



LUND
UNIVERSITY

**Multi-spectral X-ray imaging with a
laser-plasma accelerator**

MASTER THESIS

by Andrea Angella

Supervisor: Olle Lundh
Co-supervisor: Giada Cantono
Examiner: Lars Rippe

LRAP 584
Division of Atomic Physics
Lund Institute of Physics, Lund University
Lund, Sweden 2022

Abstract

This master thesis focuses on the creation of an algorithm to perform multispectral imaging of a metallic target using X-rays. The generation of an ultrashort high-power low-divergent X-ray beam is performed by means of a Ti:Sapphire laser driven wake-field plasma accelerator. The X-rays result from synchrotron radiation, and are therefore broadband. The spectrum will be characterized, along with the transmittance of a variety of different metals, using a technique called "single-photon counting". An X-ray sensitive 4 MP CCD camera is used to perform a statistical analysis of the x-ray beam spectrum, and X-ray absorption spectroscopy (XAS) is used in order to identify a "fingerprint" for each metal. It is demonstrated that the proposed system allows multi-spectral imaging of a Ross filter, being capable of detecting a series of different metals with good spatial resolution.

Popular scientific summary

It is possible to determine an object's composition by simply shining light through it and observing how the frequency components of the light source (its spectrum) change after transmission. This is because different materials absorb light in unique ways, with some frequencies being absorbed more than others. For example, when looking through a pair of red colored glasses, the world appears red because only the red frequency component of the sunlight (which is white) is transmitted, whereas all other colors are mostly absorbed. This method of studying a material is called "absorption spectroscopy", and it is an extremely powerful tool which can even allow astronomers to determine the species of gases in nebulae that are thousands of light years away from Earth.

In this thesis work, the aim is to be able to determine the composition and location in space with good resolution of different metals in a 2D image. Visible light cannot be used because metals are opaque to this type of radiation. Instead, X-rays, a more energetic type of light, are produced by means of a technique called laser wakefield acceleration. An intense laser pulse is shot into a gas target, which turns into a plasma, i.e. a state of matter where electrons are separated from their nuclei. The pulse excites a plasma wave, and electrons are trapped behind the laser pulse as it propagates through the target, following in its wake like surfers on a sea wave, accelerating to extreme energies while simultaneously wiggling up and down. As they change direction in this wiggling motion, they lose energy in the form of X-rays. This radiation has many frequencies (it is broadband) and is incredibly short, less than 40 fs. To put this in perspective, light can travel about 300 000 km in a single second (more than 7 round trips around the Earth's equator!), but only travels about 12 microns in 40 fs: less than the width of a human hair.

To measure the spectrum of the X-rays, it is inconvenient to use conventional methods based on spectrometers which rely on physical dispersive elements such as prisms or crystals, because with these the spectrum can be measured only for single points in the sample, which would be very inconvenient for imaging. Instead, the technique hereby proposed uses nothing more than a CCD camera. When X-rays strike into the pixels of a camera, they generate a number of electrons which is proportional to their energy (and therefore also their frequency). This means that if a few X-ray photons are "sprinkled" here and there on the sensor, it is possible to derive the distribution of the light source frequencies by creating a histogram of these single-photon events, which is nothing but the spectrum! This technique is called single-photon counting, and an algorithm has been developed by the author for this purpose.

After calculating the spectrum before and after the sample in this way, another algorithm has been developed to automatically detect the presence and type of a metal in an image by comparing the two spectra and using the K-edge of a metal, a dip in the transmission spectrum which occurs at different energies for different metals, that can be used as a sort of fingerprint.

CONTENTS

List of Abbreviations	vii
1 Introduction	1
1.1 Historical background	2
1.2 X-ray spectroscopy and multi-spectral imaging	2
1.3 Thesis objective	3
1.4 Thesis outline	4
2 LWFA and X-ray generation	5
2.1 Principles of LWFA	6
2.1.1 Electromagnetic waves fundamentals	6
2.1.2 Gaussian beam	9
2.1.3 High intensity laser pulses	11
2.1.4 Laser-plasma interaction	12
2.1.5 Ionization of gases	14
2.1.6 Electron acceleration	16
2.2 X-ray betatron radiation	19
3 X-ray absorption spectroscopy and single-photon counting	23
3.1 X-ray absorption spectroscopy	24
3.1.1 Motivation for single-photon counting	24
3.1.2 K-edge detection	26
3.2 Energy deposition in CCD sensors	27
3.2.1 CCD sensor basics	27
3.2.2 Quantum efficiency	31
3.3 Single-photon counting	33
3.4 Noise and dynamic range	34
3.4.1 Shot noise	35
3.4.2 Thermal noise	36
3.4.3 Read-out noise	37
3.4.4 Dynamic range	38
4 Experimental setup	41
4.1 Laser system and target chamber	42
4.1.1 The Lund High-Power Laser Facility	42
4.1.2 Experimental setup	42
4.2 X-ray camera	45
4.3 Ross filter	46

5	Measurements and data analysis	49
5.1	The single-photon counting algorithm	50
5.1.1	Cluster search algorithm	50
5.1.2	Spectral measurement	54
5.2	Energy calibration	58
5.3	Metal detection algorithm	62
5.4	Increasing the spatial resolution	63
5.4.1	Overlapping sectoring	64
5.5	MATLAB code	65
6	Results and conclusion	67
6.1	Final results	68
6.2	Possible improvements	70
6.3	Conclusion and outlook	72
7	Acknowledgements	75
	Bibliography	77

*Dedicated to my mom and my dad who are the best
parents a son could ever ask for.*

LIST OF ABBREVIATIONS

a.u.	Arbitrary Units
CCD	Charge-Coupled Device
CMOS	Complementary Metal-Oxide Semiconductor
CPA	Chirped Pulse Amplification
CT	Computed Tomography
FWHM	Full Width Half Max
HWHM	Half Width Half Max
IR	InfraRed
LWFA	Laser WakeField Acceleration
MOSFET	Metal-Oxide Semiconductor Field Effect Transistor
MOS	Metal-Oxide Semiconductor
nMOS	n-type Metal-Oxide Semiconductor
pMOS	p-type Metal-Oxide Semiconductor
QE	Quantum Efficiency
RGB	Red Green Blue
SNR	Signal-to-Noise Ratio
UV	Ultra Violet
XAS	X-ray Absorption Spectroscopy

1

INTRODUCTION

This chapter introduces some background about plasma based acceleration, X-ray absorption spectroscopy and other important historical milestones, followed by a statement on the motivation behind this master thesis.

Contents

1.1	Historical background	2
1.2	X-ray spectroscopy and multi-spectral imaging . . .	2
1.3	Thesis objective	3
1.4	Thesis outline	4

1.1 Historical background

This thesis is made possible thanks to the work of many brilliant scientists of the past, whose discoveries still affect our world deeply to the present day.

To begin with, X-rays were discovered in 1895 by the German scientist Wilhelm Conrad Röntgen, who was later awarded the Nobel prize in 1902 [1]. X-rays is a form of highly penetrative electromagnetic radiation which has been used extensively in countless applications: medicine, astronomy, microscopy, airport security and many others. About half a century later, the first optical laser, based on a ruby crystal, was developed by Theodore H. Maiman [2] at Hughes Research Labs, which marked the beginning of a new era. Since this historical event, a great number of innovations flourished in the following years. Among the sheer amount of them, there is the invention of the Nd:YAG laser in 1964 by Joseph E. Geusic and Richard G. Smith at Bell Labs, and that of the Ti:Sapphire laser in 1982 by Peter F. Moulton of MIT's Lincoln Laboratory [3], both of which are used in today in the Terawatt laser laboratory at Lund University.

The principle of laser wakefield acceleration (LWFA) is used in this thesis work to generate the X-rays to be used for multispectral imaging purposes. It was introduced for the first time in 1979, pioneered by Toshiki Tajima and John M. Dawson [4]. The extremely high laser pulse intensities required for this process, however, would hardly be achievable if not for the work of Donna Strickland and Gerard Mourou in 1985 [5], Nobel prize winners in 2018.

Lastly, multispectral imaging could not be done if not for the invention of digital camera sensors such as the CCD, developed at Bell labs by George Smith and Willard Boyle in 1969 [6], who were later awarded the Nobel prize in 2009.

1.2 X-ray spectroscopy and multi-spectral imaging

X-ray absorption spectroscopy (XAS) is a broadly used method to investigate the properties of matter at the atomic level. The first time that XAS was

utilized using radiation coming from synchrotron accelerators was in 1974 [7], and since then it has found applications in the most heterogeneous fields of science. In this thesis work, XAS is used in conjunction with laser wakefield acceleration to identify the species of various metallic samples. Unlike for a visible light color camera, where a Bayer filter array (a mosaic of red, green and blue filters) is placed on top of the sensor to detect different frequencies, with X-rays this is not possible due to the lack of appropriate band-pass filters. For this reason, spectral measurements are done in this thesis work using a technique called single-photon counting. In a CCD camera, X-ray photons generate a number of electrons that is proportional to their energy, therefore each pixel can be used as a calorimeter, and the spectrum of the incoming radiation can be retrieved by creating a histogram of single photon events. This has been done in previous studies [8,9], but in this thesis work the technique is used in conjunction with multispectral imaging. Different metals absorb X-rays in uniquely distinctive ways, so it is possible to reconstruct a virtual multispectral image of an illuminated sample with a large number of spectral channels (as many as the number of metals that are detected in the image).

1.3 Thesis objective

The ultimate goal of this thesis is twofold:

1. Firstly, a method of measuring the spectrum of an X-ray source generated as betatron radiation from a laser wakefield accelerator has to be developed. The method proposed in this thesis is based on single-photon counting, which requires nothing more than use of a CCD detector (appropriately sensitive to X-ray radiation) and a specialized algorithm.
2. Secondly, an algorithm to perform X-ray absorption spectroscopy has to be developed to detect the presence and species of some types of metals in a sample, achieving relatively high spatial resolution without any need of scanning the sample with a "pencil beam".

The multispectral imaging technique presented in this thesis wants to propose itself as a valid, fast and low-cost method to achieve these tasks.

1.4 Thesis outline

The main body of the thesis is split into 5 chapters.

Chapter 2 starts with an introduction to all the physics basics that are needed to understand the principle of laser wakefield acceleration, used to generate X-rays which represent the light source for the multispectral imaging.

Chapter 3 introduces some background about X-ray absorption spectroscopy and how CCD cameras work in their possible use as calorimeters. Furthermore, the main idea behind how both the tasks described in the previous section can be achieved is established with single-photon counting and a metal detection algorithm based on the identification of the K-edge of different metals.

Chapter 4 introduces in great detail the entire experimental setup that was used, from the laser target chamber where the X-rays are generated, to specifications of the CCD camera.

Chapter 5 contains a description of both the single-photon counting algorithm and the metal detection algorithm. It shows in detail how the spectral measurements were achieved, and how the algorithm is able to produce the final results.

Chapter 6 concludes the thesis with images of the final results, and considerations on possible improvements that could have been made. It also includes some possible future applications that the technique could have in different scientific fields.

2

LWFA AND X-RAY GENERATION

This chapter introduces the theoretical basics of Laser WakeField Acceleration (LWFA) and the physics behind the generation of X-rays as betatron radiation.

Contents

2.1	Principles of LWFA	6
2.1.1	Electromagnetic waves fundamentals	6
2.1.2	Gaussian beam	9
2.1.3	High intensity laser pulses	11
2.1.4	Laser-plasma interaction	12
2.1.5	Ionization of gases	14
2.1.6	Electron acceleration	16
2.2	X-ray betatron radiation	19

2.1 Principles of LWFA

2.1.1 Electromagnetic waves fundamentals

The propagation of electromagnetic waves in space as a function of position \mathbf{r} and time t can be described with four vector fields: the electric field strength $\mathbf{E}(\mathbf{r}, t)$, the magnetic field strength $\mathbf{H}(\mathbf{r}, t)$, the electric displacement field $\mathbf{D}(\mathbf{r}, t)$ and the magnetic flux density $\mathbf{B}(\mathbf{r}, t)$ (sometimes also called "magnetic induction"). These quantities are related via the well-known Maxwell's equations:

$$\nabla \times \mathbf{E}(\mathbf{r}, t) = -\frac{\partial \mathbf{B}(\mathbf{r}, t)}{\partial t} \quad (2.1)$$

$$\nabla \times \mathbf{H}(\mathbf{r}, t) = \mathbf{J}(\mathbf{r}, t) + \frac{\partial \mathbf{D}(\mathbf{r}, t)}{\partial t} \quad (2.2)$$

$$\nabla \cdot \mathbf{D}(\mathbf{r}, t) = \rho(\mathbf{r}, t) \quad (2.3)$$

$$\nabla \cdot \mathbf{B}(\mathbf{r}, t) = 0, \quad (2.4)$$

where $\mathbf{J}(\mathbf{r}, t)$ is the electric current density (measured in A/m²) and $\rho(\mathbf{r}, t)$ is the volumetric charge density (measured in C/m³) of any external localized charges. If no sources are present, such as in a vacuum, both $\mathbf{J}(\mathbf{r}, t) = 0$ and $\rho(\mathbf{r}, t) = 0$.

Equation 2.1 is called Faraday's law of induction, equation 2.2 is Ampère's law, equation 2.3 is Gauss's law and equation 2.4 is Gauss's law for magnetism. The equations are hereby given in their differential form.

All the components of the electric field strength and the magnetic field strength must satisfy the wave equation:

$$\nabla^2 u(\mathbf{r}, t) - \frac{1}{c_0^2} \frac{\partial^2 u(\mathbf{r}, t)}{\partial t^2} = 0, \quad (2.5)$$

where $c_0 = \sqrt{\varepsilon_0 \mu_0} \approx 3 \times 10^8$ m/s is the speed of light in vacuum, and $u(\mathbf{r}, t)$ is a scalar field that represent any component of a vector field, e.g. the z-component $E_z(\mathbf{r}, t)$ of $\mathbf{E}(\mathbf{r}, t)$, if $\mathbf{E}(\mathbf{r}, t) = (E_x(\mathbf{r}, t), E_y(\mathbf{r}, t), E_z(\mathbf{r}, t))$.

$\varepsilon_0 \approx 8.854 \times 10^{-12}$ F/m indicates the permittivity of free space and $\mu_0 = 4\pi \times 10^{-7}$ H/m is the magnetic permeability of free space. The electric displacement

field and the magnetic flux density are related to the electric field strength and the magnetic field strength via the so-called constitutive relations:

$$\begin{aligned} \mathbf{D}(\mathbf{r}, t) &= \varepsilon_0 \mathbf{E}(\mathbf{r}, t) + \mathbf{P}(\mathbf{r}, t) = \varepsilon \mathbf{E}(\mathbf{r}, t) \\ \mathbf{B}(\mathbf{r}, t) &= \mu_0 \mathbf{H}(\mathbf{r}, t) + \mu_0 \mathbf{M}(\mathbf{r}, t) = \mu \mathbf{H}(\mathbf{r}, t), \end{aligned} \quad (2.6)$$

where $\mathbf{P}(\mathbf{r}, t)$ and $\mathbf{M}(\mathbf{r}, t)$ are, respectively, the macroscopic polarization density and the magnetization density of the medium in which the light is propagating. ε and μ indicate, respectively, the electric permittivity and magnetic permeability of the medium in which the light is propagating. In free space, $\mathbf{P}(\mathbf{r}, t)$ and $\mathbf{M}(\mathbf{r}, t)$ are zero, and $\varepsilon = \varepsilon_0$, $\mu = \mu_0$, so that the constitutive relations simplify to:

$$\begin{aligned} \mathbf{D}(\mathbf{r}, t) &= \varepsilon_0 \mathbf{E}(\mathbf{r}, t) \\ \mathbf{B}(\mathbf{r}, t) &= \mu_0 \mathbf{H}(\mathbf{r}, t). \end{aligned} \quad (2.7)$$

The power of the electromagnetic wave flows in the direction of the Poynting vector $\mathbf{S}(\mathbf{r}, t)$, which is defined as:

$$\mathbf{S}(\mathbf{r}, t) = \mathbf{E}(\mathbf{r}, t) \times \mathbf{H}(\mathbf{r}, t), \quad (2.8)$$

and has units of W/m^2 . The optical intensity $I(\mathbf{r}, t)$ of the electromagnetic wave can be obtained from the Poynting vector, as it is simply equal to the magnitude of its time average over a suitable time interval that is significantly larger than the period of the wave's oscillation:

$$I(\mathbf{r}, t) = \langle \mathbf{S}(\mathbf{r}, t) \rangle. \quad (2.9)$$

The magnitude of the electric field and magnetic field are related via a quantity called the electromagnetic impedance $\eta = \sqrt{\mu/\varepsilon}$ as follows:

$$\begin{aligned} |\mathbf{H}(\mathbf{r}, t)| &= \frac{1}{\eta} |\mathbf{E}(\mathbf{r}, t)| \\ |\mathbf{E}(\mathbf{r}, t)| &= \eta |\mathbf{H}(\mathbf{r}, t)|, \end{aligned} \quad (2.10)$$

so that the magnitude of the Poynting vector can also be written as:

$$|\mathbf{S}(\mathbf{r}, t)| = \frac{|\mathbf{E}(\mathbf{r}, t)|^2}{2\eta}. \quad (2.11)$$

In free space, the electromagnetic impedance of free space is $\eta_0 \approx 377 \Omega$, and this means that in this case the magnitude of the electric field strength is higher than that of the magnetic field by 377 times, in SI units. This fact will be important in section 2.1.6, when talking about the influence of the magnetic field on the Lorentz force.

