_{\mathrm{mean\ energy}}. \tag{5.15}$$

In the case of the time axis the scaling constant is calculated as

$$C_{\text{time}}[\text{fs/pix}] = C_{\text{calibration}}[\text{fs} \cdot \text{MV/µm}^2]/A_{\text{streaking}}[\text{MV/µm}] \cdot C_{\text{camera}}[\text{µm/pix}], \quad (5.16)$$

where C_{camera} is the pixel size $A_{\text{streaking}}$ is the amplitude of the streaking electric field measured for every shot and $C_{\text{calibration}}$ is a calibration constant. Additionally the phase of the electric field streaking the electrons defines the sign of the axis. This phase $\phi_{\text{streaking}}$ is measured for every shot and compared to a reference phase ϕ_{ref} to obtain the direction of the axis,

$$\text{dir}_{\text{time}} = \text{sign}(\cos(\phi_{\text{streaking}} - \phi_{\text{ref}})). \quad (5.17)$$

While an absolute time reference is not available in this case, time zero is assigned for simplicity to the position of the centre of mass across the time axis pix_{tCOM} calculated similarly to pix_{ECOM}. This leaves the final transformation as

$$t = \text{dir}_{\text{time}} \cdot C_{\text{time}} \cdot (\text{pix}_t - \text{pix}_{\text{tCOM}}). \quad (5.18)$$

Finally the amplitude of the XTCAV images $\text{Im}_{\text{XTCAV}}(t, E)$ is only proportional to the electron density $D(t, E)$ and needs to be calibrated. Using the total electron-bunch charge (#electrons), measured for every shot at the electron dump, it is possible to obtain the calibrated electron density distribution,

$$D(t, E) = \#\text{electrons} \frac{\text{Im}_{\text{XTCAV}}(t, E)}{\int_{-\infty}^{+\infty} \int_{0}^{+\infty} \text{Im}_{\text{XTCAV}}(t, E) dE dt}. \quad (5.19)$$

5.4.2 Obtaining non-lasing references

Once the electron density distribution $D(t, E)$ is known for a given XFEL shot, the next step is to find a reference to compare it to. One could think about installing an XTCAV diagnostic before the undulators to use as a non-lasing reference and another one after the undulators to obtain the lasing images. However this is not possible for two reasons. First, XTCAV is an intrusive diagnostic that alters the electron bunch and would completely spoil the lasing process if installed before the undulator. Second, even if XTCAV did not alter the bunch, due to normal bunch propagation effects, the image taken before the undulator would be different than the image taken after the undulators even if the bunch did not lase.

As a solution, a single XTCAV device is used after the undulators and a non-lasing reference run is recorded prior to the experimental lasing runs. In order to disable the lasing process during the reference runs, the trajectory of the electrons is shifted laterally during their propagation through the undulators. This prevents the electrons from overlapping spatially with any generated x-rays, impeding the amplification of any

spontaneous emission, and avoiding the SASE lasing process. This process is commonly known as *kicking* the electron bunch.

The events from the non-lasing reference run cannot just be averaged together to obtain a single universal reference as the reference run is comprised of a very broad variety of different images resulting from the single-shot fluctuations of the acceleration and compression processes of the XFEL. Instead, images are clustered together in groups of typically 5 based on similarity, producing at the end a single average reference image for each group, and a pool of reference images overall.

Whenever a lasing electron distribution $D_{\text{lasing}}(t, E)$ needs to be analysed, a non-lasing reference $D_{\text{non-lasing}}(t, E)$ is selected from the pool. In this case the reference cannot be chosen based on the similarity of the image, because the lasing image should always be different than any other non-lasing image. Nevertheless, by assuming that the lasing process does not reorder the electrons in time (at least macroscopically, without considering micro-bunching), we can calculate the time-dependent electron current profile of the lasing image integrating across the energy axis,

$$I(t) = \int_0^{+\infty} D(t, E) dE, \tag{5.20}$$

and use the select the reference image from the pool of references that presents the most similar electron current profile. Figure 5.16c shows similarity between the electron current profile for a lasing image and the selected non-reference image used for the reconstruction.

5.4.3 Obtaining the x-ray power

There are two different methods to obtain the x-ray power. The first method is the most intuitive and is based purely on the energy transferred between the electron bunch and the x-ray pulse, calculating the time-dependent x-ray power as the time-dependent energy lost by the electrons. We first obtain the time-dependent centre of mass for the energy (COM) of the electron distributions,

$$E_{\text{COM}}(t) = \frac{\int_0^{+\infty} E \cdot D(t, E) dE}{\int_0^{+\infty} D(t, E) dE}. \tag{5.21}$$

And then calculate a first approximation to the time-dependent x-ray power as

$$P_{\text{COM naive}}(t) = (E_{\text{COM non-lasing}}(t) - E_{\text{COM lasing}}(t)) \cdot I_{\text{lasing}}(t). \tag{5.22}$$

While this method allows the calculation the absolute power directly, it is often found that it is not consistent with the total pulse energy measured with the gas detectors

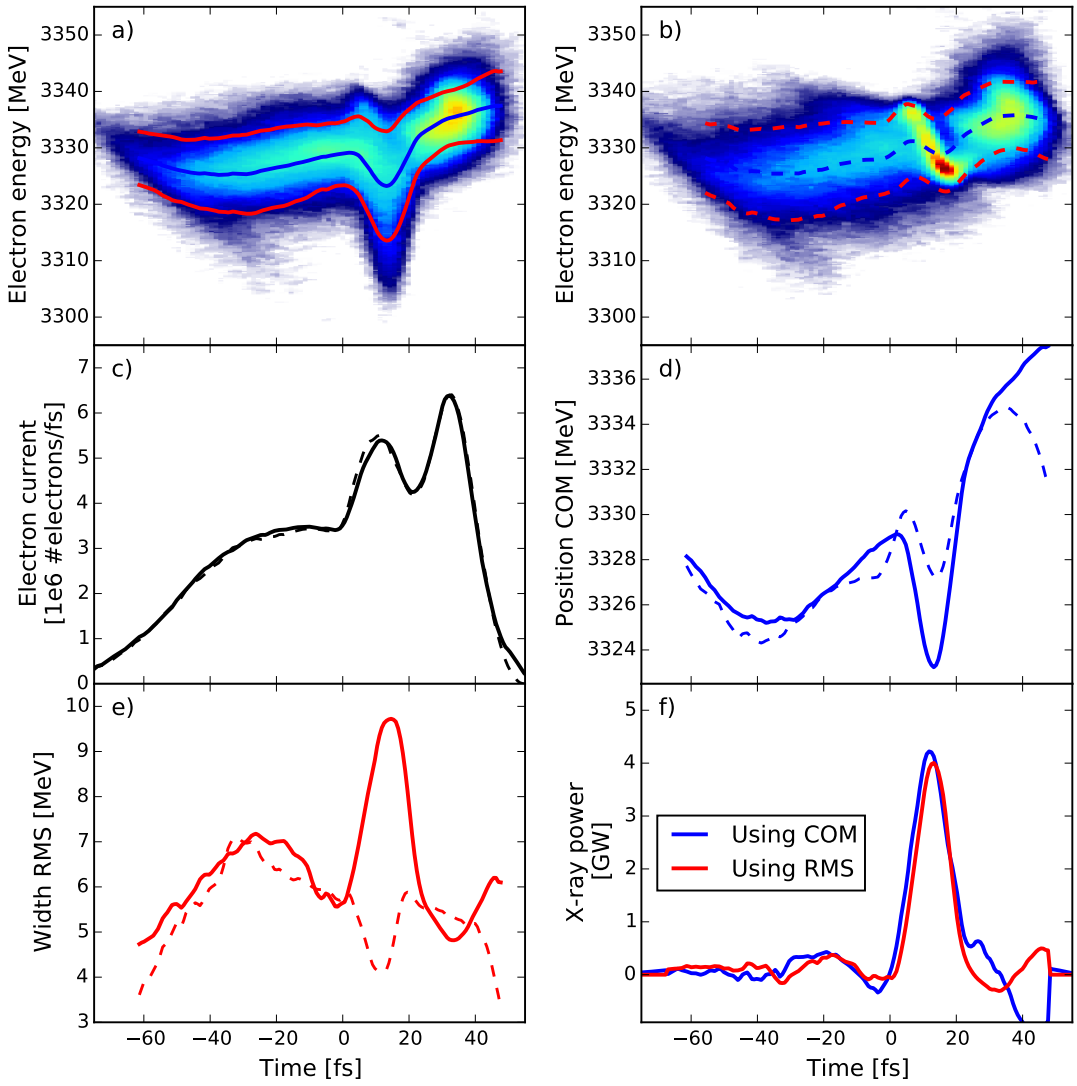


Figure 5.16: XTCAV reconstruction for a single bunch comparing properties of the lasing XTCAV image (solid lines) and the non-lasing reference (dashed lines). (a) Lasing XTCAV image. (b) Corresponding non-lasing reference. (c) Time-dependent electron current. (d) Time-dependent centre of mass (COM) of the energy. (e) Time-dependent standard deviation (RMS) of the energy. (f) Reconstruction of the time-dependent x-ray power using the COM method (blue) and the RMS method (red). The COM (blue) and COM ± RMS (red) profiles are overlaid on the corresponding XTCAV images (a,b).

PE_{GD}, in particular that

$$PE_{\mathrm{GD}} \neq \int_{-\infty}^{+\infty} P_{\mathrm{COM\ naive}}(t)dt, \tag{5.23}$$

typically due to a small offset across the energy axis between the lasing and non-lasing images. In order to satisfy that equation, an additional constant is added to the calculation that corrects for the offset,

$$P_{\mathrm{COM}}(t) = (E_{\mathrm{COM\ non\text{-}lasing}}(t) - E_{\mathrm{COM\ lasing}}(t) + C_{\mathrm{correction}}) \cdot I_{\mathrm{lasing}}(t), \tag{5.24}$$

where the correction is calculated as

$$C_{\mathrm{correction}} = \frac{PE_{\mathrm{GD}} - \int_{-\infty}^{+\infty} P_{\mathrm{COM\ naive}}(t)dt}{\int_{-\infty}^{+\infty} I_{\mathrm{lasing}}(t)dt} = \frac{PE_{\mathrm{GD}} - \int_{-\infty}^{+\infty} P_{\mathrm{COM\ naive}}(t)dt}{\#\mathrm{electrons}}. \tag{5.25}$$

Figures 5.16a-b show the COM position (blue) overlaid on the corresponding XTCAV images. Figure 5.16d compares the COM profiles for the lasing image and its reference. We can observe that the main difference occurs at around 15 fs, where the non-lasing profile (dashed line) shows a clear higher value than the lasing profile (solid line), corresponding to the energy lost by the electrons and gained by the x-rays. There is also a small mismatch above 30 fs, which in this case corresponds to an artifact of the reference. The reconstruction of the power profile is shown in blue in 5.16f. It shows a 10 fs FHWM pulse peaking just above 4 GW, consistent with a total energy of approximately 55 µJ measured with the gas detector. We can also observe a region where the power is negative, this is a result of the reference mismatch obtained in the COM profiles.

The second method to obtain the instantaneous power of the x-ray pulses relies on the broadening of the time-dependent energy spread of the electrons. We can obtain this distribution by calculating the root mean square (RMS) energy separation from the COM for each time slice,

$$E_{\mathrm{RMS}}(t) = \sqrt{\frac{\int_0^{+\infty}(E - E_{\mathrm{COM}}(t))^2 \cdot D(t,E)dE}{\int_0^{+\infty} D(t,E)dE}}. \tag{5.26}$$

It can be shown that the x-ray power is in this case related to the energy spread through a proportionality relationship [128],

$$P_{\mathrm{RMS}}(t) \propto (E^2_{\mathrm{RMS\ lasing}}(t) - E^2_{\mathrm{RMS\ non\text{-}lasing}}(t)) \cdot I^{2/3}_{\mathrm{lasing}}(t) := P_{\mathrm{RMS\ proportional}}(t). \tag{5.27}$$

While the method in this case does not provide absolute power values, the total pulse energy measured with the gas detector PE_{GD} can be used again to calculate the

proportionality constant and the absolute x-ray power,

$$C_{\text{prop}} = \frac{PE_{\text{GD}}}{\int_{-\infty}^{+\infty} P_{\text{RMS proportional}}(t)dt} \tag{5.28}$$

$$P_{\text{RMS}}(t) = C_{\text{prop}} \cdot P_{\text{RMS proportional}}(t). \tag{5.29}$$

Figures 5.16a-b show the RMS deviation at each side of the COM (red) overlaid on the corresponding XTCAV images. Figure 5.16e compares the RMS profiles for the lasing image and its reference. The main difference occurs again at around 15 fs, where the lasing profile (solid line) doubles the RMS value of the non-lasing profile (dashed line). There is again a small mismatch above 30 fs, which in this case corresponds to a small difference with the reference. The reconstruction of the power profile is shown in red in 5.16f. It shows a similar 10 fs FHWM pulse peaking at around 4 GW and slightly shifted with respect to the COM reconstruction.

5.4.4 Double bunch configuration

In the two-bunch configuration, the electron density obtained from the XTCAV images shows the contribution from each of the bunches (figure 5.17b). In this particular case we can observe two electron bunches that are overlapping in time, but with an offset of about 50 MeV in energy. The low-energy bunch seems to be lasing at around -25 fs. However, for the high-energy bunch is not clear whether it is lasing at -5 fs, at -40 fs, or at both locations. In order to solve this, we can look at the corresponding non-lasing reference image (figure 5.17a) to realize that the feature that looked like lasing at around -40 fs is also present in the reference, meaning that it is not a lasing feature, but a feature produced by the the acceleration process of the electrons. This is consistent with the fact that this data was obtained using the double-slotted foil, which should only allow each of the bunches to lase at one individual location.

Using a similar method as with the single bunch, the x-ray power can be obtained (figure 5.18). The only additional task is to separate the two electron bunches into two different images. This is typically achieved by thresholding the image until it is formed of two independent pixel "islands", each of them corresponding to each of the bunches (figure 5.18a-b). After the bunches are separated, the exact same procedure of calculating the current, COM and RMS is followed to obtain the x-ray power (figures 5.18c-j), with the only difference that the gas detector information, which still gives a single value of the total x-ray energy, has to be used to simultaneously normalize the power profiles for both pulses.

In this case the power reconstructions obtained from each of the methods (figures 5.18i-j) are not as similar as in the single bunch case. Overall it seems that the RMS method gives more reasonable results, as the comparison with the non-lasing profiles is more accurate. These inaccuracies may be a consequence of the reference pool not being

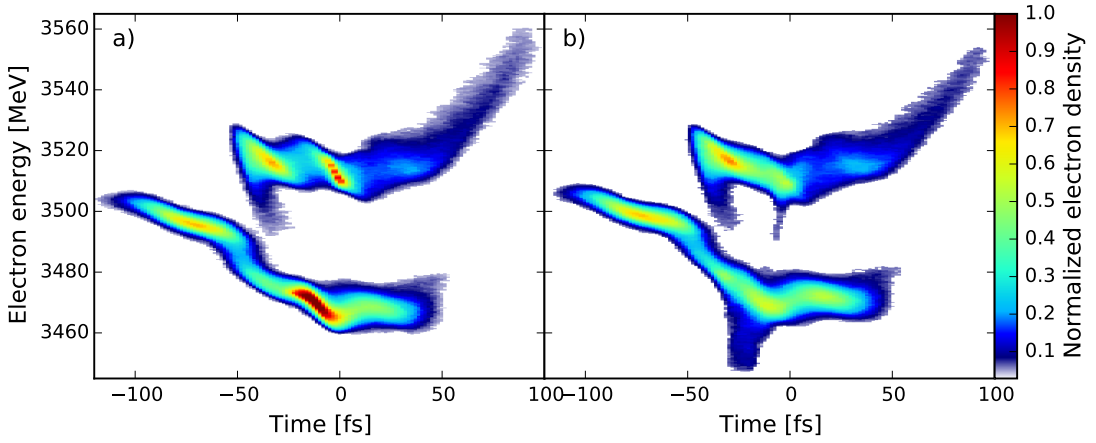


Figure 5.17: XTCAV images in a double bunch mode spoiled using the double-slotted foil (a) when the bunches are not lasing and (b) when the bunches are lasing. A decrease in the electron bunch energy can be observed in the part of each of the electron bunches that undergoes lasing, which is the region with the highest electron density before lasing. The images were obtained for two different but similar events. The normalization factor applied to the coloured axis showing the electron density in the time-energy phase space is kept constant across both images.

large enough: while the size of the reference pool is similar to the single bunch case (about 400 references) this data was obtained while scanning the electron energy. As a consequence the number of available references taken for a specific region of electron energy values is smaller. In general we have observed that the RMS method seems to be less sensitive to inaccurate references, and should be the preferred default option when the agreement between both methods is not acceptable.

Additionally it is not always necessary to perform a full reconstruction to obtain useful information from XTCAV like the time delay. When looking at the XTCAV images and their current profiles (figures 5.18a-d) we can observe that the electron bunches always lase at a peak of the electron current, the peak corresponding to the unspoiled part of the bunch. By looking to a series of XTCAV images and some references is possible to identify the peak in the electron current that is lasing, and then extract its location just by analysing the electron current. By subtracting the location of the relevant peak in the electron current for each of the bunches it is possible to directly obtain the time delay with ∼1 fs of accuracy.

5.4.5 Limits of the method

During the analysis of XTCAV images it is important not to forget that we are measuring electron properties, and that transferring those properties into the x-ray pulses may not be straightforward.

As the electron bunch lases, the generated x-rays travel faster and may be amplified by regions of the electron bunch longer than the pulse itself. This means that typically the actual x-ray pulse may be shorter that what is measured with XTCAV. This it not

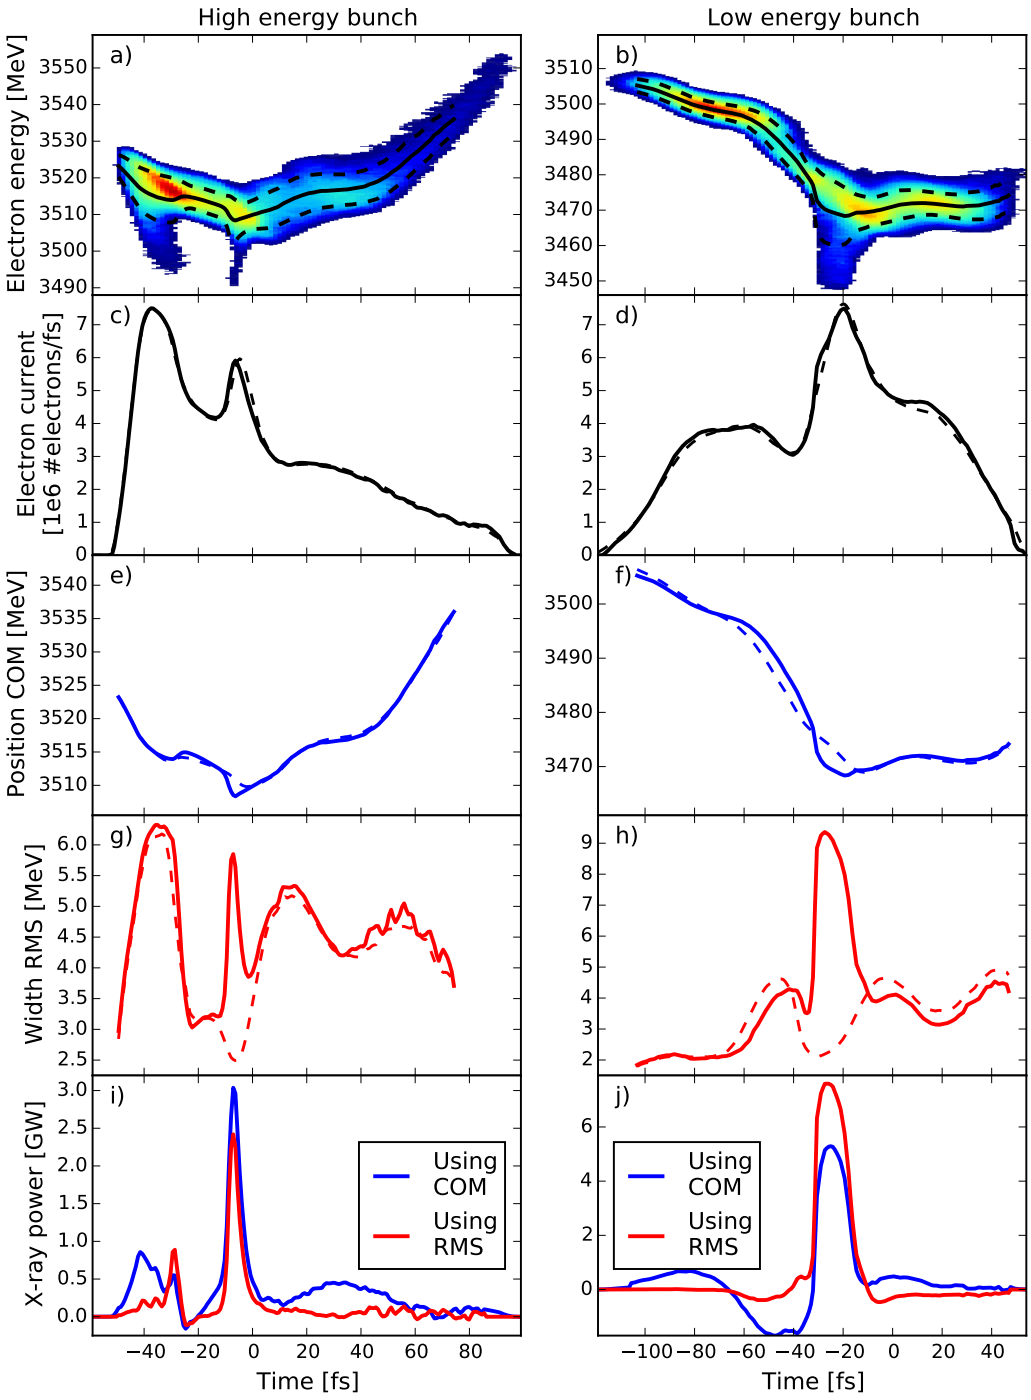


Figure 5.18: XTCAV reconstruction for a double bunch configuration performed for each of the bunches independently. (a,b) Separated bunches from the XTCAV image. (c,d) Time-dependent electron current for each of the bunches (solid lines) compared to the non-lasing reference (dashed lines). (e,f) Time-dependent COM for each of the bunches (solid lines) compared to the non-lasing reference (dashed lines). (g,h) Time-dependent RMS for each of the bunches (solid lines) compared to the non-lasing reference (dashed lines). (i,j) Reconstruction of the time-dependent x-ray power using the COM method (blue) and the RMS method (red) for each of the bunches. The COM (solid black) and COM ± RMS (dashed black) profiles are overlaid on the corresponding XTCAV images (a,b).

normally a problem when generating hard x-rays because as the energy of the electrons increase, their speed asymptotically approaches the speed of light, however, it may be a problem when generating soft x-rays. For example, when generating x-rays at 560 eV with 3520 MeV electrons, the electrons travel at 0.9999999895 the speed of light. However, this small difference still gives a slippage of 0.03 fs per metre, which can grow to a few fs if we consider propagation through typically tens of meters of undulators. Additionally, the wiggling electron trajectories add some distance to the electron paths.

Using a similar reasoning, when two different electron bunches sit at different energies, they also propagate at different overall speeds, arriving at XTCAV at slightly different times. This effect is however less significant as, for example, a difference of 50 MeV at 3500 eV leads to an slippage of only 0.001 fs per metre.

Additionally, in some configurations (section 3.2) different bunches or different parts of the same bunch are controlled to lase at different undulators [95], with magnetic chicanes between the undulators. These magnetic chicanes alter both the timings and the spectral properties of the electron bunches, making it more difficult to recover the x-ray power and to extract accurate timing from XTCAV.

5.4.6 Availability of the algorithm at LCLS

As part of the work for this Ph.D., custom software tools were built by the author to automate the XTCAV algorithms. This includes fetching all the relevant variables, generating background images for subtraction, filtering the noise of the images, separating the bunches, generating the non-lasing references, calculating the statistics on the images, finding the best reference and performing the reconstruction. These software tools are now part of the official analysis framework *psana* at LCLS [151], and were included as part of one of our publications [13] in collaboration with the LCLS data recording and analysis teams.

Using the software provided, obtaining the x-ray power for an XFEL event is as simple as,

```
from psana import *
from xtcav.ShotToShotCharacterization import *

ds = DataSource('exp=xpptut15:run=302:smd')
XTCAVRetrieval = ShotToShotCharacterization();
XTCAVRetrieval.SetEnv(ds.env())

# Code to fetch the event object goes here
# ...

XTCAVRetrieval.SetCurrentEvent(evt)
time, power, ok = XTCAVRetrieval.XRayPower()
```

5.5 Looking at static data in glycine

Much information about a sample can be obtained by looking at static data, even before considering any time-resolved results. Static data is of extreme importance to learn about the background signal and to obtain references of the expected time-resolved signal. Furthermore, before working around a resonance, it is convenient to explore the signal obtained for different values of the photon energy. This exploration can be performed explicitly by means of a manual scan of the photon energy or implicitly by the photon energy fluctuations which can be recorded shot to shot.

In this section we will present two examples of useful static data for a glycine sample (NH_2-CH_2-COOH), obtained with unspoiled bunches generating 75 fs pulses with energies between 50 and 550 µJ per pulse.

5.5.1 Non-resonant processes

As the sample contained oxygen, nitrogen and carbon atoms, to avoid any resonances we tuned the XFEL to 560 eV, 20 eV above the oxygen K-edge. Using this photon energy, we collected data at different values of the retardation potential. First, we used 200, 300 and 380 V to measure carbon-edge, nitrogen-edge and oxygen-edge non-resonant Auger spectra (blue lines in figures 5.19a-c). In each of the processes of interest, an electron was removed by direct photoionization from one of the carbons, nitrogen or oxygens core shells, leaving the molecule in an extremely unstable cationic excited state (lifetime $<$ 2-5 fs) leading to the auto-ionization of a second Auger electron using the "spare" energy obtained by the subsequent repopulation of inner-shell hole, and leaving the system in a dicationic state.

The results were compared to theoretical data obtained by *ab initio* ADC calculations (red lines) [141, 12]. In the case of the oxygen and the carbon edge, the contributions from each of the two atoms of each kind were added in equal proportion. The theoretical data was offset in the horizontal axis to fit the data in order to correct for possible inaccuracies of the ionization potential in the calculation. The overall widths and shapes follow the same patterns, except for those not discernible due to the low resolution of the MTOF at these energies.

An additional retardation potential of 450 V was used to optimize the MTOF resolution for the oxygen-edge Auger electrons (blue curve on figure 5.19 d), the region where the time-resolved experiment was going to be carried out. This configuration presents a lower collection efficiency for low-energy electrons, making the relative peak amplitudes change. However, a more detailed structure can now be observed and compared in more detail to theoretical spectra (red lines).

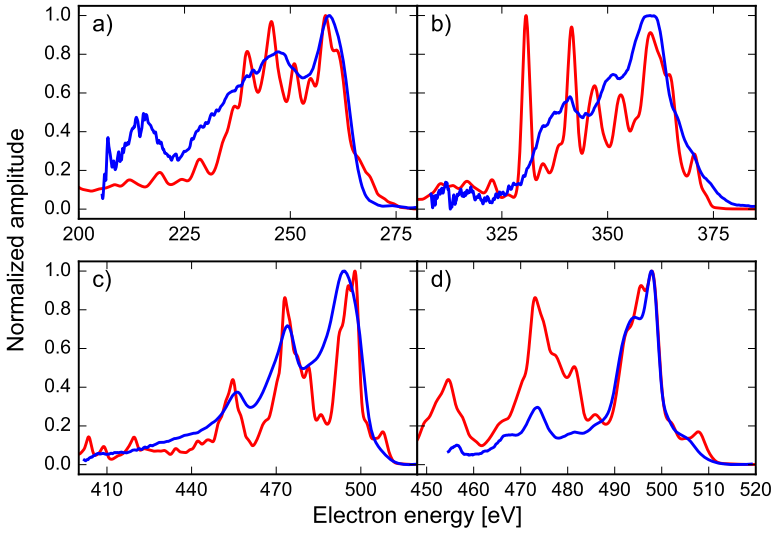


Figure 5.19: Experimental Auger spectra (solid blue lines) for a glycine sample from (a) the carbon, (b) nitrogen and (c,d) oxygen K-edges, obtained with the MBES with retardation potentials of 200, 300, 380, 450 V respectively. The ionization photon energy was 560 eV in all cases. Red lines correspond to predicted theoretical data calculated by P. Kolorenč.

5.5.2 Resonant processes

Much richer information can be extracted from looking at resonant data, in this case near the 1s oxygen-edge. The XFEL photon energy was scanned from 527 to 546 eV by tuning the XFEL to 539 eV and scanning the electron-bunch energy a range of -25 to 25 MeV. Data from the MTOF with a retardation potential of 450 V was recorded for every shot, normalized to the pulse energy, and sorted against the photon energy measured with the x-ray grating spectrometer. The resulting electron spectra for different values of the photon energy are shown in figure 5.20.

Data taken in this way is typically very rich in features. In this case there are three different types of contributions that can be clearly observed in the figure:

- The photoelectron contribution, shown as two wide peaks crossing the plot diagonally in figure 5.20a, starting at 496 and 511 eV for a photon energy of 527 eV, and increasing at the same rate as the photon energy. The high-energy peak would correspond to a binding energy of 16 eV, being consistent with spectra in the literature for the sample [152].
- The non-resonant Auger contribution for photon energies above 539 eV corresponding to the oxygen K-edge for the sample. This contribution does not change with photon energy (orange to red profiles in figure 5.20b) and is identical to the spectrum shown in figure 5.19d. This value of 539 eV for the 1s ionization threshold was consistent with the values of 538.4 eV and 540 eV for each of the oxygen atoms of the molecule found in the literature [44]. As the cutoff for this region is located at 507 eV, this would yield a value for the double ionization potential (DIP) of

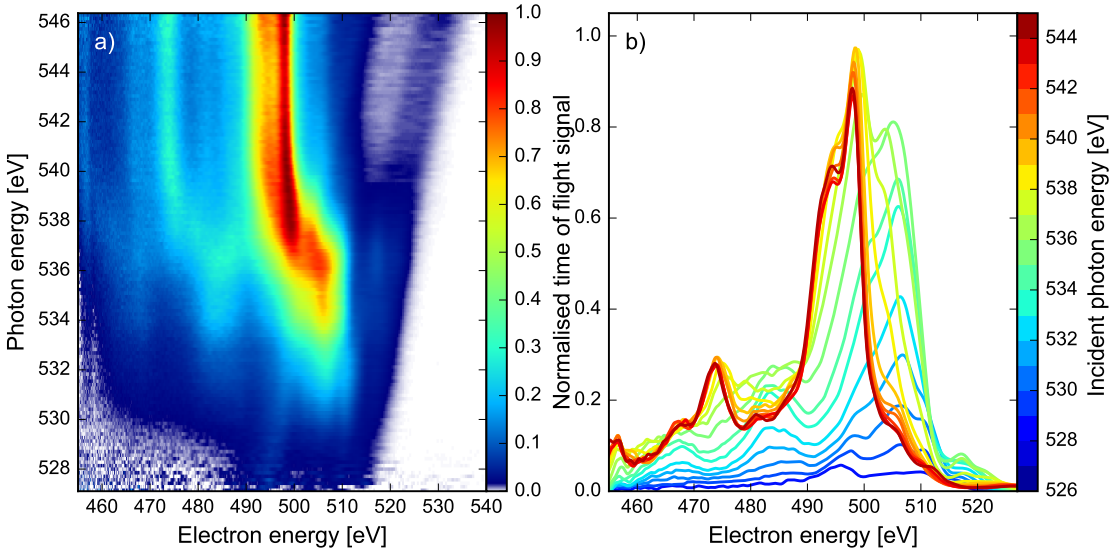


Figure 5.20: Auger electron spectra for a glycine sample at different photon energies below and across the oxygen K-edge. (a) High dynamic range colour map. (b) Profiles for selected photon energies. Figure extracted from [12].

32 eV, slightly different from the only theoretical value of 26.43 eV we found in the literature [153].

- The rest of the signal for photon energies below 539 eV corresponds to processes where the molecule is not core-ionized, but core-excited, leaving the molecule in a high-energy neutral state. After Auger emission, with a cutoff at 525 eV for the participator Auger processes, the molecule is left in a cationic state. The value of the cutoff is consistent with those found in the literature for the possible excited stages: 525 eV for the π^* resonance of the oxygen atoms, and 527 eV for the 3p Rydberg resonance [44, 154].

More information about the calculated Auger spectrum data can be found in reference [12].

5.5.3 Possible water contamination

After revisiting the data for this thesis, concerns about a possible water contamination of the sample were raised. These concerns were based first on the similarity noticed between Auger electron spectra for water found in the literature and the experimental Auger spectra measured in the sample. Figure 5.21 shows that, in fact, the experimental data is much more similar to the water spectrum obtained from the literature than to the calculated Auger spectrum. However, this evidence alone does not provide proof of contamination, as there are always similarities in the Auger spectra of different oxygen-containing molecules. Also the energy-dependent collection efficiency of the magnetic bottle, which could cause significant distortions in the relative amplitude of electron

peaks, could not be calibrated. In order to explore this issue, we tried to find more evidence of water contamination in our data.

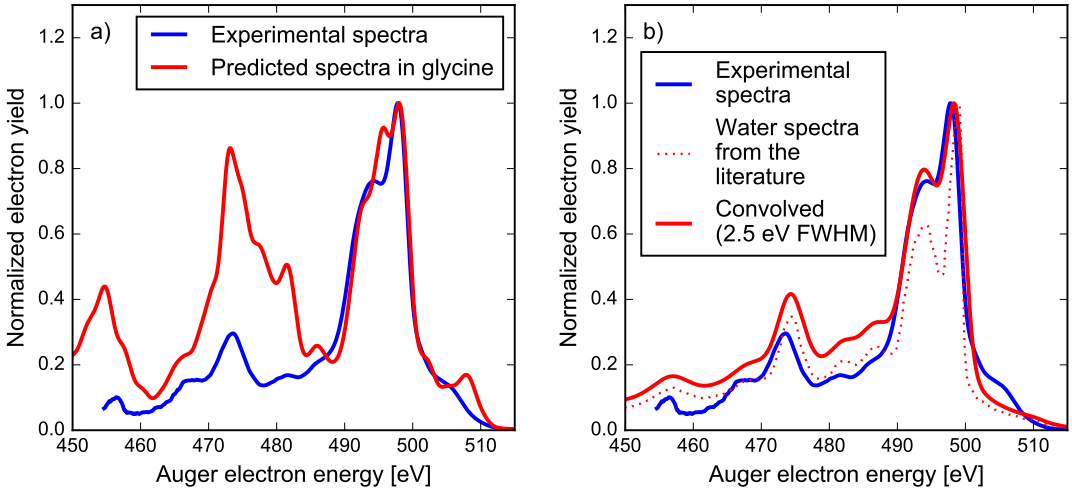


Figure 5.21: Experimental Auger electron data compared to (a) the calculated oxygen-edge Auger spectrum in glycine and (b) experimental oxygen-edge Auger spectrum in water obtained from reference [155].