When dealing with monochromatic waves, i.e. electromagnetic waves which have a fixed frequency ν (or equivalently a fixed angular frequency $\omega = 2\pi\nu$), the time dependence of the electric field strength $\mathbf{E}(\mathbf{r}, t)$ and magnetic field strength $\mathbf{H}(\mathbf{r}, t)$ can be discarded, and the two fields can be expressed in terms of complex phasors $\mathbf{E}(\mathbf{r})$ and $\mathbf{H}(\mathbf{r})$ which are related to their time-dependent counterparts by:

$$\mathbf{E}(\mathbf{r}, t) = \text{Re} \{ \mathbf{E}(\mathbf{r}) \exp(j\omega t) \} \quad (2.12)$$

$$\mathbf{H}(\mathbf{r}, t) = \text{Re} \{ \mathbf{H}(\mathbf{r}) \exp(j\omega t) \}. \quad (2.13)$$

Other complex phasors can be defined analogously for the other fields discussed above. These vectors are sometimes called complex amplitudes. With these quantities, Maxwell's equation in the absence of external sources can be expressed in the frequency domain as:

$$\nabla \times \mathbf{E}(\mathbf{r}) = -j\omega \mathbf{B}(\mathbf{r}) \quad (2.14)$$

$$\nabla \times \mathbf{H}(\mathbf{r}) = j\omega \mathbf{D}(\mathbf{r}) \quad (2.15)$$

$$\nabla \cdot \mathbf{D}(\mathbf{r}) = 0 \quad (2.16)$$

$$\nabla \cdot \mathbf{B}(\mathbf{r}, t) = 0, \quad (2.17)$$

where j denotes the imaginary unit. If we consider the complex amplitudes, we can also define the complex Poynting vector as:

$$\mathbf{S}(\mathbf{r}) = \frac{\mathbf{E}(\mathbf{r}) \times \mathbf{H}^*(\mathbf{r})}{2}, \quad (2.18)$$

where the asterisk on the magnetic field strength denotes the complex conjugate. Inserting the complex amplitudes of the electric field and magnetic field strength into the wave equation (see equation 2.5), allows the latter to be written in the frequency domain as well (sometimes referred to as the Helmholtz equation):

$$\nabla^2 U(\mathbf{r}) + k^2 U(\mathbf{r}) = 0, \quad (2.19)$$

where $k = nk_0 = n\omega/c_0$ is the wavenumber of the light, i.e. the absolute value of the wavevector \mathbf{k} which is orthogonal to both $\mathbf{E}(\mathbf{r})$ and $\mathbf{H}(\mathbf{r})$, and is normal to the wavefront of the electromagnetic wave. $n = c_0/c$ is the refractive index, where $c \leq c_0$ is the speed of light in a medium. $U(\mathbf{r})$, of course, is the complex amplitude of $u(\mathbf{r}, t)$.

Solutions to the Helmholtz equations include the plane wave, spherical wave and Gaussian beam.

2.1.2	Gaussian beam
--------------	---------------

The transverse electromagnetic (TEM) plane wave is the simplest solution to the wave equation:

$$\begin{aligned}\mathbf{E}(\mathbf{r}, t) &= \mathbf{E}_0 \cos(\omega t - \mathbf{k} \cdot \mathbf{r}) = \text{Re} \{ \mathbf{E}_0 e^{-j\mathbf{k} \cdot \mathbf{r}} e^{j\omega t} \} \\ \mathbf{H}(\mathbf{r}, t) &= \mathbf{H}_0 \cos(\omega t - \mathbf{k} \cdot \mathbf{r}) = \text{Re} \{ \mathbf{H}_0 e^{-j\mathbf{k} \cdot \mathbf{r}} e^{j\omega t} \},\end{aligned}\quad (2.20)$$

where \mathbf{E}_0 represents the maximum amplitude of the electric field strength (the direction of this vector indicates the wave's polarization) and \mathbf{H}_0 represents the maximum amplitude of the magnetic field strength (whose direction is transversal to \mathbf{E}_0). However, it is ideal and cannot actually exist in the real world, as its wavefronts extend infinitely with equal intensity. The Gaussian beam, on the other hand, which is a solution of the wave equation in the paraboloidal approximation (valid in the slowly varying envelope approximation of the Helmholtz equation, also called the paraxial Helmholtz equation), better describes the propagation of a laser beam through space, as the amplitude of the wave tends to zero as one gets farther away from the axis of propagation. If we assume an electromagnetic wave propagating in vacuum along the z-axis, the expression for the intensity of a Gaussian beam can be written as:

$$I(r, z) = \frac{2P}{\pi W(z)^2} \exp\left(\frac{-2r^2}{W(z)^2}\right), \quad (2.21)$$

where P is the total power of the electromagnetic wave, r is the radial distance from the axis of propagation and $W(z)$ is the beam radius at position z , defined as the radius for which the intensity drops to $1/e^2$ of the maximum intensity. The minimum beam radius is achieved in the focus of the Gaussian beam for $z = 0$, where the beam radius is $W(0) \equiv W_0$, called the "beam waist".

The beam radius can be calculated as:

$$W(z) = W_0 \sqrt{1 + (z/z_R)^2}, \quad (2.22)$$

where z_R is the Rayleigh length, defined as:

$$z_R = \frac{\pi W(z)^2}{\lambda}, \quad (2.23)$$

with $\lambda = n\lambda_0$ being the wavelength of the wave inside the medium in which the beam is propagating. The Rayleigh length is a measure for defining the so-called near-field (when $z < z_R$) and far-field (when $z > z_R$) of the beam. The transverse area is also doubled at one Rayleigh length away from the focal plane.

The Rayleigh length is a measure of how divergent is the Gaussian beam. Within the Rayleigh length, the wavefronts of beam appear approximately planar, and the divergence of the beam is low (the beam appears collimated in this region). However when one moves to the far-field of a Gaussian beam, the wavefronts will appear spherical. For a monochromatic beam of wavelength λ propagating along the z -axis, the complex amplitude of the electric field strength $\mathbf{E}(r, z)$ as a function of longitudinal position z and axial distance r can be written as:

$$\mathbf{E}(r, z) = E_0 \frac{W_0}{W(z)} \exp\left(-\frac{r^2}{W(z)^2}\right) \exp\left[j\left(kz - \varphi_G(z) + \frac{kr^2}{2R(z)}\right)\right] \hat{\mathbf{e}}_{\perp}, \quad (2.24)$$

where E_0 is the peak amplitude of $\mathbf{E}(r, z)$, $\hat{\mathbf{e}}_{\perp}$ is the unit vector that points along the direction of polarization of the electromagnetic wave (which is just the direction in which the electric field oscillates; the subscript \perp is to remind that the electric field oscillates perpendicularly to the direction of propagation). Notice that $\hat{\mathbf{e}}_{\perp}$ could also be a function of z if the wave isn't linearly polarized) and $R(z)$ is the radius of curvature of the Gaussian beam wavefronts, defined as:

$$R(z) = z[1 + (z_R/z)^2]. \quad (2.25)$$

The quantity $\varphi_G(z) = \arctan(z/z_R)$ in equation 2.24 is also called the Gouy phase, and measures the extra phase shift that the Gaussian beam acquires as it propagates through space when compared to a simple plane wave with the same frequency. A schematic diagram of a Gaussian beam can be seen in figure 2.1.




Figure 2.1: Schematic diagram of a Gaussian beam. In addition to the aforementioned quantities, $b = 2z_R$ represents the depth of focus and $\theta = \lambda/(\pi W_0)$ is the angular divergence of the beam.

2.1.3 High intensity laser pulses

Laser wakefield acceleration relies on the generation of extremely intense laser pulses. To increase the intensity of a laser, one can both decrease the area of the beam (by focusing) and increase the power (which is done most often by trying to generate an ultrashort pulse duration rather than increasing the energy per pulse, as this can be done over many more orders of magnitude). There are many different ways to generate short laser pulses with high peak power [10]: Q-switching, gain switching, mode locking and cavity dumping. It was only in 1985, thanks to the invention of chirped pulse amplification (CPA) by Donna Strickland and Gerard Mourou [5] (which granted them the Nobel prize in 2018) that scientists were able to reliably obtain laser pulses with powers even past the petawatt level [11]. A diagram of how CPA works is given in figure 2.2. Most laser pulses can be assumed to have a Gaussian temporal




Figure 2.2: Schematic diagram of how CPA works. [12] An ultrashort laser pulse is stretched in time with a pair of gratings, so that it can be safely amplified without damaging the optics. The gratings are arranged in such a way that they force different wavelengths to travel different optical path lengths so that they separate in time. The pulse is then re-compressed using another pair of gratings which are arranged to obtain the opposite effect.

profile (that is, their optical power over time has a Gaussian shape). Every laser pulse has a time duration τ_P (often defined by the full-width half-max (FWHM)) and a corresponding spectral bandwidth B_P . These two quantities are inversely related, but one needs additional information to obtain one from the other. This is because a laser pulse can be up-chirped (when the frequency increases as time increases), down-chirped (when the frequency decreases as time increases) or unchirped (when all frequencies are overlapped in time, also called "transform-limited"). For a transform-limited Gaussian pulse, the bandwidth and pulse duration are related via:

$$B_P \approx \frac{0.44}{\tau_P}. \quad (2.26)$$

Since laser pulses have a defined bandwidth which can be very significant if the pulse duration is short, they cannot be considered as monochromatic waves. However, thanks to Fourier analysis, any signal in time can be decomposed into an infinite sum of monochromatic frequency components, which means that the description of the electric and magnetic field in the harmonic case presented in section 2.1.1 does not lose generality.

2.1.4 Laser-plasma interaction

Francis F. Chen defines a plasma as follows [13]:

A plasma is a quasineutral gas of charged and neutral particles which exhibits collective behavior.

The term quasi-neutral indicates that a plasma appears neutral on a macroscopic scale (as the number of positively charged ions matches the number of negatively charged electrons), however it implies that there can be local inhomogeneities in charge distribution. The collective behavior of a plasma determines the motion of the electrons and ions to be influenced by other regions of the plasma itself, which could also be far away.

A peculiar property of plasma called "Debye shielding" describes its ability to shield itself from local accumulations of charge. The Debye length λ_{DE} can be thought of as the radius of the smallest sphere for which a plasma can be enclosed without exhibiting this shielding effect. Within this distance, it is therefore possible to significantly separate the negative and positive charges without excessively disturbing the plasma. The Debye length is defined as:

$$\lambda_{DE} = \sqrt{\frac{\varepsilon_0 k_B T_e}{q_e^2 n_e}}, \quad (2.27)$$

where $k_B \approx 1.38 \times 10^{-23} \text{ JK}^{-1}$ is the Boltzmann constant, T_e is the temperature of the electrons in the plasma, $q_e \approx 1.6 \times 10^{-19} \text{ C}$ is the elementary charge and n_e is the electron density.

Note that this length increases when temperature increases, as a higher temperature forces the electrons to spread out more; and it decreases when the electron density increases, as higher electron density causes a stronger collective behaviour and therefore a stronger shielding effect. The average number of electrons that are contained within a sphere of radius λ_{DE} (a

"Debye sphere") is denoted as N_D . Since the volume of a sphere is $(4/3)\pi r^3$, this can be written as:

$$N_D = n_e \frac{4}{3} \pi \lambda_{DE}^3. \quad (2.28)$$

A plasma where $N_D \ll 1$ is called "cold", while a plasma with $N_D \gg 1$ is called "hot" (also said to be ideal). In a hot plasma, almost all the atoms are fully ionized, whereas a cold plasma is generally only partially ionized. A closely related quantity to N_D is the plasma parameter g , which is defined as [14]:

$$g \equiv \frac{1}{n_e \lambda_{DE}^3}. \quad (2.29)$$

The condition of ideal plasma can therefore be also satisfied if $g \ll 1$, since g and N_D are inversely proportional.

It has been established that the collective behavior of a plasma will try to restore any disturbance that could cause the charges to be separated, and the swiftness of the plasma to do this depends on a quantity called the plasma frequency ω_p , which has the following expression:

$$\omega_p = \sqrt{\frac{n_e q_e^2}{\epsilon_0 m_e}}, \quad (2.30)$$

where $m_e \approx 9.11 \times 10^{-31}$ kg is the rest mass of the electron. The plasma frequency is the natural frequency that the electrons exhibit as they oscillate relatively to the positive ions. Higher plasma frequency means higher velocity in restoring the neutrality of the plasma if disturbances are caused on length scales greater than the Debye length. The dispersion relation for electromagnetic waves with angular frequency ω and wavenumber $k = (2\pi)/\lambda$ propagating in plasma is given by:

$$\omega^2 = \omega_p^2 + c^2 k^2. \quad (2.31)$$

Since the plasma frequency is proportional to the electron density n_e , a denser plasma will respond faster to external perturbations. If both the frequency of the light going into the plasma and the plasma frequency are fixed, equation 2.31 must be solved for k . However, if $\omega < \omega_p$, there exist no real solutions for k , i.e. the wave cannot propagate in the plasma because the shielding is faster to rearrange the electrons and the ions than the wave is able to displace them, the plasma basically acts as a mirror and the wave evanesces exponentially into the plasma, penetrating it only up to a shallow skin depth.

The electron density for which $\omega_p = \omega$, which marks the threshold between the ability of the light to penetrate the plasma, is called the critical density of the plasma n_c . A plasma whose electron density is higher than

the critical density is called "over-dense" (light cannot propagate), while a plasma where the electron density is lower than the critical density is called "under-dense" (light can propagate, and the plasma will behave as a nonlinear refractive medium). Typically, n_e is in the order of 10^{19} cm^{-3} in LWFA applications, which means that $\omega_p \approx 180 \text{ THz}$ and the corresponding plasma wavelength $\lambda_p = (2\pi c_0)/\omega_p \approx 10 \mu\text{m}$. The laser used in the experiment has a wavelength of 800 nm , therefore its frequency is substantially greater than the plasma frequency and can easily penetrate the plasma.

2.1.5 Ionization of gases

In the previous section, the basics of how light interacts with plasma have been discussed. Therefore, a logical question that needs to be answered is how such a plasma can be created.

While there are various ways to generate a plasma, the method used in LWFA revolves around supplying a material, usually a gas, with such high optical power that its atoms become fully ionized (their nuclei are stripped of all of their electrons). The ionization energy for a certain atom is the energy that is needed to expel one of its bound electrons from the nucleus. Since different electrons occupy different orbitals, the ionization energy is different for each one of them (with the innermost electrons requiring the most amount of energy to be released, since they reside closer to the positive protons in the nucleus). A photon with energy $E_{ph} = h\nu$ (where $h \approx 6.626 \times 10^{-34} \text{ m}^2 \text{ kg s}^{-1}$ is Planck's constant) can be absorbed by an atom and release an electron via photoelectric effect if $E_{ph} > E_{ion}$, where E_{ion} represents the ionization energy of its outermost electron. Complete ionization of the atom is achieved if the energy exceeds that of the binding energy of its innermost electron. Table 2.1 shows the ionization energy of the innermost electron of the first 10 elements of the periodic table, rounded to the second decimal digit.

Even the ionization of the single electron of hydrogen requires a frequency of $\nu = (13.6 \cdot 1.6 \cdot 10^{-19})/h \approx 3.3 \text{ PHz}$, which corresponds to a wavelength of around 91.3 nm , which is within the extreme ultraviolet (XUV) range. This is extremely difficult to achieve, as even the highest frequency commercially available excimer lasers do not have wavelengths that are shorter than 100 nm [16]. In spite of this, if enough optical intensity is supplied, an atom can be ionized also using photon energies that are lower than the ionization energy through the mechanism of multiphoton ionization. When a photon with an energy lower than the ionization energy gets absorbed by

Z	Ion	E_{ion} [eV]
1	H ⁺	13.60
2	He ²⁺	54.42
3	Li ³⁺	122.45
4	Be ⁴⁺	217.72
5	B ⁵⁺	340.23
6	C ⁶⁺	489.99
7	N ⁷⁺	667.05
8	O ⁸⁺	871.41
9	F ⁹⁺	1103.12
10	Ne ¹⁰⁺	1362.20

Table 2.1: Ionization energies of ions. [15] The ionization energy of the first 10 elements of the periodic table, for their innermost electrons in the 1s orbital. Z represents the atomic number.

an atom, it can excite it to a short-lived virtual energy level. If another photon is absorbed within the lifetime of the virtual level, the atom can be further excited (even multiple consecutive times) until it is ionized. Another way to ionize an electron is to deform its potential by applying a strong enough electric field (for example in the form of light) to the point that the electron can cross the potential barrier in a process called tunnel ionization. A schematic diagram of the different ways to ionize an atom can be seen in figure 2.3. It is often the case that the pulses that are used for LWFA have such high




Figure 2.3: Schematic diagram of 3 possible ionization mechanisms. Multiphoton ionization (b) and tunnel ionization (c) do not require the energy of the incident photons to exceed the ionization energy as in direct photoionization (a).

optical intensity that even the leading tail of the Gaussian pulse is able to ionize the gas (often a mixture of hydrogen or helium with a small percentage of nitrogen, in "ionization injection"), so that when the strongest part of the pulse arrives, it will interact with an already ionized hot plasma.

2.1.6 Electron acceleration

It is finally time to explain how LWFA actually utilizes a plasma to build up tremendously large fields that can accelerate electrons in a compact laboratory setup.