First of all we examined the records of the residual gas analyser (RGA) monitoring the background gases in the chamber. Around the time when the static data was taken, the RGA data showed that none of the background gases had a background pressure of more than $4 \cdot 10^{-10}$ mbar (figure 5.22). While water appears indeed as the background gas with the highest pressure in figure 5.22, this pressure corresponds to a density of approximately 10^7 molecules per cubic centimetre, and it should not be able to compete with the expected $10^{10} - 10^{11}$ glycine molecules per cubic centimetre near the tip of the oven.

Additionally, we looked at the ion spectra measured with the sample at the interaction region with the same magnetic bottle just before recording the static Auger electron data (figure 5.23). While this data shows evidence of the existence of some ions that could potentially originate from water (H_2O^+,OH^+,O^+,O^{2+}), these ions could also originate from glycine (NH_2-CH_2-COOH). On the other hand, the presence of many glycine fragment ions containing carbon and nitrogen, and carbon and oxygen together, including the COOH peak at 45 atomic mass units, suggest a significant presence of glycine compared to water at the interaction region.

However, the final piece of evidence supporting water contamination hypothesis is the photoelectron spectra. As already shown in figure 5.20, the photoelectron spectra of the sample consists of two main photoelectron peaks. These two photoelectron peaks have binding energies of approximately -33 and -15 eV (blue curves in figure 5.24) with a relative amplitude of 2.5 approximately favouring the first peak. Figure 5.24 shows that indeed the water photoelectron spectra is much more similar to the experimental signal

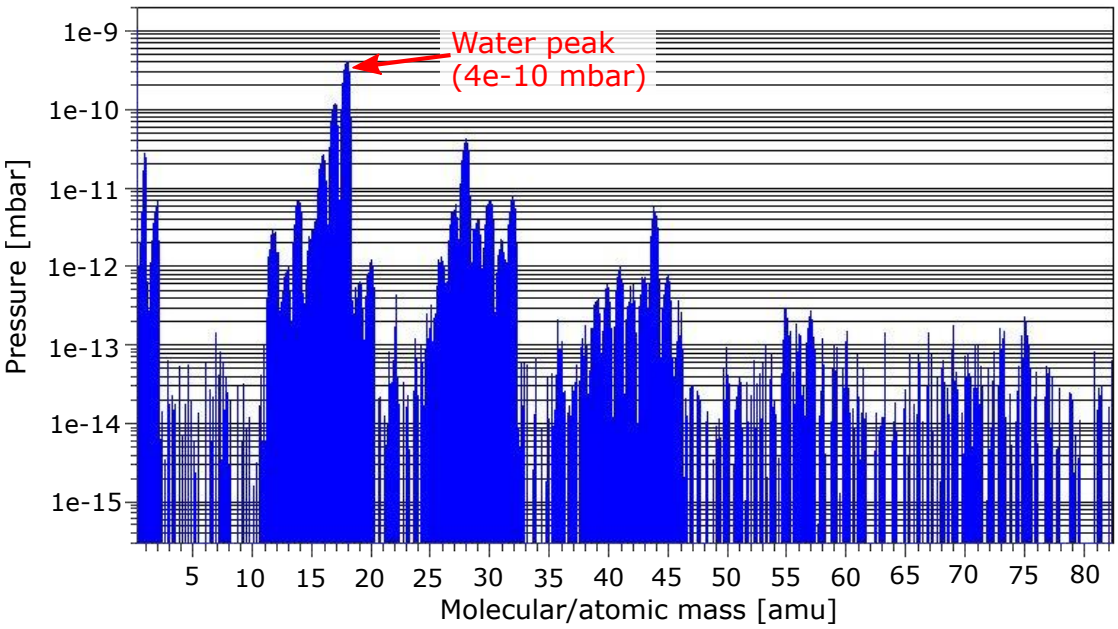


Figure 5.22: Signal from the residual gas analyser during the experiment. The dominant background gas with a molecular mass of 18 amu is indeed water. However, the low background pressure of $4 \cdot 10^{-10}$ mbar, corresponding to a density of 10^7 molecules/cm^3, should be negligible compared to the expected 10^{10} molecules/cm^3 for the glycine sample.

than the glycine photoelectron spectra found in the literature. This lack of similarity was not noticed in the original analysis as the literature data that we had found at the time only extended down to a binding energy of -25 eV, and there was agreement within that region. However it became obvious after the extended glycine photoelectron spectrum was found in the literature [156].

All things considered we are inclined to believe the presence of water contamination in our sample, even though we cannot explain the apparent inconsistencies between the ion data (suggesting higher glycine densities) and the electron data (suggesting higher water densities). A possible source for the water contamination at the interaction region could be water evaporation from the sample in the oven. This would be consistent with the low RGA water levels detected in the background, which are measured far from the interaction region. However, the glycine sample was delivered from the manufacturer in a sealed container, was not left exposed to the atmosphere for more than a few minutes and was kept in vacuum and above 100 degrees for three days before the data was taken. This should have effectively removed all water absorbed by the glycine sample.

Another possible explanation for the differences between the ion and the electron signal measured at the interaction region is that the exact size and location of the volume from where the magnetic bottle extracts the particles are different in the electron and in the ion mode. While the oven was aligned by optimizing the ion signal, it could be that the detector in the electron mode was receiving electrons from a different region where the fraction of water molecules compared to glycine molecules was much larger.

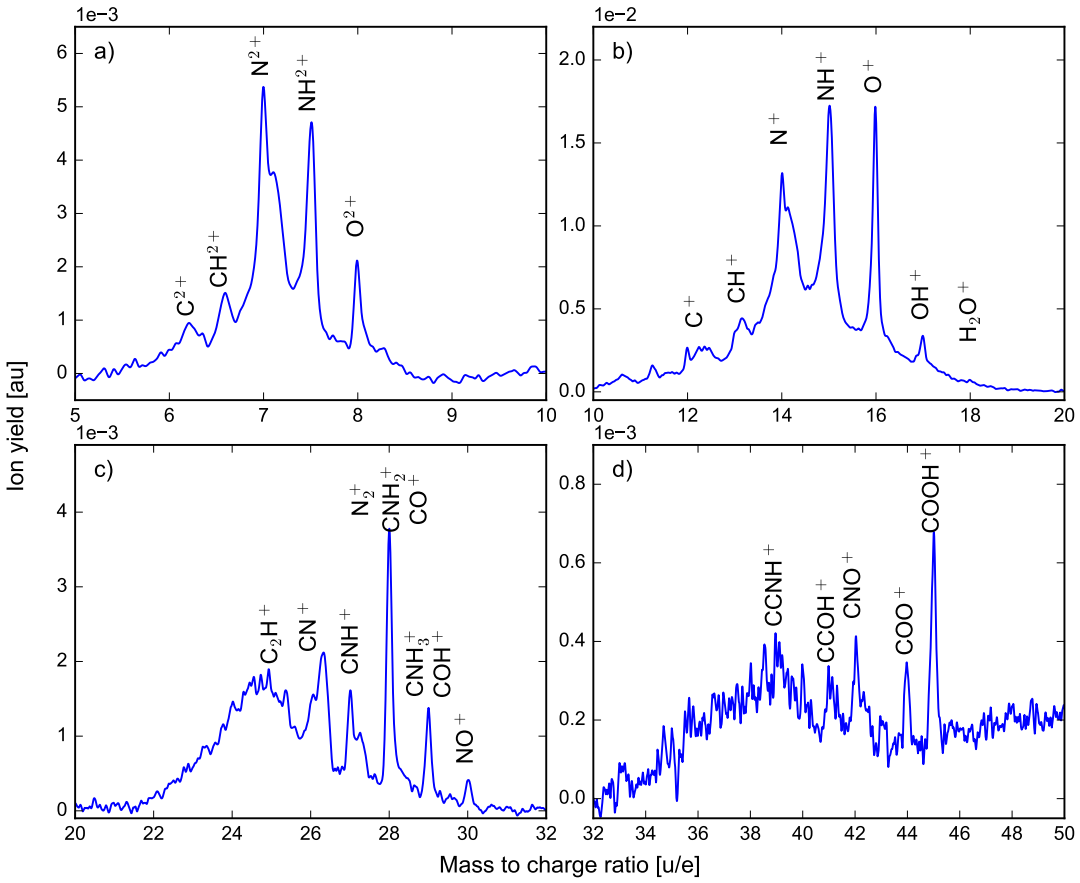


Figure 5.23: Ion spectra for different mass-to-charge intervals (a, b, c, d) measured at the interaction region using 560 eV photon energy to induce the fragmentation. This data was obtained minutes before recording the Auger-electron static data.

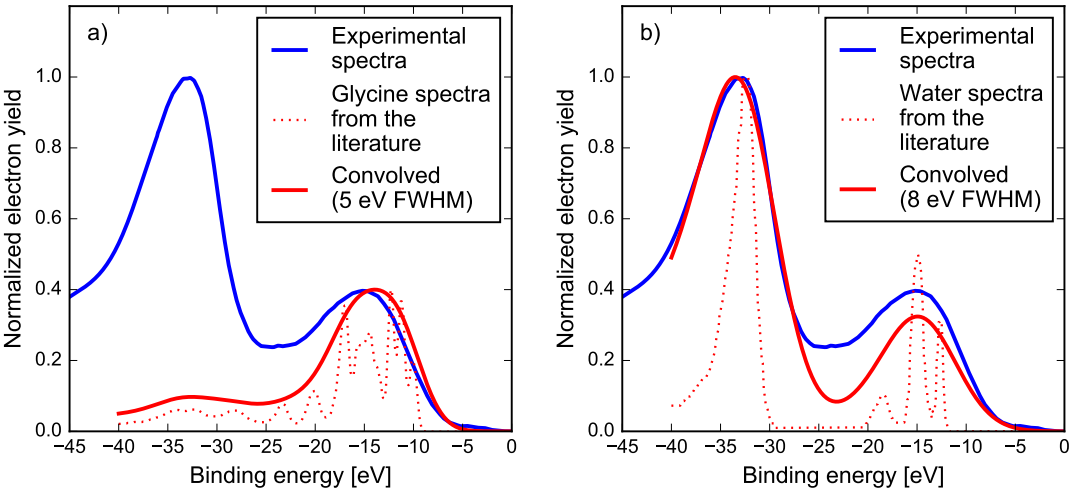


Figure 5.24: Experimental photoelectron data obtained at 520 eV compared to: (a) the experimental photoelectron spectra of glycine from reference [156] (photon energy: 99 eV) and (b) the experimental photoelectron spectra of water obtained from reference [144](photon energy: 1487 eV).

However, the Auger spectra measured for the sample also contained carbon-edge and nitrogen-edge Auger electrons in similar amounts (figure 5.10) that could not have been generated in water, proving that the magnetic bottle in electron mode was definitely measuring glycine electrons.

Details about the source of this inconsistency are still being discussed with the experts in the magnetic bottle detector, and a final conclusion still needs to be reached.

5.6 Searching for a time-resolved signal in O_2

As mentioned in section 5.1 this experiment originally aimed to measure non-trivial electron dynamics in glycine at the carbon K-edge. However, low glycine densities together with a reduced data collection time due to XFEL tuning difficulties, did not allow us to perform the experiment. As an alternative, we attempted to measure a hole opening process in O_2, in particular in an inner-valence orbital with a binding energy of 24.2 eV, making use of a 519.8 eV and 504.8 eV photon energies for the pump and probe pulses respectively. As O_2 is a small molecule we were expecting the hole to be static, not presenting any non-trivial dynamics in the hole survival probability. Accordingly, we just attempted to measure the hole opening process itself showing as a step function.

As shown in figure 5.25, the expected signal in this case is much smaller than what we showed in figure 5.4 for isopropanol. Adding the problem of the XFEL fluctuations, finding such a low signal becomes a difficult (sometimes impossible) task that requires techniques such as data binning or background subtraction.

While this particular attempt to measure few-fs time-resolved dynamics in O_2 was also unsuccessful, many aspects of the analysis may be useful for future experiments.

5.6.1 Data binning

The best way to overcome the limitations of the XFEL fluctuations is to record as much information as possible about the x-ray properties for each XFEL event. In this particular experiment, after performing the different calibrations described in the previous sections, we had access to the information shown in figure 5.26 for every single event. The spectral information is obtained from the x-ray spectrometer profile as in figure 5.26b. In this profile we can clearly identify the contribution from each of the x-ray pulses, and fit a double Gaussian curve. From the centre of the fitted Gaussians we can obtain the central photon energy of each pulse, and by comparing the area of each of the Gaussians (A_1, A_2) we can obtain the relative pulse energy of the pulses. This information can be used together with the total x-ray energy E_{gd} measured by the gas detector to calculate the absolute pulse energy for each of the pulses (P_1, P_2),

$$P_i = \frac{A_i}{A_1 + A_2} E_{\mathrm{gd}}. \tag{5.30}$$

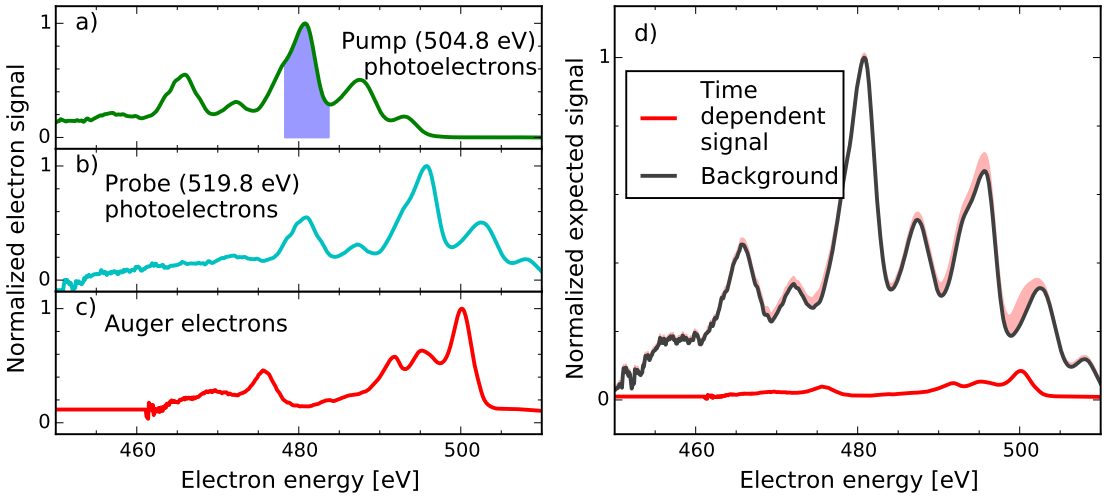


Figure 5.25: Three different electronic signals are measured in the energy range of interest. (a) Photoelectrons produced by valence ionization from the pump, with the photoelectrons ionized from the orbital of interest highlighted in blue. (b) Photoelectrons produced by valence ionization from the probe. (c) Auger electrons produced by the two-photon process presenting the time-dependent dynamics. (d) Expected signal at the detector. The black solid line corresponds to the photoelectron background signal, while the red solid profile and the red highlighted region corresponds to the expected maximum Auger signal assuming a perfect resonant absorption of the probe. The data shown corresponds to experimental data measured during the experiment.

On the other hand, the time delay information can be obtained from the individual single-shot XTCAV images (figure 5.26a) as described in section 5.4.4.

The importance of monitoring these parameters can be emphasized by looking at their distributions. Figure 5.27 shows histograms for the values of the variables obtained for $3.3 \cdot 10^5$ events during a period of approximately 10 hours. It is worth noting that this is not a measurement of the fluctuations of the XFEL alone, as many changes were applied to the configuration of the machine to make coarse adjustments to the time delay and corrections to the photon energy. In this case it was critical to keep the photon energy of the probe pulse close to the target value of 519.4 eV, so we applied XFEL tuning corrections whenever the photon energy moved away from that point, leaving us with a smooth Gaussian distribution centred at around 520.5 eV (figure 5.27b). The small difference between the target value and the central value was due to the refined spectrometer calibration achieved in the post analysis stages, however, the large bandwidth of the fluctuations showing a FWHM of approximately 4 eV was enough to allow for small inaccuracies in the calibration. On the other hand, the photon energy of the pump pulse did not require to be fixed to a particular value, so less attention was put to keep its distribution centered around a value. As a result the photon energy distribution for the pump pulse is more structure due to systematic changes (figure 5.27a). Similarly we recorded most of our data near 0 time delay to see the hole opening

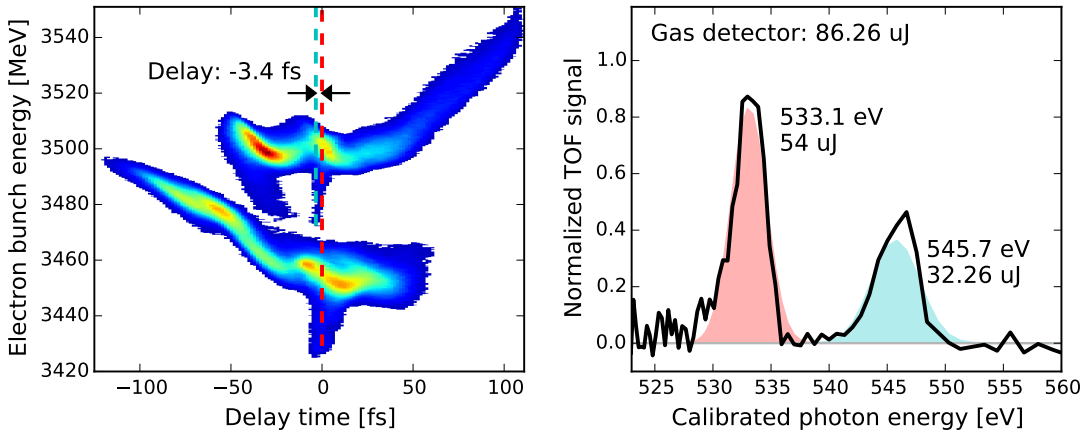


Figure 5.26: Diagnostics measured for each individual event. (a) XTCAV image used to extract the delay. (b) X-ray spectrometer used to extract the spectral shape, and in particular the photon energy and relative x-ray energy for each pulse. The absolute energies can be obtained using the total x-ray energy measured with the gas detector.

(figure 5.27f), but also recorded data for longer time delays centred at around 15 fs to measure any potential dynamics of the hole created. While scanning the delay in finer steps would have yielded a more uniform distribution of delays, adjusting the delay while keeping the photon energy fixed was a very time consuming process.

In the case of the total x-ray energy obtained from the gas detector (figure 5.27c) the distribution is also similar to a Gaussian peaking at 5 µJ, but with some asymmetry due to the fact that the total x-ray energy cannot be negative. Looking at the distributions of the individual pulses (figure 5.27d-e) we observe a similar effect with a slightly more pronounced asymmetry, and surprisingly peaking below 2 µJ for each of the pulses, which may seem inconsistent with the 5 µJ obtained from the gas detector. The enhanced asymmetry, as well as the apparent inconsistency in peak power, are signs of negative correlation between the pulse energies of the pump and the probe pulses. In general it is important to remember that XFEL emission is a very complicated process and that whenever a fluctuation or change happens it is very likely to affect more than one property simultaneously. We will see in chapter 7 that we may take advantage of this entanglement, however for now, we will just discuss some of the correlations existing in this dataset (figure 5.28).

Figure 5.28a shows a significant positive correlation between the photon energy of the pulses. The origin of this correlation is related to global fluctuations in the electron-beam energy: as the energy of both bunches increase, the photon energy of both pulses increase. This effect is partially blurred due to the adjustments performed to keep the energy of probe pulse fixed.

The correlation plot between the pulse energies (figure 5.28b) shows that, as expected, they are not independent. First we can observe unpopulated regions in the low-energy areas for each of the pulses. This is an artifact introduced by the fitting process as

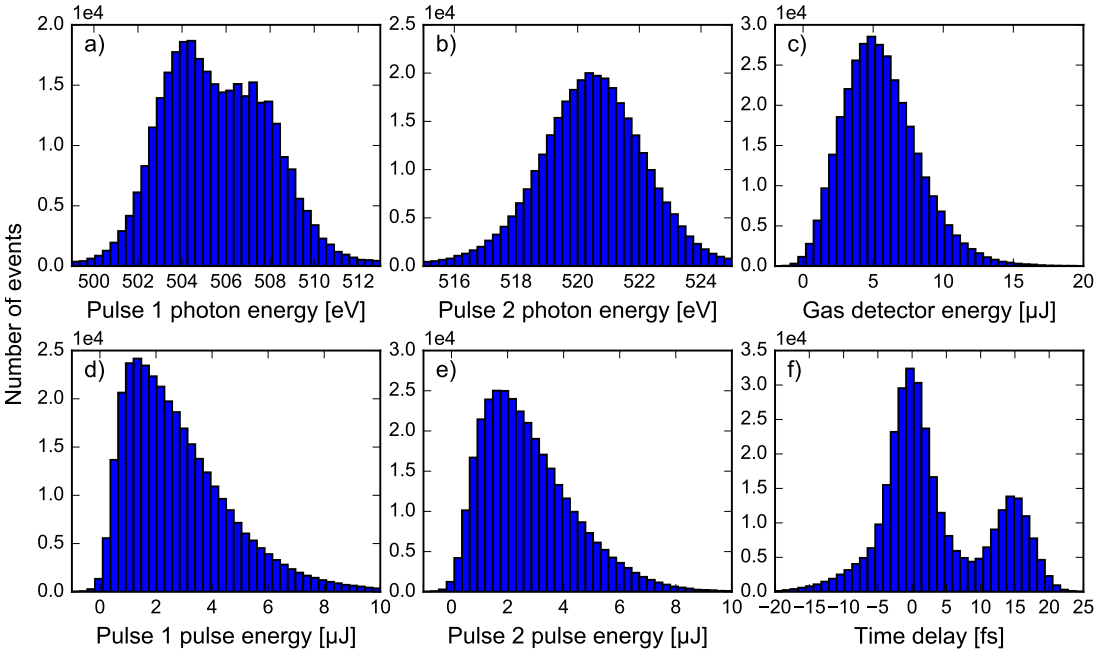


Figure 5.27: Distributions of different x-ray properties for a dataset containing $3.3 \cdot 10^5$ events obtained within 10 hours in a double bunch configuration. Pulse 1 corresponds to the pump pulse, and pulse 2 corresponds to the probe pulse.

typically when one of the pulses is much more energetic than the other the double Gaussian fit does not work well and the event is discarded, leaving this triangle-shaped distribution. On the other hand, looking at the regions inside the triangle, we can observe that the contour lines present signs of negative correlations of the type $P_2 = a - P_1$, showing that when one of the pulses lases very strongly the other one tends to be much weaker. We suspect this is due to having different optimal conditions for the lasing of each of the bunches that cannot be met for both of them simultaneously.

Additionally, each of the individual pulses lases more efficiently around a particular photon energy. This effect can be observed very clearly for the probe pulse in figure 5.28d, and it is also present for the pump pulse (figure 5.28c), although in this case it seems to peak at more than one photon energy, due to the uncontrolled changes in its central photon energy.

Finally, when looking at the correlation between the delay and the photon energy of the probe pulse (figure 5.28f) we observe that it is essentially non-existent, as care was taken to stabilize the photon energy of this pulse when changing the delay. However, the correlation of the delay with the photon energy of the pump pulse (figure 5.28e) is significant: the data taken at 15 fs is centred near 507 eV, while most of the data taken at 0 fs is centred at 504 eV. This effect is related to the fact that the time delay is adjusted by changing the position of the double-slotted foil, and this in turn has an effect on the energy separation between the pulses. Although this can be avoided by

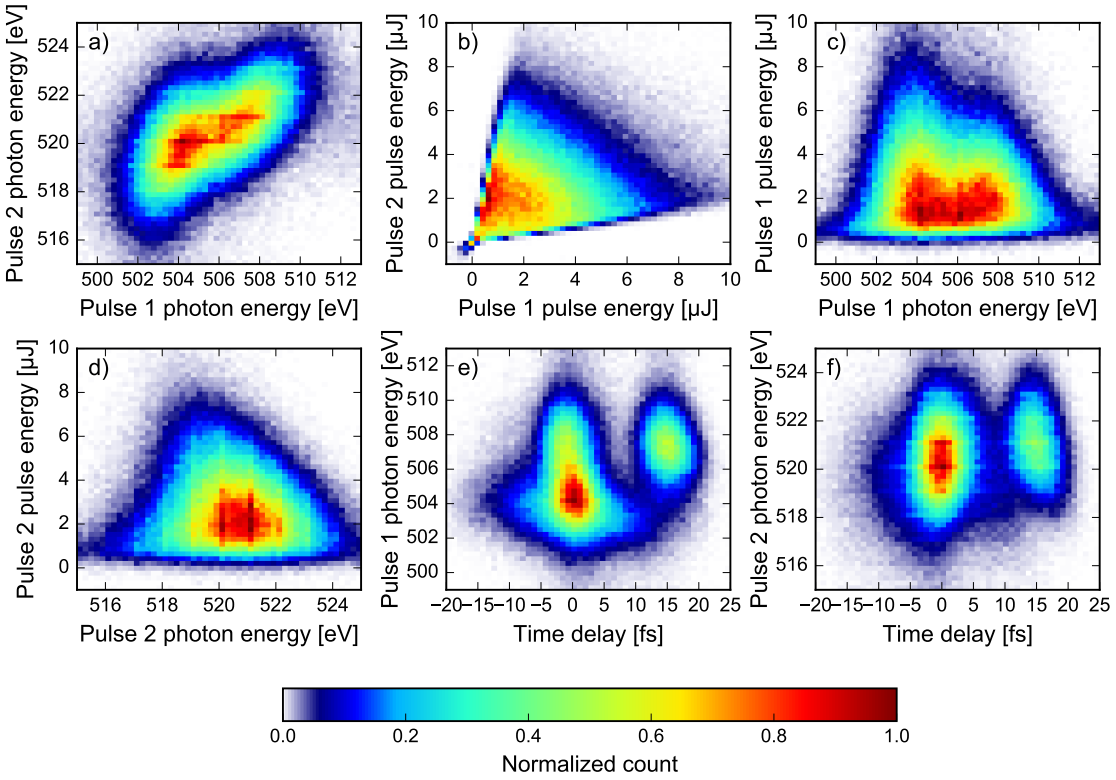


Figure 5.28: Correlations between different pairs of x-ray properties for the same dataset used to plot the distributions in figure 5.27. Pulse 1 corresponds to the pump pulse, and pulse 2 corresponds to the probe pulse. The figure shows complex correlations between the different pairs of variables, evidencing that the fluctuations are not independent. These correlations need to be taking into account together with the histograms to determine when the number of recorded events is sufficient.

other means, it involves a much more delicate and risky retuning when changing the delay, and was avoided for most of the 10 hours when the dataset was recorded.

All these effects highlight the importance of not only monitoring the distribution of the x-ray properties, but also monitoring the correlations in order to decide when enough data has been accumulated for a given experiment. Otherwise, for example, the statistics of the time delay may look high-enough for all delays of interest, and the statistics of the photon energy may similarly look sufficient, but in reality we lack certain photon energies for certain values of the delay.

Figure 5.29 shows an example of binning electron time-of-flight data in 5 dimensions (time-delay, photon energy for the pump, photon energy for the probe, pulse energy for the pump and pulse energy for the probe) to attempt to see the expected time-dependent signal (figure 5.25d). In particular we sorted the data in 2 fs time-delay bins, keeping only the events with a probe photon energy between 519 and 521 eV and then applied different limits on the other variables. The first row (figures 5.29a-c) corresponds to applying a very loose limit to the photon energy of the pump (500 eV - 507 eV) and no limits to the pulse energies. Figures 5.29a shows the average electron time-of-flight

electron spectra for the events inside each delay bin. We can definitely observe some apparent time delay effects, both in the amplitude and the position of the different photo electron peaks. However, if we look at the average pump and probe energies as a function of the delay (figure 5.29b) we observe strong correlations for the data across different temporal bins that would account for the changes in amplitude observed in figure 5.29a. This can also be observed by looking at the average x-ray spectrometer profiles as function of the delay (figure 5.29c), which in this case not only manifest a clear systematic change in amplitude, but also a systematic change in the photon energy of the probe, which accounts for the changes in the positions of the photoelectron peaks.

To avoid these effects we can tighten the photon energy limits for the pump (502 eV - 505 eV), and introduce some limits in the pulse energies of the pump (0.92 µJ - 3.7 µJ) and the probe (1.4 µJ - 3.7 µJ) as in figures 5.29d-f. While this removes most of the problems related to the changes in the photon energy of the pump, it still contains some delay-dependent changes in the pulse energy of the pump and the probe, with the corresponding modulation showing in the electron data. However, in the process we have already filtered down from an initial 119000 events (8500 per time bin) to 16800 events (1200 per time bin), greatly decreasing our signal-to-noise.

We can attempt to remove more of the fake time-dependent effects by applying even tighter limits to the pulse energies of the pump (0.92 µJ - 1.4 µJ) and the probe (1.6 µJ - 2.1 µJ) as in figures 5.29g-i. This effectively removes the pulse energy changes (figures 5.29h-i), but leaves only a total of 770 events (55 per time bin), which is not enough to see any signal above the noise.

While binning is a powerful technique, it becomes very difficult to use when sorting data as a function of more than one or two variables for the current low-repetition-rate XFELs. This is informally known as the *curse of dimensionality* since the number of possible bins grows exponentially with the number of dimensions, and so the amount of data needs to grow similarly to keep a constant number of events per bin. We will see another example of binning in section 6.1.

5.6.2 Background subtraction

In the previous section we binned our data based on delay, pulse energy, and photon energy, however, from a strict point of view it is only the photon energy of the probe that needs to be fixed to be able to probe the process, and the time delay that needs to be binned to show time-dependent effects. Considering a fixed photon energy for the probe pulse ($E_2 = 519.4$ eV) and a fixed value of the delay, the total expected signal for a given event with fluctuations in the pump and probe pulse energies (P_1 and P_2) and pump photon energy (E_1) has the following contributions in the first approximation,

$$\text{Signal}(P_1, P_2, E_1) = \text{Only}_{\text{pump}}(P_1, E_1) + \text{Only}_{\text{probe}}(P_2) + \text{Signal}_{\text{pump-probe}}(P_1, P_2). \quad (5.31)$$

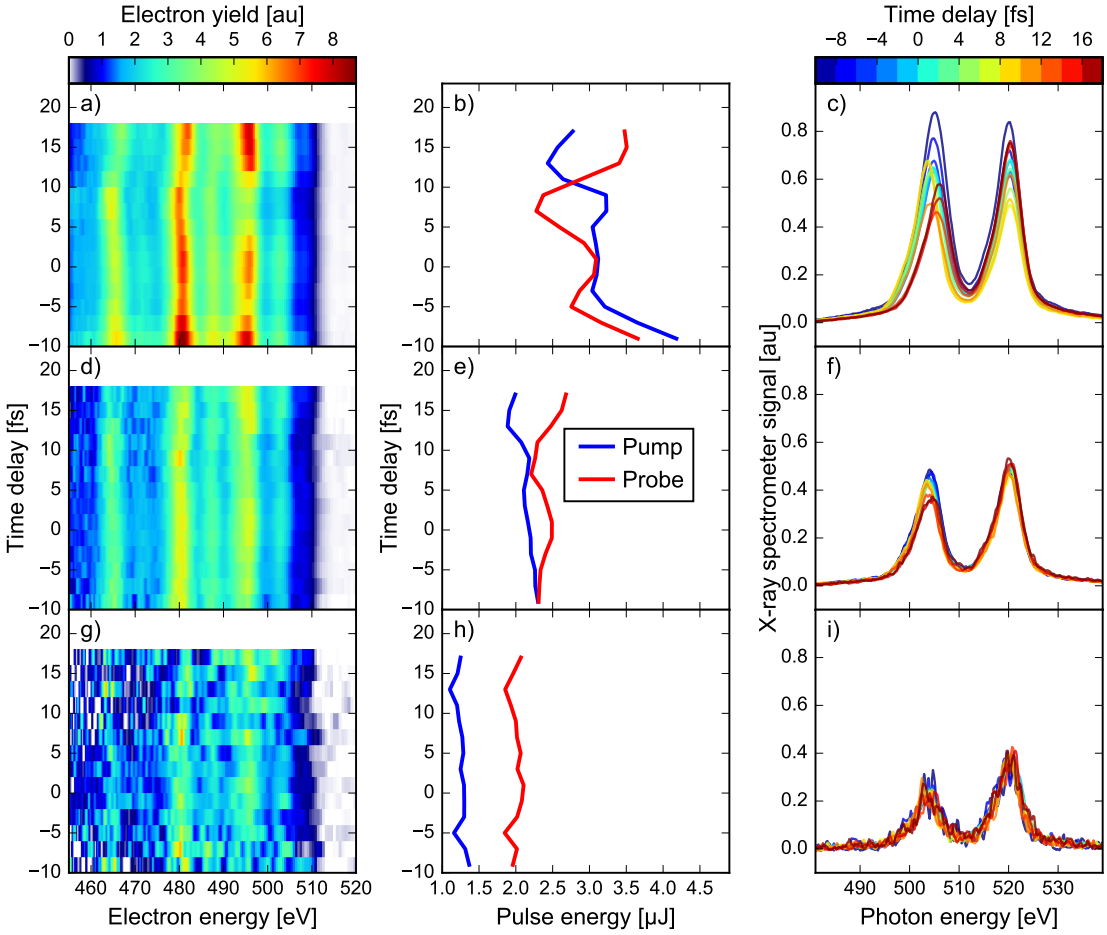


Figure 5.29: Using binning to analyse time-dependent data. As the size of the bin used to filter the data decreases (top to bottom), the systematic effects of the fluctuations are removed better from the data. However, the number of events within the bin also decreases, amplifying the noise. The first column corresponds to the average electron yield for events sorted in 2-fs delay bins. The second column shows the mean pulse energy of the pump and the probe pulses for the events contained in each delay bin. The third column shows the mean spectrometer profiles for each of the delay bins. Each of the rows shows results for bins with different limits in photon energy (P) and pulse energy (E) for each of the pulses used for filtering. First row: pump (500eV<P<507eV and 0μJ<E<18μJ), probe (519eV<P<521eV and 0μJ<E<18μJ). Second row: pump (502eV<P<505eV and 0.92μJ<E<3.7μJ), probe (519eV<P<521eV and 1.4μJ<E<3.7μJ). Third row: pump (502eV<P<505eV and 0.92μJ<E<1.4μJ), probe (519eV<P<521eV and 1.6μJ<E<2.1μJ). The mean number of events per time slice for each of the rows are 8500, 1200 and 55, respectively.

Since we are only interested in the last term corresponding to the pump-probe process, we can attempt to remove the first two terms from our signal. By ignoring two photon effects coming from individual pulses, we can rewrite the expression respectively as

$$\begin{aligned} \text{Signal}(P_1, P_2, E_1) = &\ \text{Photoelectrons}_{\text{pump}}(P_1, E_1) \\ &+ \text{Photoelectrons}_{\text{probe}}(P_2) \\ &+ \text{Auger}_{\text{pump-probe}}(P_1, P_2). \end{aligned} \tag{5.32}$$

Furthermore, as the shape of the photoelectron spectrum is very similar for the pump and probe pulses, except for an offset due to the difference in photon energy, the expression can be rewritten as

$$\begin{aligned} \text{Signal}(P_1, P_2, E_1) = &\ \text{Photoelectron}_{\text{spectrum}}(E - E_1) \propto P_1 \\ &+ \text{Photoelectrons}_{\text{spectrum}}(E - E_2) \propto P_2 \\ &+ \text{Auger}_{\text{spectrum}} \propto P_1 \cdot P_2. \end{aligned} \tag{5.33}$$

This makes the subtraction of the first two terms easier, as they are linear in the individual pulse energies, allowing us to scale any photoelectron references using the pulse energy information. Once the first two terms are removed, it should be as easy as normalizing the remaining pump-probe signal on the pump and probe pulse energies to obtain a fluctuation-free signal. This signal could in principle be averaged together for all the events within each delay bin by filtering only on the photon energy of the probe and regardless of the variation in the pump photon energy or the variation in the the pump or probe pulse energies.