A charge q in motion at velocity \mathbf{v} in an electric field \mathbf{E} and magnetic field \mathbf{B} is subjected to the Lorentz force:

$$\mathbf{F} = \frac{\partial \mathbf{p}}{\partial t} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (2.32)$$

where \mathbf{p} is the momentum of the charge q . If an electron is considered as the charge, then the (relativistic) momentum is:

$$\mathbf{p}_e = \gamma m_e \mathbf{v}, \quad (2.33)$$

where γ is the Lorentz factor:

$$\gamma = \frac{1}{\sqrt{1 - \frac{v^2}{c_0^2}}}, \quad (2.34)$$

with v being the speed of the electron.

To better describe how the electrons in a plasma will react to a high intensity laser pulse, it is easier to express the electromagnetic fields \mathbf{E} and \mathbf{B} in terms of scalar potential Φ and vector potential \mathbf{A} , as follows:

$$\begin{aligned} \mathbf{E} &= -\nabla\Phi - \frac{\partial \mathbf{A}}{\partial t} \\ \mathbf{B} &= \nabla \times \mathbf{A}. \end{aligned} \quad (2.35)$$

In vacuum, the gradient of Φ is zero, as there are no charges to vary the electric potential. The vector potential has a peak amplitude A_0 , whose normalized value is called the laser strength parameter a_0 :

$$a_0 = \frac{q_e A_0}{m_e c_0^2} = \frac{q_e E_0}{m_e \omega_0 c_0^2}. \quad (2.36)$$

When this value is greater than 1, the motion of the electron will be relativistic and the laser-plasma interaction is non-linear. If the parameter $\beta = v/c_0$ is defined as the normalized speed of an electron, relativistic motion corresponds to $\beta \rightarrow 1$. The value of the laser strength parameter can be calculated with

the following equation:

$$a_0 \approx \sqrt{7.3 \times 10^{-19} [\lambda(\mu\text{m})]^2 I_0 (\text{W}/\text{cm}^2)}. \quad (2.37)$$

Inserting $a_0 = 1$ in the expression above and solving for the required critical peak laser intensity I_0 for relativistic electron motion at a wavelength of 800 nm (as in the experiment) yields $I_0 \approx 2.14 \times 10^{18} \text{ W}/\text{cm}^2$. This can be easily achieved by tabletop pulsed lasers that make use of CPA.

A non-relativistic electron ($\beta \ll 1$) in an electromagnetic wave will not experience any forward acceleration, but will only wiggle up and down in the direction of the electric field, without any net energy gain. However, a relativistic electron ($\beta \rightarrow 1$, achieved by having $a_0 > 1$) will oscillate with a velocity that approaches c_0 . This means that the magnetic contribution in the Lorentz force (equation 2.32), which for a slow electron was quite small (as $|\mathbf{B}| = |\mathbf{E}|/c$, see equations 2.7 and 2.10), will become significant. Solving the Lorentz equation for a relativistic electron reveals a longitudinal velocity component which is able to "push" the electron in the forward direction. It can also be shown that this longitudinal motion scales with a_0^2 (as opposed to the transverse one, which scales with a_0), so that having $a_0 \gg 1$ causes the longitudinal motion to dominate over the transverse one. The longitudinal displacement of the electron equals the transverse one for $a_0 \approx 2.55$.

Even if an electron is pushed to a new position, it does not actually gain any energy if the accelerating electromagnetic field is a plane wave, according to Lawson-Woodward theorem. However, a focused pulse with a Gaussian spatial profile in an under-dense plasma will cause a plasma wave to be excited due an effect called the ponderomotive force. This force has the following expression:

$$\mathbf{F}_p = -mc_0^2 \nabla a_0^2 / 2, \quad (2.38)$$

where m is the mass of a particle (note that a plasma is comprised of both electrons and ions). The ponderomotive force effect occurs when a particle interacts with an inhomogeneous electromagnetic field (such as a Gaussian beam), which causes the particle to move towards where the field is weaker. This is confirmed in the formula above by remembering that the gradient of the strength parameter squared ∇a_0^2 is a vector that points in the direction of steepest ascent of a_0 , but the minus sign causes the ponderomotive force to point towards the direction of its steepest descent, i.e. where the field is weaker. Note from the expression of a_0 in equation 2.36, the strength parameter has the mass of the particle in the denominator. Since a_0 is squared in the expression for the ponderomotive force, the latter is actually inversely proportional to the particle mass. This means that the lightweight electrons are pushed outwards

from where the pulse enters the plasma, creating a region that is completely depleted of them, similarly to how a bullet shot in water would leave a pocket of air behind its path. The much heavier positive ions, however, will be basically unaffected, since the mass of the proton is about 1836 times larger than that of the electron. As the electrons are pushed out, a strong charge inhomogeneity is created, and Debye shielding causes the electrons to quickly try to restore their original position. However, much like in a spring in an underdamped harmonic oscillation, the electrons overshoot past their equilibrium position, causing a series of consecutive regions of inhomogeneity, which constitute a so-called plasma wave. The wavelength of the plasma wave λ_p is proportional to the electron density according to the following:

$$\lambda_p [\mu\text{m}] \approx 33(n_e[10^{18} [\text{cm}^{-3}]]^{-1/2}). \quad (2.39)$$

The plasma wave trails after the laser pulse, propagating in the plasma at a normalized speed β_p that is close to c_0 :

$$\beta_p = \frac{v_g}{c_0} = \sqrt{1 - \frac{n_e}{n_c}}, \quad (2.40)$$

where v_g is the group velocity of the plasma wave. The equation above shows that a faster wave will be achieved for lower electron densities n_e . The regions where there exists a rarefaction of the electron density are called "blow-out regions" or "bubbles", due to their characteristic appearance. Figure 2.4 shows a schematic of how the plasma wave would look like as seen from a side view.

These strong demarcations between electrons on the outside and ions in the center of the bubbles cause extremely high electric fields to build up, which can be in the order of 100 GV/m [17]. As the electrons are pushed back towards the center of axis of propagation of the plasma wave, it can happen that some of them acquire trajectories that make them end up in a region that is just between the back and the center of a bubble, where this enormous electric field is able to trap them and accelerate them with the plasma wave. This mechanism of injection of the electrons is called "self-injection" and occurs when the intensity of the electric field is so high that it causes the wave to "break" on the front. There are many other injection mechanisms that are possible, such as ionization injection (using mixtures of light and heavier gases), density gradient injection (such as shock-front injection using razor blades) and colliding-pulses injection (the beating between counter-propagating pulses can cause trapping of electrons) [18], however further investigation of the topic of controlled injection is out of the scope of this thesis.

The injected electrons will be repelled by the electrons in the back of

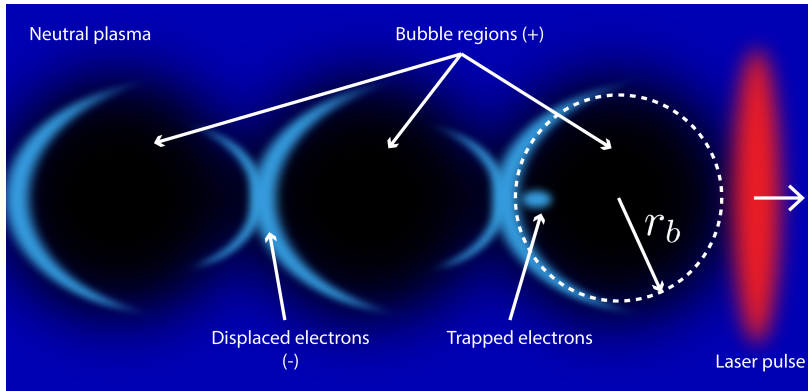


Figure 2.4: Schematic diagram of a laser-excited plasma wake. The plasma is represented in dark blue. The blow-out regions that are depleted of electrons are shown in black. The electrons, which are pushed out of the path of the laser due to ponderomotive force, are shown in light blue. In this example, a bunch of electrons have been trapped in the first bubble and will be accelerated in the axial direction of the laser, which points in the direction of the white arrow on the right of the image. However, since the electrons move faster than the phase velocity of the plasma wave, they will eventually stop accelerating after they reach a "dephasing length". This happens when the electrons reach the middle of the bubble, as the accelerating electric field there becomes zero. The radius of the bubbles is denoted with r_b .

the bubble, but attracted to the ions at its center, and since the entire plasma wave co-propagates with the laser pulse at almost the speed of light, these trapped electrons will be accelerated enormously in the span of a few millimeters, in a process that is akin to surfers riding a wave in the ocean. The electrons can be accelerated to energies in the hundreds of MeV [17].

2.2 X-ray betatron radiation

When the electrons are injected in the back of the bubble, they will only follow straight trajectories along the propagation axis of the laser if they are trapped in the middle of the bubble. However, it is often the case that some electrons are trapped off-axis. This will cause the electrons to wiggle around the optical axis as they are accelerated forward.

When a charge that is moving at relativistic velocity changes direc-

tion, it loses energy to the electromagnetic field in the form of electromagnetic radiation. This is called synchrotron radiation, if the wiggling is obtained using magnetostatic means such as in undulator devices. However in the case of LWFA, where the transverse motion of the electrons is obtained by strong focusing forces within the plasma, the emission is called betatron radiation [19].

Figure 2.5 shows the principle behind the betatron emission effect.

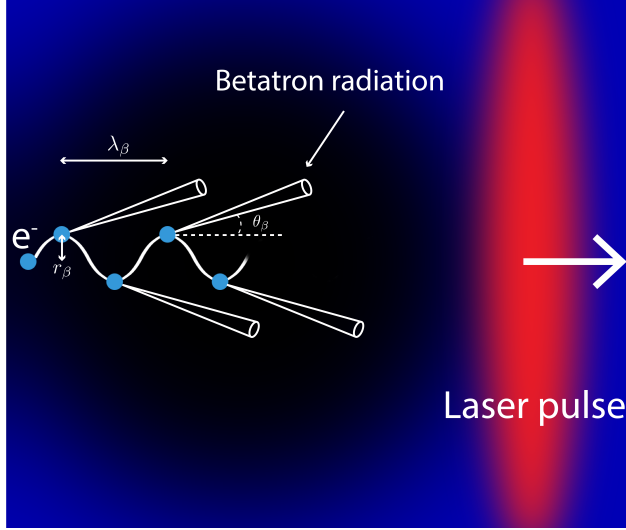


Figure 2.5: Schematic diagram of the betatron radiation emission principle. The electrons, indicated with light blue dots, move both axially and transversely within the bubble. As they wiggle at relativistic velocities, they release photons in the X-ray frequencies as betatron radiation, at a divergence angle θ_β . The radiation is incoherent because it is the product of many emissions along different points of the electron trajectory, and it is originated by different electrons in the trapped electrons bunch. The vertical motion of the electron is greatly exaggerated for clarity.

As the electrons oscillate while being accelerated, they do so in a spatial period called the betatron wavelength $\lambda_\beta = \lambda_p \sqrt{2\gamma}$, with corresponding betatron frequency $\omega_\beta = \omega_p \sqrt{2\gamma}$.

An interesting parameter for betatron radiation is the betatron strength parameter α_β , which is:

$$\alpha_\beta = r_\beta k_\beta \gamma, \quad (2.41)$$

where r_β is the amplitude of the electrons transverse oscillations and $k_\beta = \omega_\beta/c_0$ is the betatron wavenumber. The amount of generated radiation is linearly proportional to the betatron strength parameter. When $\alpha_\beta \gg 1$, as is often the case in LWFA, the betatron emission appears broadband, characterized by

a critical angular frequency ω_c , after which the radiation quickly diminishes:

$$\omega_c = 3\alpha_\beta\gamma^2\omega_\beta. \quad (2.42)$$

Typically this is in the hard X-ray region of the electromagnetic spectrum, with energies in the tens of keV. Furthermore, the betatron emission is incoherent, and is produced with a divergence $\theta_\beta = \alpha_\beta/\gamma$, which often results in divergence angles in the order of tens of milliradians, and extremely small source sizes in the order of a few microns. The peak X-ray energy of the betatron radiation is given by:

$$E [\text{eV}] = 1.45 \times 10^{-21} \gamma^2 n_e [\text{cm}^{-3}] r_\beta [\mu\text{m}]. \quad (2.43)$$

A thorough description of the characteristics of the betatron radiation spectrum is given in [20]. As a practical example: assuming a betatron radius $r_\beta = 1 \mu\text{m}$, Lorentz factor $\gamma = 200$ and electron density $n_e = 10^{19} \text{cm}^{-3}$ results in a betatron wavelength $\lambda_\beta \approx 0.21 \text{mm}$ and peak X-ray energy of 5.8 keV. The photon flux of X-rays can be increased in practice by increasing the pressure of the gas that is used as target for the laser pulse, as the higher gas density causes stronger transverse oscillations of the electrons, and thus more betatron emission.

3

X-RAY ABSORPTION SPECTROSCOPY AND SINGLE-PHOTON COUNTING

This chapter introduces X-ray absorption spectroscopy and how it can be implemented with single-photon counting, using CCD cameras as arrays of calorimeters. Details about important parameters for CCD sensors are given, with a focus on the most relevant noise sources present in the experiment.

Contents

3.1	X-ray absorption spectroscopy	24
3.1.1	Motivation for single-photon counting	24
3.1.2	K-edge detection	26
3.2	Energy deposition in CCD sensors	27
3.2.1	CCD sensor basics	27
3.2.2	Quantum efficiency	31
3.3	Single-photon counting	33
3.4	Noise and dynamic range	34
3.4.1	Shot noise	35
3.4.2	Thermal noise	36
3.4.3	Read-out noise	37
3.4.4	Dynamic range	38

3.1 X-ray absorption spectroscopy

Absorption spectroscopy is used in countless applications, such as astronomy, remote sensing, combustion physics, medicine, chemistry and many others. In this spectroscopic technique, a sample is illuminated with an electromagnetic wave. As the wave propagates through the material, some photons will be absorbed, with an absorption coefficient that depends on frequency, so that the transmitted light doesn't have the original spectrum anymore, but a filtered version of it. There is a diversified nomenclature of absorption spectroscopy, depending on the region of the electromagnetic spectrum that is used for illumination. In the case of this thesis work, one talks about "X-ray absorption spectroscopy" (XAS).

XAS is widely used across many different scientific fields, such as nanoscience, chemical and biochemical research, molecular and condensed matter physics and more. Because of all these areas of application, there are currently more than 100 beamlines or endstations worldwide which support this technique using synchrotron or free-electron laser facilities [21]. One of these synchrotron facilities is here in Lund, at the MAX IV laboratory.

3.1.1 Motivation for single-photon counting

As mentioned previously, XAS relies on being able to spectrally characterize X-rays both before and after a sample. However, traditional spectrometers based on prisms or diffraction gratings are notoriously difficult to construct for X-rays. In this case, prisms are ill-suited since the refractive index of all elements tends to 1 as frequency increases, and diffraction gratings require the grating spatial period to be in the same order of magnitude as the wavelength of light that wants to be dispersed, which can be smaller than an angstrom for hard X-rays. For this reason, crystals are often used instead to disperse X-rays.

This way of doing XAS allows for point measurements of the spectrum, and using it for imaging would be very time-consuming. Single-photon counting is an alternative technique for performing XAS, allowing to do multispectral 2D imaging as well. Another reason why this is an attractive technique to measure a spectrum is that it doesn't rely on any type of dispersive element such as prisms or diffraction grating as in traditional spectrometers.

Single-photon counting makes use of the intrinsic dispersion of the X-ray beam itself to spatially separate the photons onto a CCD camera sensor. Some studies [22, 23] have shown similar experiments, however in these cases the X-ray beam was very narrow and was scanned through the entire object. This gives an excellent spatial resolution, however it also requires much longer time. In this experiment the X-ray source has a broadband spectrum, therefore no scanning is needed, and many frequencies are detected simultaneously in a single shot by virtue of the fact that photons are scattered with a uniform distribution across the sensor of the camera. The method presented in this thesis illuminates the entire target, and then uses a sequence of many images to obtain better statistics. Figure 3.1 shows the distribution of all photons on one of the background images (no filter used), which shows that photons of different energies are distributed pretty uniformly throughout the sensor. Further details on how the algorithm can obtain the energy of the photons and their coordinates on the sensor will be given in section 5.1.1 of chapter 5.




Figure 3.1: Scatter plot of photon events in a background image. Photons with different energies are scattered uniformly across the CCD sensor. Colors represent counts on the pixels (note that a noise of 300 counts was subtracted, which is why the colorbar start from as little as 1 count). The color map is discretized to only 7 colors for clarity. Only photons with less than 15 keV of energy are plotted, corresponding to around 1282 counts (see section 5.2 of chapter 5 for details on how the calibration was done).

3.1.2 K-edge detection

X-ray absorption spectroscopy in this thesis is done with the final aim of identifying the material that the sample is composed of, i.e. species identification. To do this, the proposed method is to look at the energy absorption peak due to K-edge absorption. The K-edge refers to the sudden increase in absorption of a material when the energy of the incident light becomes slightly higher than the binding energy of its innermost electrons. When this happens, the electrons in the innermost shell of the atoms in the sample (which is called 1s shell, or orbital, in spectroscopic notation and K-shell in X-ray notation) can be ionized via photoelectric effect.