The first step to remove the background signal is to obtain single-pulse photoelectron references. In this particular experiment, due to the wild fluctuations of the pulse energy, about 50% of the events contained a single pulse (either the pump or the probe), so no additional data was recorded for this purpose. We identified these events using the x-ray grating spectrometer and used them to build a pool of references with single-pulse photoelectron spectra for different photon energies (figure 5.30). We defined 200 photon-energy bins between 500 eV and 524 eV and sorted all of the single-pulse events to their corresponding bin, with their spectral profile and electron yield profile normalized to the pulse energy. By averaging the electron yield profiles within each bin we obtained the single-pulse references (figure 5.30b), which correspond to shifted copies of the photoelectron spectrum. Simultaneously, we obtained the 200 averaged x-ray grating spectrometer profiles corresponding to each of the references (figure 5.30a) showing the change in corresponding photon energy, and a very uniform amplitude due to the pulse energy normalization. There is a clear modulation in the amplitude of the photoelectron spectra as a function of the photon energy that we believe to be a symptom

of the energy-dependent collection efficiency of the magnetic bottle used to detect the electrons. We also considered this effect to be a signature of a single-pulse two-photon Auger signal but this hypothesis was discarded as the signal did not show a quadratic dependence on the pulse energy.

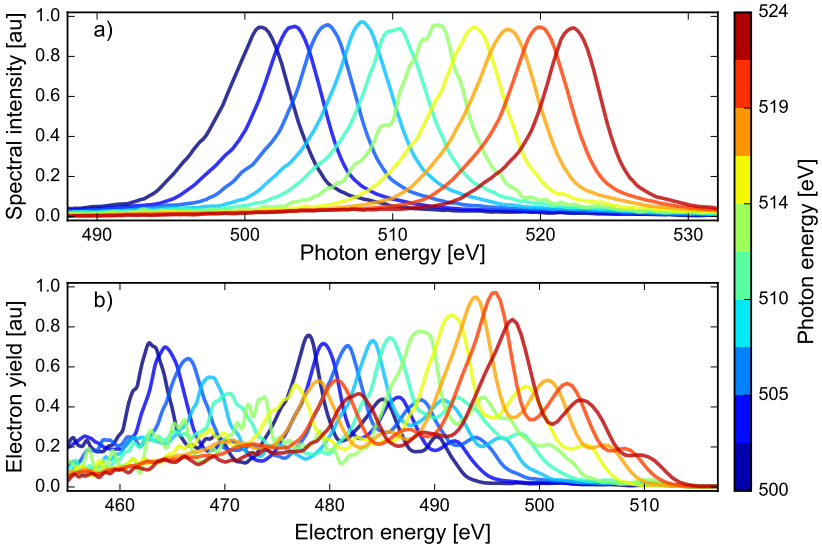


Figure 5.30: Electron yield references obtained from sorting single-pulse events at different photon energies. (a) Average normalized x-ray grating spectrometer profiles for the events contributing to each of the references. (b) Average normalized electron yield profiles created for the same groups of events in (a) used as references.

With the pool of normalized references available, we can build a single-pulse background for any individual x-ray shot just from knowing the corresponding values of the pulse energies and photon energies for the pump and the probe. For example, to build the background for the pump pulse, we find the reference in our pool with the closest photon energy value to the photon energy of the pump, and then multiply it by the pulse energy of the pump. In cases where we are averaging multiple events in the same bin, we just need to calculate the contribution to the background from each of the events in the bin. Figure 5.31 shows the result of applying this process to double-pulse data in a five-dimensional bin similar to those used in figure 5.29. Using the single-pulse references, we can interpret the average signal for the events within the double-pulse bin (green) for both the x-ray grating spectrometer (figure 5.29a) and the electron yield (figure 5.29b) in terms of the contribution of the pump pulse (red) and the probe pulse (blue), and show that in fact when we add them together (grey), the signal matches well the average signal in the bin.

If we now apply this process to the same time-dependent data used for figures 5.29a-c (also in figure 5.32a) we can independently remove the contribution of the photoelectrons created by the pump (figure 5.32d) to remove some of the electron energy changes with delay, and leave it as in figure 5.32b. Alternatively we can remove the probe

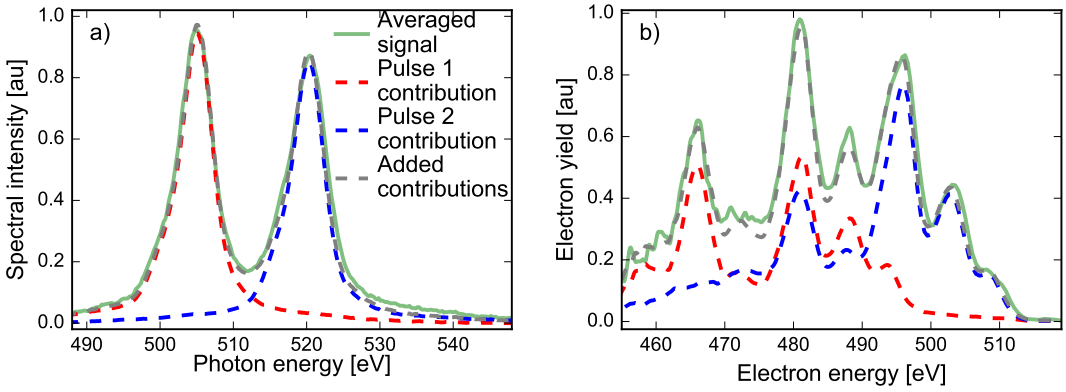


Figure 5.31: Average signals (solid green) for (a) the x-ray grating spectrometer and (b) the time-of-flight spectrometer obtained for all of the double-pulse events within a pulse energy, photon energy and delay bin. Background contributions from the low-energy pulse (dashed red) and from the high-energy pulse (dashed blue) are obtained by averaging multiple single-pulse references according to the distribution of the data in the bin. The background contribution of each of the pulses added together (dashed grey) matches the average double-pulse signal (green).

photoelectron background (figure 5.32e) resulting in figure 5.32c, or we can remove both photoelectron contributions simultaneously to obtain figure 5.32f.

The signal that is left in figure 5.32f could be normalized on both pump and probe pulse energies and should correspond to the time-dependent Auger signal, however it does not match well the Auger spectrum expected from the literature (figure 5.9a). Further analysis showed that the signal appearing in figure 5.32f is due to artifacts in the background references due to having extracted the single-pulse references from double-pulse data. Even though it is not visible in the single-shot events, we can observe really small tails in the blue and red profiles of figure 5.31a indicating that there is a small amount of "the other pulse" in the supposedly single-pulse data. Additionally, the single-pulse signals used to build the references may not be exactly linear on the pulse energy due to the presence of two-photon processes or non-linearities in the detectors. This could be solved by creating a bidimensional pool of references sorted in both pulse energy and photon energy, however again because of the *curse of dimensionality* this requires much more reference data. Many attempts were made to overcome these problems, but in the end the accuracy of the methods was never sufficient to extract without doubt the expected small time-dependent Auger contribution (figure 5.25) from the background.

We will present another example of reference subtraction in section 6.1.

5.7 Conclusions and outlook

While the work presented in this chapter did not lead to successfully measuring few-femtosecond time dynamics due to technical difficulties during the beamtime, it has contributed to set the basis that will allow performing such experiments in the future.

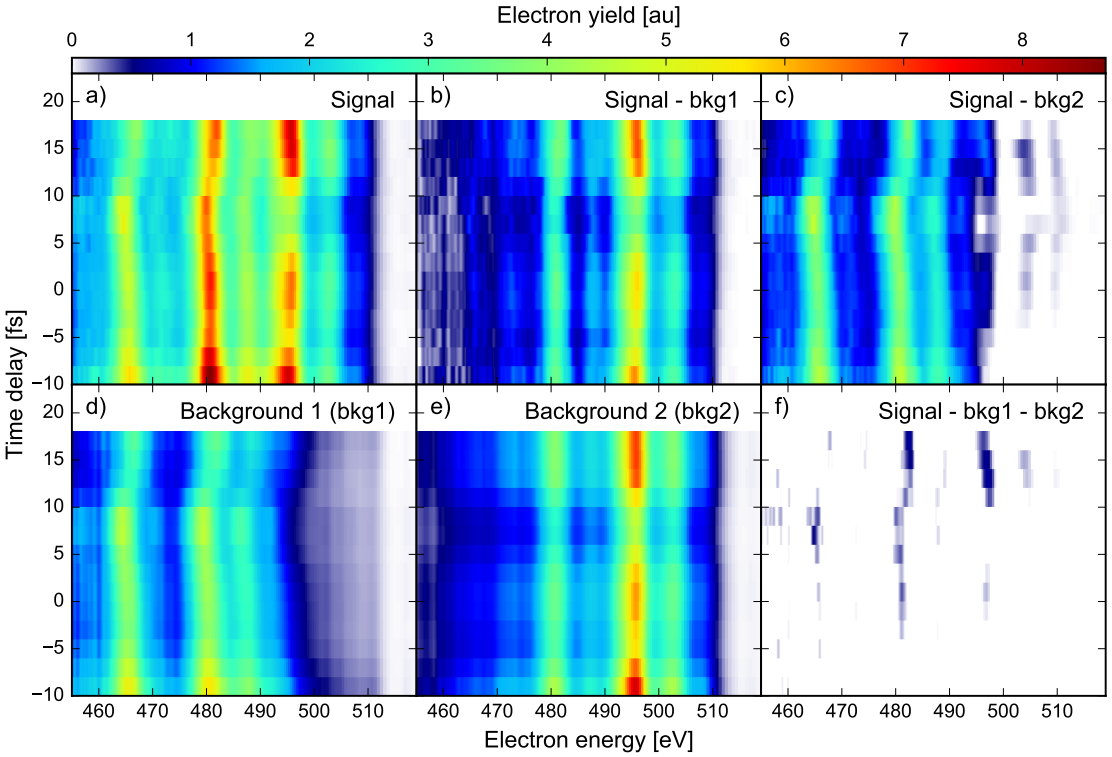


Figure 5.32: Background subtraction for the time-dependent data. Using the single-pulse references, accurate single-pulse backgrounds of the individual contributions of the pump and probe pulses can be created and subtracted from the time dependent signal to help revealing two-photon signals. (a) Delay-dependent electron yield obtained from binning the data (Same as 5.29a). (d) Calculated background corresponding to the low-energy pulse from single pulse data. (e) Calculated background for the high-energy pulse. (b) Time-dependent data from (a) after subtracting the single-pulse contribution of the low-energy pulse. (c) Time-dependent data from (a) after subtracting the single-pulse contribution of the high-energy pulse. (f) Time-dependent data from (a) after subtracting the single-pulse contribution of both the low-energy pulse and the high-energy pulse. However, the signal that is left does most likely not correspond to two-photon processes, but to inaccuracies in the single-pulse references.

Improvements to the standard calibration procedures of the magnetic time-of-flight spectrometers have been developed. New algorithms to analyse XTCAV images, and obtain temporal properties including the time delay have been established and made available to the LCLS user community. The importance of single-shot x-ray characterization has been identified as the way to deal with the large XFEL fluctuations in parameters such as time delay, photon energy and pulse energy, that are of critical importance for few-femtosecond experiments. Single-shot non-intrusive x-ray characterization has been shown to be possible by making use of an x-ray grating spectrometer, a gas detector and XTCAV combined. Furthermore, different analysis techniques, such as multidimensional data binning, and sophisticated reference subtraction have been developed to allow the extraction of small signals in future experiments.

Using all the evidence and information learned from this previous experiment, we have written a proposal (led by Dr. Thomas Barillot) to attempt to measure the electron dynamics in isopropanol included at the beginning of this chapter at LCLS. Isopropanol presents the great advantage of being in a liquid state at room temperature with a vapour pressure of approximately 50 mbar. This should allow us to achieve sample densities two or three orders of magnitude larger than for glycine, making it an ideal sample to measure the time dynamics. While the proposal has not been accepted for a beamtime yet, due to perhaps different priorities in the agenda that favour crystallography experiments, the feedback obtained from the reviewers of the proposal have been supportive of the feasibility of the experiment based on the evidence provided, and will most likely be accepted within the next proposal rounds. This should allow us to finally show for the first time that it is possible to measure few-femtosecond electron correlation dynamics using an XFEL.

Chapter 6

Ultrafast experiments on non-linear x-ray optics and timing

Following chapter 5, in this chapter the author will present the results of two further experiments relevant to ultrafast science. The first experiment will discuss some early inconclusive results of the development of an alternative method to generate more coherent few-fs pulses using XFELs, obtained by applying some of the techniques discussed in chapter 5 for the data analysis. The second experiment concentrates on obtaining accurate timing between x-ray pulses and optical pulses, assessing the performance and the limitations of current systems.

6.1 Resonant Raman scattering in an XFEL

In this experiment we attempted to measure for the first time stimulated resonant inelastic x-ray scattering (SRIXS) in a molecule using an XFEL. This type of scattering is considered one of the critical building blocks of non-linear spectroscopy, especially when applied to molecules. Also, as a lasing mechanism, it provides a way to use the XFEL as a pump to generate x-ray pulses with a shorter pulse duration and an enhanced coherence [64] compared to the XFEL emission.

The signals expected from these early experiments on SRIXS are several orders of magnitude smaller than the XFEL energy. This required developing and extensively applying a new binned-reference analysis technique. However, before going into details of SRIXS and the technique it is worth introducing one of its predecessors: self-stimulated x-ray lasing in neon.

6.1.1 Prior work: self-stimulated x-ray lasing and SRIXS in an atom

In order to achieve x-ray lasing (XRL), it is necessary to pump an electron transition of several hundreds or thousands of eV of energy to induce population inversion. This can be achieved by core-exciting or core-ionizing an atom or molecule, however this kind of

highly-energetic core-excited/ionized states usually have enough energy to decay through very fast (few fs) non-radiative processes like Auger decay. Taking advantage of the high brilliance of XFELs, Rohringer et al. [64] achieved efficient XRL in 2012, results that were reproduced many times after that [157], including in collaboration with the author [10].

The experimental setup of the original experiment [64] is shown in figure 6.1a. The free-electron laser (LCLS) was tuned to the proposed pump photon energy (960 eV, above the 870 eV neon K-edge in this case) generating 40 fs pulses that were tightly focused (1-2 µm) inside a gas cell. The light coming out of the gas cell was spectrally analysed using an x-ray grating spectrometer consisting of a slit, a curved grating and an x-ray CCD. Apart from the transmitted XFEL radiation, the spectrometer allowed the observation of an additional emission line at 849 eV (figure 6.1b).

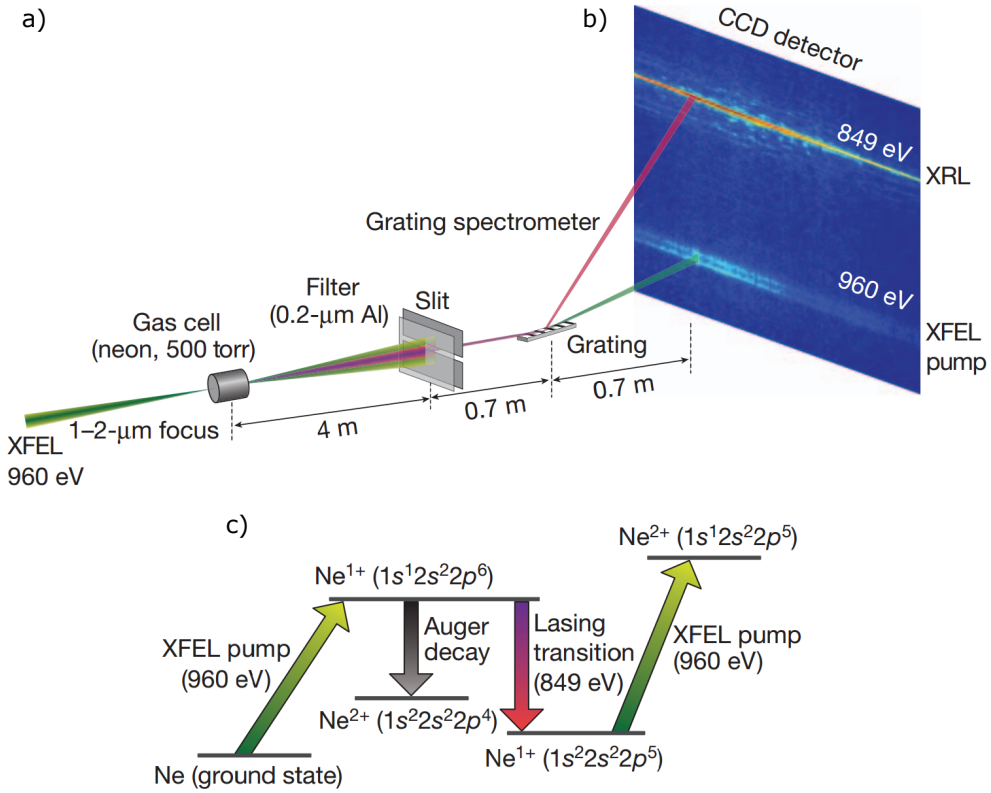


Figure 6.1: Experiment on self-stimulated x-ray lasing in neon. (a) Experimental setup. The XFEL SASE pulses tuned at 960 eV are focused into a high-density gas cell containing neon. (b) Spectrometer image showing the XFEL spectra (960 eV) with some SASE structure and the sharp neon lasing line (849 eV). The spatial axis shows the divergence of the beam after the focus. (c) Diagram of the process. Following pump ionization, the system can decay by Auger emission or by the lasing emission. Figures extracted from [64].

The process that generates this line is illustrated in figure 6.1c. The x-ray photons induce photoionization from the inner shell (1s) of neon, leaving the atoms in a short-lived excited state. This excited state has two different decay channels. The first channel

is a non-radiative Auger decay, that leaves the atom in a doubly ionized state. The second channel consists of the radiative spontaneous emission of a photon of 849 eV corresponding to a 2p to 1s decay. While the non-radiative decay is the dominant process, with a lifetime of only 2.4 fs, there is a small chance ($< 1.8\%$) that the atom decays through spontaneous emission. This small seed of emitted photons can then induce the emission of more photons of the same characteristics and with the same phase through stimulated emission in a process that can compete with the Auger decay. In this way a small seed of spontaneously emitted photons generated at the front of the gas cell travel together with the x-ray pulse undergoing exponential amplification through the gas by a factor of between 10^7 and 10^8 [157]. As a result of the process, known as self-stimulated XRL, the generated pulse presents full temporal coherence and a much shorter pulse duration of 5 fs [64].

Finding the XRL signal in this case was relatively straightforward as it appeared in a background-free area (figure 6.1b). The lack of background makes it easy to, for example, study the dependence of the signal with the photon energy as we did in one of the follow-up experiments [10]. In particular we scanned the energy of the XFEL across the neon K-edge (870 eV), to observe the appearance of the lasing line (figure 6.2b) for values above 870 eV.

Additionally, by zooming in the scanned image (figure 6.2c), we could also observe small signs of emission at slightly different wavelengths for incident photon energies below the neon K-edge. This emission cannot be a result of the previously described self-stimulated XRL, as the photon energies in that region do not allow 1s photoionization. Instead this signal corresponds to a different process: stimulated resonant electronic x-ray Raman scattering, also known as stimulated resonant inelastic x-ray scattering (SRIXS) [157, 158, 10]. This process is similar to x-ray lasing with a few key differences. The first difference is that the sample is core-excited instead of core-ionized. In the case of neon, the core excitation occurs into the 3p and 4p Rydberg states at 867.5 and 869.2 eV, respectively. The second key difference is that the process can occur even for incident photon energies that are slightly offset from the resonant transition, inducing a small offset in the energy of the emitted photons, characteristic of Raman processes. In the case of SRIXS in neon induced by SASE pulses consisting of multiple spectral spikes, this is observed as small shifts in the lasing line (figures 6.2a and 6.2c). Finally, a third key difference is that the amplification is seeded mainly by photons provided by the tail of the XFEL spectrum, which are approximately 10^4 to 10^5 larger in number than the spontaneously emitted photons at the onset of the amplification process [157].

6.1.2 Stimulated resonant x-ray Raman scattering in a molecule

Having achieved resonant inelastic scattering in neon, the next critical step was to demonstrate it in a molecule. However, the presence of multiple atoms increases the

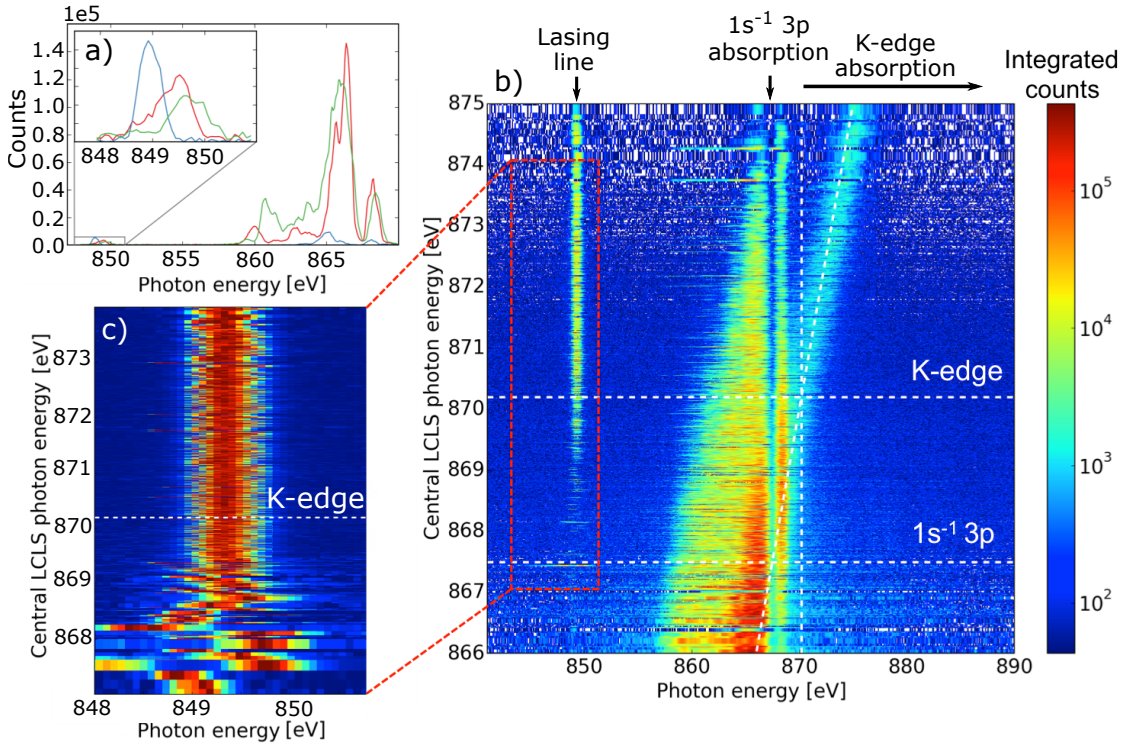


Figure 6.2: Lasing in neon as function of the photon energy. (a) Three sample spectra obtained for three single shots with a central photon energy of 868 eV and 1 mJ of pulse energy. A small contribution of stimulated signal is visible at 849 eV. (b) Single-shot spectra (horizontal axis) sorted as function of the incoming LCLS photon energy (vertical axis), in a logarithmic colourscale. A lasing line appears at 849 eV for central photon energies above the neon K-edge (870.2 eV). (c) A zoomed view of the normalized spectra of the lasing line shows signs of lasing below the K-edge, with slightly different photon energies, that would be consistent with SRIXS with the $1s^{-1}3p^{1}$ excited state. Figure extracted from [10].

.

number of competing channels, yielding much lower amplification factors. In particular, in the calculations performed for the CO molecule, the most optimistic experimental estimates predicted an amplification factor of less than 2 [159, 11, 10]. With such low amplification it is not possible to observe the amplification of the spontaneous emission within the dynamic range of our previous experimental setup and it is required to stimulate the transition by providing an external seed pulse.

For our first attempt in 2014 [10] at LCLS we used two independent SASE pulses (40-60 fs) generated in two different undulator sections from the same bunch, with a similar experimental setup to that of the neon experiment. To trigger the mechanism shown in figure 6.3a, a strong 0.5 - 1 mJ pump pulse was tuned to the CO π^* transition (534 eV) to excite the molecules and a weak Stokes pulse was tuned at the expected energy of the lasing transition (525 eV) to stimulate it. Under the applied experimental conditions, however, the predicted amplification factor was 1.05, mainly due to the long pulse durations and the large bandwidth of the SASE pump pulse, much larger than the width of the π^* transition. As a consequence it was impossible to observe any clear sign

of SRIXS in CO. In order to improve the amplification, a second experiment was carried out in July 2015, when the pulse durations were reduced to 25 fs and, most importantly, a self-seeded pulse (section 3.1.5) featuring a smaller bandwidth was used as pump. Even though the total energy of the pump pulse was reduced to a few hundreds of microjoules, the narrow bandwidth provided an increased pulse brilliance, theoretically allowing for a more efficient pumping and the expected total amplification of 2 mentioned earlier.

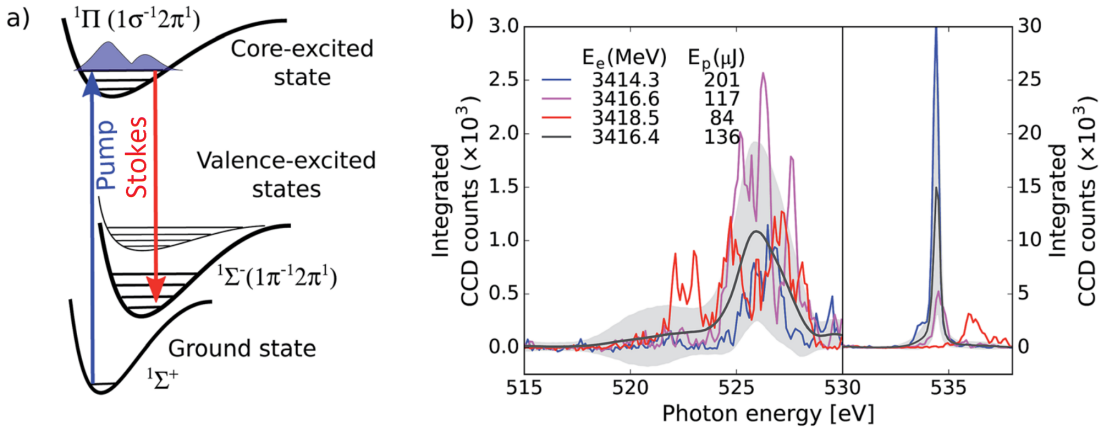


Figure 6.3: (a) Scheme of SRIXS in CO. (b) Sample single-shot spectra of the XFEL pulses used, for slightly different electron-beam energies (E_e) and total x-ray energies (E_p). A strong seeded pulse tuned to the π^* transition (534 eV) is emitted along a much weaker SASE pulse tuned at 525 eV to stimulate the emission. The solid grey line, and the the grey shaded area indicate the mean and the standard deviation of the spectra, respectively. Figure extracted from [11].

The limitations in this case were related to the particular fluctuations of the XFEL when working in the seeded mode. By looking at the correlation plot between the electron-beam energy and the total x-ray energy (figure 6.4a) we observe that it is only for a very particular electron-beam energy (3419 MeV) that the XFEL emits efficiently. This is a consequence of the self-seeding process: for electron energies that do not favour lasing at the wavelengths allowed by the monochromator (534 eV in this case), the seeding process does not work well, and normal SASE emission is produced. Indeed, looking at the dependence of the x-ray spectra as a function of the electron energy (figure 6.4b), the pump pulse (at energies above 530 eV) appears as a typical Gaussian SASE profile that increases in energy as the electron-beam energy increase for most electron-beam energies. It is only when the spectrum is well centred at 534 eV, that it shows a very narrow and strong spectral emission line, corresponding to the self-seeded process and accounting for the extra energy, in some cases of up to 0.25 mJ as shown in 6.4a. Figure 6.4b also shows the stokes pulse (at energies below 530 eV) which as expected behaves as a Gaussian SASE pulse, with the particularity that it presents a hole right at the region of interest (525 eV), which seems to be correlated to the efficient emission of the pump. In fact, looking at the dependence of the spectral shape with total x-ray energy (figure 6.4c), we observe that as the total x-ray energy increases providing

a stronger seeded pulse, the intensity of the stokes pulse decreases, and the hole in the spectrum of the Stokes pulse becomes deeper.

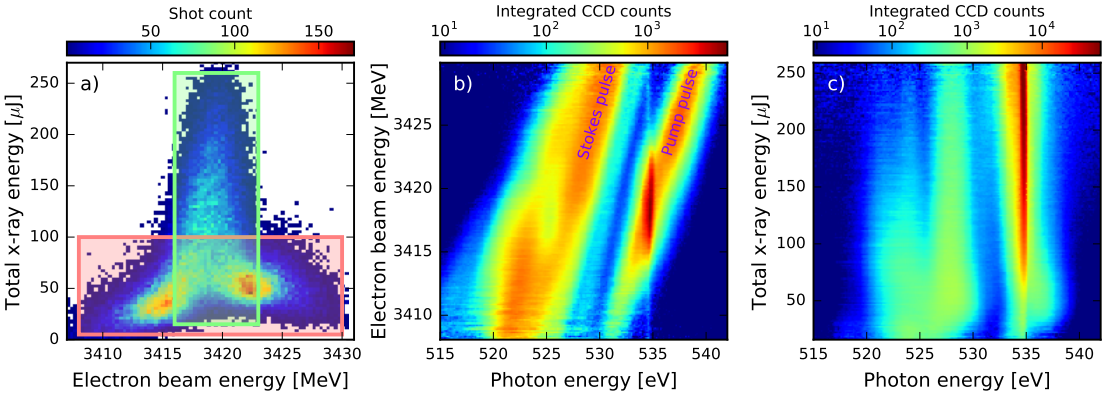


Figure 6.4: Characteristics of the XFEL emission in a double-pulse mode with a seeded pulse and a SASE pulse generated from a single bunch in different undulators. (a) Correlation plot between the electron-beam energy and the generated total x-ray energy. (b) Spectral shape as function of the electron-beam energy for the data within the pink rectangle in (a). (b) Spectral shape as function of the total x-ray energy for the data within the green rectangle in (a). Figure extracted from [11].

This situation is much less favourable than the background-free neon lasing, and the changes in the spectra shown in figure 6.4 require treating the XFEL radiation differently for any possible electron-energy x-ray-energy pairs of values. To solve this problem we recorded references without any gas in our cell for every possible experimental configuration, to use as a database of the fluctuations of the XFEL spectra. This allowed us to then, for every dataset taken with gas, choose bins in certain regions of the electron energy and x-ray energy, and build a custom reference of what the average incoming XFEL spectra should be for that bin. As an example figure 6.5 shows four averaged spectrometer profiles (coloured lines in figures 6.5b-e) for four different bins at significantly different intensities (figure 6.5a), along with the corresponding no-gas references in black which are considerably different between bins. The references have been scaled by a constant factor to account for the linear absorption due to valence ionization. This allows us to observe very clearly the effect of the strong π^* absorption at 534 eV, as well as other small regions with additional absorption. In particular, the transmission curves as function of the x-ray energy (figure 6.5f), obtained by dividing each of the spectra by the corresponding no-gas references, show an additional non-linear absorption that increases with the x-ray power in most regions of the spectra, which we believe correspond to absorption in the CO cations. It is worth noting, however, that in the region near 525-527 eV this effect is not observed, being potentially a symptom of amplification, compensating the extra non-linear absorption.

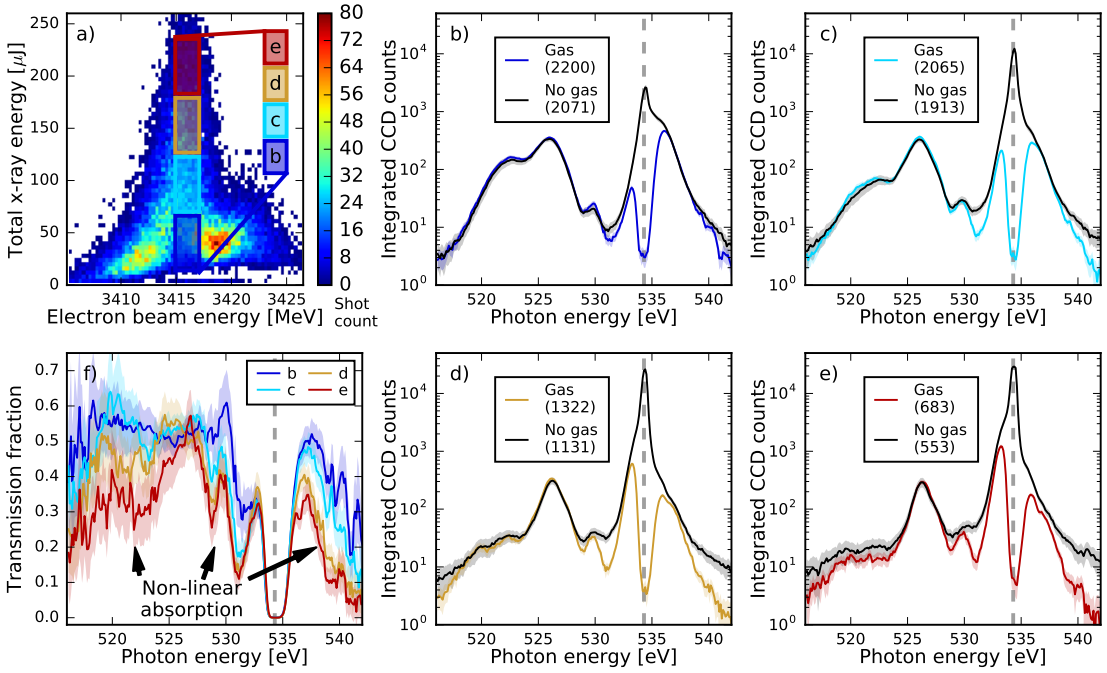


Figure 6.5: Attempt to find the SRIXS signal by looking at the intensity dependence. (a) Location of the different regions studied in the following figures b-e within the correlation plot. (b-e) Spectra obtained with 100 mbar of CO (blue, orange red) compared to a no-gas reference (black) scaled by a factor of 0.486 to roughly match the amplitude, compensating for the absorption due to single photon ionization. Numbers in brackets indicate the number of events in the bin. (f) CO transmission obtained by dividing the signal by the corresponding no-gas references at the different intensities. A strong absorption is observed at the π^* resonance around 534.3 eV (vertical dashed lines in figures b-f). Figure (f) also shows an overall increase of the absorption with the intensity probably corresponding to resonant states in the CO ions. The shaded regions correspond to 99% confidence intervals. Figure extracted from [11] and modified.

We attempted to find any changes in the shape of the spectra around the 525 eV region for a different dataset with increased statistics following a similar binning strategy (figure 6.6). In this case we looked only at bins at the highest available x-ray energies, as well as at a low x-ray energy region to use as a control bin. This however did not show any significant signs of change in the shape of the spectra due to the amplification (figure 6.6a-f) as the only differences lay within the error bars of our experimental data, and were not stronger than the differences observed in the control bin (figure 6.6g).

6.1.3 Conclusion

While we could not observe the expected SRIXS signal, the binned reference comparison method that we developed allowed us to determine unequivocally that the expected amplification of 2 was not occurring. In fact, based on the accuracy of the method, we expect that we could have observed differences down to 5-10%, proving the potential of the technique. There are however many reasons why the actual amplification could have

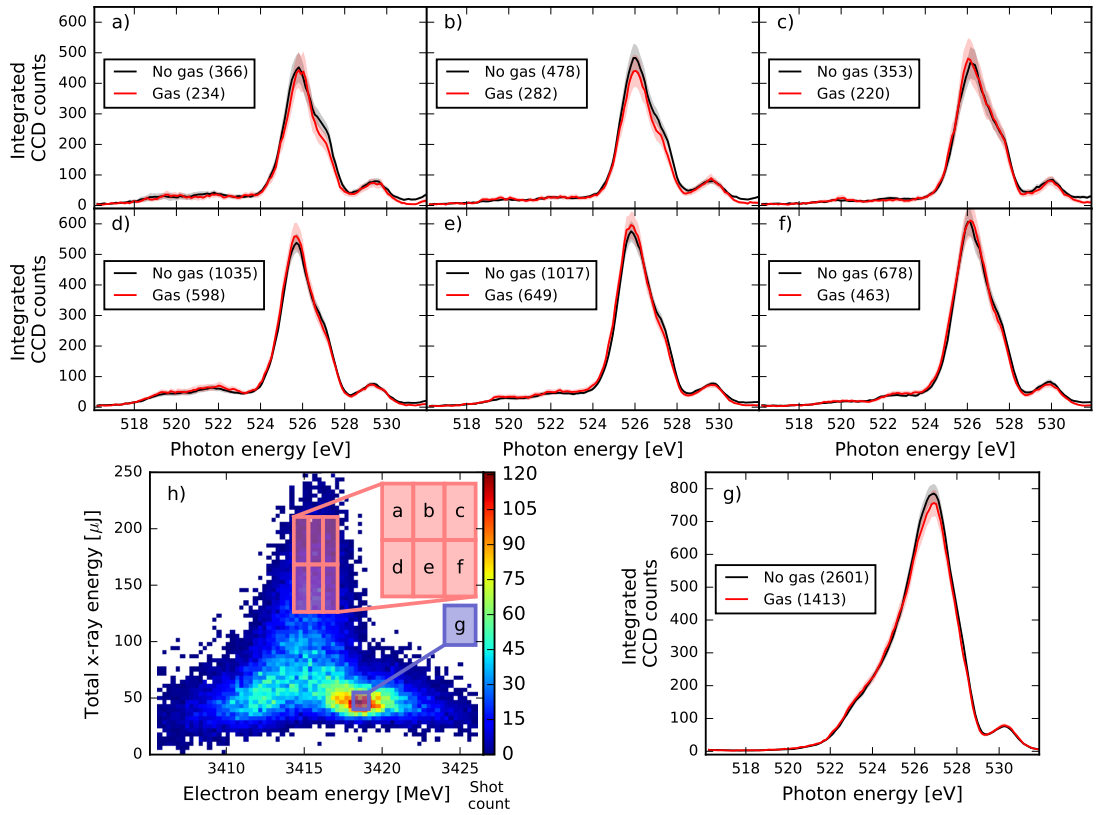


Figure 6.6: Attempt to find the SRIXS signal in different pulse-energy electron-beam-energy regions. (a-g) Experimental signal with 100 mbar of CO (red) near the region of interest (525 eV) compared to no-gas references (black) for different x-ray energy and electron-beam energy bins as shown in (h). The shaded regions correspond to 99% confidence intervals. The no-gas references were adjusted by a factor of 0.537 to match the amplitude of the signal, compensating for the absorption due to single photon ionization. Numbers in brackets indicate the number of events in the bin. Figure extracted from [11].

been even smaller than calculated including inaccuracies in the nominal focal dimensions, nominal reflectance of the focusing and transport mirrors, or spatial overlap. These parameters are typically measured only every few years at LCLS beamlines, as measuring each of the them is an experiment in itself. Further details and discussion about this experiment is out of the scope of this thesis and can be found in the relevant reference [11].

6.2 Synchronization between an optical laser and an XFEL at SACLA

So far we have mainly discussed how to perform time-resolved experiments using x-ray pulses for both the pump and the probe, but, as we already mentioned in section 5.6 there are many experiments that can benefit from using XFEL pulses and optical-based laser pulses together in pump probe experiments. However, XFELs and optical lasers are sources of completely different nature, which can only be synchronized to a few

hundreds of femtoseconds of accuracy, as allowed by the electronics or optoelectronics triggering the processes. As a result, there is always a timing uncertainty associated to the arrival time of each of the pulses, which is typically characterized using time tools [97, 98, 99, 109], previously introduced in section 3.3.7. Some of these time tools claim few-fs or even sub-femtosecond accuracies when measuring the arrival timing at the time tool, however this accuracy may not be necessary transferable to what happens at the interaction region in an actual pump probe experiment.

We designed and implemented an experiment at SACLA using the available time tool [109] to allow us to gain some insights in the accuracy of the timing for mixed x-ray optical pump-probe experiments. The results presented in this section are also in preparation for publication [16].

6.2.1 Experimental method

The experimental setup used for the experiment is depicted in figure 6.7. It consists of two main branches: the time-tool branch and the experiment branch.

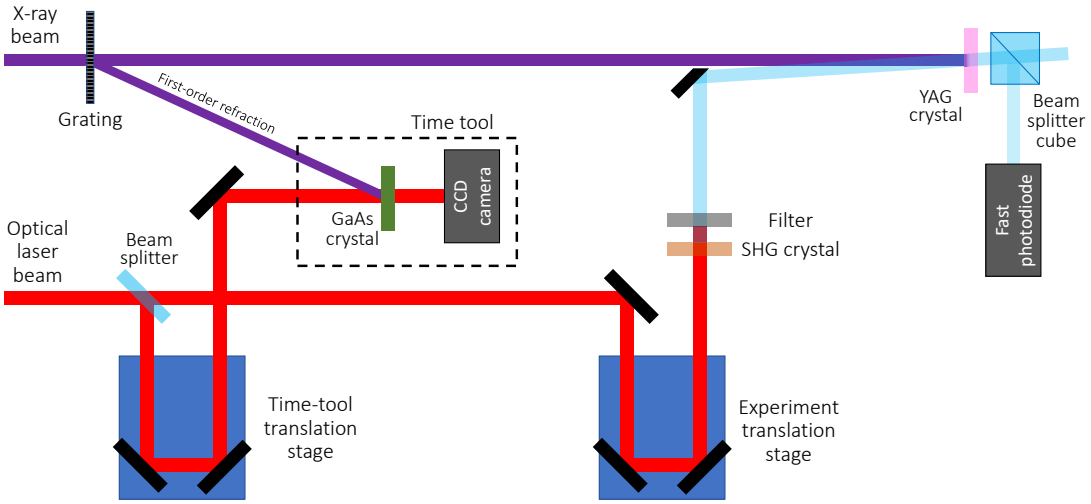


Figure 6.7: Experimental setup used to provide absolute timing calibration to an optical-pump x-ray-probe experiment at SACLA. The relative timing between the optical and the x-ray pulses is measured independently in two different experiments. The first experiment is a spatially encoded time-tool with a GaAs crystal fed with a first order refraction of the x-ray beam and a fraction of the optical 800 nm pulse. The relative timing information is extracted by recording the spatially resolved transmission of the optical pulse after the GaAs crystal. The second experiment, situated at the interaction region, measures the time-dependent transmission of the second harmonic of the 800 nm pulse through a Ce:YAG screen, which is modified by the x-ray pulses propagating collinearly.

The time-tool branch

The time-tool branch is fed with approximately 2% of the x-ray energy picked off from the main beam (10.5 keV, 500 µJ) using the first diffraction order of a transmission

grating. This beam is focused at 45 degrees into a 10 µm thick gallium arsenide (GaAs) screen using an elliptical mirror. At the same time, a small fraction of the 800-nm optical beam (< 1 mJ) is separated using a beam splitter, driven through a translation stage to control the timing, and focused at normal incidence into the GaAs screen with a cylindrical lens. The effect of the x-ray beam on the GaAs is to suddenly increase the density of free carriers by ionization and a subsequent cascade process. This effectively switches the behaviour of the GaAs crystal from semiconductor to metallic, in a process called metallization. As a consequence of this process, the optical properties of the GaAs crystal change suddenly, decreasing the amount of transmitted optical beam through the GaAs crystal whenever the x-ray pulse arrives at the screen before the optical pulse. By having the x-ray beam and the optical beam arrive to the GaAs screen at 45 degrees angle, different spatial points of the GaAs screen map into different relative arrival times at a rate of 2.36 fs/µm. This allows one to obtain an image of the back of the GaAs crystal and project it to obtain a time-dependent optical absorption profile that can be used to extract the relative timing. More details about the experimental details of this time tool can be found in reference [109].

Figure 6.8 summarizes the extraction process of the relative arrival times from the time-tool profiles. Figure 6.8a shows examples of profiles obtained in the time tool for five different consecutive events (blue, green, red, cyan, purple). Each of these profiles shows a particular pixel where the amplitude drops abruptly when moving from right to left due to the arrival of the x-ray pulse. This can be verified by comparing the profiles to an averaged profile obtained with the x-ray beam blocked (grey), which does not contain any such edges. Extracting the timing involves measuring the location of the edge in the absorption, which is performed in three steps. First, the transmission is obtained by dividing each of the single-shot profiles by the reference profile and smoothing the signal (figure 6.8b). Then the derivative of the transmission is calculated to localize the position of the maximum change in transmission (figure 6.8c). Finally the peak in the derivative can be fitted to extract its exact location, which can then be used as a measurement of the relative timing. To validate the process the transmission profiles for all events in an experimental run can then be plotted together, sorted as a function of the relative arrival time obtaining a perfect diagonal edge (figure 6.8d).

Figure 6.8d, already highlights the need for time tools, as it shows that even for the same experimental run, some events were separated by more than 600 pixels, corresponding to 1560 fs. More precisely, this jitter was characterized to be approximately Gaussian with a standard deviation of 270 fs. We also observed significant drift in the timing by pooling all of the time tool data that we collected for the duration of our experiment and plotting it together (figure 6.9). We found this drift to have two components. The first component (figure 6.9a) corresponds to long term changes over few-hour periods that can shift the timing by values of up to 0.5 ps. This is expected as there are many external factors that can affect the timing over long periods of time. We also observed

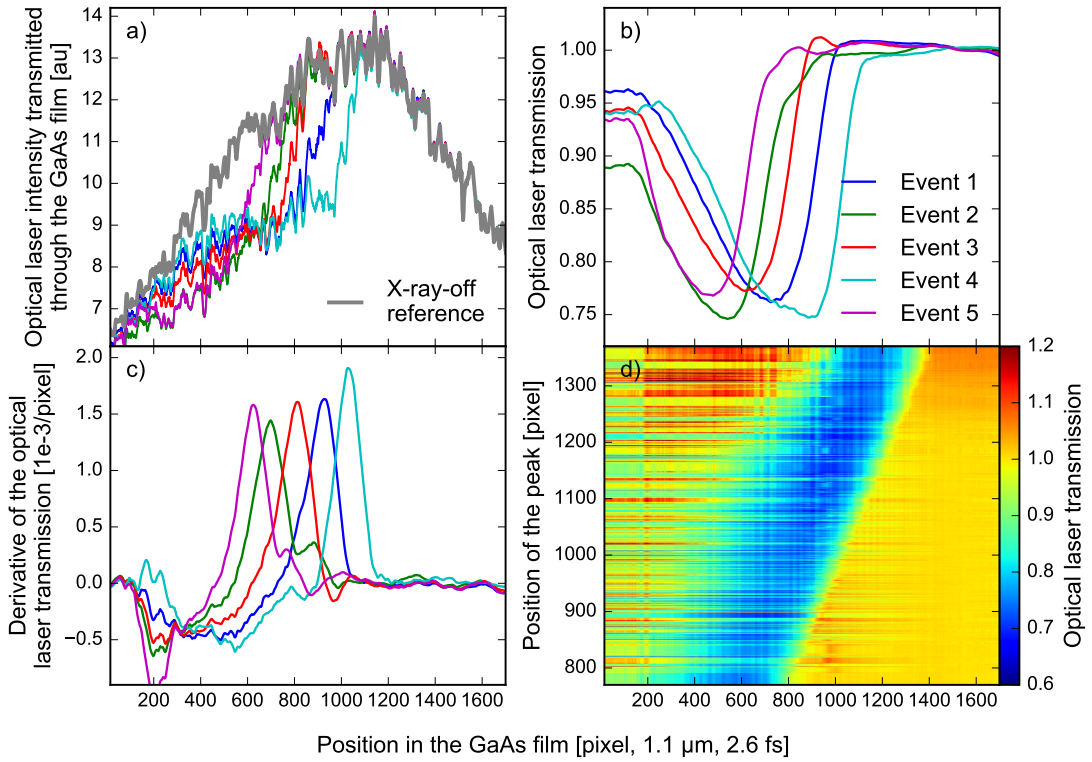


Figure 6.8: Retrieval of the timing information from the time tool. (a) Time-tool profiles measuring the spatial profile of the optical pulse transmitted through the GaAs crystal for 5 sample consecutive events with different x-ray arrival times (blue, green, red, cyan, purple). The grey profile corresponds to a reference with the x-ray beam blocked. The edge in the profiles corresponds to the arrival time of the x-ray pulse. The colours assigned to the individual events are consistent across figures a-c. (b) Time-dependent absorption of the optical pulse obtained dividing the profiles by the reference and smoothing the signal. (c) Derivative of the profiles in (b) used to transform the edge into a peak at the x-ray pulse arrival time. (d) Transmission profiles for a run including 4200 events sorted as function of the location of the edge (peak in the derivative) spanning a full range of approximately 600 pixels corresponding to 1560 fs.

a much faster periodic drift with periods of about 10 minutes (figure 6.9b) shifting the centre of the distribution of the jitter by up to 250 fs. We suspect this fast drift is related to some of the feedback loops stabilizing the synchronization of SACLA.

The experiment branch

The previous published results estimating the accuracy of the time tool at SACLA [109] essentially relied on a second identical time tool set up in the experimental chamber to prove that timing information in the time-tool branch can be correlated to timing in the experimental chambers. This approach presents two problems however. First, while the timing is measured in the experimental chamber, it is done in a configuration where the x-ray beam and the optical beam interact at 45 degrees, which is very different from the normal pump-probe configuration with collinear or quasi-collinear beams. Second, since the beams intersect at 45 degrees the concept of temporal overlap is not well

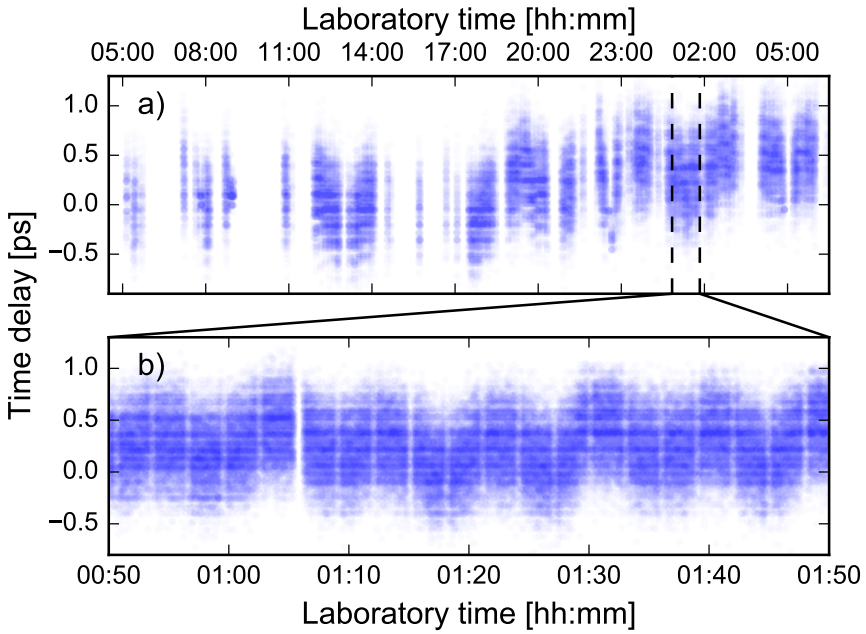


Figure 6.9: Absolute drift between the arrival time of the optical laser and the x-ray pulses measured using the time tool at SACLA. (a) For a period of 26 hours. (b) For a small time window of 1 hour. We can observe small oscillations in the timing with a period of 10 minutes approximately. The distribution of the jitter is approximately Gaussian with a standard deviation of 290 fs using all data in (b) and 270 fs when looking at a small 3-minute slice to avoid the oscillations. Figure extracted from [16].

defined (different spatial points have different timings) making it impossible to extract a reference for absolute time-zero, which is needed for time-resolved experiments.

In order to solve those problems we used a cross-correlation timing diagnostic at the experimental interaction region using the collinear beams (figure 6.7). This diagnostic consisted of a 300 µm Ce:YAG screen situated in the experimental interaction region where both the x-ray beam and the optical beam were delivered at normal incidence. The optical beam was also driven through a translation stage to control the relative timing, and through a second-harmonic-generation (SHG) crystal and a filter, to produce 400 nm light which was being used at the time as an optical pump for a parallel experiment in the beamline. As an x-ray pulse arrives at the Ce:YAG crystal, a metallization process similar to the one occurring in the time tool takes place, depleting the transmission of the optical 400 nm pulse when arriving after the x-ray pulse. In this case there is no spacial-time encoding in the YAG screen, so the absorption can be measured by a single photodiode, which for each event records a signal proportional to the amount of optical light transmitted through the Ce:YAG crystal.

While the method uses an identical configuration to that used for actual pump-probe experiments with collinear beams (with the exception of beam expanders to reduce the intensity of the beams), it cannot provide single-shot timing measurements, and it requires scanning the arrival time by means of the experimental translation stage. Figure

6.10a shows an example of a scan, where single-shot photodiode values have been plotted against different positions of the translation stage corresponding to different timings. We can clearly observe that for large values of the delay, corresponding to situations when the optical pulse arrives first, the photodiode values are much larger than for low values of the delay, for which the x-ray pulses arrive first. The transition between these to regimes is however completely blurred. This is a consequence of the jitter as for certain positions of the translation stage, the jitter allows single-shot positive, zero and negative delays, corresponding to high, intermediate and low photodiode values.

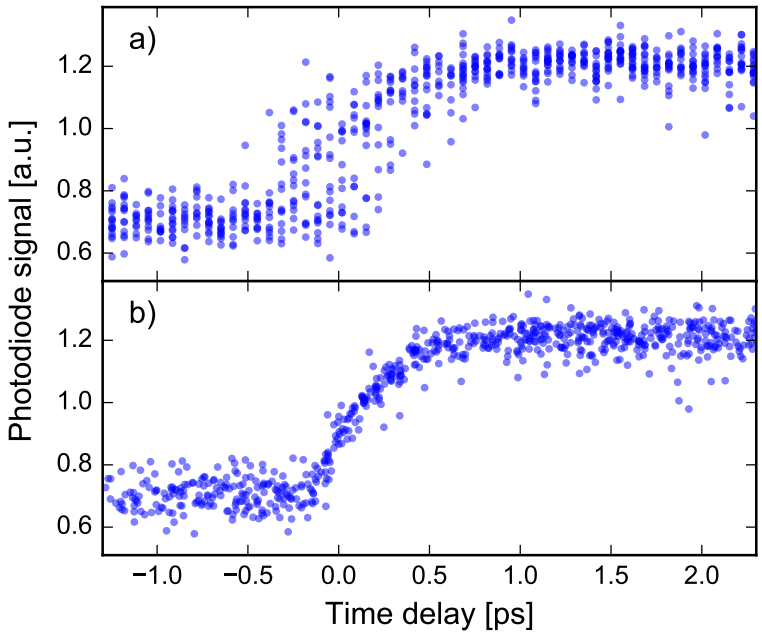


Figure 6.10: Absorption of the optical pulse in Ce:YAG as a function of the arrival time of the x-ray pulse. (a) Using timing information from the translation stage. (b) Applying the time-tool correction. Positive delays indicate the optical pulse arriving first. Figure extracted from [16].

In order to remove the effect of the jitter, we can apply a time-tool timing correction to the delay values obtained from the translation stage. In general, the single-shot relative timing between the x-ray pulse and the optical pulse can be calculated as

$$\tau_{\text{delay}} = \tau_{\text{exp-stage}} + \tau_{\text{tt}} - \tau_{\text{tt-stage}}. \tag{6.1}$$

The first term $\tau_{\text{exp-stage}}$ corresponds to the change in timing due to displacements of the experimental translation stage. This change can be expressed as

$$\tau_{\text{exp-stage}} = 2 \cdot (\text{pos}_{\text{exp-stage}} - \text{ref}_{\text{exp-stage}})/c, \tag{6.2}$$

where c is the speed of light, $\text{pos}_{\text{exp-stage}}$ is the position of the translation stage, $\text{ref}_{\text{exp-stage}}$ is a reference position fixed at the beginning of the experiment, and the factor of 2 accounts for the double change in optical path for a given translation.

The second term $\tau_{\text{exp-stage}}$ corresponds to single-shot corrections obtained from the time tool calculated as

$$\tau_{\text{tt}}[\text{fs}] = -1 \cdot (\text{pix}_{\text{tt}} - \text{ref}_{\text{tt}}) \cdot 2.6\text{fs/pixel}, \tag{6.3}$$

where pix_{tt} is the position of the edge in the timing tool in pixels, ref_{tt} is a reference pixel fixed at the beginning of the experiment, and the -1 factor is used to be consistent with the sign criteria for the delay.

The third term $\tau_{\text{tt-stage}}$ corresponds to changes in the time-tool translation stage. While this does not have an effect in the timing at the interaction regions, it systematically shifts the timing measured at the time tool, so it needs to be subtracted from the time tool values. It is calculated similarly to the first term,

$$\tau_{\text{tt-stage}} = 2 \cdot (\text{pos}_{\text{tt-stage}} - \text{ref}_{\text{tt-stage}})/c. \tag{6.4}$$

Using the corrected values, it is possible to reorder the data points in figure 6.10a to resolve the transition between high-absorption and low-absorption as shown in figure 6.10b.

Additionally we recorded photodiode references to calculate the dimensionless x-ray-dependent transmission (T) of the optical laser through the Ce:YAG from single-shot photodiode values (PD) as

$$T = \frac{PD - \text{ref}_{\text{only-xray}}}{\text{ref}_{\text{long-delay}} - \text{ref}_{\text{only-xray}}}, \tag{6.5}$$

where $\text{ref}_{\text{only-xray}}$ is a reference taken with the optical beam blocked, accounting for the room background signal and the signal coming from x-ray fluorescence in the Ce:YAG crystal, and $\text{ref}_{\text{long-delay}}$ is a reference obtained for long values of the delay, to normalize the transmission to one in cases where there is not x-ray induced absorption.

6.2.2 Stability of the timing

With both experiments (time tool and Ce:YAG absorption) working simultaneously, we repeated a series of time scans similar to the one in figure 6.10b under the same experimental conditions but at different times. Some of the previous work [109], already studied the time stability but only for periods of up to 1 hour. Since the data acquisition for XFEL experiments is typically much longer, in the order of a few tens of hours, we repeated the scans first after 16 hours of the initial scan, and then twice more every three hours.

Figure 6.11a shows the data points for each of the transmission scans in different colours, overlaid with the average absorption curve obtained for each scan by calculating the median photodiode values in 100 fs bins. There is a clear shift between the first scan

(blue) and the following scans (red, green, purple), which was quantified by fitting the curves and plotting the offset of each subsequent scan with respect to the first one (figure 6.11b). This effect shows a positive change of 60 fs approximately after the first scan and then a return towards the initial value in steps of 40 fs and 15 fs. Similar shifts can be observed in the distributions of the delay values (figure 6.11c) within the fixed highlighted area of figure 6.11a.

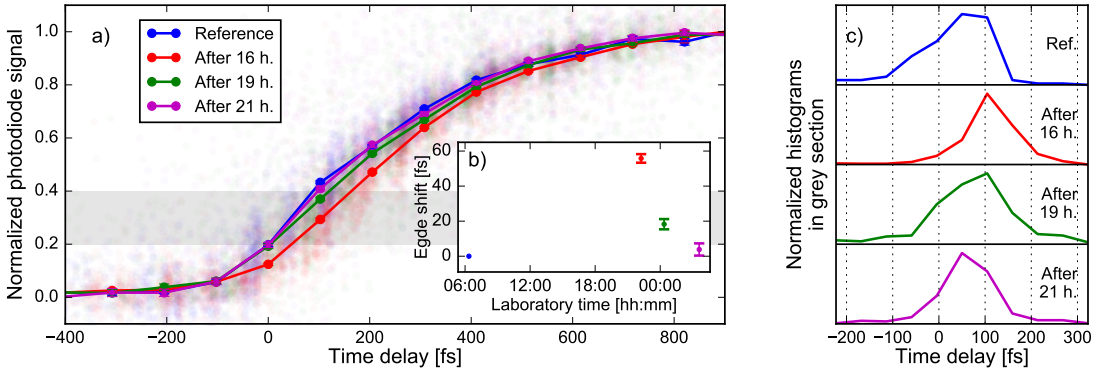


Figure 6.11: (a) 4 different photodiode scans were performed at different times and corrected using the time tool. (b) Time shift measured in the scans when compared to the first scan (blue) used as reference. (c) Histogram of the timing values lying within the shaded region in (a) where the same shift can be observed. Figure extracted from [16].

As these results are shown applying the timing tool correction, the observed drift must be independent of the global XFEL synchronization drift that we observed in figure 6.9. We believe the origin of this drift is related to the long optical paths (in the order of 25 meters) that the beams follow from the corresponding beam splitter/grating until the GaAs crystal in the time-tool branch, and the Ce:YAG screen in the experimental branch. In fact, a small change of less than 20 µm in the 25 meters of optical path would already induce a timing difference of more than 60 fs.

Due to the low number of data points, it is difficult to make many more assumptions about this drift or to create a model. Considering that the first point and the last point show a similar timing and that both of them were taken at similar times of the day, we are inclined to think that there is a daily periodicity to this trend probably related to the large changes of temperature occurring throughout the day, which could change the optical path by means of expansion and compression of the optical tables. Taking that into account, and assuming some kind of sinusoidal behaviour, we would expect to have changes of up to 120 fs throughout the course of a day.

Finding solutions to this problem would be necessary to truly provide XFELs with optical-pump x-ray-probe few-fs capabilities. The optimal way to solve this problem would be to design future time tools to be physically closer (less than a few meters) to the final experiment, which would scale this drift down probably by an order of magnitude. Alternatively, care should be taken to perform auxiliary timing experiments, similar to this one, to monitor the drift and collect enough data points similar to figure 6.11b, that

would allow one to interpolate over time, and use that to add another correction term in the calculation of the "true" timing.

6.2.3 The position of time zero

As our experiment is performed in the same experimental configuration as any pump-probe experiment would be, it also allows us to gain information about the absolute timing between the x-ray pulse and the optical pulse and, in particular, about the exact location of zero time delay. For many experiments, the location of time zero is just estimated using a similar approach by taking the middle point in figure 6.10b. This is typically the case when the experiment attempts to measure "long" few-ps dynamics, being able to afford a systematic error of a few hundreds of femtoseconds in the absolute delay estimation, or in experiments where the signal is strong enough to provide the true location of temporal overlap from the experimental data. Such large systematic errors are however not acceptable for few-fs spectroscopy experiments for which the expected signals are typically weak compared to the background, and the dynamics occur within a few tens of fs. To solve this problem we developed a model of the time-dependent transmission in the Ce:YAG crystal, in order to determine which point in 6.10b corresponds exactly to zero time delay.

We started by propagating a 7-fs-FWHM Gaussian x-ray pulse with unit amplitude through the 300 µm thickness of the Ce:YAG. Note that the effect of this arbitrary unit amplitude will be suppressed later by a free parameter of the model. The real part of the refractive index of Ce:YAG was considered to be exactly $n_{\text{x-ray}} = 1$ for the x-ray pulse. The intensity of the x-ray pulse was also attenuated by absorption with an exponential decay for different z positions inside the crystal,

$$I_{\text{x-ray}}(t, z) = I_{\text{x-ray}}(t, 0) \cdot e^{-c_{\text{x-ray}} z}, \tag{6.6}$$

where $c_{\text{x-ray}} = 1/(58.2\mu m)$ is the absorption coefficient extracted from reference [160] by weighting the atomic data according to the doping ratios.

At the same time a 100-fs-FWHM Gaussian optical pulse with unit amplitude was propagated inside the crystal at different time delays τ_{delay} with respect to the x-ray pulse. The unit amplitude in this case is irrelevant as we are only interested in the transmission coefficient. The optical pulse propagated with a refractive index of n_{optical} of 1.9 extracted from reference [161], by comparing the value of 1.86 index for YAG at 400 nm, and taking into account the resonance present at 450 nm in Ce:YAG. In this case the absorption of the optical pulse depends on the propagation of the x-ray pulse, in particular on the energy deposited by the x-ray pulse before the arrival of the optical pulse. As both pulses travel at different speeds, we performed a discretization of the thickness of the crystal in 120 different spatial slices distributed proportionally to the amount of x-ray absorption at different depths inside the crystal. This allowed us to have

a higher density of slices in the first 58.2 µm of the crystal where most of the energy is deposited. Once the slices were chosen we propagated the electric field of the optical pulse through each of the slices with the following iterative formula,

$$E_{\text{optical}}(t, z + dz) = E_{\text{optical}}(t, z) \cdot e^{-c_{\text{int}} \cdot A_{\text{x-ray}}(t, z+dz/2) \cdot dz/2}, \tag{6.7}$$

where dz is the width of each slice, $A_{\text{x-ray}}(t, z+dz/2)$ corresponds to the time-dependent accumulated x-ray energy at each spatial slice and c_{int} is a free parameter of the model representing the strength of the influence of the x-ray energy in the optical pulse absorption. The factor of two accounts for the calculation being performed in terms of electric field instead of intensity. The time-dependent accumulated x-ray energy at each slice was calculated as

$$A_{\text{x-ray}}(t, z) = \int_{-\infty}^{t} I_{\text{x-ray}}(t' - \tau_{\text{offset}}(z), z) dt', \tag{6.8}$$

where $\tau_{\text{offset}}(z)$ corresponds to the time delay between the x-ray pulse and the optical pulse at that z position inside the crystal which is calculated using the absolute time delay τ_{delay} plus the timing difference due to the different propagation speeds,

$$\tau_{\text{offset}}(z) = \tau_{\text{delay}} + \frac{z}{c}(n_{\text{x-ray}} - n_{\text{optical}}). \tag{6.9}$$

After propagating the optical pulse across the full width of the crystal, the transmission coefficient was calculated by dividing the amount of optical energy at the back of the crystal by the amount of incident energy,

$$T = \frac{\int_{-\infty}^{\infty} (E_{\text{optical}}(t, width))^2 dt}{\int_{-\infty}^{\infty} (E_{\text{optical}}(t, 0))^2 dt}. \tag{6.10}$$

The free parameter c_{int} of the model was adjusted to match the 0.24 value of the transmission for very negative time delays. Then the experimental points were sorted in 80 fs bins, calculating the median absorption value for each bin, and the model was fitted into the data simultaneously adjusting an offset parameter, to match the data left and right, as well as a Gaussian convolution parameter to match the exact width of the transition in the data, minimizing the mean-squared error of the differences between the model and the data. The result is shown in figure 6.12.

The first interesting effect observed from the fit is that the absolute time zero is towards the lower side of the transition, as opposed to the centre of the transition as might be naively assumed. This is due to the difference in propagation speeds in the Ce:YAG of the x-ray pulse compared to the optical pulse, in particular, due to the fact that the x-ray pulse travels almost at twice the speed of the optical pulse. For positive, but not very large time delays, when the optical pulse enters the crystal slightly before than the x-ray pulse, the x-ray pulse can still overtake the optical pulse, and induce some absorption. In fact, figure 6.12 shows that the optical pulse is partially absorbed even

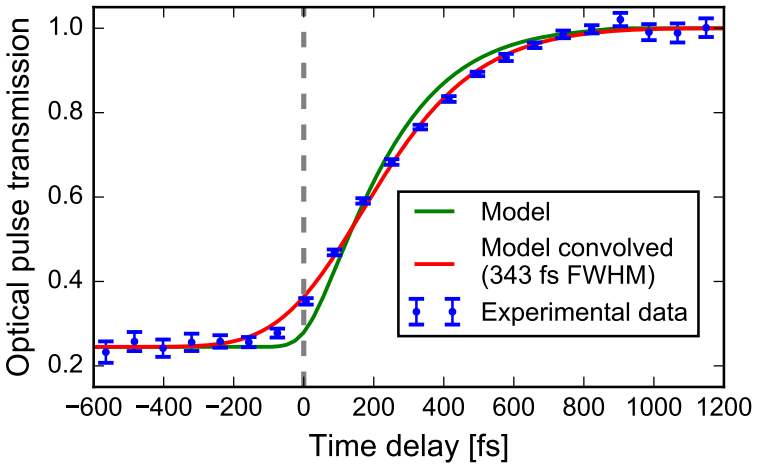


Figure 6.12: Comparison between the experimental data sorted in 80 fs bins (blue) and the fitted absorption model convolved with a 343 fs FWHM Gaussian (red) in Ce:YAG. The width of the convolution was a free parameter of the fitting procedure. Error bars are calculated as the standard deviation of all experimental points above 900 fs divided by the square root of the number of events included in each bin. The green profile shows the behaviour of the model without the convolution. Figure extracted from [16].

for long positive delays of up to 800 fs. On the other hand, for negative delays, with the x-ray pulse entering the crystal first, the time order of the pulses is maintained, and the transmission remains constant at the lower value.

The second interesting effect is that, even though we used realistic parameters for our model, we still have to convolve the results with a 343 fs FHWM Gaussian in order to fit the experimental data. The first reason for this could be inaccuracies in the time delay retrieval from the time tool: while the time tool images have a resolution of 2.6 fs/pixel, the actual precision of the measurements is probably in the order of a few tens of fs. This however cannot account for a 343 fs FWHM uncertainty. Another possible reason, which we think accounts for most of this effect, may be related to the x-ray and optical intensity fluctuations, or photodiode shot noise. We can observe in figure 6.10b that even for very long time delays or very short time delays, the photodiode values present a significant variance due to intensity fluctuations or shot noise. The effect of the fluctuations may be similar to having timing uncertainty as for example a slightly less/more intense x-ray pulse would induce a very similar effect in the transmission as if it had arrived slightly later/earlier than the optical pulse. We attempted to model the effect of the fluctuations using a Monte-Carlo approach, however the information we had on the distributions of the intensity fluctuations was very limited and the model did not fit the data well. A third possible reason for the need of the convolution could be related to the fact that we have completely neglected the rise time in the carrier density in the Ce:YAG crystal due to the x-ray absorption, however we expect this number not to be more than a few tens of fs [162, 163]. Finally, bleaching in the YAG absorption due to the high x-ray

intensity could increase the x-ray penetration length producing longer time dynamics as measured in the experiment.

While the model developed here still has some of those limitations, it shows evidence towards the true time-zero being approximately 200 fs away from the middle point of the change in the transmission and should be taken into account for any experiment requiring absolute timing accuracy below a few tens of fs. Further experiments recording the single-shot x-ray and optical pulse intensities could be performed to provide more experimental data to improve the model further and obtain more accurate results.