The energy at which the K-edge occurs is specific to each element, generally increasing with increasing atomic number Z . This provides a sort of "fingerprint" for each element. Table 3.1 shows the K-edge of elements which occur within the range of the X-ray energies used for the experiment. If the

Z	Element	K-edge [eV]
21	Scandium	4492.8
22	Titanium	4966.4
23	Vanadium	5465.1
24	Chromium	5989.2
25	Manganese	6539.0
26	Iron	7112.0
27	Cobalt	7708.9
28	Nickel	8332.8
29	Copper	8978.9
30	Zinc	9658.6

Table 3.1: K-edge energy of different elements. [24] The K-edge energy of elements which are relevant for the work presented in this thesis, as the energies are within the spectrum of the X-rays used. The elements all belong to the 4th group of the periodic table and are transition metals.

spectrum of the X-rays prior to entering the sample is known, by comparing it with the one after the absorption (the filtered spectrum) it is possible to identify the dip in transmission due to the K-edge and relate the energy at which this occurs with an element (or a combination of them), effectively allowing its identification. More details on the metal detection algorithm that is able to automate this process is given in section 5.3 of chapter 5.

3.2 Energy deposition in CCD sensors

Before explaining in detail how single-photon counting can be used to measure the X-ray spectrum of a light source, the way in which the energy is deposited on the CCD camera sensor must be understood. As is explained in the following sections, the main idea is that pixel counts in the X-rays are not only proportional to the optical flux of the incoming light, but also to the energy of each individual photon. This means that each CCD pixel can effectively be used as a rudimentary calorimeter. These devices are widely used in particle physics facilities to measure the energy of particles produced after collisions in accelerators.

3.2.1 CCD sensor basics

CCD stands for Charge Coupled Device, and it is an integrated circuit which uses specialized metal oxide semiconductor capacitors (MOS) arranged in an array, usually rectangular, with each transistor called "photosite" or more commonly "pixel".

Even though CCD sensors have been steadily decreasing in market share (the global sales of CCD sensors were surpassed by complementary metal oxide semiconductor (CMOS) sensors already in the early 2000s, with the CMOS representing 89% of the total sensor sales in 2017 [25]), they are still widely used in scientific applications due to their yet superior performance, offering higher dynamic range and lower levels of noise compared to the however cheaper CMOS sensors [26].

A schematic of how a MOS used as pixel for a CCD sensor works can be seen in figure 3.2. A nMOS has a p-type semiconductor substrate, i.e. pure silicon doped with acceptor atoms (atoms with one less electron in their valence shell compared to silicon), such as boron. If a positive voltage difference (called bias) is applied between the gate (indicated in red in the figure) and the p-type substrate (called bulk, indicated in gray), the holes, which are majority carriers and are present in the doped silicon, will be attracted to the bottom part of the bulk, which is at a lower potential (normally ground). As the holes go deeper into the bulk, they leave behind a region that is void of mobile carriers called "depletion region" (indicated in green). This region does




Figure 3.2: Schematic diagram of a nMOS pixel in a CCD sensor. [27] A positive voltage is applied to the gate of a MOS capacitor, such that a depletion region is created within the p-type doped silicon used as substrate. In the case of visible light, each photon that is absorbed by the silicon ions in the depleted region and cause the generation of a single electron-hole pair. The one-to-one relationship between photons and electron-hole pairs holds regardless of the energy of the visible photon. High energy photons such as X-rays, instead, can generate a great number of free electrons that is proportional to their energy. In both cases, the photo-generated electrons can then be retrieved and proportionally converted to a voltage.

not extend indefinitely, as an electric field builds up between the negatively charged fixed acceptor ions and the holes injected in the gate, which opposes the holes diffusion. In typical nMOS operation, thermally generated electrons accumulate at the interface between oxide (which is an insulator, indicated in light blue) and semiconductor, creating a so-called "n-channel" that can be used in transistors as a conductive element. This is why the device is called nMOS in spite of the substrate being p-doped. This device, when other contacts are added by the sides of the gate, called "source" and "drain", is appropriately called metal oxide semiconductor field effect transistor (MOSFET). However, if the temperature of the nMOS is kept low (in the order of tens of degrees below 0°C) and if the voltage applied to the gate is relatively high, the device goes in a non-equilibrium state called "deep depletion", where the depletion region extends further than usual and no electron channel is formed at the oxide-semiconductor interface.

If the nMOS is exposed to light, there is a chance that photons pass through the gate and silicon dioxide layers and are absorbed by the atoms within the depletion region, provided that their energy is high enough to exceed the energy band gap of the material (that is, the energy difference between the top of its valence band and the bottom of its conduction band). In the case of visible light, each photon can generate a single electron-hole pair, and the photo-generated holes will rush to the bulk region while the photo-generated electrons will be attracted to the gate and collected at the oxide-semiconductor

interface, creating a so-called n-channel (indicated in yellow).

The amount of photo-generated electrons is proportional to the intensity of the light impinging on the nMOS, therefore it can be used as a meter to measure it, which is the way that normal photographs are taken. However, each visible photon only generates one electron-hole pair (there is a one-to-one relationship), so there is no dependence between counts and energy in this case, and the spectrum of the light could not be inferred using single-photon counting.

This is quite different from what happens for an high-energy X-ray photon. Because the energy of X-rays is thousands of times larger than the band gap of silicon, any atom which absorbs the X-ray photon will be ionized, releasing a high-energy electron. This electron can collide with other atoms in the substrate lattice, consequently ionizing them and triggering an avalanche effect which ultimately leads to the generation of a great number of free electrons, proportional to the energy of the original photon. This is illustrated in figure 3.3. The generated electrons form a "cloud" of a certain diameter (see




Figure 3.3: Avalanche ionization. After an X-ray photon (black arrow) has been absorbed, it causes the release of an high energy electron (indicated in light blue on the left), because the ionization energy of silicon is only around 8.15 eV, while the photon can be in the keV range. The ionized electron can impact on other atoms and ionize them as well, leading to a cascading effect called avalanche ionization. Yellow lines illustrate the motion of electrons schematically. Note that impact ionization could cause the release of more than only 2 electrons in reality. The process stops when the final electrons have reached an energy which puts them in the conduction band of the silicon substrate.

figure 3.4) which can actually overflow into different neighboring pixels. When an X-ray photon is absorbed, the number of electron-hole pairs formed (the number of "counts") and the size of the cloud they form in the substrate is linearly proportional to the energy of the incident photon, as can be seen in

figure 3.5.




Figure 3.4: Diameter of the detected photons as a function of energy for the X-ray sensitive BR-DD CCD sensor by Andor. [28] Photons with different energies will deposit it over different finite-size areas on the sensor. This is because each X-ray photon will excite the release of a "cloud" of electrons as it is absorbed. The relationship between energy and diameter of the sensor area where the energy was deposited is linear.




Figure 3.5: Number of photo-electrons per detected X-ray. [28] The number of photo-electrons has a linear dependence on the energy of the X-ray photons. If the energy is expressed in eV, the slope of the line is $1/3.65$.

The relationship between photo-generated electrons and photon energy which only holds for high energy photons such as X-rays is the key idea behind single-photon counting, as it allows to directly measure the spectrum of light by creating a histogram of the photon events. CCD sensors are comprised of an array of nMOS pixels, and a schematic configuration of this can be seen in figure 3.6.

Every nMOS element in the CCD sensor has a fixed well-depth (also called full-well capacity), which is the maximum number of electrons that it can hold before reaching total saturation, i.e. the amount of photo-generated electrons on

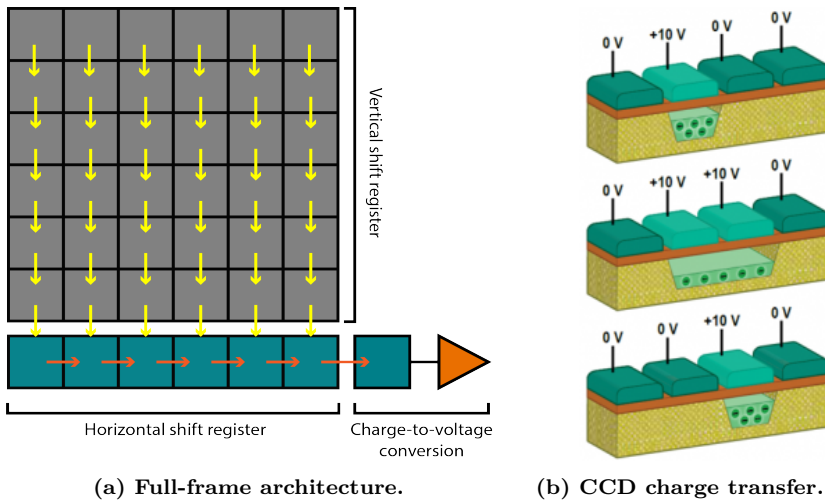


Figure 3.6: Full-frame architecture of a CCD sensor. A common CCD architecture is the full-frame depicted in (a). The grid with gray rectangles is the photosensitive CCD sensor matrix, where each square represents one individual MOS pixel. The photo-generated charge is shifted down row by row (indicated with yellow arrows), by appropriately switching the voltage on the gates of each MOS, as can be seen in (b). At the bottom of the sensor there is an horizontal array (shown in cyan in (a)) which can store the current row, and is shielded from light so that it is not exposed as the image is captured. The electrons in this register are shifted horizontally (indicated with orange arrows in (a)) to a charge-to-voltage converter and then stored. The operation repeats, alternating one vertical row shift and one horizontal array shift, until the entire array of pixels is read out. The frequency at which each pixel is read is called the read-out rate of the CCD.

the pixel is so high that it becomes impossible to distinguish further increments in energy (the pixel will appear completely white in the final image).

3.2.2 Quantum efficiency

As mentioned previously, not all photons can be absorbed and generate electron-hole pairs. Only photons with an energy that exceeds the band gap energy of the semiconductor material that comprises the bulk of each pixel can force the electrons in the silicon to "make the jump" from the valence band to the conduction band. The band gap of silicon is about 1.12 eV, which corresponds to a wavelength of around 1.1 μm , which is in the infrared region of the electromagnetic spectrum.

Even if X-rays possess vastly superior energy to this, not all energies are absorbed with the same efficiency. This is because photons can be detected by the CCD only if they are absorbed within the depletion region, but the photons need to cross both the gate and oxide layer of the CCD as well. The number of absorption events per incident photon is called "quantum efficiency" (QE), and it is a function of photon energy. In the case of the CCD sensor used for this master's thesis work, the quantum efficiency can be seen in figure 3.7. As can be seen, the efficiency is always under 100% for this sensor, being

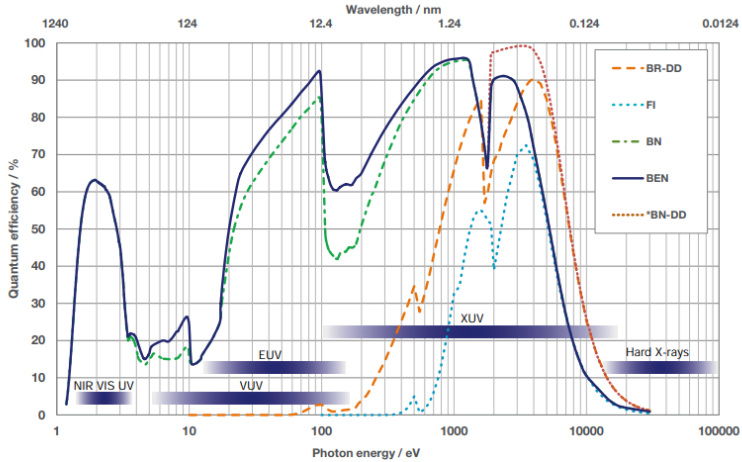


Figure 3.7: Quantum efficiency of various CCD sensors manufactured by Andor. [30] The iKon L SO camera used in the experiment is equipped with a BR-DD sensor (DD stands for "deep depletion"), whose QE is indicated with an orange dashed line in the figure.

particularly high in the XUV/Soft X-rays range (between around 0.2 keV and 10 keV it stays above 20%), but quite low for lower UV photons and higher energy X-rays (the so-called "hard X-rays"). This is because photons in the UV range are heavily absorbed by silicon, which means that they are more easily absorbed before even reaching the depletion layer, effectively being "wasted", as they won't contribute to the photo-generated electrons that are actually read out by the camera (this problem can be compensated if one were to use a back-illuminated CCD sensor). On the other hand, photons whose energy is very high (hard X-rays) will sometimes pass through the entire pixel without being absorbed at all. This is why each photosite of CCD is set into deep depletion, with the purpose of increasing the likelihood of absorption for these higher energy photons, by incrementing the possible volume of interaction.

The quantum efficiency of the CCD sensor acts as a filter in the spectral domain, which means that to obtain the real spectrum from a

measurement, one needs to divide the measured spectrum by the QE response first.

3.3 Single-photon counting

Single-photon counting is a technique that can be used to measure the spectrum of a light source without the use of physical dispersion mediums such as prisms or diffraction gratings.

If a camera is placed at a distance from a light source such that the photon flux impinging on its sensor is sufficiently low, it is possible that each of its pixels will be struck by at most one photon. This condition, which is ideal for the technique, is called "single-photon regime". Since for X-ray photons there is a dependence between photo-generated electrons and energy, the more energetic photons will appear brighter on the final image compared to the lower energy ones (again, this wouldn't be the case for lower energy photons in the visible range). The pixel depth of a camera is defined as the number of bits that are used to store the value of brightness for any pixel. For example, the 16-bit camera used in the experiment can store $2^{16} = 65536$ different brightness levels for each pixel of the sensor. There is a well-defined relationship that binds the pixel value that will be recorded on the final image with the amount of energy deposited on each pixel, which will be explored further in section 5.2 of chapter 5.

Because of this fact, pixel brightness counts can be used to measure the spectrum of the light source, by counting all the single-photon events that occurred on the image and collecting their energies (obtained after converting the pixel values with an appropriate calibration) in a histogram.

This process is not as trivial as one may think, because it is not sufficient to simply compute the total histogram of the entire image, which is done pixel by pixel. The first reason for this can be understood by examining again figure 3.4. As it can be seen, each photon that deposits its energy on the sensor is associated with an area of detection which is non-zero and increases linearly in diameter as the energy increases.

This means that it can be the case (and it often is) that a photon may strike the CCD on a peripheral region of a pixel, so that the area of energy deposition overflows into neighboring pixels. When such an event occurs, the histogram algorithm must take into account the fact that the

entire cluster of pixels which have absorbed the energy of the photon must be counted as resulting from the energy deposition of a single photon. This causes any attempt to construct a standard histogram for spectral measurement to fail, as a regular histogram is done on individual pixels, not on pixel clusters. For example, consider a photon of 10 keV distributing its energy onto two pixels, depositing 4 keV in one and 6 keV in the other. A regular histogram calculated pixel by pixel would erroneously conclude that two weaker photons struck the sensor, rather than one with a higher energy.

A second reason why the single-photon counting algorithm is not straightforward is that it can be the case that different photons could strike pixels that are adjacent (or very close), so that their energy deposition areas overlap. This will result in a large cluster of connected pixels with pixel values that are well above noise level in the final picture. If the algorithm was to count such a large cluster in the histogram for the spectrum measurement, it would erroneously attribute it to a single-photon event, which would result in an overestimation of its energy and an overall blue-shifting effect of the spectrum.

Even if the large cluster is recognized to be as the result of more than a single photon absorption, it is not possible to retrieve the individual energies the photons, but only that of their sum. To give an intuitive explanation, it would be like having two equal glasses containing different amounts of water and mixing them into a bowl. The content of the bowl only reveals the sum of the amounts of water in the glasses, yet it is impossible just by looking at the water in the bowl to find out what the original amount of water was in each individual glass. This means that larger clusters of photons must be discarded while creating the histogram in order to more accurately measure the spectrum of the light source.

A thorough description of the algorithm that was devised for this thesis work is presented in chapter 5.

3.4 Noise and dynamic range

Noise is an unwanted signal that is inevitably added to any real physical system. In cameras, noise can take many different forms. Arguably the most relevant types of noise in scientific cameras are shot noise, thermal noise (also called "dark current noise") and read-out noise (cheaper consumer cameras sometimes need to also take into account "fixed-pattern noise"). The amount of noise in an image is actually of little relevance to determine its quality, as what is

actually important is how much the desired signal stands out from it. The ratio between the power of the wanted signal P_S and that of the noise P_N is called signal-to-noise ratio (SNR), and is a figure of merit of a system:

$$\text{SNR} = \frac{P_S}{P_N}. \quad (3.1)$$

Equivalently, the SNR can also be defined as the signal (for instance the mean number of photo-generated electrons S) divided by its standard deviation σ (which is the noise). The total effective noise of a system σ_{eff} can be calculated as the root mean square of the individual noises:

$$\sigma_{eff} = \sqrt{\sigma_{\text{shot}}^2 + \sigma_{\text{th}}^2 + \sigma_r^2}, \quad (3.2)$$

where σ_{shot} is the shot noise, σ_{th} is the thermal noise and σ_r is the read-out noise.

3.4.1 Shot noise

Shot noise is the noise due to the particle nature of light itself. Photons are light quanta, which means that they travel in discrete packets (although light also exhibits wave-like behavior). Because photons can either be absorbed by the CCD or not, the photo-generation of electrons in the sensor is not continuous, but also discrete. Furthermore, the arrival of the photons on the sensor itself is intrinsically random (it should be highlighted that shot noise is actually due to the particle nature of light itself, not to the detector). If the number of photons is sufficiently high, the effect of any fluctuations over the total number of photons received per unit time is insignificant. However, when the intensity of the light is low, the number of photons per unit time is also low, and any fluctuation in the stream of photons has a much greater importance. To give an intuition about this, it is much easier for a human to notice a change in the number of people that exits a room if the group of people is small, rather than trying to spot the same change within a large crowd. The SNR of shot noise (if we consider the signal to be the photons) can be calculated as:

$$\text{SNR}_{\text{shot}} = \frac{\bar{N}}{\sqrt{\bar{N}}} = \sqrt{\bar{N}}, \quad (3.3)$$

where \bar{N} is the average number of photons that we would expect to receive on the camera over a certain time. As the formula suggests, shot noise is quite irrelevant when the illumination is high, as the signal to noise ratio increases.

However, when the number of photons is low, shot noise becomes significant. In single-photon counting, the amount of photons is purposely set to be low to reduce the possibility of more than one photon to hit the same spot on the CCD sensor, which means that the shot noise is quite significant.

However, because variations in the total number of photons hitting the sensor over time do not influence the spectrum, but just provide better or worse statistics from shot to shot (no pun intended), shot noise doesn't actually matter from a spectral measurement standpoint, if considered upon the entire sensor rather than on each pixel. Nevertheless, shot noise on each pixel could lead to multiple photons hitting the same pixel, which is undesirable.

3.4.2 Thermal noise

Another type of noise which actually does not depend on the amount of light hitting the sensor is thermal noise. Thermal noise is caused by electrons which are thermally generated within the depletion region of each pixel, rather than being photo-generated. When the camera counts the amount of charge in each pixel, the thermally generated electrons which have merged with the photo-generated ones (only the latter having an actual dependence on the photon energy, in the case of X-ray photons) will be completely indistinguishable from the others, effectively distorting the final result.

Since the rate of thermal generation only depends on the temperature T , the electrons can actually be generated in complete darkness as well, which is why thermal noise is also called "dark current noise". Thermal noise can never be completely eliminated, as this would only occur if the system were set at a temperature of absolute zero, which is impossible [31]. The expression for thermal noise is given here:

$$\sigma_{\text{th}} = \sqrt{I_D(T)t}, \quad (3.4)$$

where $I_D(T)$ is the temperature-dependent dark current and t is the integration time of the camera (which includes the exposure time plus the read-out time). This means that dark noise can be greatly reduced either by choosing a smaller integration time t or by cooling the camera sensor (the camera used in the experiment is cooled at -70°C for this exact purpose). The exposure time was set at 100 ms for the experiment. Notice that choosing a shorter exposure time would be quite irrelevant for thermal noise, as the majority of the thermal generation occurs during the read-out of the camera which is 2 orders of magnitude greater.

As previously mentioned in section 3.2.1, the pixels in a CCD sensors are read sequentially row by row. The rate at which this is done is called "read-out rate", and in the experiment this was set to be 1 MHz. Since the number of pixels is $2048 \cdot 2048 \approx 4.2 \cdot 10^6$, reading at this rate means that the charge generated in the pixels over the entire chip needs 4.2 seconds to be fully collected. Because the rows on the bottom part of the sensor are read first compared to the ones at the top, the former will collect less thermal noise, as their total integration time is smaller. If the temperature is high enough so that the contribution of dark current noise over the total effective noise is significant, this will show up in an image as a gradient over the vertical direction of the image. Figure 3.8 shows a sequence of dark frames (images taken without any light on the sensor), where the effect of this noise gradient can be seen at different temperatures. As is evident from the montage, at room temperature the effect of thermal noise is completely dominant over all other noise sources and the gradient effect is extremely visible (the image contrast had to be adjusted to compare it with the others, but notice that the line plot on the right shows pixel values well over 10000 counts), rendering the CCD camera effectively unusable at this temperature. The gradient is still noticeable up to around -40°C , below which it is basically nullified. At -70°C , the noise due to thermal generation can be considered irrelevant, being well below $1\text{ e}^-/\text{px}$ over the total read-out time (see table 4.1).

3.4.3 Read-out noise

Finally, the last relevant type of noise scientific CCD cameras is read-out noise. When the charge is transferred after the image has been acquired (see figure 3.6), it is converted into a voltage by means of a charge-to-voltage converter, which is an amplifier that produces an output voltage proportional to the injected input charge. The amplification process is noisy, and adds up a certain amount of disturbance to each pixel, which constitutes the read-out noise. The faster the read-out rate, the more noise (in the forms of electrons per pixel) is added. There is therefore a trade-off between read-out noise and thermal noise, as a higher read-out rate implies less thermal noise but more read-out noise. In the experiment, the camera was set at a read-out rate of 1 MHz which offers a good balance between integration time and read-out noise limitation. Read-out noise contributes much more to the total noise seen on the sensor than thermal noise. In fact, even at the minimal temperature of -70°C , figure 3.8 reveals that the camera is still reading an approximately constant value of noise around 300 counts, which is due to the read-out amplification process. The 300 counts

isn't noise, but just an electronic bias offset value that is clamped at that value. The actual read-out noise is the standard deviation around this offset, which was calculated at 1 MHz to be about 2.5 counts. This corresponds to a noise level σ_r of about $(2.5 \cdot 11.7)/3.65 = 8.125 \text{ e}^-/\text{px}$, which is slightly larger than the one advertised by the manufacturer in the datasheet (see table 4.1). To know how the calibration factor 11.7 was obtained, see section 5.2 in chapter 5.

3.4.4 Dynamic range

The ratio between the full-well capacity (in electrons) and the total effective noise (in electrons, considered for all the pixels in the chip) is called the dynamic range of the camera. This value determines the ability of the sensor, for a given exposure time, to simultaneously capture images of scenes in which the illumination is very different, from dark to bright.

When taking a picture of a sunset with a cheap camera, it will be the case that the image will either appear with a properly exposed sky but too dark in the foreground, or with a properly exposed foreground but a blown-out saturated sky. This happens because the dynamic range of the scene (the difference in luminosity between the darkest part of the image and the brightest one) is greater than the dynamic range offered by the camera.

The camera used for the experiment is a scientific-grade camera where read-out noise dominates over the rest, so its dynamic range d_R can be calculated as:

$$d_R = \frac{150000}{6.8} \approx 22059 : 1, \quad (3.5)$$

at a read-out rate of 1 MHz. Notice that the value of 6.8 comes from the datasheet of the manufacturer (table 4.1). The equation above means that the sensor can detect regions of illumination that are more than 22000 brighter than the absolute darkest ones without saturating. Dynamic range is often written in terms of dB, and doing this for the camera used in the experiment yields:

$$[d_R]_{\text{dB}} = 20 \log(22059) \approx 86.87 \text{ dB} \approx 14.48 \text{ stops}, \quad (3.6)$$

where 1 stop = 6 dB is a typical unit used when talking about dynamic range in photography. This dynamic range is exceptionally high, to put this in perspective it is about as high as that of an Arri Alexa LF [32], a popular camera that has been used to film many Oscar winning movies and famous TV series [33].

If the illumination on the pixel exceeds the pixel well depth, then it can happen that charge will overflow into neighboring pixels, in an effect called "blooming". Another similar effect called "streaking" which actually can happen even in dark frames is particularly noticeable at high temperatures (see figure 3.8, at room temperature and -10°C especially), where it can be seen that there are a few long vertical lines over the image. This happens when some defective pixels, so-called "hot pixels", tend to produce more noise than the other ones. As the high temperature makes thermal noise more prominent, it is possible that as the charge is shifted downwards at read-out, these pixels will keep accumulating unwanted electrons, therefore transferring them to the pixels in the column above, creating the vertical streaks.

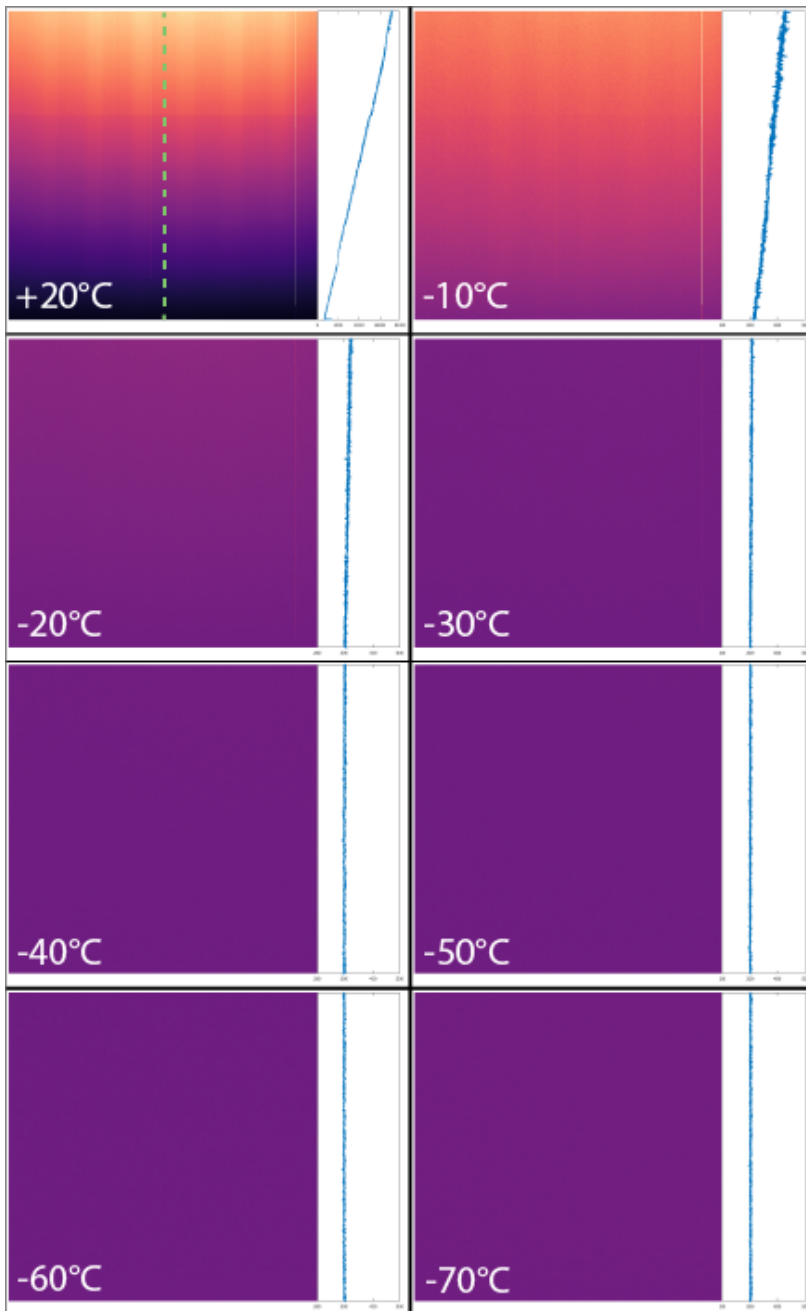


Figure 3.8: Dark frames for different sensor temperatures. The CCD was exposed for 100 ms without any light. The 8 images represent the exposure of the CCD, while the plots on the right of the images are line plots (tilted by 90°) over a vertical column bisecting each image (indicated in the first image with a dashed green line), which show the noise gradient at higher temperatures.

4

EXPERIMENTAL SETUP

This chapter illustrates the main elements of the experimental setup used for the measurements: laser system, target chamber, X-ray camera and Ross filter.

Contents

4.1	Laser system and target chamber	42
4.1.1	The Lund High-Power Laser Facility	42
4.1.2	Experimental setup	42
4.2	X-ray camera	45
4.3	Ross filter	46

4.1 Laser system and target chamber

4.1.1 The Lund High-Power Laser Facility

As of 2022, the Atomic Physics department at Lund University possess a high-power laser system that is capable of delivering ultrashort laser pulses with a peak power of almost 40 TW. For comparison, this is about equivalent to the average power consumption per second of all of the human world in 2019 [34].

The Terawatt laser uses a Ti:Sapphire crystal with central wavelength of 800 nm (in the near IR), and each pulse has 1.5 J of energy on target with a time duration of about 40 fs, at a repetition rate of 10 Hz [35]. It uses chirped pulse amplification (CPA) to achieve these feats. Because the optical power is so high, after the compression of the pulse in the CPA process the laser propagates in vacuum, and reflective optics is used in order to avoid optical breakdown that would otherwise occur in refractive elements such as lenses.

4.1.2 Experimental setup

Figure 4.1 shows a sectioned view of the target chamber that is used to perform the laser wakefield acceleration experiments and generate the X-ray radiation used for the multispectral imaging that is the subject of this thesis. The whole chamber is kept at a pressure in the order of 10^{-6} mbar, which can be considered as "high vacuum". The 800 nm laser pulse (1), shown in red, enters the chamber and is redirected by two mirrors to a 45° dichroic mirror (2) which can reflect the IR beam with a reflectivity well above 99% and only allows a very small leakage. The reflected beam, which is still collimated, is sent to an off-axis parabolic mirror (3) which focuses it on the center of the chamber, where a nozzle (4) can be synchronized to trigger the release of a pressured jet of gas, which is then ionized (see chapter section 2.1.5 of chapter 2 for details about this process). The target is normally a mixture between a light gas such as hydrogen or helium and a low concentration of a heavier gas such as nitrogen. As the laser pulse passes through the plasma, electrons are accelerated due to LWFA and are deviated using a large permanent dipole magnet (5) which disperses them in different orbits depending on their




Figure 4.1: Experimental setup - Target chamber. Cutaway side-view render of the vacuum chamber used to perform LWFA and generate the X-rays.

energy. This dipole is placed on a motorized stage so that it can be moved out of the way if needed. The electrons (shown in light blue) then strike a phosphor screen (6) which emits visible photons so that the signal can be picked up by a camera (7). The obtained image can be used to measure the energy spectrum of the accelerated electrons, since faster high-energy electrons will be deviated less than the slower low-energy ones. The electrons, as explained previously in chapter 2, produce X-rays as betatron radiation, which is collinear with the IR beam. The X-ray beam mixed with the IR beam is shown in purple in the figure. To isolate the X-rays, an aluminum foil (9) with a thickness of $3\mu\text{m}$ is used to block the IR beam, while allowing most of the X-rays through. Before this, at about 40 cm distance from the nozzle, a $125\mu\text{m}$ thick copper filter with a 5.5 mm hole (8) is used to select a solid angle that was calculated to result in the X-rays, which have a divergence of about 40 mrad, to not intersect any part of the 2.4 m long pipe (10) which leads to the X-ray camera used for the single photon counting. This was done in order to avoid any type of X-ray fluorescence from the edges of the pipe. Copper fluorescence from the filter itself is possible, but would result in a much smaller contribution compared to fluorescence inside the pipe, as the fluorescence signal is radiated approximately isotropically and therefore diminishes according to the inverse square law. As will be evident later in section 5.1.2 of chapter 5, this countermeasure was only partially effective.

There are some additional diagnostics that are possible using the setup above. First, a small pick-up mirror (11) is used to extract a small portion of the IR beam and direct it to a motorized translation stage (12). With the aid of a few mirrors, this beam (called the probe beam) is intersected on the nozzle perpendicularly to the main focused laser and sent to a camera (13). By changing the position of the translation stage (12), it is possible to take side-view shadowgrams of the plasma wake at different moments in time (for different pulses). Secondly, the leakage from the dichroic mirror (2) can also be imaged on a camera (15), to serve as a pre-interaction diagnostic of the quality of the IR beam. Finally, a remotely controlled flip-mirror (16) can be used to redirect the focused IR beam (shown in orange, to highlight the fact that the nozzle targeting cannot be done simultaneously) on a camera (17) which can be used to examine the focal spot. Lastly, another mirror and camera (which are not shown in this image for clarity) is used to image the plasma wake from a top-down viewpoint, and an additional phosphor screen with camera (also not shown) can be inserted into the beam after moving the dipole away so that the dispersion of the accelerated electron beam can also be measured.

Figure 4.2 shows the experimental setup with the closed target chamber. The IR laser (1) enters the chamber (2) via vacuum tubes (not shown)




Figure 4.2: Experimental setup - Target chamber and X-ray camera. Render of the target chamber with the extension tube used to allow the X-ray camera to be sufficiently far away for single-photon counting. Note that the IR beam, shown in red, is also coming from a vacuum and the chamber is not open as it appears in the figure, which is shown as such just for clarity.

and after the X-rays are generated (see figure 4.1) they are sent into a 2.4 m extension tube (3). The tube terminates with an opening that is covered by a $50\ \mu\text{m}$ film of kapton (not shown), which maintains the vacuum pressure. The X-ray beam exits the tube and propagates in air for about 10 cm, before entering the X-ray camera (4) whose sensor is hidden behind a $250\ \mu\text{m}$ beryllium window to avoid visible light interference. The sensor is also kept in vacuum by means of a flexible tube (not shown) that connects the space between the beryllium window and the sensor to the extension tube (3). The total distance between the sensor and the nozzle in the center of the chamber is about 3.10 m. In the 10 cm of air between the end of the vacuum tube and the camera, a Ross filter can be inserted (not shown in figure 4.2). A description of this filter is provided in section 4.3. Figure 4.3 shows a diagram with a more precise description of the relative distances between the different elements of the setup.




Figure 4.3: Experimental setup - Schematic side diagram. A schematic diagram that explicitly shows the distances between the various components that are relevant for the X-ray experiment. The IR beam isn't shown for clarity.

4.2 X-ray camera

The X-ray camera used for the single-photon imaging is the iKon-L SO by Andor, which is part of the Oxford Instruments group. The datasheet for the camera is available in [30], with the most important parameters reported in table 4.1, and an image of the camera can be seen in figure 4.4.

An explanation of the parameters indicated in the table is provided in the previous chapter.




Figure 4.4: Andor iKon-L SO X-ray camera. The X-ray camera used for single-photon counting. The camera has a 27.6×27.6 mm square CCD sensor. In the experiment, the camera had a flange with a $250 \mu\text{m}$ beryllium window covering the sensor, which allows it to be kept in vacuum to prevent condensation.