6.2.4 Conclusions and outlook

Using experimental data from SACLA, we have shown evidence concerning the stability of the time tool measurements of the single-shot jitter in the relative arrival time of the pulses in optical-pump x-ray-probe experiments. While these devices claim down to sub-femtosecond accuracy in some cases [99], we have found that the reliability of these values compared to an absolute cross-correlation reference may degrade by tens of femtoseconds in periods of a few hours due to drift in the optical paths. Considering that most experiments require much longer data acquisition times, on the order of tens of hours, care should be taken to record timing references often enough to track the drift. This may however be unfeasible for few-fs experiments, as it would require recording references every few minutes for the observed drift rates of up to 20 fs per hour.

Furthermore, the procedures currently used to obtain the absolute timing reference, i.e. the time-zero position, are also typically only accurate to a few hundreds of femtoseconds. This is due to open questions about the response time of the materials used to obtain the timing references, as well as other effects, such as pulse energy fluctuations, and propagation effects that make extracting an accurate value very challenging. While such an absolute reference may be not necessary for experiments showing a well contrasted time-resolved signal that may allow the extraction of the position of time-zero from the signal itself, this is not usually the case, and a time-zero reference accurate to the few-fs would be required.

These results show that performing optical-pump x-ray-probe experiments to measure few-fs dynamics is still very challenging using current XFEL facilities. This could be improved however by building few-fs beamlines specifically designed to improve the timing stability, as well as by finding and choosing thin materials with fast response times to perform the cross-correlation reference measurements.

Chapter 7

X-ray characterization, fluctuations and machine learning

We have learned in the previous chapters that, due to the fluctuating nature of XFELs, successful time-resolved experiments depend upon a reliable single-shot characterization of the generated x-ray pulses. This will remain true for the next generation of XFELs working at higher repetition rates, however, many of the detectors necessary for the characterization will not be able to cope with the enhanced repetition rate due to data bandwidth or heat load constraints. Some examples of detectors that will present these limitations are XTCAV for the measurement of the delay, time tools based on solid targets for the optical timing, and some x-ray grating spectrometers for spectral characterization.

In this chapter the author proposes a method based on machine learning to overcome that limitation. The author demonstrates that it is possible to predict complex single-shot x-ray information using electron-beam parameters and other properties measured with fast detectors as the only single-shot input. This can be achieved with the only requirement being that some training events containing information from both, the complex diagnostics, and the single diagnostics are available. The results presented in this chapter are already available in the form of a manuscript [14] containing all the experimental details and a full discussion of the results. In this chapter the author will summarize the experimental methods and results, providing some additional evidence of the correlations between electron-bunch parameters and x-ray properties that make the approach possible.

7.1 The problem of the fluctuations

We already analysed in section 5.6.1 the distributions and the correlations of the x-ray properties for a typical double-bunch pump-probe dataset (figures 5.27 and 5.28). In that section, we justified the large width of the distributions mainly by the fact that the data had been recorded within a period of 10 hours, interleaved with many corrections to

the tuning of the machine that changed the value of the parameters. However, looking at data from a single 13-minute experimental run (figures 7.1 and 7.2) for which no changes were made to the XFEL configuration we can still observe similar effects. In this case, the large fluctuations in the pulse energy, the photon energy and the time delay are purely due to the single-shot jitter of the XFEL operating in a nominally "stable" manner.

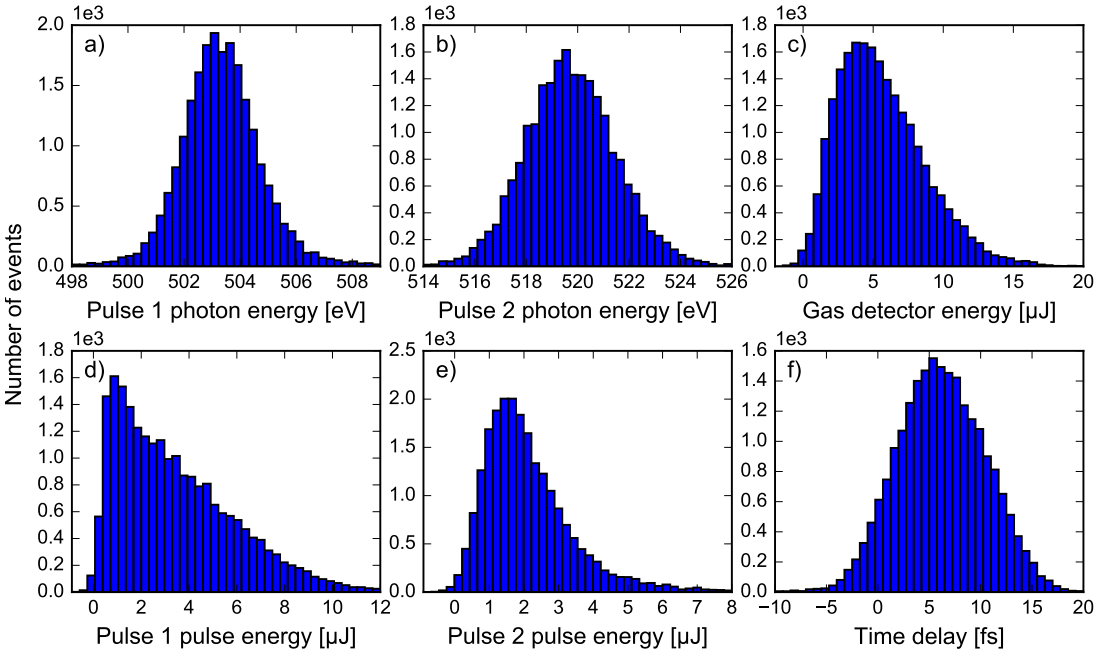


Figure 7.1: Distributions of different x-ray properties for a single run containing $2.2 \cdot 10^4$ events obtained within 13 minutes in a double-bunch configuration. Recording and monitoring these distributions is one of the keys to understand the XFEL fluctuations, which can vary between different modes and across different days.

In the distributions of this particular dataset (figure 7.1) we observe symmetric Gaussian fluctuations with FWHM of up to 3.5 eV, 4.5 eV and 12.5 fs in the photon energy of the pulses and the delay, respectively, and asymmetric fluctuations in the pulse energies yielding values between 0 and 12 μJ for each pulse. These distributions are however not typically reproducible, as we have observed changes of up to 50% in the width of the photon energy, and delay distributions, and even larger changes of the mean pulse energy distribution under similar conditions on different occasions.

The correlations between the variables are typically more consistent across different days. This is the case of the positive correlation between the photon energy of the pulses (figure 7.2a), as they typically fluctuate together depending on the electron-bunch energy. It is also the case of the correlation between the pulse energy and the photon energy for each pulse (figures 7.2c-d), as there is always a particular photon energy at which the XFEL lases more efficiently. On the other hand, the correlation of the photon energy with delay is not always consistent. In figures 7.2e-f we observe a slightly negative

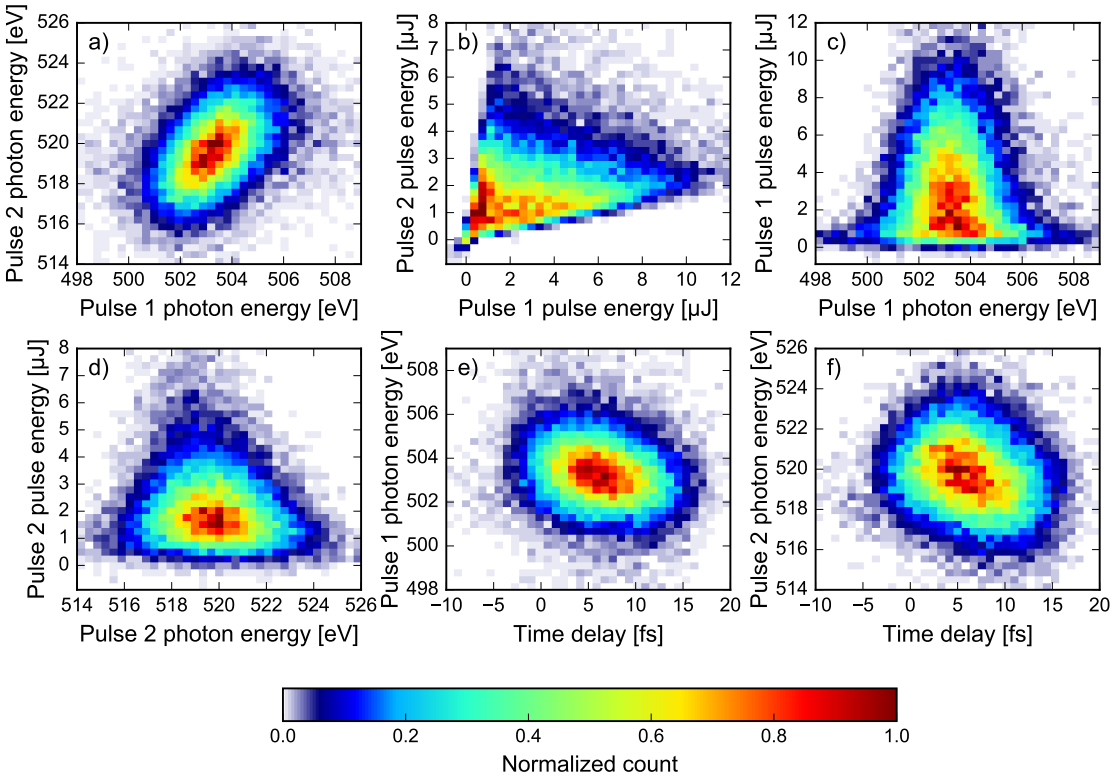


Figure 7.2: Correlations between different pairs of x-ray properties for the same dataset used in figure 7.1. The fluctuations between different variables can be highly correlated. For example, (a) shows that the photon energy of the first pulse typically fluctuates in the same direction as the photon energy of the second pulse. On the other hand, the pulse energies (b) present negative correlation, as when one pulse is particularly strong, the other tends to be particularly weak.

correlation, while in the figure 5.28 including the complete dataset we observed a slightly positive correlation.

These results show definite evidence that when doing few-femtosecond experiments, one should not rely on the nominal values given for the x-ray parameters, and always seek to measure the single-shot properties for all the events, even if changes are not made to the XFEL configuration. However, as mentioned earlier in this chapter, this may become impossible in future high-repetition-rate XFELs, as some of the devices involved will not work at the full repetition rate. In general, devices that require intercepting the electron or the photon beam with a solid target will be exposed to thermal damage. Also, devices required to store single-shot two-dimensional images may not have enough bandwidth or may just generate too much data to realistically store it for every single shot.

However, as we have mentioned many times throughout this thesis, XFELs are highly correlated machines where changes in any electron-beam parameters across the accelerator and undulator stages have a impact on the final x-ray properties. In fact, most of the fluctuations that we observe in figures 7.1 and 7.2 are a result of fluctuations on the

electron-beam parameters. Figures 7.3, 7.4 and 7.5 show the correlations between single-shot x-ray properties and single-shot electron-beam properties measured for a different dataset.

We observe that, as we expected, the electron beam after the third LINAC (L3), is very strongly correlated with the final photon energy of the pulses (figures 7.3p-q). The same happens with the position of the electron beam at dispersive regions in the LINAC to undulator (LTU) tunnel (figures 7.4a-b,d-e), which is also sensitive to changes in the electron-beam energy. In the case of the delay, we observe the maximum correlation to happen with the beam position and the peak current at the second beam compressor (BC2) (figures 7.3o and 7.5f). This is consistent with the fact that the delay is set by the double-slotted foil which is located precisely in the dispersive region of the second beam compressor and is sensitive to fluctuations at that stage.

Apart from those simple correlations, figures 7.3, 7.4 and 7.5 show many other subtle correlations between other electron-beam properties and the x-ray properties. Some of these correlations are just a consequence of how the feedback loops stabilizing the XFEL output were operating at that particular time, and cannot be understood in terms of a physical mechanism without doing a full realistic model of the entire XFEL. For example, the correlation between the electron-beam energy after L3 and the time delay (figure 7.3r) shows clear signs of multidimensional correlations, where a third unknown variable may control the range of variation of the delay at each electron-beam energy. This is also the case for the photon energy of the first pulse, and the position of the bunch at BC2 (figure 7.3m), a correlation which surprisingly does not show at all for the second pulse (figure 7.3n). Such complex correlations cannot easily be understood in terms of linear Pearson coefficients between pairs of variables. Furthermore there are many other variables that show clear signs of weaker correlation with the different x-ray properties, which, while they cannot be easily interpreted, are real for the particular dataset.

These correlations indicate that it should be possible to infer information about the x-ray properties from information about the electron beam. This approach has the advantage that the beam-position monitors (section 3.3.9), used to measure the electron-beam properties, are not so limited in terms of repetition rate, as they only measure the amplitude of the pulsed electric field generated by the passing electron bunch. However, none of those individual correlations are strong enough on their own to allow the extraction of accurate values of the x-ray properties. Instead, a single multivariate model exploiting all of the correlations together would be required to make predictions of x-ray properties based on all the electron-beam variables. Furthermore, many of the subtle correlations change their behaviour over time due to drift in the XFEL. In order to follow that drift, additional variables, containing information about environmental changes in the experiment that may have an impact on the drift (such as time, temperatures, chamber pressures,...), would have to be included as additional inputs. Finally, the model would have to be of empirical nature, as characterizing all

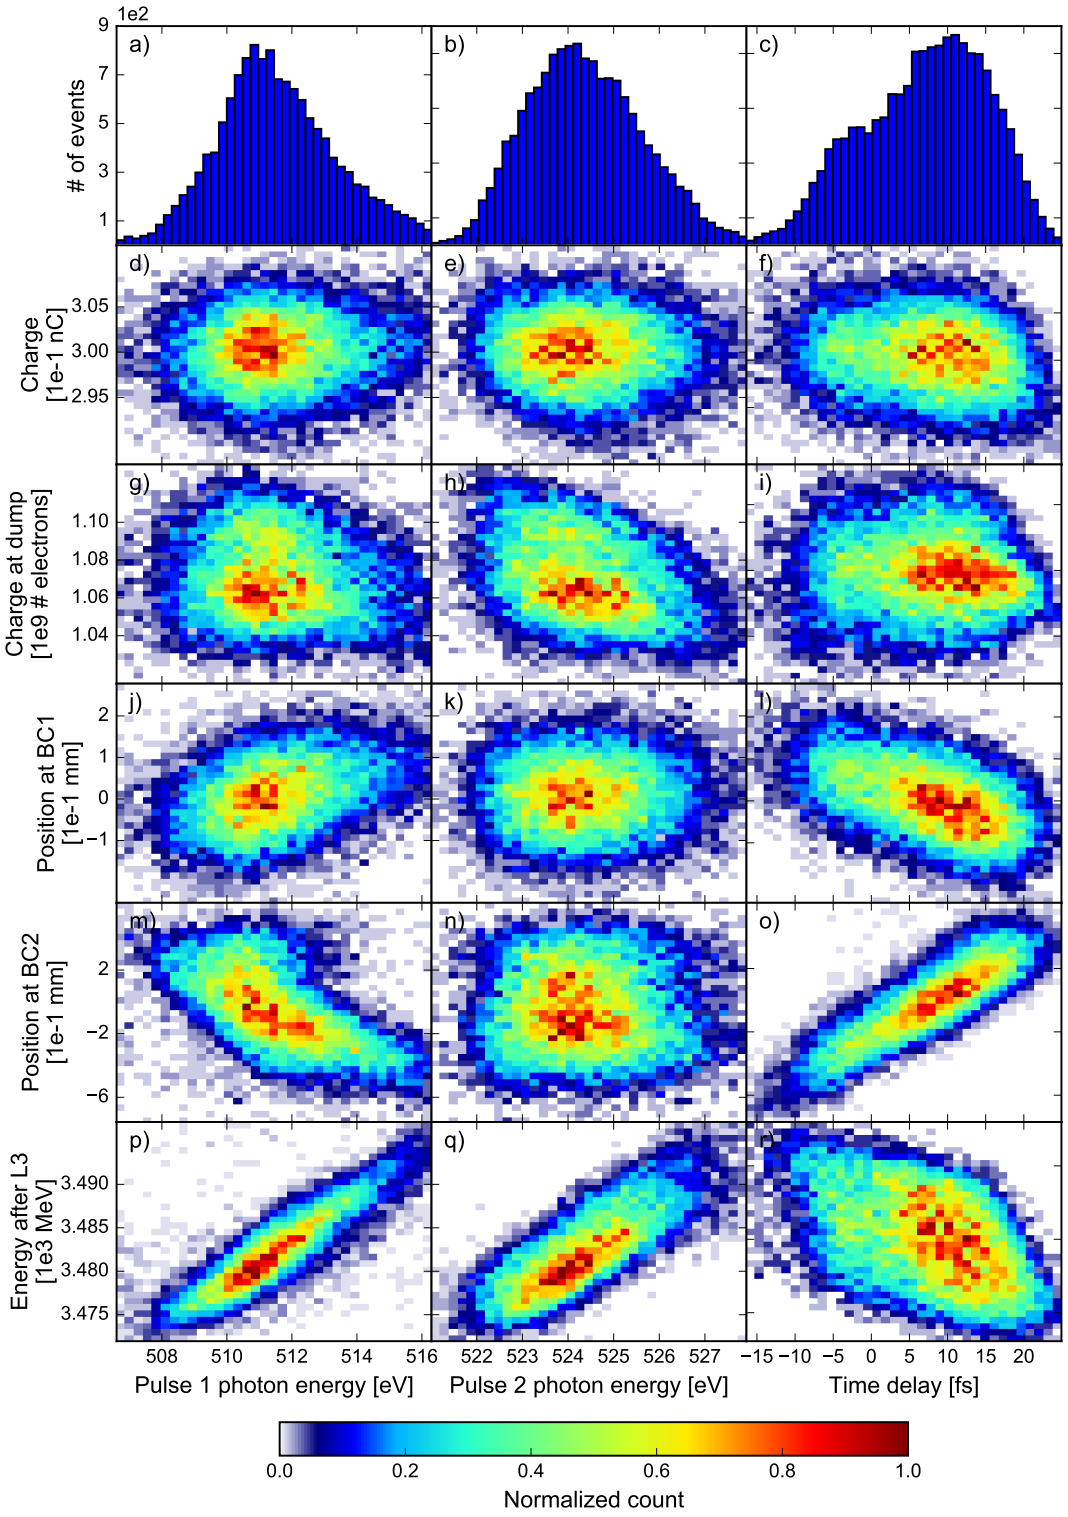


Figure 7.3: Correlation between x-ray properties and electron-beam parameters. (a, b, c) Distributions of the x-ray properties in a double-bunch mode. (d-r) Correlation plots of the x-ray properties with the different electron-beam parameters. Some of the plots (o, p, q) show strong linear correlations, however, the majority of them (f, g, h, i, j, l, m, n, r) show much more complex correlation patterns, with signs of multivariate correlations where more than two electron-beam variables are simultaneously involved. Considering that there are many more variables involved (Figures 7.4 and 7.5), the situation become difficult to handle for analytical models and can benefit from applying machine learning approaches. BC: Beam compressor. L3: third LINAC.

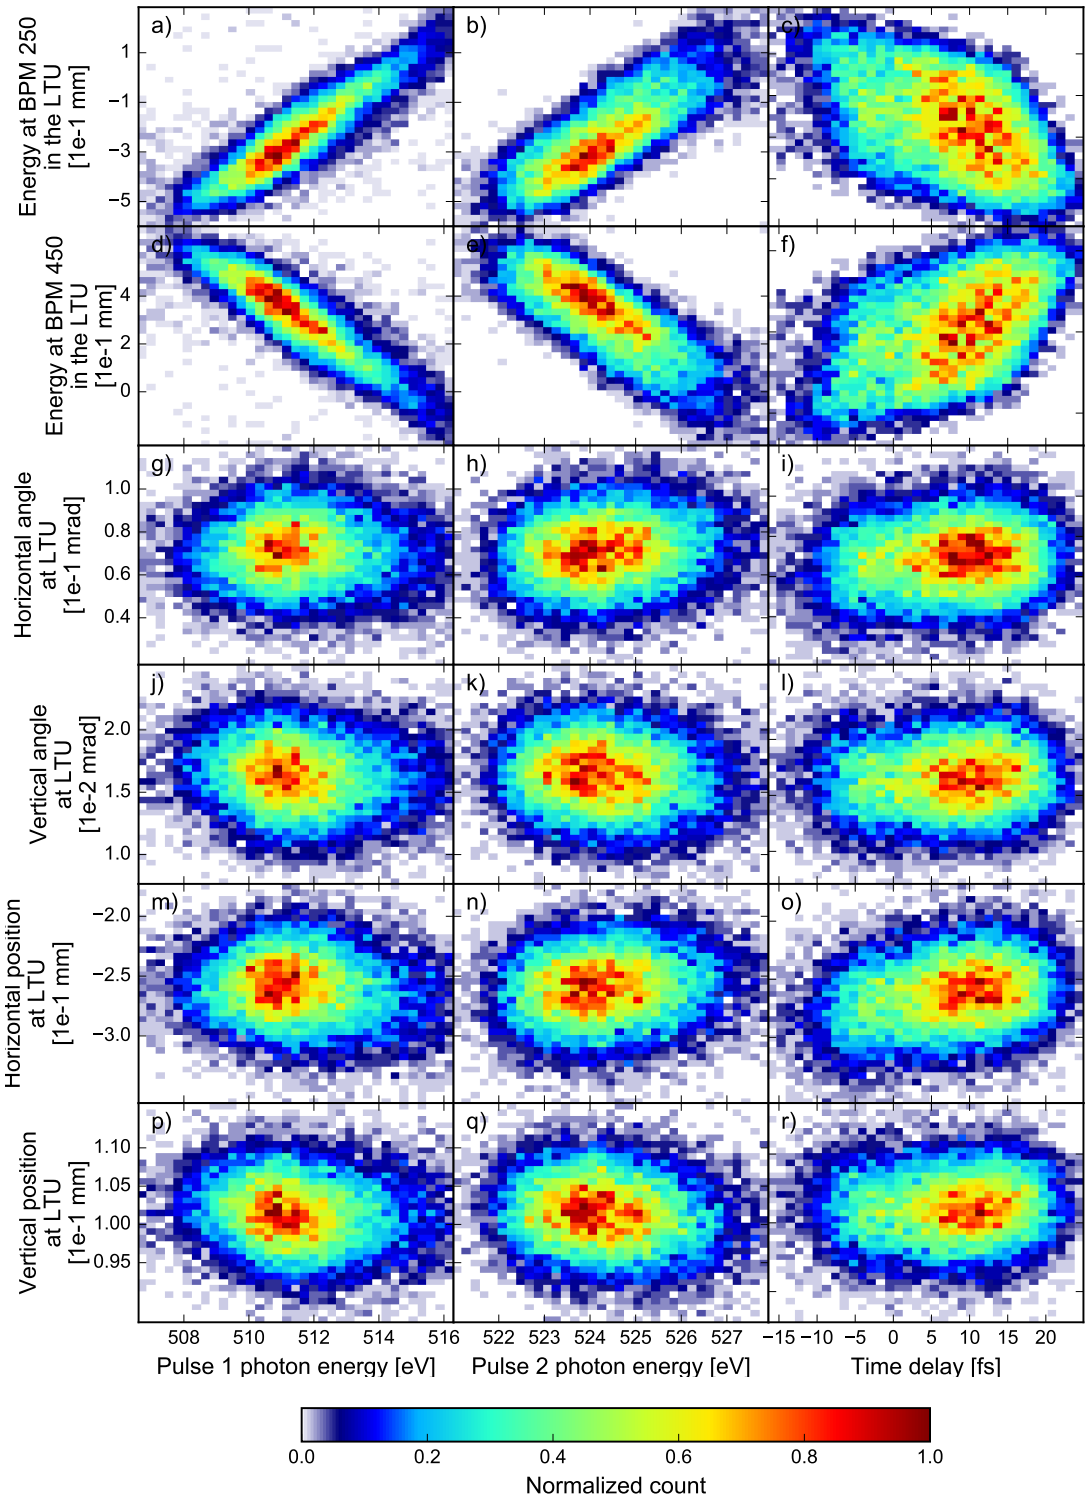


Figure 7.4: (a-r) Continuation from figure 7.3. In this case we observe linear correlations in (a), (b), (d), (e), and more complex correlations in (c) and (f). BPM: Beam-position monitor. LTU: LINAC to undulator tunnel.

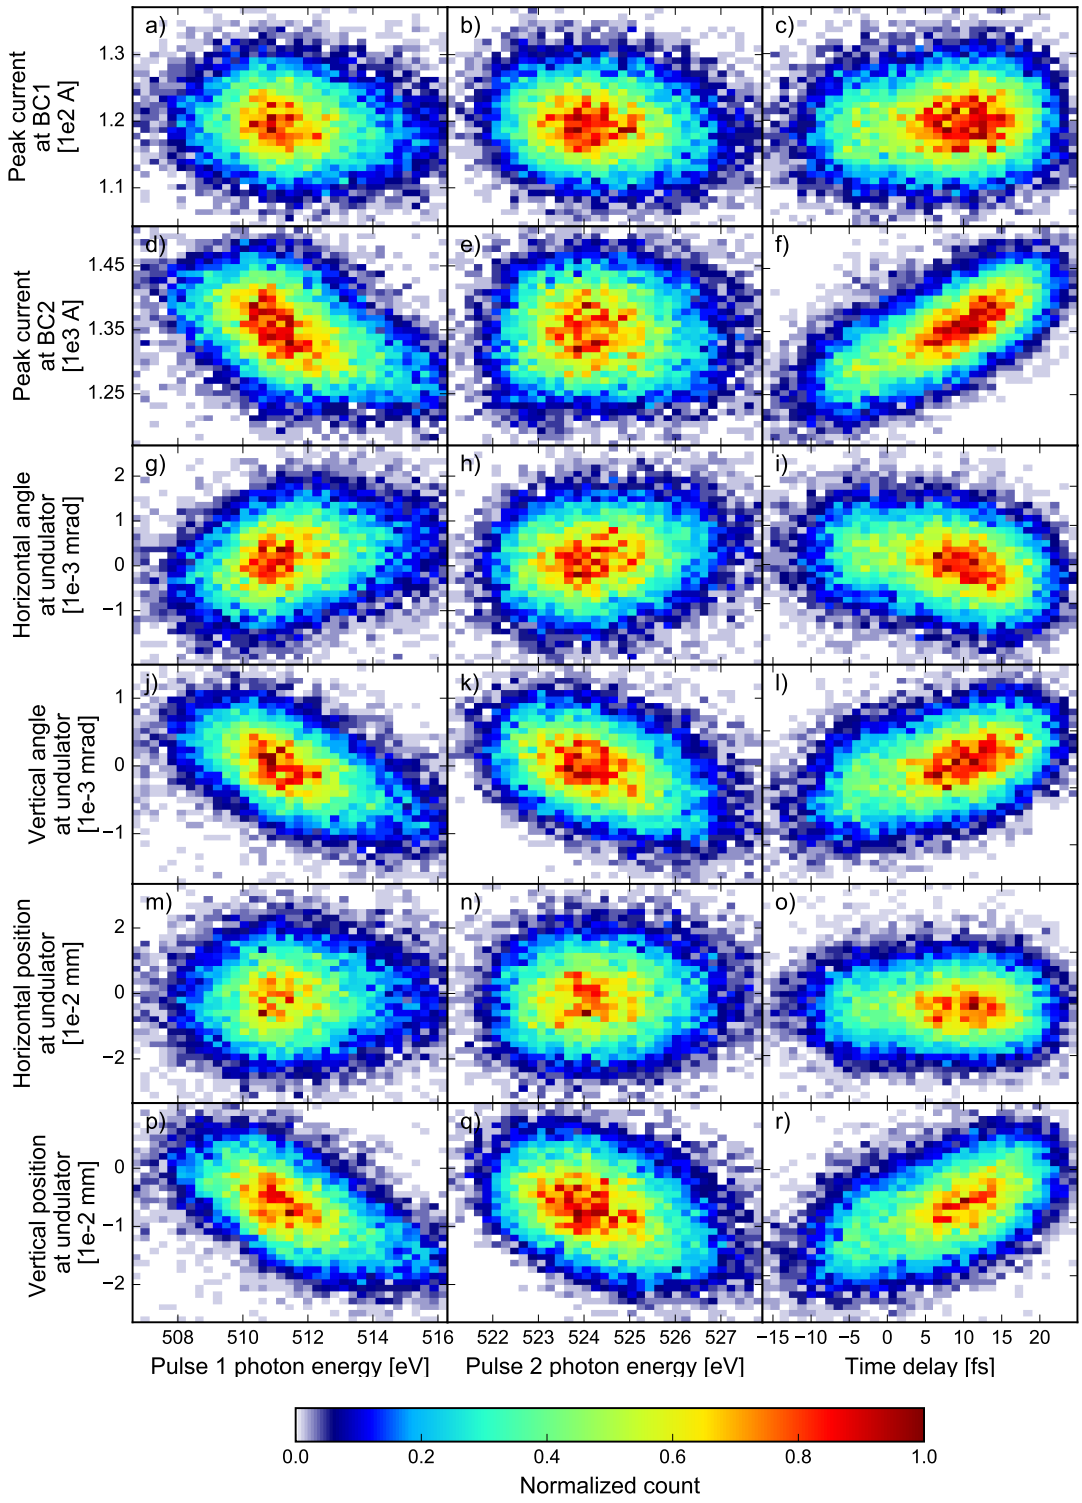


Figure 7.5: (a-r) Continuation from figure 7.4. In this case we observe linear correlations in (f) and more complex and subtle correlations in (d), (e), (g), (i), (j), (k), (l), (p), (q) and (r). BC: Beam compressor.

possible electron-beam fluctuations in a few-km long few-fs precision machine to fit a full analytical model is practically impossible.

Considering the large amount of tentative input variables, the empirical character of the problem and the short-lived stability of some of the fluctuations, manually designing and fitting a model becomes very inefficient. As an alternative, it is much more practical to use machine learning techniques to solve this task.

7.2 Introduction to machine learning

The great majority of the computer software behind most of the technology developed within the past 40 years relies on the use of specific algorithms to solve each specific problem. An *algorithm* is a finite series of well-defined actions designed specifically to solve a task. The concept of algorithm is not a result of the development of computer science, in fact, algorithms such as the Euclidean algorithm to compute the *greatest common divisor* of two numbers have existed for thousands of years. However, with the development of computers as machines with a fixed set of instructions, algorithms specifying how to solve particular tasks by making use of those particular instructions in a certain way became the building blocks of computer software.

There are however many tasks that we can perform as humans that cannot be easily written as a set of instructions. While we can identify objects, faces or actions from a picture very easily, we do not understand how we do it well enough to write an algorithm that would mimic that behaviour. Creative processes like writing a book (or Ph.D. thesis), composing a melody or painting an artwork are intrinsically very complex, as we make use of all our past experiences, and this cannot be automated in a closed algorithm easily. Finally, non well-defined problems such as "find something interesting in this dataset" cannot be written in terms of classical algorithms. While humans can easily solve all these problems for small amounts of data, with the increasing amounts of stored data, it is becoming unfeasible for humans to analyse and interpret all that data.

Machine learning [164] is the branch of computer science dedicated to the development of generic algorithms, not designed to solve individual problems, but instead, to provide a generic framework that lets the computer find the optimal way to solve any given problem. In this sense, machine learning is based on *meta-algorithms*, which define how the computer will *learn* how to build its own algorithms based on large amounts of data. Currently, LCLS generates data at a rate of 1 GB/s, with this number being expected to increase by three orders of magnitude for LCLS-II. This makes XFEL science a textbook example of a field that can benefit from applied machine learning.

Machine learning is very closely related to statistics, as it attempts to synthesize experimental data to create higher-level concepts and understand the process that allows the generation of that data. In this sense, classical machine learning is typically divided in two categories [165]: *supervised learning*, which focuses on finding relationships between many variables to learn how to make predictions on another variable,

and *unsupervised learning*, which finds relationships between many variables with the aim of understanding the probability distribution generating the data. Also, as machine learning attempts to solve problems that are "easy" to solve for humans, current machine learning also includes categories that make use of human-like strategies, such as *reinforcement learning* based on action-reward approaches, and *deep learning* which allows the computer not only to learn how to solve a specific problem, but to do so by creating its own abstract concepts, yielding machine learning models much more robust to extrapolation [165].

7.2.1 Supervised learning

In this thesis we will only make use of supervised learning. Returning to the statistics analogy, supervised learning is analogous to statistical regression. For a vector of n_f input variables $\mathbf{x} = (x_1, x_2, ..., x_{n_f})$ and a vector of n_t output variables $\mathbf{y} = (y_1, y_2, ..., y_{n_t})$, supervised learning attempts to find a mathematical model f that allows the value of the output variables to be obtained as a function of the input variables $\mathbf{y} = f(\mathbf{x})$. The input variables are known as *features* and the output variables are known as *targets*. Finding the required mathematical model can be expressed as an optimization problem,

$$f = \arg\min_{f^*} E_D[\|\mathbf{y} - f^*(\mathbf{x})\|^2], \tag{7.1}$$

where the optimal function f that minimizes the expected value E across the data generating distribution D of the norm of the difference between the targets $\mathbf{y}$ and the computed targets $f(\mathbf{x})$ is selected among all the possible functions f^*. Often, instead of choosing a function that can change its analytical form in the optimization process, universal function approximators are used, and in particular, *parametric models* for which a function with a specific analytical form $f_{\mathbf{a}}(\mathbf{x})$ can adapt to any output by means of the tuning of a fixed set of parameters $\mathbf{a}$. When using this type of model, the optimization problem concentrates on finding the optimal set of parameters $\mathbf{a}$ so

$$\mathbf{a} = \arg\min_{\mathbf{a}^*} E_D[\|\mathbf{y} - f_{\mathbf{a}^*}(\mathbf{x})\|^2]. \tag{7.2}$$

This equation still presents the problem that we do not typically know the data generating distribution D, but just some *training examples* that have been drawn from the distribution. For a set of n_e training examples consisting of a list of features $\{\mathbf{x}^1, \mathbf{x}^2, ..., \mathbf{x}^{n_e}\}$ and their corresponding target values $\{\mathbf{y}^1, \mathbf{y}^2, ..., \mathbf{y}^{n_e}\}$, the optimization problem can be approximated as

$$\mathbf{a} \approx \arg\min_{\mathbf{a}^*} \frac{1}{2n_e} \sum_{i=1}^{n_e} \|\mathbf{y}^i - f_{\mathbf{a}^*}(\mathbf{x}^i)\|^2 = \arg\min_{\mathbf{a}^*} J(\mathbf{a}^*), \tag{7.3}$$

where $J(\mathbf{a})$ is known as the *cost function* or *loss function*. Once the parameters $\mathbf{a}$ of the model are optimized by minimizing the cost function, the function can be used to make predictions, using only measured features as inputs,

$$\mathbf{y}_{\text{predicted}} = f(\mathbf{x}_{\text{measured}}). \tag{7.4}$$

Any additional parameters of the model that are not optimized in this process are known as *hyperparameters*. These hyperparameters can be qualitative or quantitative, and typically control the analytical family of functions included in the model and the number of parameters of the model. Models with a larger number of parameters have a higher *capacity*, as they are much more flexible when approximating more complex functions. Models with lower capacity than the true underlying function that generates the data tend to underfit the data, while models with larger capacity tend to overfit the data. In this sense the capacity of a machine learning model for a particular problem should be ideally tuned following A. Einstein's quote: *"Everything should be made as simple as possible, but no simpler"*. Apart from the capacity of the model, different aspects of the optimization process, such as the optimization speed, or the particular training algorithm used, are also controlled by hyperparameters. In this sense, any decision that a machine learning practitioner makes can be seen as a hyperparameter .