Sensor type	BR-DD
Sensor size	2048×2048 pixels (4 MP)
Pixel size	$13.5 \times 13.5 \mu\text{m}$
Image area	27.6×27.6 mm
Min. temperature	-70°C
Pixel well depth	$150000 e^-$
Read-out noise	$6.8 e^-/\text{px} @ 1 \text{MHz}$
Dark current noise	$0.066 e^-/\text{px}/\text{s} @ -70^\circ \text{C}$
Digitization	16 bit

Table 4.1: Key parameters of the Andor iKon L SO CCD camera. [30] The main parameters that are relevant in the datasheet provided by Andor about their product.

4.3 Ross filter

The target used to assess the performance of the metal detection algorithm (see section 5.3 of chapter 5) is a Ross filter. A Ross filter is comprised of a series of stripes of different metals which are arranged in a grid, such that they overlap on each other in some locations. The metal stripes are mounted on a frame holder, in this case a 2.5 mm thick piece of painted aluminium in which an square aperture was drilled. Figure 4.5 shows the Ross filter that was constructed for the experiment.



Figure 4.5: Ross filter used in the experiment. The Ross filter contains different metallic stripes. The copper stripe is $25\ \mu\text{m}$ thick, the titanium stripe is $6\ \mu\text{m}$ thick and the zinc stripe is $5\ \mu\text{m}$ thick. Originally, the filter included also a nickel strip, however it eventually got damaged so it was discarded. The stripes are held in place on the opening with kapton tape.

When imaged on the camera, the filter was rotated 90° counterclockwise. The opening of the filter's frame measures $30 \times 23\ \text{mm}$. This was done to accommodate the possibility of using metal stripes that were cut from square foils of $25 \times 25\ \text{mm}$ size, which were actually discarded later after being damaged. However, the sensor size of the camera is $27.6 \times 27.6\ \text{mm}$. Because the filter is placed at a few centimeters from the camera, and both are placed more than $3\ \text{m}$ from the source, the X-rays are approximately collimated in this region, which means that a small part of the aluminium frame holder is actually imaged on the sensor as well (the photon flux through it is however almost insignificant, as the transmission through such a thick piece of metal is very low). Figure 4.6 shows how a render of how the Ross filter is placed when it is imaged on the camera.




Figure 4.6: Ross filter placed between the end of the vacuum tube and the X-ray camera. The filter was rotated 90° counterclockwise (it doesn't matter since the sensor has a square aspect ratio). The filter can be inserted or removed quickly using a flip mount mechanism (not shown). The X-ray beam is shown in blue.

5

MEASUREMENTS AND DATA ANALYSIS

This chapter contains a description of how the algorithm developed to perform single-photon counting works, and the process used to calibrate the system and detect the different metals in the Ross filter.

Contents

5.1	The single-photon counting algorithm	50
5.1.1	Cluster search algorithm	50
5.1.2	Spectral measurement	54
5.2	Energy calibration	58
5.3	Metal detection algorithm	62
5.4	Increasing the spatial resolution	63
5.4.1	Overlapping sectoring	64
5.5	MATLAB code	65

5.1 The single-photon counting algorithm

5.1.1 Cluster search algorithm

As previously mentioned in chapter 3, it is not possible to construct a normal histogram of the entire image in order to obtain the spectrum of the X-rays, because photons will frequently deposit their energies in more than only one pixel. By examining again figure 3.4, we can extrapolate that even for a photon energy of 20 keV, which is greater than what we would expect from the betatron radiation spectrum for the experiment, the diameter of energy deposition would be about $12\ \mu\text{m}$. Because the size of a pixel in the CCD camera used for the experiment is $13.5 \times 13.5\ \mu\text{m}$, we know that all photons which we can reasonably expect will always deposit their energies in an area that is smaller than a single pixel.

In spite of this, it is possible (and quite common) that the energy from a single photon is deposited over adjacent pixels because the photon may hit the pixel in its peripheral region. However, combining this fact with our knowledge of the diameter of the electron cloud generated subsequently to the photon's absorption we can establish that only a finite amount of configurations of neighboring excited pixels may exist for the excitation of a single photon. These bunches of neighboring pixels represent all the possible ways in which the electron cloud can overflow into adjacent pixels. They will be from now on referred to as "valid clusters", and they are shown in figure 5.1. The reason why they are referred to as valid cluster can be understood by looking at figure 5.2. All the cluster shapes which differ from those in figure 5.1 are the result of more than a single photon and therefore need to be discarded by the algorithm, as previously explained.

Although each single-photon event will always appear on the image as one of the valid clusters in some orientation, no one-to-one correspondence between number of photons and type of cluster exists. This is because it is always possible for two or more photons to strike the sensor so close to each other that they still generate one of the valid clusters (a "false positive"). This is an unavoidable problem, and each time it occurs there will be an overestimation of the photon energy (if more than one photon is detected as being a single photon, the sum of the energies of all the photons will be

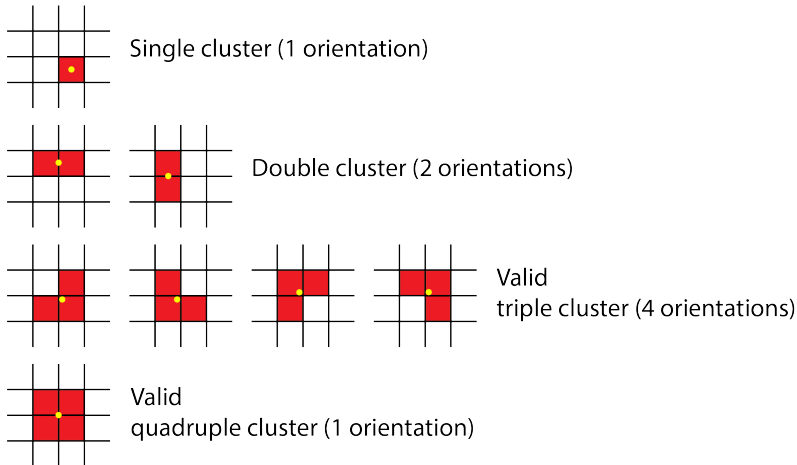


Figure 5.1: Valid cluster types. The clusters that are possibly generated by the interaction with a single photon are shown. The red color represents the pixels that have had a sufficient energy deposition so that their pixel counts exceed a threshold value above the surrounding noise. The yellow dots represent the likely position where a photon may have hit the pixel to generate the cluster.

summed and incorrectly attributed to a single high energetic photon, resulting in a blue shifting of the spectrum). However, this is why it is important that the photon flux isn't too high, as the occurrence of such events where more photons strike in the same spot becomes statistically less probable the less photons hit the sensor, as they will be more spread out from each other.

This creates a trade-off, as less photons also implies less statistics to construct the spectrum. The number of statistics can be increased by taking multiple images and combining the information from each one of them together, as will be shown in section 5.1.2. As mentioned previously, the histogram to construct to measure the spectrum needs to be performed on valid clusters, not on individual pixels. It therefore follows that it is necessary to construct an algorithm which is able to identify all elements above a threshold (to differentiate them from noise), discard any invalid cluster, and calculate the histogram using exclusively the remaining valid elements.

The issue of finding the clusters is a popular problem in programming, called the "number of islands" problem. It is a variation of a problem in graph theory which is concerned with finding all the connected components in an undirected graph. The image from the camera is represented as a matrix of pixels, but each pixel could also be represented by a vertex in a graph. Pixels which are illuminated by X-ray light may have neighbors which are also illuminated, and if this happens the analogous situation in a graph would be




Figure 5.2: Example of a background image (no filter used) in single-photon regime. A close-up of the image shows some of the different shapes that the clusters can take when photons hit the sensor. Only the clusters whose shapes belong to the ones described in figure 5.1 can be the result of single photon events, and therefore are referred to as "valid clusters". All other shapes surely originate from more than a single photon event and are therefore referred to as "invalid clusters". The colors represent pixel counts, and only pixels with more than 320 counts are shown to differentiate them from the background noise level, which is around 300 counts for all images taken throughout the experiment.

that neighboring vertices are connected by arcs. The analogy between an image (matrix of pixels) containing clusters and a graph containing connected components can be more easily understood by looking at diagram 5.3. The connected components (clusters) of the matrix are found in the following way:

1. The image is scanned pixel by pixel (in practice, its matrix is scanned cell by cell with two nested "for" loops) row by row from left to right.
2. Each time that a cell is analyzed, the algorithm checks whether the cell has been visited before (it keeps track of this with a Boolean matrix, i.e. a matrix which only contains 0s (if cell is unvisited) and 1s (if cell is visited)). If the cell has been previously visited, then the next cell is analyzed. However, if the cell is unvisited, the algorithm checks if its pixel value is above a threshold.
3. If the cell's pixel value is above threshold, then the cell is considered a cluster (momentarily with dimension 1). Whether each cell has been

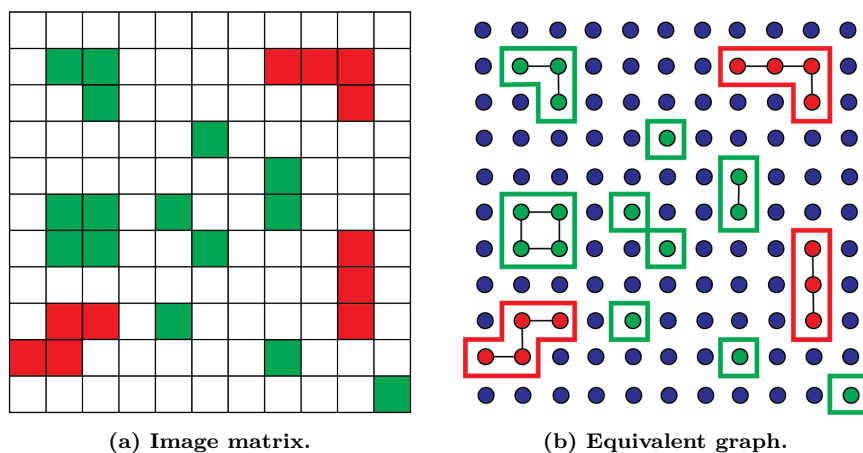


Figure 5.3: Duality between a matrix with clustered cells and an undirected graph with disjointed connected components. Subfigure (a) shows a grid which represents one of the matrix images acquired by the camera. Each cell represents a pixel, and some groups of neighboring cells above threshold can be identified as valid clusters (green) or invalid clusters (red). Equivalently, the cells of the matrix can be represented as vertices of a graph, indicated with blue dots (b). Each cluster can be thought of as a connected components of the graph, and is indicated in color (green if valid, red if invalid). Lines between vertices represent undirected arcs of the graph.

identified as part of a cluster is also recorded into another binary matrix. A counter is initiated, which keeps track of the size of the cluster, and a recursive function is called. This function pauses the main "for" loop, and then proceeds to search the four adjacent cells to the current cell to the north, east, south and west. It "asks" these cells the same 3 questions: "is the cell unvisited?", "is the cell above threshold?" and "does the cell not belong to any previously identified cluster?". If the answer to all questions is "yes", then the function calls itself in a recursive manner, searching the adjacent cells to this new one, and keeps doing it until it exhausts the full size of the cluster. If the answer is "no", the function "breaks" the sub-loop and resumes the previous instance of the function. This method is called depth-first search (DFS) because every time a new element is searched with the recursive function, each branch of subsequent investigated cells is explored as deeply as possible before backtracking onto a new one. The reason why the algorithm also keeps track of whether a cell is already part of a previously identified cluster is to avoid infinite recursion. All of this is done while keeping track of the size of the cluster with a counter, and also while checking that each adjacent cell is actually within bounds of the total matrix.

4. Each time that a cluster is fully discovered, before resuming the scanning

of the matrix, another algorithm is delegated to finding out whether it is a valid cluster or not. It does in the first place by checking its size. If the size of the cluster is greater than 4, the cluster is automatically deemed invalid. All single and double clusters are accepted. Triple and quadruple clusters require more attention, since not all triple and quadruple clusters are acceptable (only "L" shaped triple clusters and "square" shaped quadruple clusters are valid, see figure 5.1). A sub-algorithm can tell the shape of these clusters and only accepts the valid configurations.

5. The coordinates (row and column number), sizes and pixel counts (considered as the sum of all pixel counts within the cluster, after subtracting the average noise) of all valid clusters are recorded in an $N \times 4$ matrix where N is the number of clusters. The coordinates only identify the location of the pixel having the largest number of counts within the other ones in the cluster. Because the maximum distance between elements of a cluster is 1 row and 1 column (diagonal elements in a triple or quadruple cluster are the maximally spaced pixels in a valid cluster), the error associated with this method is minimal.

The threshold value was 320 counts, as the image background noise level was 300 counts with a standard deviation of about 2.5 counts. This allowed the threshold to be set about 8 standard deviations above the noise average, which should effectively minimize the amount of false positives in the single-photon counting. The proportions between the different cluster sizes across 9 different background images can be seen in figure 5.4. As evident from the diagram, clusters with high dimension (higher than 4) represent a small percentage of the total, which means that we are approximately in single-photon regime.

5.1.2 Spectral measurement

Once all the invalid clusters have been discarded and the matrix that contains the coordinates of all valid clusters has been generated, the spectrum can finally be calculated by simply taking the histogram of the 4th column of the $N \times 4$ matrix, which contains the pixel counts of each cluster. Coupling this with logical indexing in MATLAB also allows to selectively plot the histogram considering only single, double, triple or quadruple clusters individually, as can be seen in figure 5.5. The general shape of the spectra of the different clusters is the same (this can be confirmed by normalizing the spectra by number of clusters). This is expected, and corroborates the hypothesis that all 4 types of cluster represent valid statistics, representative of single-photon events. Furthermore, the spectra are in accordance with those expected from theory of betatron emission [20].




Figure 5.4: Pie diagram showing proportions among differently sized clusters. The amounts of each cluster size (valid and invalid) up to dimension 10 that were detected across 9 different background images can be seen in percentage of the total. The data contains a total of 418045 clusters. Higher cluster dimensions were ignored because they occurred in insignificant amounts.

As previously mentioned, to obtain more statistics one could increase the photon flux, to collect more photons. This is a trade-off, because going too high with the number of photons decreases the amount of valid clusters and also increases the number of false positives. For this reason, it is better to instead take multiple images at a low photon flux, and then sum up the histograms for each image. The spectrum for multiple images can be saved in a matrix, where each row stores the histogram of each image obtained by using all valid cluster types. As multiple images are taken, although the experiment conditions are kept as constant as possible, there still arise quite significant shot-to-shot fluctuations, which are especially visible in the total number of photons per image. However, these fluctuations only affect the histograms as a scale factor, yet do not actually influence the spectral shape whose energy distribution maintains itself rather constant over multiple images, as both figure 5.6 and 5.7 show. It should be noted that the energy spectrum presented in figure 5.5 differs from the spectrum that was actually generated inside the chamber. Instead, it arises from the transmission of the original spectrum through $3\ \mu\text{m}$ of aluminium (the filter used in the chamber to block the IR beam), $50\ \mu\text{m}$ of kapton tape (at the end of the extension tube, used to seal the vacuum), $250\ \mu\text{m}$ of beryllium (in front of the sensor, to keep it in

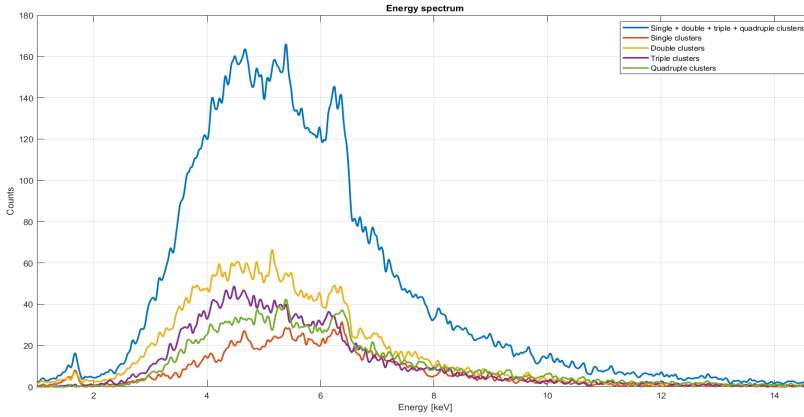


Figure 5.5: Energy spectra obtained using different sizes of valid clusters. The energy spectra for different clusters of a single background image (no Ross filter used) between 0 keV and 15 keV. Each submatrix of the $N \times 4$ matrix corresponding to a differently sized cluster can be used for a histogram. The sum of all histograms is plotted in blue.

vacuum in order to avoid condensation) and 10 cm of air. It also doesn't take into account the QE of the BR-DD detector. To obtain the original spectrum, one needs to divide the spectrum by the transmission spectra of all the three filters and also by the quantum efficiency. However, as will be shown in section 5.3, doing this is actually irrelevant for the final goal of detecting the species of an element in X-ray absorption spectroscopy, as all spectra that arise after transmission through these filters differ from the original one in the chamber by simple scalar division of a common factor (the product of all the 3 filters transmission spectra and the QE). When a spectrum is divided by another this scalar factor cancels out.

It can be seen that there consistently is a spectral peak slightly above 6 keV. This peak is actually not due to betatron radiation, but was thought to be originated from chromium X-ray fluorescence within the extension tube where the beam propagated, which contained some traces of it. To combat this, a drilled copper plate with a thickness of 250 μm was inserted into the chamber, selecting a solid angle of the X-ray beam small enough to allow propagation in the tube without intersecting its inner surface. The height of the peak was greatly reduced, but the thickness of the copper plate was proved to be insufficient to completely eliminate it, as the results show. Regardless, the presence of the peak doesn't affect the metal detection algorithm described in section 5.3, as the peak appeared consistently across all shots, causing the spectrum to be relatively time invariant.




Figure 5.6: Energy spectra of 9 different background images. The energy spectra for different images is combined visually in a 2D image. The figure was simply obtained by showing the $9 \times M$ matrix of the combined spectra as an image, where M is the number of bins used to construct the histogram. The spectra are normalized to their mean as to show more clearly their similarity in shape. The spectrum is shown between 2 keV and 15 keV.




Figure 5.7: Line plots of energy spectra of 9 different background images compared. The energy spectra for different images is plotted on simultaneously in the same graph. The different spectra align very well, meaning that in spite of variations in the photon flux among different background images, the spectral shape remains the same. The spectra are normalized to their mean as to show more clearly their similarity in shape.

5.2 Energy calibration

As previously mentioned, the algorithm which calculates the histograms of the valid clusters actually does so using the pixel counts. However, there is a scaling factor between pixel counts and energy in eV which must be measured to calibrate the system.

The calibration was done by following these steps:

1. A background spectrum is obtained by just illuminating the camera with X-rays, with the method described in the previous section.
2. A thin sheet of copper foil (25 μm in this case) is then placed in front of the camera, so that the beam passes through it. The spectrum is calculated again, but now it will appear distorted by the fact that the absorption coefficient depends on frequency.
3. The filtered spectrum is divided by the background spectrum, which yields the transmission spectrum of copper (see figure 5.8).
4. The transmission spectrum of copper shows a dip in absorption (the K-edge), which can be compared with tabulated data (see figure 5.9).
5. The number of pixel counts at which the measured K-edge is found then corresponds to tabulated K-edge energy of copper, which gives the calibration scaling factor.

The dip in the measured transmission spectrum of the copper plate was seen to occur at about 768 pixel counts. Since the K-edge of copper is at 8978.9 eV, the calibration conversion is calculated as:

$$1 \text{ count} = 11.7 \text{ eV}. \quad (5.1)$$

It is possible to calculate the spectral resolution in energy with the Rayleigh criterion. If the X-ray spectrum was perfectly monochromatic at an energy E_0 , it would appear as a Gaussian peak centered around E_0 with a standard deviation equal to the total effective noise σ_{eff} (see section 3.4 in chapter 3) appropriately converted in units of energy, with an intensity I (in arbitrary units):

$$I = \exp \left[-\frac{(E - E_0)^2}{2\sigma_{eff}^2} \right]. \quad (5.2)$$

The total effective noise is approximately equal to the read noise, which was measured at 2.5 counts, corresponding to an energy of $2.5 \cdot 11.7 \approx 29$ eV. At a distance of one standard deviation from the mean of a normal distribution, the intensity drops to $1/\sqrt{e}$ of the maximum. The half width half max (HWHM) with some calculations can then be seen to occur at about $1.18\sigma \approx 34.5$ eV. It can be concluded that two energy peaks can be distinguished if they are separated by a full width half max (FWHM) in energy of $E_{\text{FWHM}} = 2.35\sigma \approx 69$ eV.

The actual energy resolution is actually inferior to this, because another effect other than noise contributes to energy broadening. Concretely, the choice of threshold influences the spectral resolution as well. Since the threshold for considering a photon event was arbitrarily set at 8 standard deviation above the electronic offset bias (at 320 counts), if a pixel exhibits anything less than $320 - 300 = 20$ counts above the bias it will not be detected by the algorithm, even if a small part of the energy actually was deposited in the pixel. This means that in the worst-case scenario, the energy of a photon can be underestimated by as much as $[(20 \cdot 2) + 20/\sqrt{2}] \cdot 11.7 \approx 633$ eV. This result arises when considering a photon which hypothetically hits a pixel in proximity of a corner, so that the photogenerated electron cloud partially overflows in the 3 closest neighbors of the corner with a sub-threshold amount of counts. This level of uncertainty is higher than that found in previous studies [8] of about 200 eV, but it could be improved by reducing the threshold (with the downside of introducing more false positives). This value is in accordance with the experimental measurements, as the K-edge in the middle panel of figure 5.8 appears broadened by about that amount, compared to the much sharper falling edge of figure 5.9.

This resolution is much worse than that of spectrometers based on crystals, which offer resolutions in few of meV range [38], but as will be apparent in chapter 6, it is enough to effectively detect metals using the K-edge method.

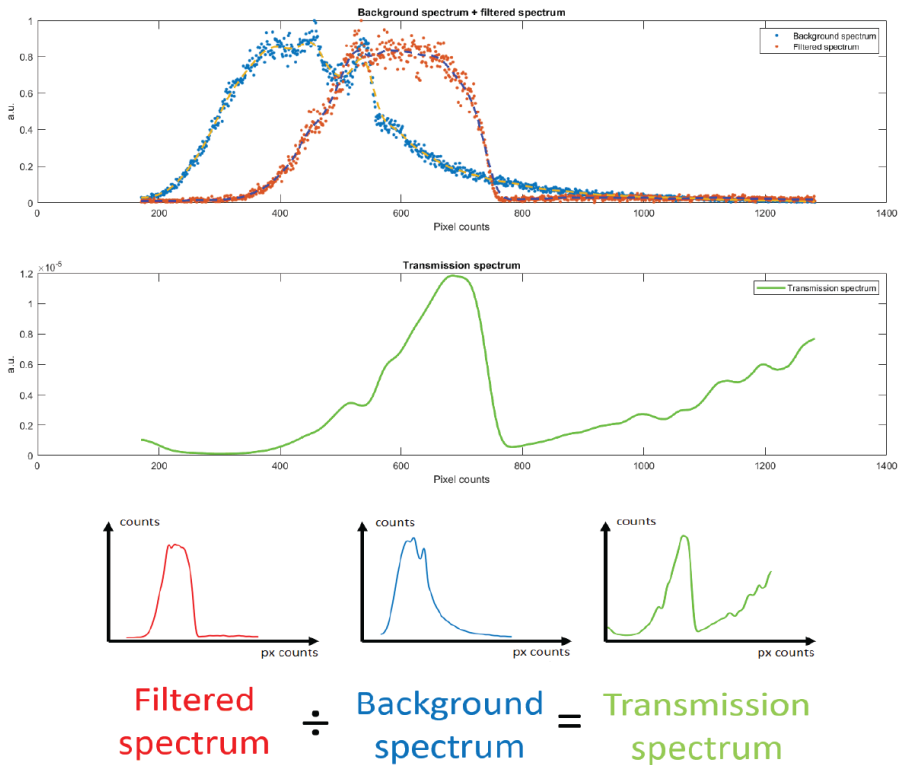


Figure 5.8: How the transmission spectrum for a metal is obtained. This is the crucial idea behind the whole experiment. The filtered spectrum (red) can be divided by the background spectrum (blue) to give the transmission spectrum (green) of whatever material is in front of the camera. This was done with a known material (copper) to calibrate the system, but the same process can be used to detect any other element with a K-edge within the energy band of the X-rays. The top two plots show the actual experimental data used for the calibration, notice how there is a dip in the transmission spectrum, corresponding to the K-edge of copper. Notice that the filtered spectrum and the background spectrum have been normalized in the first panel above.

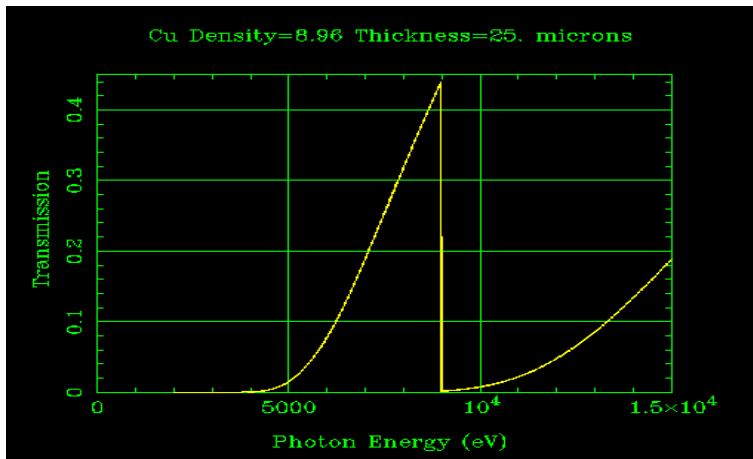


Figure 5.9: Theoretical transmission spectrum of 25 μm of copper. [37]

The measured transmission spectrum is compared with a tabulated version which is readily available online. Copper was the metal of choice for the calibration, because its K-edge dip is very sharp (the transmission almost reaches zero at the K-edge energy in this case) and therefore would be more easily visible in the measured spectrum compared to other metals. The reason why the peak appears sharper here than that in the bottom of figure 5.8 is because the energy resolution is much better in the tabulated case compared to the single-photon counting measurement.

5.3 Metal detection algorithm

Once the spectra have been calibrated (the pixel counts are converted to eV), it is in principle possible to do the same procedure as in figure 5.8 and obtain the transmission spectrum of any unknown material in the path of the X-ray beam. Afterwards, one could check at which energy an eventual K-edge occurs, and relate it to some element, effectively being able to identify it.

However, doing this by hand is a time-consuming procedure that can be automated. An explanation to how this was done in practice is as follows:

1. The transmission spectrum for an unknown element is derived with respect to energy (see figure 5.10). This is done because any sudden dip in the transmission spectrum will show up as a sudden negative oscillation of the derivative.
2. If any part of the derivative goes below a certain threshold, the program is alerted that there is a K-edge in the transmission spectrum (which means that some metal is in the path of the X-rays), and retrieves the energy at which any local minimum occurs within all below-threshold regions. Note that this method can detect more K-edges simultaneously (this can occur if different metals are overlaid on each other).
3. The algorithm now verifies if the value (or values) of energy for which the derivative is below threshold belongs to any interval within 200 eV of the nominal K-edge energies of different elements (see table 3.1 in chapter 3). Each interval is disjoint from the others so that it is impossible to have any ambiguity in the detection.
4. If there is a match between a measured K-edge and one of the intervals, the metal corresponding to the interval is considered to be detected.
5. An array is created which contains as many entries as there are detection intervals. If no metal is detected, this array is 0 for all its entries. If a metal is detected, its corresponding index in the array is set at a value which is proportional to the amount of "overshoot" below threshold. This is done to not only limit the detection to a binary answer (detected or undetected element), but rather to give a value of confidence (if the K-edge is very prominent, the overshoot past the threshold is higher, therefore the confidence of detection for the element should also be higher).

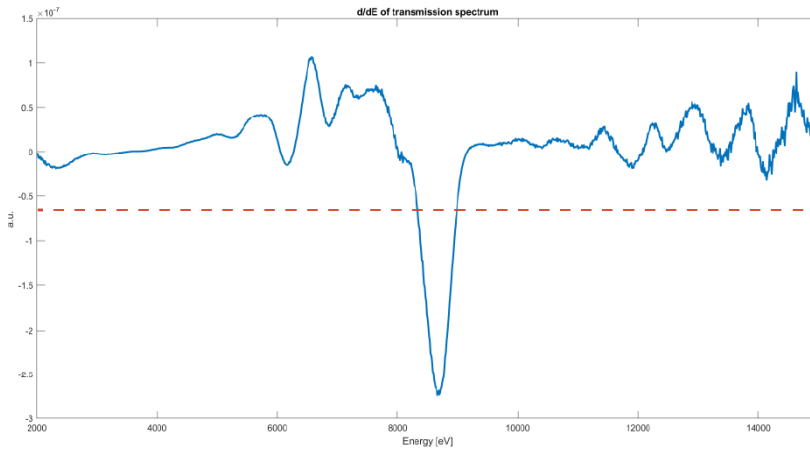


Figure 5.10: Derivative with respect to energy of the transmission spectrum of an unknown element. When a material absorbs X-rays at its K-edge, an absorption dip will be visible in its transmission spectrum. This dip shows in the derivative with a negative oscillation. If the derivative goes below a threshold (dashed red line) then the algorithm understands that there is an element in the path of the X-ray beam and identifies it based on the energy for which the dip occurs.

5.4 Increasing the spatial resolution

It is now possible to detect the presence of any element with a K-edge within the energy band of the X-rays. However, doing this over an entire image is not particularly interesting, as it offers no spatial resolution. This problem can be solved by sectoring the image in a multitude of quadrants, such that they cover the entire image. Then, by detecting the metals at each quadrant, an array of confidence for each metal can be created for each location, with an increasing spatial resolution which linearly increase for increasing subdivisions of the matrix. This allows to create different channels for different detected metals for the same image, that is, different images which only show the detection of a particular element. For example, it is possible to create a "copper channel": an image with as many pixels of resolution as the number of quadrants, which has different shades of gray for each quadrant depending on the confidence of detection of copper at that particular location. See section 6.1 in the next chapter to see these metal channels in practice.

Subdividing the image comes at a cost, however, as the size of each quadrant will collect a smaller amount of photons, reducing the quality of the spectrum calculated over the quadrant. This is why multiple images

are taken. For example, if the full image is 2048×2048 pixels, and each quadrant is 205×205 pixels, it will take about 100 images to obtain the same amount of statistics in a quadrant as 1 single shot calculated over the entire image. The spectrum in each quadrant is actually computed after processing the clusters of all the images. This can be done because the coordinates of all the clusters are stored in memory by the algorithm. When collecting multiple images, the cluster coordinate matrices of all images can be merged to increase performance, and the image subdivision into quadrants is done "virtually" in the sense that all calculations are done on the matrix that holds the coordinates of all valid clusters and not on the sensor itself before each shot. This means that the size and number of quadrants can be adjusted even after collecting the coordinates of all the clusters.

5.4.1 Overlapping sectoring

Partitioning the matrix in a grid where each sector is dissociated from the others (as in "binning") would require a tremendous amount of images in order to obtain a reasonable amount of statistics. Considering that an acceptable spectrum can be obtained with a single shot over a full image, even just a single 10×10 partition would require 100 images to be combined, with less than ideal results.

Instead, a better way to obtain much higher resolution is to virtually scan the image at discrete locations which are closer than the size of a quadrant. In practice, this means that each quadrant overlaps with other ones. An illustration of the different methods of sectoring the image is given in figure 5.11.

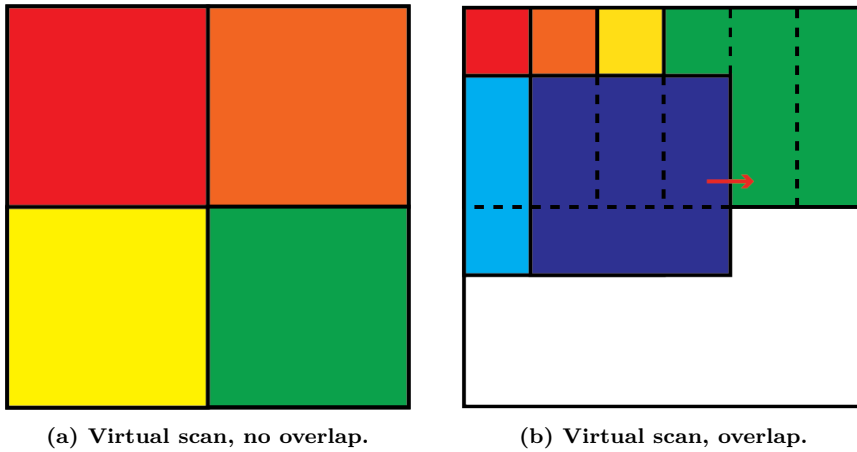


Figure 5.11: Differences between two different sectoring methods of the matrix. Subfigure (a) shows a virtual scan in which each quadrant is disjoined from the others (indicated with different colors to indicate that they're analyzed at different times). If the size of a square quadrant is $1/n$ that of the full image, then the resulting channel will only have a lateral resolution of \sqrt{n} . In the example, a quadrant with an area of $1/4$ th of the total results in a disappointing 2×2 image. Subfigure(b) shows a virtual scan where each quadrant can overlap with the others (not all quadrants are shown for clarity, the quadrants are incrementally scanned in the direction of the red arrow, row by row). Note that the size of each quadrant in (b) is the same as in (a), but now we can fit many more quadrants, allowing higher resolution (at a cost, see section 6.1 in the next chapter) while collecting the same amount of statistics per quadrant (as this only depends on the quadrant's area). In the example, (b) would result in a 4×4 image.

5.5 MATLAB code

The code that was developed by the author for this thesis work is available online as a GitHub repository at the following link: <https://github.com/Andratos95/Single-Photon-Counting-Algorithm>.

6

RESULTS AND CONCLUSION

This chapter concludes the thesis with images of the final results and discusses possible improvements to the technique and future developments.