For the results shown in this thesis, we have used linear and quadratic models, support vector regressors and artificial neural networks.

7.2.2 Linear and polynomial regression

The simplest model in supervised machine learning is the linear model (figure 7.6a), for which each of the targets is calculated as a linear combination of the features,

$$y_i = a_0^i + \sum_{j=1}^{n_f} a_j^i x_j, \tag{7.5}$$

where a_j^i are the parameters of the model. As the problem is completely equivalent to linear regression, the optimal values of the parameters can be very easily calculated using the analytical solution provided by the Moore–Penrose pseudoinverse [165].

This linear model can be very easily extended into a polynomial model, able to represent non-linear relationships across multiple features, by transforming the initial vector of features, e.g.

$$\mathbf{x} = (x_1, x_2, x_3), \tag{7.6}$$

into a vector containing artificial features consisting of all possible products between the original features up to a certain degree. For example, to obtain a quadratic model,

$$\mathbf{x}_{\text{aux}} = (1, x_1, x_2, x_3, x_1^2, x_2^2, x_3^2, x_1x_2, x_1x_3, x_2x_3). \tag{7.7}$$

The degree of the chosen polynomial acts as a hyperparameter that allows tuning the capacity of the model. In figure 7.6b we observe polynomial models of orders 1, 2 and 8 fitting the same experimental data. In this case, the linear model clearly underfits the data not being able to model most of the behaviour. On the other extreme, the 8th degree model overfits the data, creating very much more complex structures than necessary to reproduce the experimental points. The balance in this case is found with the quadratic model, which fits the data reasonably.

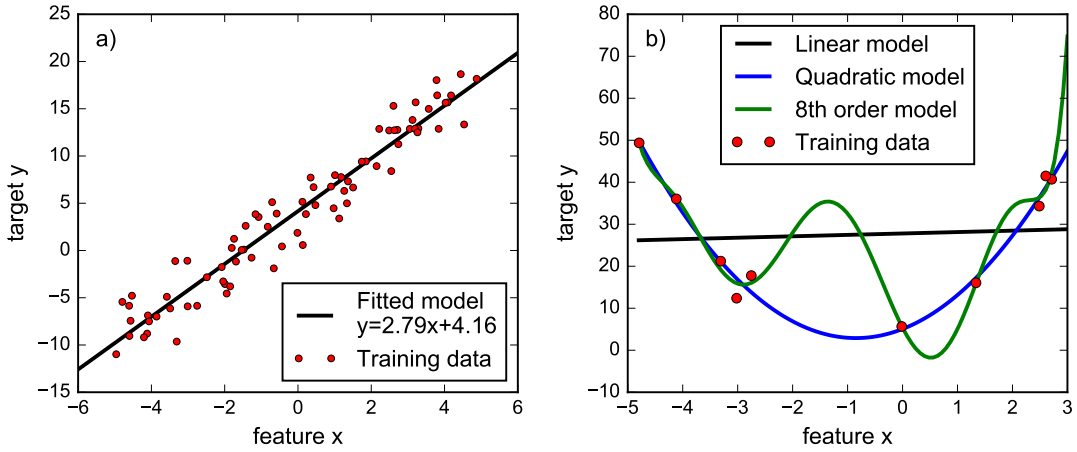


Figure 7.6: (a) Experimental data fitted using linear regression. (b) Finding the optimal degree for the polynomial model. For a human reader it is apparent that the model that best explains the data without overfitting is the quadratic model.

The main limitation of the polynomial models comes with the very large growth rate of the number of artificial features (and model parameters) as a function of the polynomial degree. In fact the number of features roughly scales as the number of k-multicombinations of n elements, where k is the degree of the polynomial and n the number of original features. For example, for 100 input features, the linear, quadratic and cubic models would have 101, 5151 and 176851 parameters respectively. Since the number of parameters has to be equal to or smaller than the number of training examples for the optimization problem not to be underdetermined, working with models of higher degree than linear or quadratic often becomes impractical.

7.2.3 Support vector regressors

Support vector regressors (SVR) [166] attempt to find a relationship between the features and each of the targets with the following analytical form,

$$y_i = a_0^i + \sum_{j=1}^{n_e} a_j^i k(\mathbf{x}, \mathbf{x}^j), \tag{7.8}$$

where, in contrast to the polynomial models, the summation is performed across the examples in the training set. This allows the number of parameters a_j^i of the model, to

grow with the number of training examples, dynamically adapting the capacity of the model to the number of training points. k is a kernel function, which for the linear SVRs is simply the dot product between the input features and each of the training examples. Instead, a Gaussian kernel was used,

$$k(\mathbf{u}, \mathbf{v}) = \frac{1}{\sqrt{(2\pi\gamma)^{n_f}}} e^{-\frac{1}{2}\gamma\|\mathbf{u}-\mathbf{v}\|^2}, \tag{7.9}$$

which allows the SVR to represent non-linear behaviour across the input features. γ in this case is treated as a hyperparameter.

While the SVR models require storing the training examples to evaluate the function, the cost function of the SVR is designed to converge to a solution where most of the parameters a_j^i are equal to zero, allowing one to, in practice, store only a few of the training examples corresponding to non-zero coefficients. These examples are denominated *support vectors*. There are two hyperparameters that allow tuning this behaviour, C and ϵ. As shown in figure 7.7, ϵ determines the maximum distance that the true value of a target is allowed to be away from the predicted value without contributing to the cost function. C on the other hand defines the trade-off between the flatness of f (corresponding to more zero coefficients) and the contribution to the loss function for predictions that deviate by more than ϵ (ξ in figure 7.7) [166].

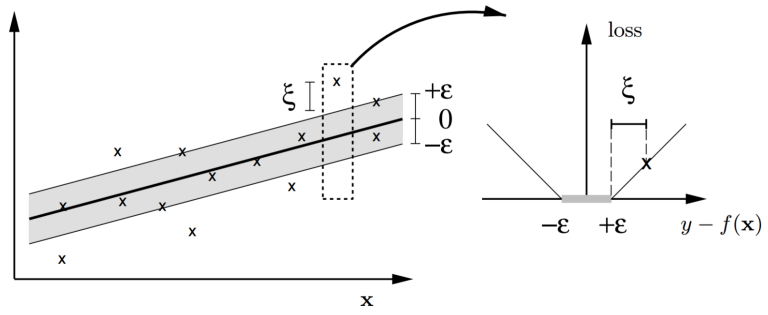


Figure 7.7: The cost function of a linear support vector regressor defines a margin ϵ around the model within which errors in the predictions are not penalized and do not contribute to the cost. Data points further away from the predictions by more than the margin only contribute by the amount by which they surpass the margin ξ. SVR cost functions defined in this manner are always convex with respect to the parameters of the model, facilitating the numerical optimization. Figure extracted from [166].

SVR models are commonly used in cases where the number of training examples is small, as they always allow finding of the perfect balance between the number of training examples and the capacity of the model, solving some of the problems of the polynomial models. They do however become impractical to use when the number of training examples is very large, as they need to store and use a significant part of those examples every time the function is evaluated.

7.2.4 Artificial neural networks

Artificial neural reworks (ANN) [167] attempt to find a balance between the polynomial models (very large growth of number parameters due to all possible non-linear combinations) and the SVRs (need to store many training examples). The architecture of a typical fully connected ANN is shown in figure 7.8. It is made of different layers, each of them containing different cells computing a single output value, using all the output values from the previous layer as inputs. The output of all cells at any given layer $\mathbf{l}_{i+1}$, can be expressed using matrix notation in terms of the output of the previous layer $\mathbf{l}_i$ as

$$\mathbf{l}_{i+1} = \phi(\mathbf{b}_i + \mathbf{A}_i \mathbf{l}_i), \tag{7.10}$$

where $\mathbf{A}_i$ is a matrix of parameters with dimensions corresponding to the number of cells in the current and the previous layer and $\mathbf{b}_i$ is a vector of parameters with the same size as the current layer. ϕ is denominated the *activation function*, which is the only non-linear element in the ANN. Choosing the shape of the activation function is considered a hyperparameter on its own. Typical activation functions are the sigmoid function or the hyperbolic tangent, however, in this case we chose the commonly used [165] rectified linear activation (ReLU) function,

$$\phi_{\mathrm{ReLU}}(x) = \begin{cases} x & \text{if} \quad x \geq 0 \\ 0 & \text{if} \quad x < 0 \end{cases} . \tag{7.11}$$

This activation function presents the advantage that it is linear *almost everywhere*[1], making it easier to avoid saturation or numerical overflow problems when training.

As shown in figure 7.8 the input to the first layer corresponds to the features, and the output of the last layer corresponds to the targets. Typically the activation function is removed from the last layer in order for the targets to be able to take any value within the real numbers. The other intermediate layers are denominated *hidden layers*, with their corresponding *hidden cells*. The number of hidden layers and hidden cells per layer are hyperparameters of the model. Increasing these numbers linearly increases the number or parameters of the model also linearly, allowing much more fine control of the capacity of the model. Since ANNs designed this way are linear almost everywhere, their output can be interpreted as a sophisticated piecewise function formed of many linear regions (hyperplanes in an n_f-dimensional space) that increase in number as the number of hidden cells increase [168].

For a given number of model parameters, ANNs allow one to find much more complex non-linearities between the features and the targets than a polynomial model. The key difference is that while the polynomial models always make use of all possible combinations between the features, the ANN learns through the training process how to select

[1] *Almost everywhere* indicates everywhere except for some regions with measure zero.

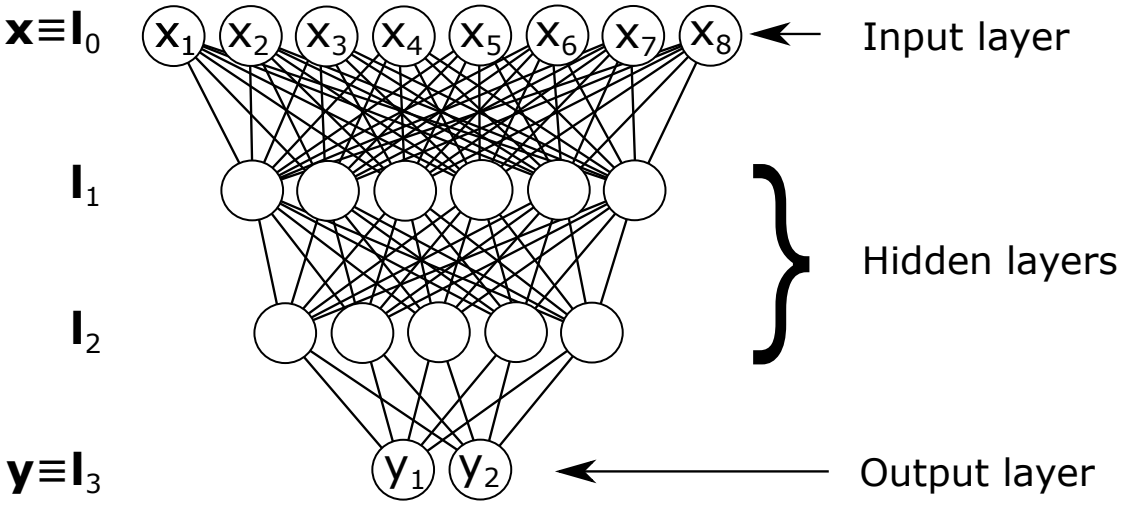


Figure 7.8: Diagram of a fully connected artificial neural network (ANN) for predicting 2 targets based on 8 features. This ANN contains two hidden layers, with 6 and 5 hidden cells, respectively. At each of the layers, the inputs are linearly combined to produced a set of outputs, which are modified using the activation function.

the key relationships between the features that allow prediction of the targets, being able to ignore other less-useful features or combination of features.

7.2.5 Iteratively training a model

In order for the polynomial, SVR and ANN models to work, the values of the optimal parameters $\mathbf{a}$ have to be found. While the minimization problem for the polynomial case has an analytical solution, in general, the optimization of the cost function has to be done iteratively, typically with variations of the gradient descent algorithm. Figure 7.9a shows the values of the cost function for the linear regression from figure 7.6a as function of the a and b parameters of the model $y = ax + b$, with the global minimum located at $a = 2.79$ and $b = 4.16$. The *gradient descent* algorithm involves iteratively updating the values of the parameters $\mathbf{a}^i \to \mathbf{a}^{i+1}$ in the opposite direction of the gradient of the cost function at that point,

$$\mathbf{a}^{i+1} = \mathbf{a}^i - \alpha \nabla J(\mathbf{a}^i), \tag{7.12}$$

where α is a hyperparameter known as the *learning rate*. This allows the algorithm to track the location of the minimum of the function as shown in figure 7.9b.

While the optimization of polynomial and SVR models can be typically expressed in terms of convex cost functions, where a single absolute minimum is present, ANNs typically present non-convex cost functions. This means that the minimization algorithm may get "stuck" in a local minimum, being unable to find the key relationships that allow making the optimal predictions. There are many variations of the gradient descent algorithm such as *stochastic gradient descent*, *mini-batch gradient descent*, and *adaptive methods* that dynamically change the learning rate [169], specifically designed to avoid

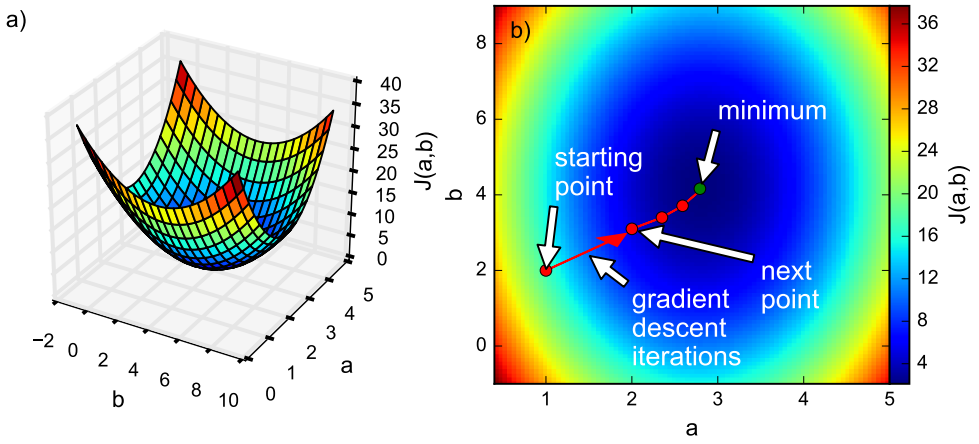


Figure 7.9: (a) Cost function as a function of the a and b parameters of a linear regression. The cost function shows a minimum at the coordinates a_0 and b_0 that solve the least-squares problem. (b) Gradient descent algorithm to minimize the cost function. By following the gradient from a random starting point, the location of the minimum is found.

converging to local minima or saddle points. However, finding a completely effective algorithm is still an open problem. In the results of this thesis we used the AdaGrad algorithm [169] combined with mini-batch gradient descent with batch sizes of 1000 training examples to train the neural networks until convergence.

Additionally, there are different techniques that can be used when training ANNs to avoid overfitting the data, including regularization [167] and dropout [170]. However, as we will show in section 7.4 overfitting was not a significant problem in our case, so none of these techniques were applied.

7.3 Proposed machine learning technique

The proposed technique based on machine learning aims to solve the problem outlined in section 7.1: to extract as much information as possible from the correlations between fast single-shot electron-beam diagnostics and x-ray properties, and to then use that information to make predictions of the x-ray properties using the fast diagnostics as the only input. A tool like this would allow researchers to keep using very valuable but slow single-shot x-ray diagnostics at high-repetition-rate XFELs. By slow single-shot x-ray diagnostics we refer to devices that are working well as single-shot diagnostics at the typical repetition rates of current XFELs (∼100 Hz), but that will not be able to operate at the full repetition rates provided by the next generation of XFELs: European XFEL (27000 events/s) and LCLS-II (1 MHz). On the other hand, even if not at the full repetition rate, the slow devices may still be operated at a lower repetition rate (for less than 0.1-1% of the events) at those facilities.

Under those conditions, the fast data stream containing the raw data would consist of two different types of events (figure 7.10). A small fraction of the events would be fully-characterized, containing single-shot information from all the detectors, just like at

current XFELs. The rest of the events would contain only single-shot information from the fast diagnostics that are scalable to high repetition rates.

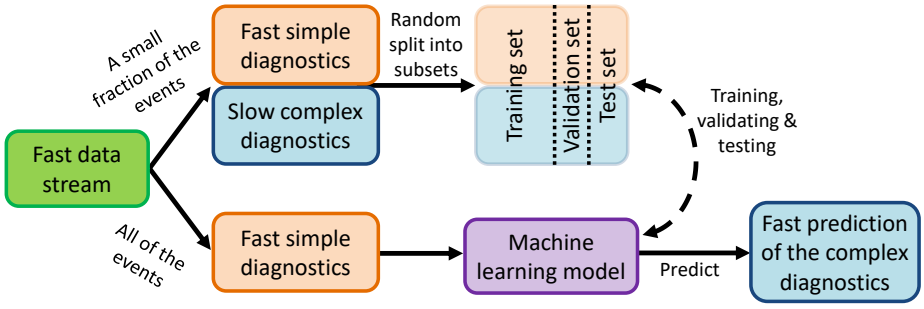


Figure 7.10: Proposed technique based on machine learning to overcome the limitation of low-repetition-rate devices at high-repetition-rate XFELs. A subset of fully-characterized events are used to infer relationships between the electron-beam parameters and the x-ray properties using a machine learning model. This model can then be used to predict the x-ray properties for the rest of the events. Figure extracted from [14].

The idea behind the technique (figure 7.10) is to use the fraction of fully characterized events to learn about the relationships between variables measured with the fast diagnostics and variables measured with the slow diagnostics to build a machine learning model. This machine learning model can be trained, validated and tested using different subsets of these events to estimate the accuracy of the model. Finally, the model can be used to predict the single-shot values of the variables that would have been measured with the slow complex diagnostics for the rest of the events, filling the data gaps, and increasing the amount of "fully characterized" events by a factor of 100 to 1000. As a proof of concept we implemented the proposed technique using data obtained in the AMO instrument [171] at LCLS [5] in two different experimental campaigns in February and April of 2015, as described in the following subsections.

7.3.1 Variables involved in the process

As shown in figure 7.11 there are three different types of relevant variables that were recorded at LCLS for our experiment: the slow variables, the fast variables and the environmental variables. The experiment was performed using the double-bunch double-slotted foil mode to generate pairs of multicolour x-ray pulses (section 3.2.3) although data was also recorded in single-pulse mode.

Slow single-shot variables from complex diagnostics

The slow variables measure single-shot properties making use of devices that will not scale well to high repetition rates. In this case, the devices were used to extract spectral information and timing information (blue in figure 7.11).

The spectral diagnostics consisted of an x-ray grating spectrometer (section 3.3.2) for the first experimental campaign, and an x-ray electron TOF spectrometer (section 3.3.3)

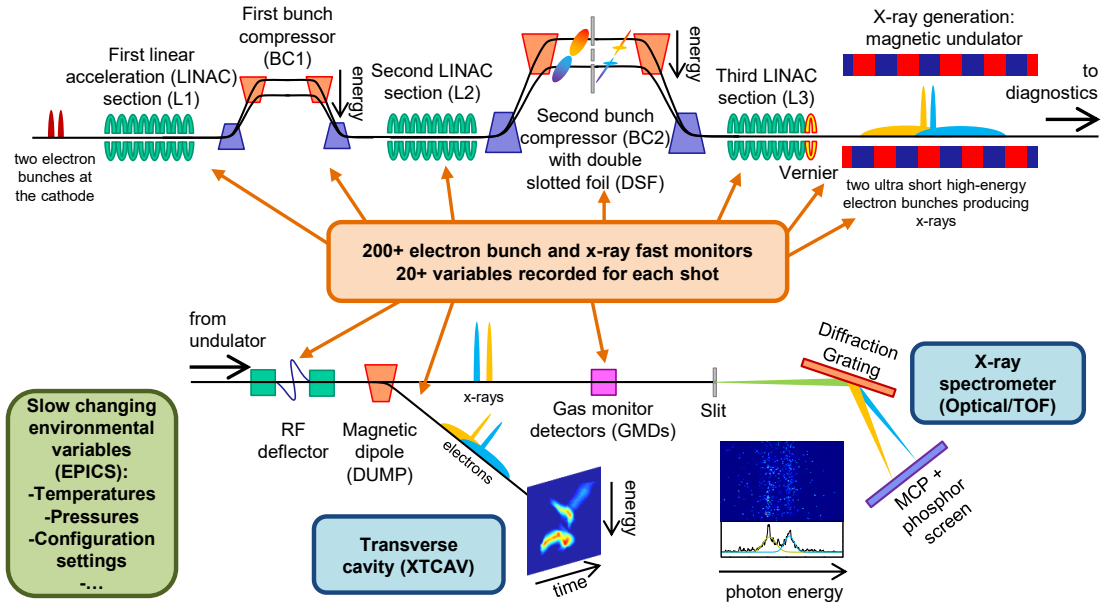


Figure 7.11: Overview of the experimental setup for machine learning predictions. More than 200 single-shot variables (orange) measuring electron-bunch and some x-ray properties with fast detectors that would scale well to MHz. At LCLS only approximately 20 of these single-shot variables are recorded by default. Complex detectors (blue) measuring single-shot x-ray properties that may not scale well to high repetition rates. Slow-changing variables (green) that only need to be measured at a low repetition rate. Figure extracted from [14].

for the second experimental campaign. In the case of the electron TOF spectrometer, the x-ray spectral information was obtained by analysing the kinetic energy of photoelectrons ionized by the x-ray pulses from the 2p shell of neon. Both spectrometers allowed us to obtain single-shot spectral profiles with resolutions of 0.2 eV and 0.5 eV, respectively, although we found the signal to noise of the x-ray grating spectrometer to be approximately 16 times higher than that of the electron TOF. We performed single-shot estimation of the photon energy by retrieving the centre of a Gaussian fit with a mean error in the estimation of 0.07 eV and 0.12 eV for the optical and TOF spectrometers, respectively.

The timing information consisted of pump-probe delay values extracted from XTCAV images (section 3.3.8) and obtained from the electron-bunch current profiles as indicated in section 5.4.4. The resolution of the XTCAV images was approximately 1.6 fs per pixel, with an expected mean error in the measurement of the delay in the order of 1 fs. While the LCLS repetition rate is 120 Hz, at the time XTCAV was working only at 60 Hz due to synchronization limitations.

The single-shot photon energies, spectral shape and time delay were used as targets for our machine learning models.

Fast single-shot variables from simple diagnostics

The fast variables consisted mainly of more than 200 single-shot electron-beam properties (orange in figure 7.11) measured at different stages of the XFEL using beam-position monitors (section 3.3.9), which should scale well to the MHz repetition rates. Typically, less than 20 of those single-shot variables are recorded by default for all events. In this case, 16 of those variables, corresponding to the ones shown in figures 7.3, 7.4 and 7.5 were used as features for the predictions.

Additionally, information from the four gas-monitor detectors (section 3.3.1) was also used as features, as this type of detector should also scale well to high repetition rates. The inclusion of these features is of extreme importance as they are the only coming from x-ray diagnostics, instead of electron-beam diagnostics, and subsequently, the only features that can be sensitive to the single-shot stochastic fluctuations originated by the SASE process, which are only correlated to non-measurable microscopic properties of the electron bunch.

Environmental EPICS variables

The third type of variables used in the process (green in figure 7.11) correspond to environmental variables typically updated only at a few Hz of repetition rate. They contain information related to the feedback loops controlling the XFEL including position of translation stages, power supplies, magnetic fields and radiofrequency fields. They also include information about temperatures and pressures at different stages or chambers, as well as the status of beam blockers or other configuration values input manually by operators.

At LCLS more than 300 of these variables are recorded routinely by the Experimental Physics and Industrial Control System (EPICS) [172]. Since the EPICS variables only account for slow environmental changes, they were treated as single-shot features for the predictions. While, of course, these features cannot really contribute to predict single-shot fluctuations, as they are not truly single-shot variables, they can help follow the drift due to, for example, temperature or pressure changes, as well as changes in the configuration of the machine.

Additionally, a timestamp for each of the events was included as an environmental feature variable. This allows the machine learning models to relate events that are closer in time, and perform time-dependent predictions to easily take the drift into account.

7.3.2 Numerical treatment of the data

The procedure allowing the creation of the machine learning models able to predict the target variables using the features as input included many different steps as shown in figure 7.12. The related software was implemented partly at the LCLS servers using Python 2.7 and the software package *psana* [13], and partly on a standard consumer computer

using Python 2.7, the *Scikit-learn* toolkit (v.0.17.1) [173] and *Tensorflow* (v0.8.0) [174]. The different steps can be summarized in two main categories: the preprocessing of the data, and the training, validation and testing of the models.

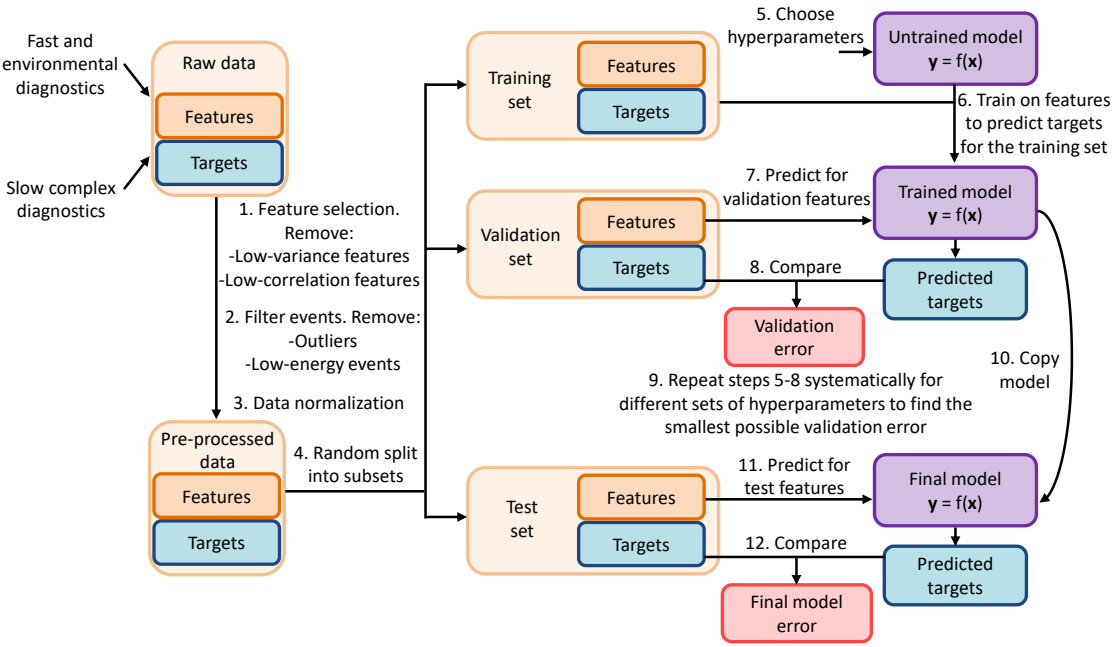


Figure 7.12: Flow chart of the preprocessing (steps 1-4), training (steps 5-6), validation (steps 5-10) and testing (steps 11-12) of the machine learning models. Figure extracted from reference [14].

Preprocessing of the data

The preprocessing of the data was performed independently for each type of prediction, but common to all the different models used for each type of prediction.

For each dataset, we performed feature selection (step 1 in figure 7.12) on the initial pool of more than 300 features, including fast variables and EPICS variables, to remove constant features and features taking a small number of discrete values (<10) with the effect of reducing the total number of features to around 90. We sorted the remaining features as function of the correlation with the variable to be predicted, typically keeping only the 40 features most highly correlated with the targets.

For each of the datasets, typically consisting of approximately $3 \cdot 10^4$ events, we removed those events for which the total x-ray energy was lower than 5 µJ accounting for approximately less than 10% of the events. We also removed the events for which the target values were separated from the median by more than four times the median absolute deviation to avoid including outliers in the predictions. This process (step 2 in figure 7.12) typically removed less than 1% of the events, except in the case of the spectral prediction, where the spectral spikes raised this value to approximately 20%.

The distributions of each of the features and the targets were normalized by subtracting the mean value and dividing by the standard deviation (step 3 in figure 7.12). While it was very relevant to do this for the features, we found that the normalization of the targets only improved performance when using artificial neural networks. In the case of the spectral predictions, we reduced the original 350 spectral components provided by the x-ray grating spectrometers to only 20 by using principal component analysis (PCA) [175]. This allowed the removal of random single photon spikes from the spectrum, improving the predictions.

The events contained in each dataset were then randomly split (step 4 in figure 7.12) into a training set (70% of the events), a validation set (15% of the events) and a test set (15% of the events).

Training, validation and testing of the models

We used a linear model, a quadratic model, an SVR and an ANN as our base models for each of the predictions. We also tried the K-neighbours and the decision tree regressor models, but none of them achieved competitive results compared to the other models. The steps of training, validation and testing of the models were done independently for each of the models and each of the predictions. The algorithm followed is summarized in figure 7.12 (steps 5-10).

For a chosen set of hyperparameters, we trained the corresponding machine learning models using the features and the targets from the events in the training set. After training a model with a given set of hyperparameters, the mean absolute error of the model was calculated by making predictions using features from the events in the validation set and comparing them to the targets from the same events. This error is called the *validation error*. In the case of the spectral shape, we used a similarity function between the vectors containing the 350 spectral components for the measured $\mathbf{V}_\mathrm{m}$ and the predicted $\mathbf{V}_\mathrm{p}$ data to assess the quality of the predictions,

$$\text{Agreement} = \frac{2\mathbf{V}_\mathrm{m} \cdot \mathbf{V}_\mathrm{p}}{\|\mathbf{V}_\mathrm{m}\|^2 + \|\mathbf{V}_\mathrm{p}\|^2}. \tag{7.13}$$

Different variations of the hyperparameters were selected using an informal grid search, training the model for each of them with the training set, and calculating the validation error on the validation set. Finally, the set of hyperparameters achieving the smallest validation error (or largest agreement), while at the same time not overfitting the training data (similar error or agreement for the validation set and the training set) were selected as the final optimal hyperparameters. The final values of the hyperparameters for each of the models (SVR and ANN) and each of the predictions are shown in table 7.1.

	SVR			ANN			
	C	ϵ	γ	Number of hidden layers	Number of hidden cells per layer	Batch size	Training steps
Single-pulse photon energy	100	0.08	0.005	2	[10,5]	1000	5000
Single-pulse spectrum	100	0.3	0.005	3	[50,50,20]	1000	5000
Double-pulse delay	100	0.4	0.005	2	[50,10]	1000	2400
Double-pulse photon energy	100	0.5	0.005	2	[20,5]	1000	2000

Table 7.1: Optimal hyperparameters found for the SVR and the ANN models for each of the predictions. The linear and quadratic models do not have hyperparameters.

Once the optimal set of hyperparameters was chosen for each model, the mean absolute error on the predictions was calculated making use of the test set, which was kept isolated during the training and validation processes (steps 11-12 in figure 7.12).

7.4 Results of the predictions

In order to prove the effectiveness of the technique, we trained the four models to predict photon energy and spectral shape in a single-pulse mode, time delay in a two-colour double-pulse mode and the individual photon energies of each of the pulses in the same double-pulse mode. The results are summarized in table 7.2 comparing the mean absolute error and agreement of the predictions to the mean absolute error of the distributions. This comparison is very relevant, as the error of the distributions would be the average single-shot error obtained in the case of taking the average value of the distribution as the nominal value of the variable for all of the events for which information from the corresponding diagnostic is not available. Each of the individual types of prediction are discussed in the following sections.

7.4.1 Photon energy prediction for a single pulse

The photon energy of the single x-ray pulses was obtained by fitting the spectral profiles measured with the x-ray grating spectrometer (figure 7.13a). In order to provide a larger dynamic range for the predictions, we scanned the photon energy by means of an electron-energy scan (triangular periodic shape, 1-minute period with over 10 periods) a full amplitude of 60 MeV. This yielded a photon-energy distribution with a mean absolute error of 5.6 eV (figure 7.13b).

In this case, all four models were able to achieve a mean error in the predictions of around 0.3 eV (table 7.2 and figure 7.13c). This mean prediction error implies a

Test set (Training set)	Initial distribution	Linear model	Quadratic model	Support vector regressor	Artificial neural network
Mean error of single-pulse photon energy [eV]	5.62	0.29 (0.28)	0.30 (0.24)	0.32 (0.27)	0.30 (0.29)
Shape agreement of single-pulse spectrum	67%	88% (88%)	94% (95%)	95% (95%)	97% (97%)
Mean error of double-pulse delay [fs]	6.82	2.07 (2.04)	1.67 (1.58)	1.67 (1.57)	1.59 (1.52)
Mean error of double-pulse photon energy [eV]	pulse 1: 1.45 pulse 2: 1.03	0.47 (0.49) 0.44 (0.44)	0.49 (0.48) 0.41 (0.39)	0.50 (0.48) 0.41 (0.39)	0.46 (0.47) 0.40 (0.40)

Table 7.2: Summary of the resulting mean absolute error or mean agreement for the predictions of each of the x-ray properties obtained for each of the four models. The first column corresponds to the mean error/agreement obtained by comparing the single-shot measured values with the corresponding mean value. The main values provided correspond to predictions on the test set, while the values in brackets correspond to predictions on the training set.

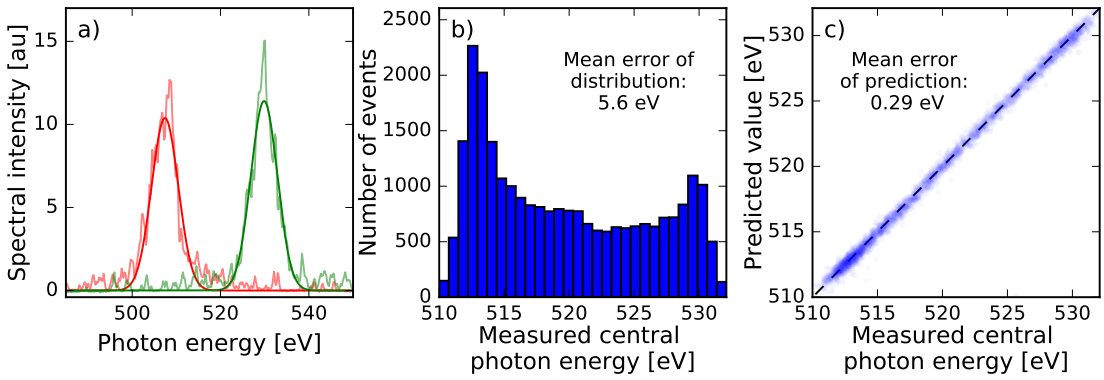


Figure 7.13: Prediction of the photon energy for a single pulse. (a) Two samples of single-pulse spectra (light red and light green) with their corresponding Gaussian fits (red and green). (b) Distribution of measured single-pulse photon energies obtained from the centre of the Gaussian fits for the dataset. (c) Comparison between the measured and the predicted photon-energy values using a linear model. Figure extracted from reference [14].

reduction of the mean error of the distribution by a factor of 20. While the scan of the electron energy artificially increased the width of the initial distribution, we also measured the mean error corresponding to the jitter when the electron-beam energy is not scanned to be 1 eV approximately. This means that an improvement of a factor of at least 3-4 would still be expected when not scanning the photon energy.

It is not surprising that the prediction of the photon energy is accurate, as one of the features included in the model was the single-shot electron-beam energy measured with a BPM. It is well known that changes in the photon energy E_{ph} and the electron-beam energy E_{e} are related in first order approximation as

$$\frac{\Delta E_{\mathrm{ph}}}{E_{0,\mathrm{ph}}} = 2\frac{\Delta E_{\mathrm{e}}}{E_{0,\mathrm{e}}}. \tag{7.14}$$

In fact, values of the electron-beam energy are routinely used to sort the data to achieve a similar effect to sorting as a function of the photon energy. However, we observed that by training our models using that electron-beam energy as the only input, the mean absolute error achieved was in the order of 0.7 eV, and in fact it took including up to 20 input features to obtain a mean error rounding to 0.30 eV. This indicates that even for well-known relationships like this one, it is preferable to use a machine learning model, which can take advantage of all the other existing correlations.

Comparing the results across all four different models (table 7.2) we observe that the error of the predictions is very similar in all cases. The fact that the linear model provides the best results of all and that the quadratic model clearly overfits the training data (much lower error for the training set than for the test set), suggests that most of the useful correlations in this case are essentially linear. This is also consistent with the fact that the optimal ANN was very small (2 hidden layers, with 10 and 5 hidden cells respectively) providing a low capacity model.

7.4.2 Spectral-shape prediction for a single pulse

Considering how well the individual photon-energy prediction worked, we attempted to build models to directly predict the spectral shape for the same dataset. In particular we built a model with 20 outputs corresponding to the first 20 principal components combining the original 350 individual spectral components using PCA [175]. The distribution of agreements for each type of model are shown in figure 7.14a.

Contrary to the prediction of the photon energy, this problem is much more non-linear, limiting the predictions of the linear model to an average of 88% agreement, with the other models achieving mean agreements above 94% (table 7.2). Specifically, the relatively deep ANN (three layers, with 50, 50 and 20 hidden cells per layer) is able to model the non-linearities with a mean agreement of 97%, with 86% of the events showing an agreement larger than 96%. Such high accuracy can be better illustrated by looking at the sample predictions for 96%, 97%, 98% and 99% (figures 7.14b-e), where

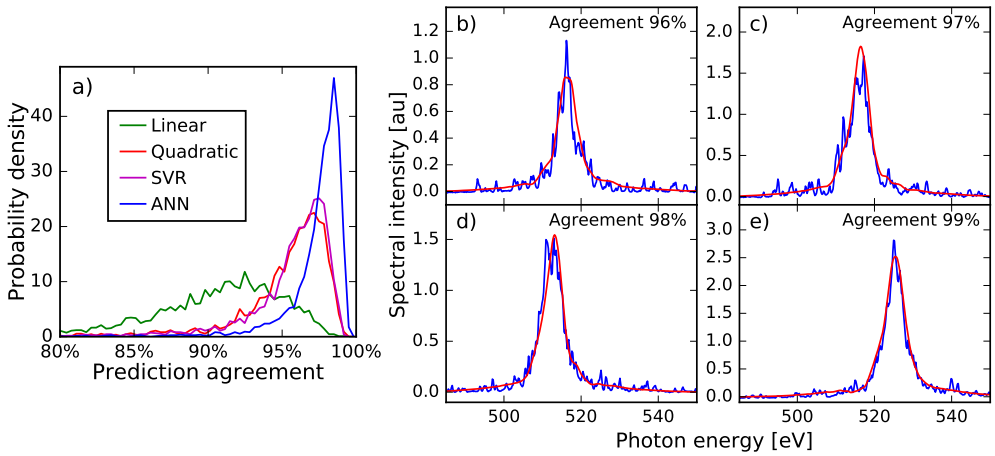


Figure 7.14: Spectral shape prediction. (a) Distribution of agreement of the predictions for the test set using the different models. (b-e) Examples of predicted (red) and measured (blue) spectra for different agreements using a neural network, presented to illustrate the accuracy corresponding to the typical agreement values for the ANN predictions. Figure extracted from reference [14].

even the lowest agreement shows a quantitative match between measured and predicted spectra, containing more information about the asymmetry of the spectrum than could have been obtained just by fitting a curve.

Comparing the results for the training set and the test set in table 7.2, we observe that none of the models seem to be overfitting the training data. This is probably a consequence of the complexity of the problem and implies that given more training, more features or more complex models, even better results could be achieved. This is typically known as a *high-bias* [164] situation. In any case, it is worth noting that the spectral shape is very sensitive to SASE fluctuations and that, unless more x-ray diagnostics are included as features, it may become very difficult for any machine learning model to gain access to the necessary correlations to accurately predict the final spectral shape. This is not a problem for this dataset, for which the pulses are driven to saturation, containing multiple SASE spikes per pulse, but would be very challenging for non-saturated shorter pulses showing individual SASE spikes.

Apart from filling data gaps when the repetition rate does not allow single-shot measurements for all the shots, this type of prediction may also be very useful for absorption experiments. In this type of experiments the spectrum transmitted through a sample has to be compared to the fluctuating incident spectrum. This is typically achieved by recording reference spectra and sorting/binning it as a function of one or two variables [11]. However, as we already mentioned in section 5.6.1 sorting as function of more than two variables to better characterize the incident spectra is typically not possible due to the *curse of dimensionality*. Alternatively, models for spectral prediction could be trained on the reference spectra to predict the single-shot incident spectral shape

for each of the events where the x-rays are transmitted through the sample, allowing one to calculate the single-shot transmission.

7.4.3 Time-delay prediction

We also applied our technique to predict the time-delay values obtained from XTCAV images as described in section 7.3.1. This is of critical importance for time-resolved experiments because, as we can observe in figures 7.15a-b, the time delay may even change sign within the same dataset due to the jitter. In fact the distribution of time-delay values for the events in a single run of our dataset (figure 7.15c) presents a FWHM of more than 25 fs, corresponding to a mean absolute error of 6.8 fs.

Table 7.2 and figures 7.15d-g show the results of the predictions for the different models. In this case, as the physical processes involved in the fluctuations of the delay are much more complex, the non-linear models achieved better results, with mean errors under 1.7 fs for all of them. In particular, by comparing the residuals around the perfect correlation line (dashed lines) in figures 7.15d-g, we observe that the neural network presents the most symmetric behaviour, which also leads to the smallest mean error of 1.59 fs.

In this case, most of the models seem to overfit the training data slightly. This could be an indication that the training of the models could benefit from having access to more training data, typical of *high-variance* [164] situations. In order to study this effect, we looked at the learning curves: the comparison between the error of the training set (dashed lines) and the error of the validation set (solid lines) as a function of the amount of training data (figure 7.15h). This figure shows that, for the available data, the ANN has not fully converged to a single value and that having access to more training data would improve the results further. In any case, it is worth noting that all the non-linear models are able to provide mean errors smaller than 1.8 fs with less than 5000 training events, with the SVR performing best in this sense, able to optimally adapt its capacity based on the amount of training data to make predictions with errors smaller than 1.8 fs with less than 2000 events. Even at low repetition rates of 100 Hz, the required amount of events for this kind of prediction could be recorded in less than a minute.

This type of prediction opens up the possibility to use XTCAV information at high-repetition-rate XFELs for all the events, even if XTCAV itself can only operate at a reduced repetition rate. Considering that we believe the accuracy of the measured values not to be smaller than 1 fs, predict the values with an accuracy in the order of 1.5 fs is very close to being as good as truly measuring the values. Furthermore, the predictions will allow recovering XTCAV timing information, duplicating the number of events for LCLS experiments performed before June 2016, when XTCAV was working

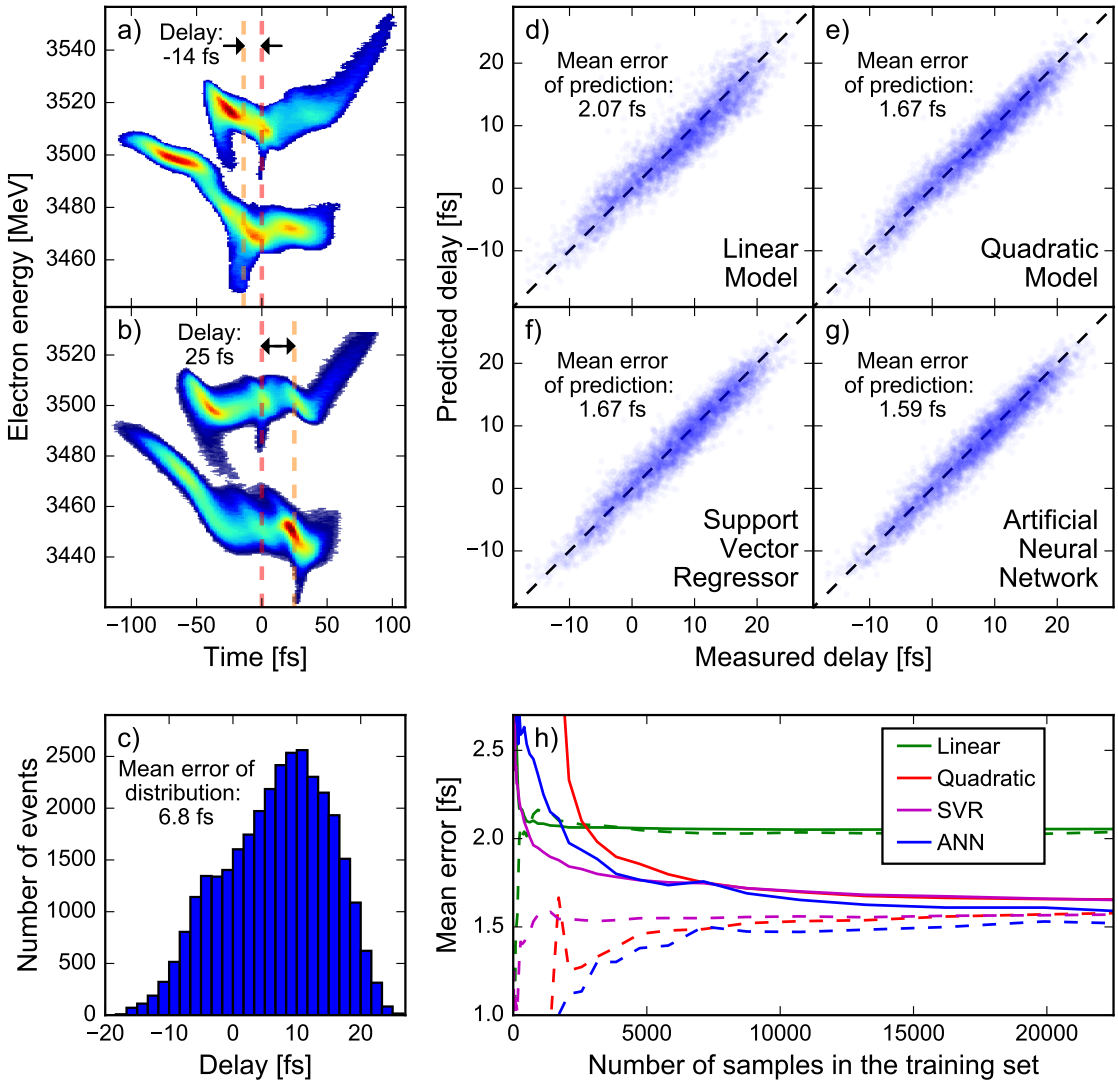


Figure 7.15: Prediction of the time delay. (a-b) Two sample XTCAV images obtained from the same dataset showing opposite values of the delay. (c) Distribution of measured time-delay values obtained from the XTCAV images for the dataset used. (d-g) Predicted values compared to the measured values for each of the four different models. (h) Learning curve showing the training error (dashed lines) and the validation error (solid lines) for each of the four models for increasing values of the number of samples in the training set. Figure extracted from reference [14].

only for alternating events, at 60 Hz, with the rest of the fast diagnostics still being recorded at 120 Hz.

7.4.4 Photon-energy prediction for a double pulse

Finally, we trained models to predict the individual photon energies in the double-pulse mode tuned around the oxygen K-edge. In this case, the electron TOF spectrometer was used to measure the spectral shape and a double Gaussian fit was applied to obtain the single-shot photon energy of each of the pulses, as shown in figure 7.16a. In this case, the electron-beam energy was being scanned by a full amplitude of only 20 MeV, which led to mean errors of the distribution of the photon energy for each of the pulses of 1.45 and 1.03 eV. Similarly to the single-pulse photon energy predictions, all four methods showed similar results, with the neural network achieving the lowest mean errors of 0.46 and 0.40 eV, respectively.

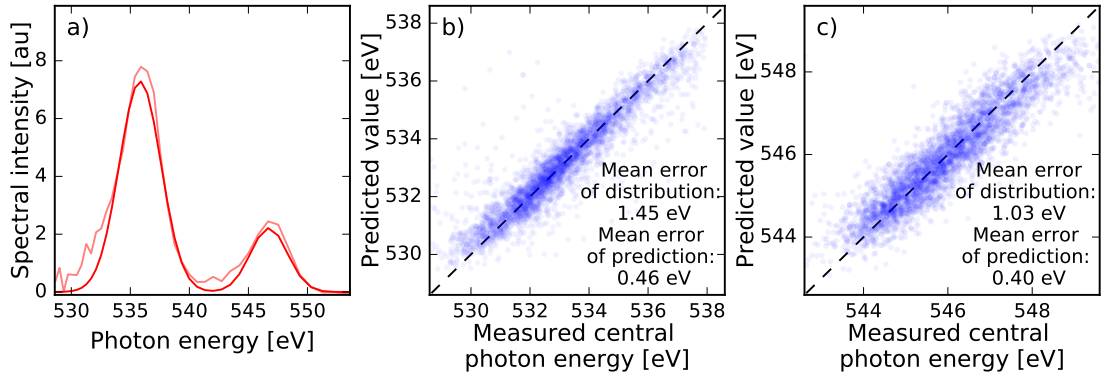


Figure 7.16: Prediction of the photon energy for a double pulse. (a) A sample of double-pulse spectra (light red) overlaid with its corresponding double-Gaussian fit. (b-c) Comparison between the measured photon energy and the predicted photon energy for the low-energy and the high-energy pulses, respectively, using an artificial neural network. Figure extracted from reference [14].

These errors are however still larger by a factor of ∼1.5 than the prediction errors obtained for a single pulse (0.3 eV). We believe that the difference is related to the lower signal to noise of the electron TOF spectrometer compared to that of the x-ray grating spectrometer, which provided less accurate values in the fitting process. Additionally, while the average total x-ray energy was approximately equal to 30 µJ in both cases, in this case it had to be divided between the two pulses, reducing the amount of signal available for each of them, and increasing the shot noise. In summary, we believe the prediction of the models is limited by the accuracy of the values provided by the fitting process.

We also attempted to predict the spectral shape in the double-pulse mode, however, the results were not as satisfactory as for a single pulse. Overall, the model could successfully predict the position of the peaks, but not the relative amplitude. The first possible reason to explain this could again be the lower signal to noise of the spectrometer

and the increased shot noise. The most plausible reason on the other hand, has to do with the stochastic nature of the SASE lasing process. As we have indicated earlier, most of the fast variables that we use as features measure electron-bunch properties and not x-ray properties, with the only exception being the gas detectors measuring the total x-ray energy. In the single-pulse spectral prediction this was not a problem, as the models could learn how to extract the information about the amplitude of the single peak from the gas detector signals. However, with two independent amplitudes to predict and a single value from the gas detectors measuring the total x-ray energy, the model is not fed with any feature allowing it to gain information about the relative amplitude of the pulses. This problem could be solved in the multiple-undulator multiple-pulse approach (section 3.2.2), for example, by placing additional gas detectors in between the undulators, which would allow us to also measure the total pulse energy of the first pulse alone.

These kind of considerations should be taken into account when designing new XFELs, in order to record as much single-shot information as reasonably possible at all different stages of the XFEL. In particular, new detectors could be strategically placed to highlight some key correlations with final x-ray properties, even if *a priori* those correlations cannot be fully understood.

7.5 Conclusions and outlook

We have demonstrated that some properties of the electron-beam trajectories, measurable using fast beam-position monitors, contain much information about important x-ray properties of the generated pulses, including spectral and temporal information. In particular, we have shown, using data from LCLS, that machine learning models can be trained to make single-shot predictions of these properties using information about the electron beam, simple x-ray diagnostics and environmental variables as the only input. The only requirement is that a small amount of experimental data ($\sim$5000 events, $\sim$1 minute), containing full information about the x-ray properties to be predicted together with the input variables, is available for training the models.

Using trained models we have successfully made accurate predictions of time delay between x-ray pulses, photon-energy values for double pulses and spectral shape. These predicted variables are currently measured at existing XFELs making use of very complex detectors (XTCAV and imaging spectrometers) that are very limited in terms of repetition rate (< 1 kHz). This limitation will make the use of these detectors very challenging at the next generation of high-repetition-rate XFELs based on superconducting LINACs working at up to MHz repetition rates. Our proposed prediction technique, recently published [14], offers a solution to this problem by recording information from those diagnostics only for a fraction of the events, recorded at a lower repetition rate, and then making predictions on the variables for the rest of the events. Furthermore, the technique could also be used to predict properties that may have to otherwise be

measured invasively, such as the x-ray spectral shape before a sample, required in absorption experiments to calculate the transmission. Additionally the technique may be used to re-analyse data from past experiments, for which some of the detectors, such as XTCAV, may not have been recording data at the full repetition rate due to technical reasons. Using the predictions would effectively allow us to double or triple the amount of events available for the data analysis of those experiments.

Now that the proof of principle has been demonstrated and published, the next critical step will be to make it available to XFEL users. While some users may be able to implement their own version of the models, ideally XFEL facilities should provide their own robust and fully-tested implementation of the machine learning algorithms. While none of the specific details have been decided yet, the possibility of providing such tools at LCLS-II has been discussed with LCLS machine and data acquisition scientists (some of them are members of our collaboration) who have been interested in the technique since the beginning of its development. Additionally, the possibility of embedding the machine learning predictions with automatic training at the European XFEL has already been discussed at the 2017 European XFEL user's meeting (25-27 January). This sets a very optimistic landscape for the future usage of the developed technique, one of the main results of this Ph.D.

Chapter 8

Conclusion

Through the chapters of this thesis the author has shown that XFELs are, and will continue to be, critical instruments for ultrafast x-ray science. Current XFELs, such as LCLS, already allow the generation of multi-colour pairs of few-femtosecond x-ray pulses tunable across a wide range (∼400 eV - 20 keV) of the x-ray spectrum and with an adjustable time delay tunable to femtosecond resolution at repetition rates of up to 120 Hz.

In contrast to optical lasers, XFELs are very large machines built up from hundreds of distinct sophisticated devices and across hundreds of metres, or even kilometres, of length. As a consequence, XFELs are prone to significant shot-to-shot fluctuations depending on the XFEL configuration, which are of up to 1% in the photon energy, 100% in the pulse energy and ∼10 femtoseconds in the relative arrival time of x-ray pulses when generating multiple few-fs pulses from the same bunch.

Time-resolved x-ray-only experiments aiming to measure nuclear dynamics in the 10 to 1000 fs timescales can usually benefit from a more stable operation by using longer, few-tens of femtoseconds pulses consisting of multiple SASE spikes saturating the emission. This usually reduces the shot-to-shot pulse energy fluctuations down to approximately 10% and the photon energy fluctuations to a small fraction of a percent. The decreased values allow in many cases (especially for non-resonant experiments) one to completely ignore the fluctuations, and work with nominal machine values of pump-probe delay, the photon energy, and the pulse energy. In these cases, the data analysis may be as simple as averaging together all data taken in an individual experimental run, and comparing it across different runs for different sets of nominal parameters.

Few-fs experiments, however, cannot afford to ignore the typical large fluctuations existent in few-fs XFEL configurations, and require a full x-ray shot-to-shot characterization to first understand the character and correlations of the fluctuations, and second, process the recorded data according to these fluctuations. This also applies to resonant

experiments, that require fixed precise values of the photon energy, and to experiments measuring highly non-linear effects, for which small fluctuations in the pulse energy can lead to much larger fluctuations in the signal. Additionally, optical-pump x-ray-probe experiments, for which the timing synchronization presents fluctuations of more than 500 fs FWHM, also require single-shot time-stamping of the XFEL shots, even for experiments in the 100-fs range.

The author has shown in this thesis that it is possible to non-invasively perform a complete single-shot characterization of the relevant x-ray properties: photon energy, pulse energy, and timing information, to enable few-fs ultrafast experiments at XFELs. This will allow end users to directly measure, for example, few-fs electron-correlation dynamics in small molecules such as glycine or isopropanol, as we have proposed for a future experiment. Making use of spectral information obtained with x-ray spectrometers (both electron time-of-flight and grating spectrometers), total x-ray energy information from gas detectors, and time-delay information from XTCAV using advanced retrieval techniques, we measured the fluctuations, studied the correlations between the different parameters and developed multidimensional binning and reference subtraction techniques to allow the extraction of small signals from data with a high degree of fluctuations. We applied these techniques in four different experiments at LCLS and SACLA to first determine whether or not the expected signals were hidden by the fluctuations, and then clean the signal from the effect of the fluctuations.

XFEL science and in general all of the proposed analysis techniques benefit from having large amounts of data. While the author has only demonstrated the necessity of single-shot x-ray characterization at current XFELs, working at a maximum of 120 Hz of repetition rate, this requirement will be carried over to the next generation of XFELs, including LCLS-II and the European XFEL, which will make use of superconducting LINACs to provide x-ray pulses at up to MHz repetition rates. Due to the increased statistics and the implementation of new techniques such as laser-dressed bunches [176, 177] potentially allowing the generation of attosecond pulses in the soft-x-ray region, the next generation of XFELs will revolutionize few-fs ultrafast x-ray science.

On the other hand, some of the complex devices required to perform the full single-shot x-ray characterization cannot be easily operated at the increased repetition rate of the next generation of XFELs, and will be unable to monitor the jitter of the x-ray properties for all x-ray pulses. This is the case of, for example, XTCAV and the "time tools", currently used to monitor fluctuations in temporal properties, and some x-ray spectrometers, used to measure spectral properties. The author has presented a technique based on machine learning that allows single-shot predictions of these variables to be made using mainly electron-beam information as single-shot input. This can be achieved by training the machine learning model using a small number (<5000) of fully characterized events that can be obtained at a low repetition rate. In the work presented here, the technique has been used to predict both spectral properties (central

photon energy and spectral shape) and temporal properties (time delay between two x-ray pulses), which are sufficient to perform few-fs x-ray-only experiments. This will open up new opportunities to apply machine learning in XFEL science, to deal with and make the most out of the large amounts of data that will be generated at the next generation of XFELs.

Considering the contents of the thesis and the conclusions presented above, the author would like to summarize the main conclusion of this thesis in a single sentence: single-shot x-ray characterization is critical to the success of any XFEL experiment and, in particular, for few-fs time-resolved experiments, and should not be taken lightly.

First of all, as x-ray sources, XFELs are not as refined as x-ray synchrotrons and, in order to propose a successful XFEL experiment, users must work with accelerator scientists and learn about some of the intricacies of the pulse generation processes, before deciding what the most appropriate configuration for the experiment is, and to design an experimental setup according to the expected fluctuations. This could be achieved by bringing together the user community with the accelerator community, as in some occasions users do not understand the limitations of XFELs, and accelerator scientists do not completely understand the needs of users.

Second, the peer-review process of XFEL articles should encourage authors to include explicit information about the assumptions made when estimating the pulse energy and the photon energy. XFEL fluctuations can have enormous effects in the experimental signals, so understanding the fluctuations, their effects, and being self-critical with the final results in terms of those fluctuations is of extreme importance and should always be discussed.

Finally, there should be an additional effort on the administration side of the facilities to allocate beamtime to perform basic but very important maintenance experiments. Important properties, such as the size of the focal spot or the overall transmission of the beamline, cannot be realistically measured during the allocated user beamtimes, having to rely in some occasions on values measured years before the user experiment. Additionally, most of the data that allowed study of the XFEL properties and even proving the machine learning predictions in this thesis were recorded in the context of other experiments and treated as a side matter. Much more accurate information and results could be obtained by dedicating more beamtime for the sole purpose of improving x-ray characterization and understanding the XFEL output and the correlations between the parameters, which could be used to provide the users with up to date exhaustive specifications of the different beamlines.

The author hopes that the results and conclusions presented in this thesis will help new XFEL users to understand the capabilities of XFELs, as well as provide new insights to experienced XFELs users on how to design their experiments to be robust against fluctuations and bring ideas on how to tackle the data analysis of fluctuating data. Furthermore, now that most of the practical aspects of the next generation of XFELs,

and in particular LCLS-II and European XFEL, are being determined, the author hopes that the contents of this thesis will provide some solutions to the new problems arising from the expected higher repetition rates. This should contribute towards the global goal of providing more versatile x-ray sources that will help push the boundaries of ultrafast x-ray science.

References

1. Remacle, F. & Levine, R. An electronic time scale in chemistry. *Proceedings of the National Academy of Sciences* **103,** 6793–6798 (2006).
2. Goulielmakis, E., Loh, Z.-H., Wirth, A., Santra, R., Rohringer, N., Yakovlev, V. S., Zherebtsov, S., Pfeifer, T., Azzeer, A. M., *et al.* Real-time observation of valence electron motion. *Nature* **466,** 739–743 (2010).
3. Marangos, J. Introduction to the new science with X-ray free electron lasers. *Contemporary Physics* **52,** 551–569 (2011).
4. Pellegrini, C. X-ray free-electron lasers: from dreams to reality. *Physica Scripta* **2016,** 014004 (2017).
5. Emma, P., Akre, R., Arthur, J., Bionta, R., Bostedt, C., Bozek, J., Brachmann, A., Bucksbaum, P., Coffee, R., *et al.* First lasing and operation of an ångstrom-wavelength free-electron laser. *Nature Photonics* **4,** 641–647 (2010).
6. Ishikawa, T., Aoyagi, H., Asaka, T., Asano, Y., Azumi, N., Bizen, T., Ego, H., Fukami, K., Fukui, T., *et al.* A compact X-ray free-electron laser emitting in the sub-angstrom region. *Nature Photonics* **6,** 540–544 (2012).
7. Ayvazyan, V., Baboi, N., Bähr, J., Balandin, V., Beutner, B., Brandt, A., Bohnet, I., Bolzmann, A., Brinkmann, R., *et al.* First operation of a free-electron laser generating GW power radiation at 32 nm wavelength. *The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics* **37,** 297–303 (2006).
8. Ackermann, W. a., Asova, G., Ayvazyan, V., Azima, A., Baboi, N., Bähr, J., Balandin, V., Beutner, B., Brandt, A., *et al.* Operation of a free-electron laser from the extreme ultraviolet to the water window. *Nature Photonics* **1,** 336–342 (2007).
9. Liekhus-Schmaltz, C. E., Tenney, I., Osipov, T., Sanchez-Gonzalez, A., Berrah, N., Boll, R., Bomme, C., Bostedt, C., Bozek, J. D., *et al.* Ultrafast isomerization initiated by X-ray core ionization. *Nature Communications* **6,** 8199 (2015).
10. Rohringer, N., Kimberg, V., Weninger, C., Sanchez-Gonzalez, A., Lutman, A., Maxwell, T., Bostedt, C., Montero, S. C., Lindahl, A., *et al.* in *X-Ray Lasers 2014* 201–207 (Springer, 2016).

11. Kimberg, V., Sanchez-Gonzalez, A., Mercadier, L., Weninger, C., Lutman, A., Ratner, D., Coffee, R., Bucher, M., Mucke, M., *et al.* Stimulated X-ray Raman scattering–a critical assessment of the building block of nonlinear X-ray spectroscopy. *Faraday Discussions* **194,** 305–324 (2016).

12. Sanchez-Gonzalez, A., Barillot, T., Squibb, R., Kolorenč, P., Agaker, M., Averbukh, V., Bearpark, M., Bostedt, C., Bozek, J., *et al.* Auger electron and photoabsorption spectra of glycine in the vicinity of the oxygen K-edge measured with an X-FEL. *Journal of Physics B: Atomic, Molecular and Optical Physics* **48,** 234004 (2015).

13. Damiani, D., Dubrovin, M., Gaponenko, I., Kroeger, W., Lane, T., Mitra, A., O'Grady, C., Salnikov, A., Sanchez-Gonzalez, A., *et al.* Linac Coherent Light Source data analysis using psana. *Journal of Applied Crystallography* **49,** 672–679 (2016).

14. Sanchez-Gonzalez, A., Micaelli, P., Olivier, C., Barillot, T., Ilchen, M., Lutman, A., Marinelli, A., Maxwell, T., Achner, A., *et al.* Accurate prediction of X-ray pulse properties from a free-electron laser using machine learning. *Nature Communications* **8,** 15461 (2017).

15. Berrah, N., Sanchez-Gonzalez, A., Jurek, Z., Obaid, R., Xiong, H., Squibb, R., Osipov, T., Lutman, A., Fang, L., *et al.* Femtosecond-resolved observation of the fragmentation of buckminsterfullerene following X-ray multiphoton ionization. *In Preparation* (2017).

16. Sanchez-Gonzalez, A., Johnson, A. S., Fitzpatrick, A., Hutchison, C. D., Fare, C., Cordon-Preciado, V., Dorlhiac, G., Ferreira, J., Morgan, R. M., *et al.* Coincidence timing of femtosecond optical pulses in an X-ray free electron laser. *In Preparation* (2017).

17. Hutchison, C. D., Cordon-Preciado, V., Morgan, R. M., Nakane, T., Ferreira, J., Dorlhiac, G., Sanchez-Gonzalez, A., Johnson, A. S., Fitzpatrick, A., *et al.* X-ray free electron laser crystal structures of dark and light states of a reversibly photoswitching fluorescent protein at room temperature. *In Preparation* (2017).

18. Simpson, E. R., Sanchez-Gonzalez, A., Siegel, T., Diveki, Z., Hutchinson, S. E., Miseikis, L., Strüber, C., Austin, D. R. & Marangos, J. P. *A beamline for attosecond pump-probe experiments: towards tracking ultrafast electron dynamics in atoms and molecules* in *SPIE Optical Engineering+ Applications* (2015), 958402–958402.

19. Simpson, E., Sanchez-Gonzalez, A., Austin, D., Diveki, Z., Hutchinson, S., Siegel, T., Ruberti, M., Averbukh, V., Miseikis, L., *et al.* Polarisation response of delay dependent absorption modulation in strong field dressed helium atoms probed near threshold. *New Journal of Physics* **18,** 083032 (2016).

20. Warren, B. E. *X-ray Diffraction* (Courier Corporation, 1969).

21. Bragg, W. H. & Bragg, W. L. The reflection of X-rays by crystals. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* **88,** 428–438 (1913).

22. Klug, H. P., Alexander, L. E., *et al.* X-ray diffraction procedures (1954).

23. Hecht, E. *Optics* (Addison-Wesley, 2002).

24. Millane, R. P. Phase retrieval in crystallography and optics. *JOSA A* **7,** 394–411 (1990).

25. Chapman, H. N., Fromme, P., Barty, A., White, T. A., Kirian, R. A., Aquila, A., Hunter, M. S., Schulz, J., DePonte, D. P., *et al.* Femtosecond X-ray protein nanocrystallography. *Nature* **470,** 73–77 (2011).

26. Chapman, H. N., Caleman, C. & Timneanu, N. Diffraction before destruction. *Phil. Trans. R. Soc. B* **369,** 20130313 (2014).

27. Kupitz, C., Basu, S., Grotjohann, I., Fromme, R., Zatsepin, N. A., Rendek, K. N., Hunter, M. S., Shoeman, R. L., White, T. A., *et al.* Serial time-resolved crystallography of photosystem II using a femtosecond X-ray laser. *Nature* **513,** 261–265 (2014).

28. Bonnelle, C. X-Ray spectroscopy. *Annual Reports Section" C"(Physical Chemistry)* **84,** 201–272 (1987).

29. Bohr, N. The spectra of helium and hydrogen. *Nature* **92,** 23 (1913).

30. Bohr, N. in *Niels Bohr, 1913-2013* 13–33 (Springer, 2016).

31. Schrödinger, E. An undulatory theory of the mechanics of atoms and molecules. *Physical Review* **28,** 1049 (1926).

32. Sakurai, J. J. & Tuan, S. F. *Modern quantum mechanics* (Addison-Wesley Reading, Massachusetts, 1985).

33. Atkins, P. W. & Friedman, R. S. *Molecular quantum mechanics* (Oxford university press, 2011).

34. Miller, K. & Green, A. Energy levels and potential energy curves for H2, N2, and O2 with an independent particle model. *The Journal of Chemical Physics* **60,** 2617–2626 (1974).

35. Einstein, A. Concerning an heuristic point of view toward the emission and transformation of light. *American Journal of Physics* **33,** 367 (1965).

36. Baltzer, P., Wannberg, B., Karlsson, L., Göthe, M. C. & Larsson, M. High-resolution inner-valence uv photoelectron spectra of the O 2 molecule and configuration-interaction calculations of Π u 2 states between 20 and 26 eV. *Physical Review A* **45,** 4374 (1992).

37. Hüfner, S. *Photoelectron spectroscopy: principles and applications* (Springer Science & Business Media, 2013).

38. Wagner, C. & Muilenberg, G. *Handbook of x-ray photoelectron spectroscopy: a reference book of standard data for use in x-ray photoelectron spectroscopy* (Physical Electronics Division, Perkin-Elmer Corp., 1979).

39. Scofield, J. H. Hartree-Slater subshell photoionization cross-sections at 1254 and 1487 eV. *Journal of Electron Spectroscopy and Related Phenomena* **8,** 129–137 (1976).

40. Yeh, J. & Lindau, I. Atomic subshell photoionization cross sections and asymmetry parameters: 1< Z< 103. *Atomic Data and Nuclear Data Tables* **32,** 1–155 (1985).

41. Bernath, P. F. *Spectra of atoms and molecules* (Oxford university press, 2015).

42. Koningsberger, D. & Prins, R. X-ray absorption: principles, applications, techniques of EXAFS, SEXAFS, and XANES (1988).

43. Center for X-ray optics, T. *X-ray database, Filter Transmission* `http://henke.lbl.gov`. Accessed: 2017-02-28.

44. Plekan, O., Feyer, V., Richter, R., Coreno, M., De Simone, M., Prince, K. & Carravetta, V. An X-ray absorption study of glycine, methionine and proline. *Journal of Electron Spectroscopy and Related Phenomena* **155,** 47–53 (2007).

45. Auger, P. The compound photoelectric effect. *J. Phys. Radium* **6,** 205–208 (1925).

46. Larsson, M., Baltzer, P., Svensson, S., Wannberg, B., Martensson, N., de Brito, A. N., Correia, N., Keane, M., Carlsson-Gothe, M., *et al.* X-ray photoelectron, Auger electron and ion fragment spectra of O2 and potential curves of O22+. *Journal of Physics B: Atomic, Molecular and Optical Physics* **23,** 1175 (1990).

47. Harris, L. Analysis of Materials by Electron-Excited Auger Electrons. *Journal of Applied Physics* **39,** 1419–1427 (1968).

48. Weber, R. E. & Peria, W. Use of LEED apparatus for the detection and identification of surface contaminants. *Journal of Applied Physics* **38,** 4355–4358 (1967).

49. Palmberg, P. & Rhodin, T. Auger electron spectroscopy of fcc metal surfaces. *Journal of Applied Physics* **39,** 2425–2432 (1968).

50. Palmberg, P. W. *Handbook of Auger Electron Spectroscopy: A Reference Book of Standard Data* (Physival Electronics Industries, 1972).

51. Drescher, M., Hentschel, M., Kienberger, R., Uiberacker, M., Yakovlev, V., Scrinzi, A., Westerwalbesloh, T., Kleineberg, U., Heinzmann, U., *et al.* Time-resolved atomic inner-shell spectroscopy. *Nature* **419,** 803–807 (2002).

52. Lakowicz, J. *Principles of Fluorescence Spectroscopy* (Springer US, 2007).

53. Sauer, M., Hofkens, J. & Enderlein, J. Basic principles of fluorescence spectroscopy. *Handbook of Fluorescence Spectroscopy and Imaging: From Single Molecules to Ensembles,* 1–30 (2011).

54. Indian, J. A new radiation. *Indian J. Phys* **2,** 387–398 (1928).

55. Einstein, A. On the quantum mechanics of radiation. *Physikalische Zeitschrift* **18,** 121–128 (1917).

56. Zewail, A. H. *Femtochemistry: Ultrafast Dynamics of the Chemical Bond:(Volumes I & II)* (World Scientific, 1994).

57. El-Sayed, M. A., Tanaka, I., Molin, Y., *et al. Ultrafast processes in chemistry and photobiology* (Blackwell Science Oxford, 1995).

58. Stolow, A., Bragg, A. E. & Neumark, D. M. Femtosecond time-resolved photoelectron spectroscopy. *Chemical Reviews* **104,** 1719–1758 (2004).

59. Urban, P. L., Chen, Y.-C. & Wang, Y.-S. *Time-resolved mass spectrometry: from concept to applications* (John Wiley & Sons, 2016).

60. Yan, Y. J. & Mukamel, S. Femtosecond pump-probe spectroscopy of polyatomic molecules in condensed phases. *Physical Review A* **41,** 6485 (1990).

61. Weiner, A. *Ultrafast optics* (John Wiley & Sons, 2011).

62. Zhao, K., Zhang, Q., Chini, M., Wu, Y., Wang, X. & Chang, Z. Tailoring a 67 attosecond pulse through advantageous phase-mismatch. *Optics letters* **37,** 3891–3893 (2012).

63. Fabris, D., Witting, T., Okell, W., Walke, D., Matia-Hernando, P., Henkel, J., Barillot, T., Lein, M., Marangos, J., *et al.* Synchronized pulses generated at 20 eV and 90 eV for attosecond pump–probe experiments. *Nature Photonics* **9,** 383–387 (2015).

64. Rohringer, N., Ryan, D., London, R. A., Purvis, M., Albert, F., Dunn, J., Bozek, J. D., Bostedt, C., Graf, A., *et al.* Atomic inner-shell X-ray laser at 1.46 nanometres pumped by an X-ray free-electron laser. *Nature* **481,** 488–491 (2012).

65. Thomson, E. *Roentgen-ray tube* US Patent 575,772. Jan. 1897.

66. And others. *Method and Apparatus for the Acceleration of Ions* US Patent 1,948,384. Feb. 1934.

67. Sokolov, A. A. & Ternov, I. M. *Synchrotron radiation* (Nauka Eds., 1966).

68. Ternov, I. Synchrotron radiation. *Physics-Uspekhi* **38,** 409 (1995).

69. Chao, A. W. Evaluation of beam distribution parameters in an electron storage ring. *Journal of Applied Physics* **50,** 595–598 (1979).

70. Winick, H., Brown, G., Halbach, K. & Harris, J. Wiggler and undulator magnets. *Physics Today* **34,** 50–63 (1981).

71. Katoh, M. Ultra-short pulses of synchrotron radiation on storage rings. *Japanese Journal of Applied Physics* **38,** L547 (1999).

72. Schoenlein, R., Chattopadhyay, S., Chong, H., Glover, T., Heimann, P., Shank, C., Zholents, A. & Zolotorev, M. Generation of femtosecond pulses of synchrotron radiation. *Science* **287,** 2237–2240 (2000).

73. Saes, M., van Mourik, F., Gawelda, W., Kaiser, M., Chergui, M., Bressler, C., Grolimund, D., Abela, R., Glover, T. E., *et al.* A setup for ultrafast time-resolved x-ray absorption spectroscopy. *Review of Scientific Instruments* **75,** 24–30 (2004).

74. Khan, S., Holldack, K., Kachel, T., Mitzner, R. & Quast, T. Femtosecond undulator radiation from sliced electron bunches. *Physical Review Letters* **97,** 074801 (2006).

75. Cavalleri, A., Rini, M., Chong, H., Fourmaux, S., Glover, T., Heimann, P., Kieffer, J. & Schoenlein, R. Band-selective measurements of electron dynamics in VO 2 using femtosecond near-edge X-ray absorption. *Physical Review Letters* **95,** 067405 (2005).

76. Bressler, C., Milne, C., Pham, V.-T., ElNahhas, A., van der Veen, R. M., Gawelda, W., Johnson, S., Beaud, P., Grolimund, D., *et al.* Femtosecond XANES study of the light-induced spin crossover dynamics in an iron (II) complex. *Science* **323,** 489–492 (2009).

77. Bressler, C. & Chergui, M. Molecular structural dynamics probed by ultrafast X-ray absorption spectroscopy. *Annual review of physical chemistry* **61,** 263–282 (2010).

78. *Europeal XFEL, "In Comparison,"* Internet: `http://www.xfel.eu/overview/in_comparison/`. [Jul. 29, 2014].

79. Schneider, J. R. *FLASH. The free-electron laser in Hamburg.* (Deutsches Elektronen-Synchrotron DESY, Germany, 2007).

80. Kim, K.-J. Three-dimensional analysis of coherent amplification and self-amplified spontaneous emission in free-electron lasers. *Physical Review Letters* **57,** 1871 (1986).

81. Emma, P., Bane, K., Cornacchia, M., Huang, Z., Schlarb, H., Stupakov, G. & Walz, D. Femtosecond and subfemtosecond x-ray pulses from a self-amplified spontaneous-emission–based free-electron Laser. *Physical Review Letters* **92,** 074801 (2004).

82. Wang, X., Freund, H., Harder, D., Miner Jr, W., Murphy, J., Qian, H., Shen, Y. & Yang, X. Efficiency and spectrum enhancement in a tapered free-electron laser amplifier. *Physical Review Letters* **103,** 154801 (2009).

83. Allaria, E., Diviacco, B., Callegari, C., Finetti, P., Mahieu, B., Viefhaus, J., Zangrando, M., De Ninno, G., Lambert, G., *et al.* Control of the polarization of a vacuum-ultraviolet, high-gain, free-electron laser. *Physical Review X* **4,** 041040 (2014).

84. Lutman, A. A., MacArthur, J. P., Ilchen, M., Lindahl, A. O., Buck, J., Coffee, R. N., Dakovski, G. L., Dammann, L., Ding, Y., *et al.* Polarization control in an X-ray free-electron laser. *Nature Photonics* **10,** 468–472 (2016).

85. Ratner, D., Abela, R., Amann, J., Behrens, C., Bohler, D., Bouchard, G., Bostedt, C., Boyes, M., Chow, K., *et al.* Experimental demonstration of a soft x-ray self-seeded free-electron laser. *Physical Review Letters* **114,** 054801 (2015).

86. Li, X., L'Huillier, A., Ferray, M., Lompré, L. & Mainfray, G. Multiple-harmonic generation in rare gases at high laser intensity. *Physical Review A* **39,** 5751 (1989).

87. Yu, L. H. Generation of intense UV radiation by subharmonically seeded single-pass free-electron lasers. *Physical Review A* **44,** 5178 (1991).

88. Allaria, E., Appio, R., Badano, L., Barletta, W., Bassanese, S., Biedron, S., Borga, A., Busetto, E., Castronovo, D., *et al.* Highly coherent and stable pulses from the FERMI seeded free-electron laser in the extreme ultraviolet. *Nature Photonics* **6,** 699–704 (2012).

89. Castagna, J., Murphy, B., Bozek, J. & Berrah, N. *X-ray split and delay system for soft x-rays at LCLS* in *Journal of Physics: Conference Series* **425** (2013), 152021.

90. Bozek, J., Castagna, J., Fang, L., Hui, Z., Kukk, E., Murphy, B. & Berrah, N. *X-ray split and delay device for ultrafast x-ray science at the AMO instrument at LCLS* in *Journal of Physics: Conference Series* **635** (2015), 012018.

91. Wöstmann, M., Mitzner, R., Noll, T., Roling, S., Siemer, B., Siewert, F., Eppenhoff, S., Wahlert, F. & Zacharias, H. The XUV split-and-delay unit at beamline BL2 at FLASH. *Journal of Physics B: Atomic, Molecular and Optical Physics* **46,** 164005 (2013).

92. Lutman, A., Coffee, R., Ding, Y., Huang, Z., Krzywinski, J., Maxwell, T., Messerschmidt, M. & Nuhn, H.-D. Experimental demonstration of femtosecond two-color x-ray free-electron lasers. *Physical Review Letters* **110,** 134801 (2013).

93. Ding, Y., Huang, Z. & Ruth, R. D. Two-bunch self-seeding for narrow-bandwidth hard x-ray free-electron lasers. *Phys. Rev. ST Accel. Beams* **13,** 060703 (6 June 2010).

94. Marinelli, A., Ratner, D., Lutman, A., Turner, J., Welch, J., Decker, F.-J., Loos, H., Behrens, C., Gilevich, S., *et al.* High-intensity double-pulse X-ray free-electron laser. *Nature Communications* **6,** 6369 (2015).

95. Lutman, A. A., Maxwell, T. J., MacArthur, J. P., Guetg, M. W., Berrah, N., Coffee, R. N., Ding, Y., Huang, Z., Marinelli, A., *et al.* Fresh-slice multicolour X-ray free-electron lasers. *Nature Photonics* **10,** 745–750 (2016).

96. Gahl, C., Azima, A., Beye, M., Deppe, M., Döbrich, K., Hasslinger, U., Hennies, F., Melnikov, A., Nagasono, M., *et al.* A femtosecond X-ray/optical cross-correlator. *Nature Photonics* **2,** 165–169 (2008).

97. Bionta, M. R., Lemke, H., Cryan, J., Glownia, J., Bostedt, C., Cammarata, M., Castagna, J.-C., Ding, Y., Fritz, D., *et al.* Spectral encoding of x-ray/optical relative delay. *Optics express* **19,** 21855–21865 (2011).

98. Harmand, M., Coffee, R., Bionta, M., Chollet, M., French, D., Zhu, D., Fritz, D., Lemke, H., Medvedev, N., *et al.* Achieving few-femtosecond time-sorting at hard X-ray free-electron lasers. *Nature Photonics* **7,** 215–218 (2013).

99. Hartmann, N., Helml, W., Galler, A., Bionta, M., Grünert, J., Molodtsov, S., Ferguson, K., Schorb, S., Swiggers, M., *et al.* Sub-femtosecond precision measurement of relative X-ray arrival time for free-electron lasers. *Nature Photonics* **8,** 706–709 (2014).

100. Cooper, B., Kolorenč, P., Frasinski, L. J., Averbukh, V. & Marangos, J. P. Analysis of a measurement scheme for ultrafast hole dynamics by few femtosecond resolution X-ray pump–probe Auger spectroscopy. *Faraday discussions* **171,** 93–111 (2014).

101. Hau–Riege, S., Bionta, R., Ryutov, D. & Krzywinski, J. Measurement of x-ray free-electron-laser pulse energies by photoluminescence in nitrogen gas. *Journal of Applied Physics* **103,** 053306 (2008).

102. Boyle, W. S. & Smith, G. E. Charge coupled semiconductor devices. *Bell Labs Technical Journal* **49,** 587–593 (1970).

103. Kita, T., Harada, T., Nakano, N. & Kuroda, H. Mechanically ruled aberration-corrected concave gratings for a flat-field grazing-incidence spectrograph. *Applied Optics* **22,** 512–513 (1983).

104. Nakano, N., Kuroda, H., Kita, T. & Harada, T. Development of a flat-field grazing-incidence XUV spectrometer and its application in picosecond XUV spectroscopy. *Applied Optics* **23,** 2386–2392 (1984).

105. Tokushima, T., Harada, Y., Watanabe, M., Takata, Y., Ishiguro, E., Hiraya, A. & Shin, S. Design of a flat field spectrometer for soft X-ray emission spectroscopy. *Surface Review and Letters* **9,** 503–508 (2002).

106. Beiersdorfer, P., Magee, E., Träbert, E., Chen, H., Lepson, J., Gu, M.-F. & Schmidt, M. Flat-field grating spectrometer for high-resolution soft x-ray and extreme ultraviolet measurements on an electron beam ion trap. *Review of Scientific Instruments* **75,** 3723–3726 (2004).

107. Wiza, J. L. Microchannel plate detectors. *Nuclear Instruments and Methods* **162,** 587–601 (1979).

108. Ceglio, N. *X-ray CCD camera* tech. rep. (1985).

109. Katayama, T., Owada, S., Togashi, T., Ogawa, K., Karvinen, P., Vartiainen, I., Eronen, A., David, C., Sato, T., *et al.* A beam branching method for timing and spectral characterization of hard X-ray free-electron lasers. *Structural Dynamics* **3,** 034301 (2016).

110. West, B. G. *Time-to-digital converter* US Patent 6,501,706. Dec. 2002.

111. Kornilov, O., Eckstein, M., Rosenblatt, M., Schulz, C., Motomura, K., Rouzée, A., Klei, J., Foucar, L., Siano, M., *et al.* Coulomb explosion of diatomic molecules in intense XUV fields mapped by partial covariance. *Journal of Physics B: Atomic, Molecular and Optical Physics* **46,** 164028 (2013).

112. Cryan, J. P., Glownia, J. M., Andreasson, J., Belkacem, A., Berrah, N., Blaga, C. I., Bostedt, C., Bozek, J., Buth, C., *et al.* Auger electron angular distribution of double core-hole states in the molecular reference frame. *Physical Review Letters* **105,** 083004 (2010).

113. Fang, L., Hoener, M., Gessner, O., Tarantelli, F., Pratt, S. T., Kornilov, O., Buth, C., Gühr, M., Kanter, E. P., *et al.* Double core-hole production in N 2: beating the Auger clock. *Physical Review Letters* **105,** 083005 (2010).

114. Bozek, J. D. AMO instrumentation for the LCLS X-ray FEL. *The European Physical Journal-Special Topics* **169,** 129–132 (2009).

115. Wabnitz, H., Bittner, L., De Castro, A., Döhrmann, R., Gürtler, P., Laarmann, T., Laasch, W., Schulz, J., Swiderski, A., *et al.* Multiple ionization of atom clusters by intense soft X-rays from a free-electron laser. *Nature* **420,** 482–485 (2002).

116. Hoener, M., Fang, L., Kornilov, O., Gessner, O., Pratt, S. T., Gühr, M., Kanter, E. P., Blaga, C., Bostedt, C., *et al.* Ultraintense X-ray induced ionization, dissociation, and frustrated absorption in molecular nitrogen. *Physical Review Letters* **104,** 253002 (2010).

117. Kruit, P. & Read, F. Magnetic field paralleliser for 2π electron-spectrometer and electron-image magnifier. *Journal of Physics E: Scientific Instruments* **16,** 313 (1983).

118. Mucke, M., Förstel, M., Lischke, T., Arion, T., Bradshaw, A. M. & Hergenhahn, U. Performance of a short "magnetic bottle" electron spectrometer. *Review of Scientific Instruments* **83,** 063106 (2012).

119. McFarland, B., Berrah, N., Bostedt, C., Bozek, J., Bucksbaum, P., Castagna, J., Coffee, R., Cryan, J., Fang, L., *et al. Experimental strategies for optical pump–soft x-ray probe experiments at the LCLS* in *Journal of Physics: Conference Series* **488** (2014), 012015.

120. McFarland, B., Farrell, J., Miyabe, S., Tarantelli, F., Aguilar, A., Berrah, N., Bostedt, C., Bozek, J., Bucksbaum, P., *et al.* Ultrafast X-ray Auger probing of photoexcited molecular dynamics. *Nature Communications* **5,** 4235 (2014).

121. Murphy, B., Osipov, T., Jurek, Z., Fang, L., Son, S.-K., Mucke, M., Eland, J., Zhaunerchyk, V., Feifel, R., *et al.* Femtosecond X-ray-induced explosion of C60 at extreme intensity. *Nature Communications* **5,** 4281 (2014).

122. Manura, D. J. & Dahl, D. A. Simion Version 8.0/8.1 User Manual. *Scientific Instrument Services* **5** (2011).

123. Eppink, A. T. & Parker, D. H. Velocity map imaging of ions and electrons using electrostatic lenses: Application in photoelectron and photofragment ion imaging of molecular oxygen. *Review of Scientific Instruments* **68,** 3477–3484 (1997).

124. Sobottka, S. E. & Williams, M. B. Delay line readout of microchannel plates. *IEEE Transactions on Nuclear Science* **35,** 348–351 (1988).

125. Jagutzki, O., Cerezo, A., Czasch, A., Dorner, R., Hattas, M., Huang, M., Mergel, V., Spillmann, U., Ullmann-Pfleger, K., *et al.* Multiple hit readout of a microchannel plate detector with a three-layer delay-line anode. *IEEE Transactions on Nuclear Science* **49,** 2477–2483 (2002).

126. RoentDek Handels GmbH. *MCP Delay Line Detector Manual. Version 11.0.1304.1* `http://www.tokyoinst.co.jp/product_file/file/RD01_man01_ja.pdf`.

127. Ding, Y., Behrens, C., Emma, P., Frisch, J., Huang, Z., Loos, H., Krejcik, P. & Wang, M. Femtosecond x-ray pulse temporal characterization in free-electron lasers using a transverse deflector. *Physical Review Special Topics-Accelerators and Beams* **14,** 120701 (2011).

128. Behrens, C., Decker, F.-J., Ding, Y., Dolgashev, V., Frisch, J., Huang, Z., Krejcik, P., Loos, H., Lutman, A., *et al.* Few-femtosecond time-resolved measurements of X-ray free-electron lasers. *Nature Communications* **5,** 3762 (2014).

129. Smith, S., Hoobler, S., Johnson, R. G., Straumann, T., Young, A., Lill, R. M., Morrison, L. H., Norum, E., Sereno, N., *et al. Commissioning and performance of LCLS cavity BPMs* in *Proceedings of the 2009 Particle Accelerator Conference, Vancouver, BC* (2009), 754–756.

130. Osipov, T., Rescigno, T. N., Weber, T., Miyabe, S., Jahnke, T., Alnaser, A., Hertlein, M. P., Jagutzki, O., Schmidt, L. P. H., *et al.* Fragmentation pathways for selected electronic states of the acetylene dication. *Journal of Physics B: Atomic, Molecular and Optical Physics* **41,** 091001 (2008).

131. Li, Z., Inhester, L., Liekhus-Schmaltz, C., Curchod, B., Snyder Jr, J. W., Medvedev, N., Cryan, J., Osipov, T., Pabst, S., *et al.* Ultrafast Isomerization in Acetylene Dication: To Be or Not To Be. *Preprint arXiv:1605.05707* (2016).

132. Hishikawa, A., Matsuda, A., Fushitani, M. & Takahashi, E. J. Visualizing recurrently migrating hydrogen in acetylene dication by intense ultrashort laser pulses. *Physical Review Letters* **99,** 258302 (2007).

133. Ibrahim, H., Wales, B., Beaulieu, S., Schmidt, B. E., Thiré, N., Fowe, E. P., Bisson, É., Hebeisen, C. T., Wanie, V., *et al.* Tabletop imaging of structural evolutions in chemical reactions demonstrated for the acetylene cation. *Nature Communications* **5,** 4422 (2014).

134. Jiang, Y., Rudenko, A., Herrwerth, O., Foucar, L., Kurka, M., Kühnel, K., Lezius, M., Kling, M. F., van Tilborg, J., *et al.* Ultrafast extreme ultraviolet induced isomerization of acetylene cations. *Physical Review Letters* **105,** 263002 (2010).

135. Osipov, T., Cocke, C., Prior, M., Landers, A., Weber, T., Jagutzki, O., Schmidt, L., Schmidt-Böcking, H. & Dörner, R. Photoelectron-photoion momentum spectroscopy as a clock for chemical rearrangements: Isomerization of the di-cation of acetylene to the vinylidene configuration. *Physical Review Letters* **90,** 233002 (2003).

136. Xu, H., Okino, T., Nakai, K., Yamanouchi, K., Roither, S., Xie, X., Kartashov, D., Zhang, L., Baltuska, A., *et al.* Two-proton migration in 1, 3-butadiene in intense laser fields. *Physical Chemistry Chemical Physics* **12,** 12939–12942 (2010).

137. Xu, H., Okino, T. & Yamanouchi, K. Ultrafast hydrogen migration in allene in intense laser fields: Evidence of two-body Coulomb explosion. *Chemical Physics Letters* **469,** 255–260 (2009).

138. Duflot, D., Robbe, J.-M. & Flament, J.-P. Ab initio study of the acetylene and vinylidene dications fragmentation. *The Journal of Chemical Physics* **102,** 355–363 (1995).

139. Zyubina, T., Dyakov, Y., Lin, S., Bandrauk, A. & Mebel, A. Theoretical study of isomerization and dissociation of acetylene dication in the ground and excited electronic states. *The Journal of Chemical Physics* **123,** 134320 (2005).

140. Kirkpatrick, P. & Baez, A. V. Formation of optical images by x-rays. *JOSA* **38,** 766–774 (1948).

141. Trofimov, A. & Schirmer, J. Molecular ionization energies and ground-and ionic-state properties using a non-Dyson electron propagator approach. *The Journal of Chemical Physics* **123,** 144115 (2005).

142. Vacher, M., Steinberg, L., Jenkins, A. J., Bearpark, M. J. & Robb, M. A. Electron dynamics following photoionization: Decoherence due to the nuclear-wave-packet width. *Physical Review A* **92,** 040502 (2015).

143. Vacher, M., Bearpark, M. J., Robb, M. A. & Malhado, J. P. Electron Dynamics upon Ionization of Polyatomic Molecules: Coupling to Quantum Nuclear Motion and Decoherence. *Physical Review Letters* **118,** 083001 (2017).

144. Nordfors, D., Nilsson, A., Svensson, S., Gelius, U., Ågren, H., *et al.* X-ray excited photoelectron spectra of free molecules containing oxygen. *Journal of Electron Spectroscopy and Related Phenomena* **56,** 117–164 (1991).

145. Rightor, E., Hitchcock, A., Ade, H., Leapman, R., Urquhart, S., Smith, A., Mitchell, G., Fischer, D., Shin, H., *et al.* Spectromicroscopy of poly (ethylene terephthalate): comparison of spectra and radiation damage rates in X-ray absorption and electron energy loss. *The Journal of Physical Chemistry B* **101,** 1950–1960 (1997).

146. Kivimäki, A., Kempgens, B., Piancastelli, M., Neeb, M., Maier, K., Rüdel, A., Hergenhahn, U. & Bradshaw, A. Angle-resolved electronic decay of the π^*, σ^* and Rydberg resonances below the OK-edge in O 2. *Journal of Electron Spectroscopy and Related Phenomena* **93,** 81–88 (1998).

147. Ellis, K., Hall, R., Avaldi, L., Dawber, G., McConkey, A., Andric, L. & King, G. High resolution threshold photoelectron and photoion spectroscopy of oxygen in the 12-50 eV photon range. *Journal of Physics B: Atomic, Molecular and Optical Physics* **27,** 3415 (1994).

148. Hikosaka, Y., Aoto, T., Hall, R. I., Ito, K., Hirayama, R., Yamamoto, N. & Miyoshi, E. Inner-valence states of O2+ and dissociation dynamics studied by threshold photoelectron spectroscopy and a configuration interaction calculation. *Journal of Chemical Physics* **119,** 7693–7700 (2003).

149. Center for X-ray optics, T. *X-ray database, X-ray properties of the elements* `http://henke.lbl.gov`. Accessed: 2015-09-18.

150. Bozek, J. *The AMO Instrument & Science Drivers* `https://portal.slac.stanford.edu/sites/lcls_public/Instrument_Document_Repository/Bozek_Prop_WS_Jun08.pdf`.

151. *XTCAV - LCLS Data Analysis - SLAC Confluence* `https://confluence.slac.stanford.edu/display/PSDM/XTCAV`.

152. Cannington, P. & Ham, N. S. He (I) and He (II) photoelectron spectra of glycine and related molecules. *Journal of Electron Spectroscopy and Related Phenomena* **32,** 139–151 (1983).

153. Maclot, S., Piekarski, D. G., Domaracka, A., Mery, A., Vizcaino, V., Adoui, L., Martín, F., Alcamí, M., Huber, B. A., *et al.* Dynamics of glycine dications in the gas phase: ultrafast intramolecular hydrogen migration versus Coulomb repulsion. *The Journal of Physical Chemistry Letters* **4,** 3903–3909 (2013).

154. Gordon, M. L., Cooper, G., Morin, C., Araki, T., Turci, C. C., Kaznatcheev, K. & Hitchcock, A. P. Inner-shell excitation spectroscopy of the peptide bond: Comparison of the C 1s, N 1s, and O 1s spectra of glycine, glycyl-glycine, and glycyl-glycyl-glycine. *The Journal of Physical Chemistry A* **107,** 6144–6159 (2003).

155. Moddeman, W. E., Carlson, T. A., Krause, M. O., Pullen, B., Bull, W. & Schweitzer, G. Determination of the K—LL Auger Spectra of N2, O2, CO, NO, H2O, and CO2. *The Journal of Chemical Physics* **55,** 2317–2336 (1971).

156. Plekan, O., Feyer, V., Richter, R., Coreno, M., de Simone, M., Prince, K. C. & Carravetta, V. Investigation of the amino acids glycine, proline, and methionine by photoemission spectroscopy. *The journal of physical chemistry A* **111,** 10998–11005 (2007).

157. Weninger, C., Purvis, M., Ryan, D., London, R. A., Bozek, J. D., Bostedt, C., Graf, A., Brown, G., Rocca, J. J., *et al.* Stimulated electronic x-ray Raman scattering. *Physical Review Letters* **111,** 233902 (2013).

158. Weninger, C. & Rohringer, N. Stimulated resonant x-ray Raman scattering with incoherent radiation. *Physical Review A* **88,** 053421 (2013).

159. Couto, R. C., Guarise, M., Nicolaou, A., Jaouen, N., Chiuzbăian, G. S., Lüning, J., Ekholm, V., Rubensson, J.-E., Såthe, C., *et al.* Coupled electron-nuclear dynamics in resonant 1 σ to 2π x-ray Raman scattering of CO molecules. *Physical Review A* **93,** 032510 (2016).

160. Henke, B. L., Gullikson, E. M. & Davis, J. C. X-ray interactions: photoabsorption, scattering, transmission, and reflection at E= 50-30,000 eV, Z= 1-92. *Atomic Data and Nuclear Data Tables* **54,** 181–342 (1993).

161. Zelmon, D. E., Small, D. L. & Page, R. Refractive-index measurements of undoped yttrium aluminum garnet from 0.4 to 5.0 μm. *Applied Optics* **37,** 4933–4935 (1998).

162. Medvedev, N., Ziaja, B., Cammarata, M., Harmand, M. & Toleikis, S. Electron Kinetics in Femtosecond X-Ray Irradiated SiO2. *Contributions to Plasma Physics* **53,** 347–354 (2013).

163. Medvedev, N. Femtosecond X-ray induced electron kinetics in dielectrics: application for FEL-pulse-duration monitor. *Applied Physics B* **118,** 417–429 (2015).

164. Murphy, K. P. *Machine learning: a probabilistic perspective* (MIT press, 2012).

165. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016).

166. Smola, A. J. & Schölkopf, B. A tutorial on support vector regression. *Statistics and Computing* **14,** 199–222 (2004).

167. Cheng, B. & Titterington, D. M. Neural Networks: A Review from a Statistical Perspective. *Statistical Science* **9,** 2–30 (Feb. 1994).

168. Montufar, G. F., Pascanu, R., Cho, K. & Bengio, Y. *On the number of linear regions of deep neural networks* in *Advances in Neural Information Processing Systems* (2014), 2924–2932.

169. Duchi, J., Hazan, E. & Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* **12,** 2121–2159 (2011).

170. Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15,** 1929–1958 (2014).

171. Ferguson, K. R., Bucher, M., Bozek, J. D., Carron, S., Castagna, J.-C., Coffee, R., Curiel, G. I., Holmes, M., Krzywinski, J., *et al.* The atomic, molecular and optical science instrument at the Linac Coherent Light Source. *Journal of Synchrotron Radiation* **22,** 492–497 (2015).

172. Dalesio, L. R., Hill, J. O., Kraimer, M., Lewis, S., Murray, D., Hunt, S., Watson, W., Clausen, M. & Dalesio, J. The experimental physics and industrial control system architecture: past, present, and future. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **352,** 179–184 (1994).

173. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., *et al.* Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12,** 2825–2830 (2011).

174. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., *et al.* Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *Preprint arXiv:1603.04467* (2016).

175. Jolliffe, I. *Principal component analysis* (Wiley Online Library, 2002).

176. Marinelli, A., Coffee, R., Vetter, S., Hering, P., West, G., Gilevich, S., Lutman, A., Li, S., Maxwell, T., *et al.* Optical shaping of X-ray free-electron lasers. *Physical Review Letters* **116,** 254801 (2016).

177. Huang, S., Ding, Y., Huang, Z. & Marcus, G. Generation of subterawatt-attosecond pulses in a soft x-ray free-electron laser. *Physical Review Accelerators and Beams* **19,** 080702 (2016).

Appendix A

List of figure permissions

The author presents below a comprehensive list of figures published elsewhere and reused with or without modification for this thesis. This list also includes figures made by the author that are already published. All other figures not listed here have been made by the author for their use in this thesis.

- Figure 2.1a: Material published more than 70 years ago. Permission for this particular request is granted for print and electronic formats.

  Source: [21]

- Figure 2.1b: Public domain figure by Matthias M modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:Bragg.svg`

- Figure 2.2: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4086520255677).

  Source: `https://commons.wikimedia.org/wiki/File:Bragg.svg`

- Figure 2.3: Public domain figure modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:Bohr_Model.svg`

- Figure 2.5: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4055470480339).

  Source: [34]

- Figure 2.6b: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4057650064822).

  Source: [36]

- Figures 2.8b, c and d: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4156710813494).

  Source: [44]

- Figure 2.9: Permissions obtained using the Copyright Clearance Center (Confirmation Number: 11635704).

  Source: [46]

- Figure 2.11a: Public domain figure by Hmilch modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:WaterCooledXrayTube.svg`

- Figure 2.11b: Public domain figure by LinguisticDemographer modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:TubeSpectrum.jpg`

- Figure 2.12a: Public domain figure by R. Bartolini.

  Source: `https://commons.wikimedia.org/wiki/File:Syncrotron_radiation_energy_flux.png`

- Figure 2.12b: Creative Commons figure by TNorth modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:Cyclotron.svg`

- Figure 2.12c: Creative Commons figure by Antivolt modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:Synchrotron.jpg`

- Figure 2.12d: Copyright © EPSIM 3D/JF Santarelli, Synchrotron Soleil. Redistribution, derivative work, commercial use, and all other use is permitted provided that the copyright holder is properly attributed.

  Source: `https://commons.wikimedia.org/wiki/File:Sch%C3%A9ma_de_principe_du_synchrotron.jpg`

- Figure 2.12f: Creative Commons figure by Bastian Holst modified by the author.

  Source: `https://commons.wikimedia.org/wiki/File:Undulator.png`

- Figure 2.13: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4061360613951).

  Source: [72]

- Figure 3.1: Modified figure with Copyright © European XFEL. Permission obtained by email from Frank Poppe (frank.poppe@xfel.eu). "Dear Alvaro, sure you get the permission-to reuse and modify. Attached you find the high-res PDF file. Just indicate European XFEL as copyright information".

  Source: [78]

- Figures 3.2, 3.3, 3.4, 3.5a and 3.5b: Permission obtained by email from Hannah Gerth (hannah.gerth@desy.de). "Dear Alvaro, I have asked in our group about your request and in the name of the whole FLASH team you have explicit permission to include figures of the mentioned report in the introductory sections of your thesis. We wish you all the best!".

  Source: [79]

- Figure 3.6: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080220454049).

  Source: [81]

- Figure 3.8: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080220579688).

  Source: [85]

- Figure 3.9: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080220690669).

  Source: [88]

- Figure 3.10: Under a CC BY-NC 3.0 license.

  Source: [89]

- Figure 3.11: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080230369354).

  Source: [92]

- Figure 3.12: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080230369354).

  Source: [92]

- Figure 3.14: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080230546582).

  Source: [95]

- Figure 3.15: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080231304593).

  Source: [101]

- Figure 3.17: Under a CC BY-NC 3.0 license.

  Source: [111]

- Figure 3.18: © IOP Publishing. Reproduced with permission. All rights reserved. Permission obtained by email from Kathryn Shaw (permissions@iop.org). "Dear Alvaro Sanchez-Gonzalez, Thank you for your request to reproduce IOP Publishing material in your thesis. We are happy to grant permission for the use you request on the terms set out below. -Terms indicate including a reference, and the copyright statement".

  Source: [117]

- Figures 3.19a and 3.19b: Reproduced with the permission of their original author Thomas R. Barillot.

  Source: Grant review presentation by Thomas Barillot (2014)

- Figure 3.20b: © 2011 IEEE. The IEEE does not require individuals working on a thesis to obtain a formal reuse license.

  Source: [125]

- Figure 3.21: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080241100397).

  Source: [98]

- Figure 3.22: Under a CC BY-NC 3.0 license.

  Source: [127]

- Figures 4.1c-d, 4.2, 4.3 and 4.4: Published in one of the author publications, but not originally made by the author. Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080241372008).

  Source: [9]

- Figures 5.1: Reused from an LCLS proposal. Original figure Bridgette Cooper, modified and reproduced by the author with permission.

  Source: Internal report by Dr. Bridgette Cooper

- Figures 5.2 and 5.3: Reused from an LCLS proposal. Made in collaboration Bridgette Cooper.

  Source: LCLS proposal

- Figure 5.15: Under a CC BY-NC 3.0 license.

  Source: [127]

- Figure 5.20: Published in one of the author publications and made by the author. Permissions obtained using the Copyright Clearance Center (Confirmation Number: 11635709).

Source: [12]

- Figure 6.1: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080250427296).

  Source: [64]

- Figure 6.2: Published in one of the author publications, but not originally made by the author. Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4080250638121).

  Source: [10]

- Figure 6.3: Published in one of the author publications, but not originally made by the author. Under a CC BY-NC 3.0 license.

  Source: [11]

- Figures 6.4, 6.5 and 6.6: Published in one of the author publications and made by the author. Under a CC BY-NC 3.0 license.

  Source: [11]

- Figures 6.9, 6.10, 6.11 and 6.12: Contents not published yet. Figures made by the author.

  Source: [16]

- Figure 7.7: Permission obtained using the Copyright Clearance Center's RightsLink service (Order number: 4067630691435).

  Source: [166]

- Figures 7.10, 7.11, 7.12, 7.13, 7.14, 7.15 and 7.16: Published in one of the author publications and made by the author. Under a CC BY license.

  Source: [14]