Contents

6.1	Final results	68
6.2	Possible improvements	70
6.3	Conclusion and outlook	72

6.1 Final results

It is finally time to show the final results that were obtained using the algorithm. Figure 6.1 shows a saturated image (in the sense that there is a large number of photons compared to an image where single-photon counting could be performed) of the Ross filter, which was obtained by summing together a series of 100 images, individually in single-photon regime. The image will serve

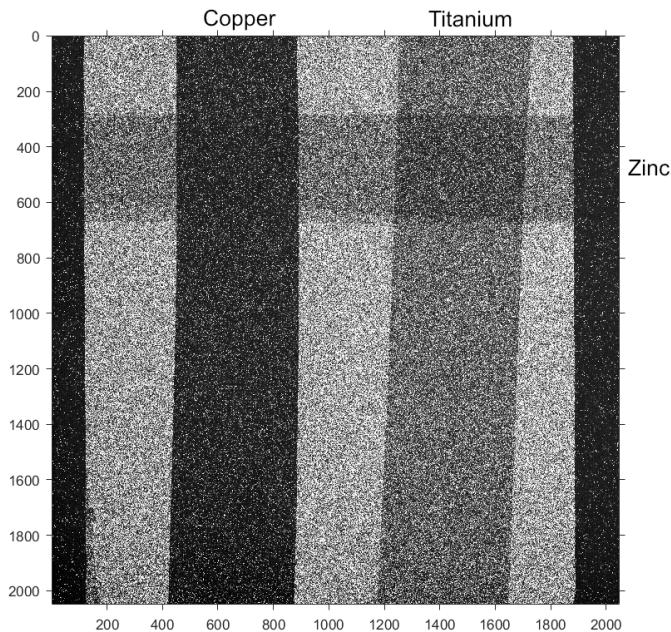


Figure 6.1: Saturated image of the Ross filter. The image was obtained by combining 100 images of the Ross filter in single-photon regime. The same 100 images were used to create the filtered spectrum in all the detection quadrants of the algorithm. The different metals for the stripes are marked on the sides of the image.

as a benchmark for judging the quality of the final result. As can be seen, the image shows the three stripes of titanium, copper and zinc (see figure 4.5 and 4.6 in chapter 4). By just observing the image without previous knowledge of what type of elements were used, it would be impossible to tell the composition of the stripes. However, the algorithm is able to use single-photon counting across all images and generate the appropriate metal channels. Figure 6.2 shows the metal channels of titanium, copper and zinc, which were obtained using 16384 overlapping quadrants with a size of 256×256 pixels, distributed

over the full sensor (as much as possible, 14 pixels were wasted on the right and bottom side of the image due to the impossibility of achieving fully uniform tessellation of the image). This results in each channel to be a 128×128 pixels image. It can be seen that copper was detected with the greatest confidence,




Figure 6.2: Metal channels. The metal channels for titanium, copper and zinc which were obtained with the detection algorithm described in section 5.3 of the previous chapter. Each channel has a resolution of 128×128 pixels, and the colorbar indicates confidence of detection.

since as previously mentioned it has a very sharp K-edge compared to the other metals. Since in this example there are exactly 3 metals, it is possible to associate each channel to a color and create an RGB image where each color represents a different detected metal. Figure 6.3 shows the final RGB image, side by side with a comparison with the "true" saturated image. It can be seen that the algorithm can be used to recognize metals with good spatial resolution.

The zinc channel looks worse than the other ones because the zinc stripe was the thinnest, at only $5 \mu\text{m}$, and because the K-edge of zinc is 9658.6 eV , which is really on the tail of the energy spectrum of the X-rays used in the experiment (see for example figure 5.5 in the previous chapter). Furthermore, the overlap between copper and zinc in the upper left portion of the image was undetected, probably due to the fact that the copper stripe, being much thicker ($25 \mu\text{m}$) than the zinc one, further reduced the amount of statistics that would've been necessary to effectively detect the zinc in that region. Nevertheless, the overlap between zinc and titanium can be clearly detected, as shown in the top right portion of the image, where the region in which the stripes overlap appears purple (since titanium is associated with red and zinc is associated with blue in the RGB image).

It should be noted that the channels and the RGB image do not fully image the entire sensor. This is because each pixel in the channels carries information which can be associated to the center of each quadrant. Because the virtual scan starts and ends with quadrants which are fully within the

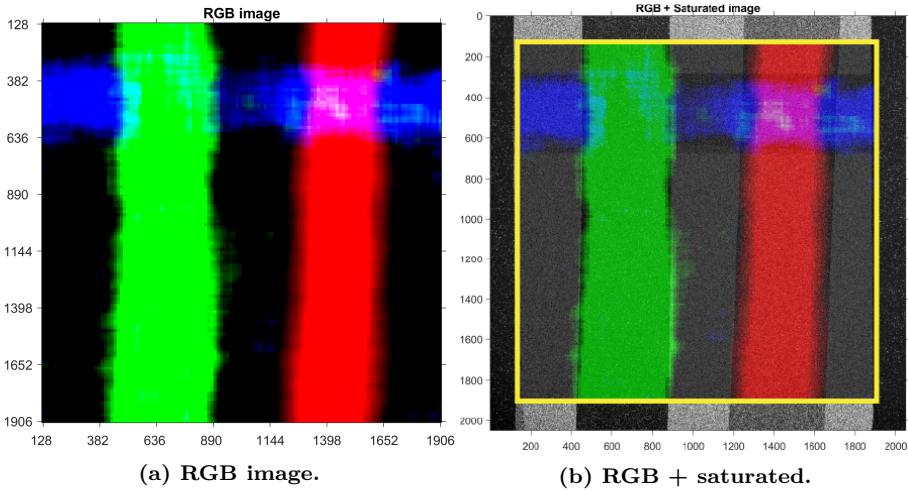


Figure 6.3: RGB image with all the combined channels and comparison with saturated image. The RGB image (a) combines all channels in one, assigning the color red to titanium, green to copper and blue to zinc. The RGB images a zoomed portion of the saturated image due to the way it was constructed, which can be seen in (b), and it spans it from cell (128, 128) to cell (1906, 1906). Note that the complete image is 2048×2048 pixels, so there is a magnification of about 1.33.

bounds of the matrix, and since each quadrant is 256×256 pixels wide, 128 pixels are "wasted" on each of the 4 sides of the square CCD sensor (to be precise, 128 are wasted from the left and top of the image, while 142 are wasted from the right and bottom due to incomplete uniform tessellation of the image with the given side resolution and quadrant size), effectively making the detected metal channels image a slightly magnified portion of the full sensor (to better understand why, see figure 5.11 (b) in the previous chapter, paying attention to how each quadrant starts already within the matrix and not outside).

6.2 Possible improvements

There are a number of things that could have been improved in the general method and the algorithm, some of which are described here:

- Since thermal noise is minimal at -70°C , a slower read-out rate could've been used to decrease the read-out noise (at the cost of increasing acquisition time). However, the camera used in the experiment had fixed

read-out rates, and the next slower setting was 0.05 MHz. This would've taken about 84 s per image, being quite inconvenient for collecting multiple images. The image acquisition method used for the experiment is automated, and all the shots to be used to construct the spectrum can be collected in less than 10 minutes at the read-out rate of 1 MHz.

- The metal detection algorithm could be different and possibly offer better results. Instead of detecting the K-edge using the derivative method presented in section 5.3 of chapter 5 (or in conjunction with it), one could for example cross-correlate the spectrum of the measured transmission spectrum of an element with a set of theoretical transmission spectra at different thicknesses. This may be able to reveal the thickness of a metal besides merely its presence. One could even think to use machine learning for this task.
- The LWFA betatron radiation offers extreme temporal resolution (each X-ray pulse has a temporal duration shorter than that of the IR pulse, therefore less than 40 fs), however this characteristic of the light source was not utilized, as the target was stationary. The method presented in this thesis could however be used to detect metals in extremely fast particle flows such as in jets, engines or even explosives, as long as the target event can be repeated in time with sufficiently low shot-to-shot fluctuations.
- A more intense source of X-rays could be used to image thicker targets, perhaps generated by vacuum tubes such as those used in airport security screening or CT scans. Also, X-rays with different energies could be used to observe other K-edges, therefore allowing detection of metals other than the ones listed in table 3.1 of chapter 3 (note however that one would also need a CCD sensor which could detect these energies).
- The performance of the algorithm was greatly increased since its inception. Originally, the counting of all clusters in a single image took about 4 minutes. After a great amount of optimization, and using MATLAB's parallel computing tool, my laptop equipped with an Intel Core i7 9th Gen processor, 16 GB of RAM and Nvidia GeForce 1660Ti GPU takes approximately 4 seconds per image, generates the histograms of background and filter in about 3 minutes and detects the metal with a resolution of 128×128 pixels in about 2 minutes. The entirety of the algorithm pipeline requires less than 10 minutes of run time to generate the final results. As MATLAB is an interpreted language, faster speed could be achieved using a compiled language such as C. Also, using a more powerful desktop computer would also shorten the run time.

- Since the quadrants are scanned and overlap on each other, the final images present smearing due to convolution. This happens because as a quadrant passes over some metal (e.g. one of the stripes), it starts detecting it even before the center of the quadrant is superimposed on its actual location. This is the same effect that arises when a rectangle is convoluted with itself, leading to a triangle with a wider base. To fix this problem, one could use some deconvolution function. However, the RGB image (and therefore the channels as well) in practice didn't really show extreme signs of smearing, as can be seen in figure 6.3 (b), where only the copper seems to be slightly smeared, therefore this idea was not pursued further. The reason why limited smearing is visible is probably due to the limited sensitivity of the threshold method used to detect a negative oscillations in the transmission spectrum's derivative, which only actually triggers a detection event only when a quadrant is well overlapping a metal.

6.3 Conclusion and outlook

In conclusion, a method to measure the spectrum of X-rays was developed, using the technique of single-photon counting. The results were in good agreement with those expected from the X-ray source in theory. The final goal of detecting metals in a sample with good spatial resolution using multispectral imaging was reached, which serves as a proof of concept of the method.

As mentioned previously, the algorithm developed for this thesis work could be used with any source of polychromatic X-rays, not only betatron radiation. This makes it a possibly valuable tool for a wide range of applications: from geology, to medicine and to food safety. The making of the algorithm from scratch took a substantial amount of work, with the MATLAB code being composed of almost 2000 lines. Better optimization of the code, coupled with GPU parallelization could lead to the execution of the single-photon counting algorithm and the data analysis to be much faster, in the sub-minute range. The real limitation of the technique is the time that is needed to acquire all the images to be combined for accurate measurement of the spectrum in small quadrants.

This could perhaps be improved in the future by using higher photon flux and then dispersing the beam onto a larger sensor, which would allow for more statistics to be collected with a single shot. X-ray dispersion could

be achieved with e.g. Fresnel zone plate. Also, the exact location of metals in an image could be obtained with more refined methods, such as with Bayesian optimization. Lastly, the technique could be used on a rotating target which would allow to identify metals in 3D space, with a tomographically reconstructed image as well.

7

ACKNOWLEDGEMENTS

There have been so many people who helped me in the creation of this thesis work. I would like to start by thanking my supervisor Olle Lundh, who has had the patience to read all of the many emails I sent him and to spend hours discussing the project with me in person. I want to thank Cornelia Gustafsson, for being so kind and available to explain to me so many things in the lab and outside; Erik Löfquist and Kristoffer Svendsen, for literally allowing me to be able to do all the measurements in time and being so nice to me; Anders Persson for teaching me how to use the milling machine and for giving me tips and tricks for the lab. I would also like to thank my girlfriend Flavia for listening to all my babbling about how cool LWFA and single-photon counting are in spite of probably thinking I had lost my mind. I would like to thank my mom Anna, my dad Stefano and my amazing brother Matteo for being always on my side and being some genuinely great friends, beyond just family. Last but not least, thank to everyone who has been part of my life here in Lund for these two amazing years: primarily my great Photonics classmates Dani and Lauro for spending countless hours at Wayne's coffee discussing about physics while having fika; all of the amazing teachers who taught me so many new things and anyone else who I may have forgotten to mention. Thank you all!

Andrea Angella

Lund, Sweden, June 2022

BIBLIOGRAPHY

- [1] Historiek. Wilhelm röntgen (1845-1923) – ontdekker röntgenstraling. <http://historiek.net/wilhelm-rontgen-ontdekker-rontgenstraling/550/>. [Online; accessed 14-May-2022].
- [2] Encyclopaedia Britannica. Theodore h. maiman. <https://www.britannica.com/biography/Theodore-Maiman>. [Online; accessed 14-May-2022].
- [3] Photonics Media. A history of the laser: 1960 - 2019. https://www.photonics.com/Articles/A_History_of_the_Laser_1960_-_2019/a42279. [Online; accessed 14-May-2022].
- [4] T. Tajima and J. M. Dawson. Laser electron accelerator. *Phys. Rev. Lett.*, 43:267–270, Jul 1979.
- [5] Donna Strickland and Gerard Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 56(3):219–221, 1985.
- [6] Nokia Bell Labs. Charge-coupled device: The breakthrough that enabled digital imaging — from dslr cameras to medical endoscopes. <http://www.bell-labs.com/about/history/innovation-stories/charge-coupled-device/#gref>. [Online; accessed 27-April-2022].
- [7] Annibale Mottana and Augusto Marcelli. The historical development of x-ray absorption fine spectroscopy and of its applications to materials science. 27, 12 2013.
- [8] C. Bonte T. Caillaud D. Descamps F. Dorchies M. Harmand S. Hulin S. Petit C. Fourment, N. Arazam and J. J. Santos. Broadband, high dynamics and high resolution charge coupled device-based spectrometer in dynamic mode for multi-keV repetitive x-ray sources. *Review of Scientific Instruments*, 80(8), Aug 2009.
- [9] S Fourmaux, S Corde, K Ta Phuoc, P M Leguay, S Payeur, P Lassonde, S Gnedyuk, G Lebrun, C Fourment, V Malka, S Sebban, A Rousse, and J C Kieffer. Demonstration of the synchrotron-type spectrum of laser-produced betatron radiation. *New Journal of Physics*, 13(3):033017, mar 2011.
- [10] O. Svelto. *Principles of Lasers*. Springer US, 5th edition, 2010.
- [11] M Perry. Crossing the petawatt threshold. *Science and Technology Review*, 12 1996.
- [12] Christoph Heyl, Cord Arnold, Arnaud Couairon, and Anne L’Huillier. Introduction to macroscopic power scaling principles for high-order harmonic generation. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 50:013001, 01 2017.
- [13] F. Chen. *Plasma Physics and Controlled Fusion*. Plenum Press, 2nd edition, 1984.
- [14] P. Gibbon. Introduction to plasma physics. *Proceedings of the 2014 CAS-CERN Accelerator School: Plasma Wake Acceleration*, 1, 02 2016.
- [15] A. Kramida, Yu. Ralchenko, J. Reader, and and NIST ASD Team. NIST Atomic Spectra Database (ver. 5.9), [Online]. Available: <https://physics.nist.gov/asd> [2022, April 9]. National Institute of Standards and Technology, Gaithersburg, MD., 2021.
- [16] M. J. Weber. *Handbook of Laser Wavelengths*. CRC Press, 1st edition, 1998.
- [17] J. Faure et al. A laser–plasma accelerator producing monoenergetic electron beams. *Nature*, 431, Jul 2004.

- [18] Faure J. Plasma injection schemes for laser–plasma accelerators. *Proceedings of the CAS-CERN Accelerator School: Plasma Wake Acceleration*, 1, Nov 2014.
- [19] M.A.C. Potenza B. Paroli. Radiation emission processes and properties: synchrotron, undulator and betatron radiation. *Advances in physics: X*, 2:978–1004.
- [20] W.P. Leemans, E. Esarey, J. van Tilborg, P.A. Michel, C.B. Schroeder, C. Toth, C.G.R. Geddes, and B.A. Shadwick. Radiation from laser accelerated electron bunches: coherent terahertz and femtosecond x-rays. *IEEE Transactions on Plasma Science*, 33(1):8–22, 2005.
- [21] Alexander S. Ditter, Evan P. Jahrman, Liam R. Bradshaw, Xiaojing Xia, Peter J. Pauzaskie, and Gerald T. Seidler. A mail-in and user facility for X-ray absorption near-edge structure: the CEI-XANES laboratory X-ray spectrometer at the University of Washington. *Journal of Synchrotron Radiation*, 26(6):2086–2093, Nov 2019.
- [22] Jacques S. Wilson M. et al. Egan, C. Dark-field hyperspectral x-ray imaging. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 470, 05 2014.
- [23] C. K. et al. Egan. 3d chemical imaging in the laboratory by hyperspectral x-ray computed tomography. *Sci. Rep.*, 5, 2015.
- [24] A. F. Burr J. A. Bearden. Reevaluation of x-ray atomic energy levels. *Reviews of modern physics*, 29:128–129, 01 1967.
- [25] IC Insights. Research bulletin: Cmos image sensor sales stay on record-breaking pace. <https://www.icinsights.com/data/articles/documents/1065.pdf>. [Online; accessed 27-April-2022].
- [26] D. Litwiller. Ccd vs. cmos: Facts and fiction. *PHOTONICS SPECTRA*, 2001.
- [27] YSC Technologies. Introduction to charge-coupled devices (ccds). <http://www.ysctech.com/digital-microscope-CCD-camera-info.html>. [Online; accessed 27-April-2022].
- [28] Andor. Understanding direct x-ray detection in scientific cameras. <https://andor.oxinst.com/learning/view/article/direct-detection>. [Online; accessed 24-April-2022].
- [29] Teledyne Princeton Instruments. Silicon-based ccds: The basics - educational notes. <https://www.princetoninstruments.com/learn/camera-fundamentals/ccd-the-basics>. [Online; accessed 29-April-2022].
- [30] Andor. ikon-m/1 so series direct detection imaging. <https://andor.oxinst.com/assets/uploads/products/andor/documents/andor-ikon-1-and-m-so-specifications.pdf>. [Online; accessed 25-April-2022].
- [31] E.A. Guggenheim. *Thermodynamics. An Advanced Treatment for Chemists and Physicists*. Elsevier Science Publishers, 5th revised edition, 1967.
- [32] Videocineimport. Technical data - alexa lf. <https://www.videocineimport.com/wp-content/uploads/2020/02/Technical-Data-ALEXA-LF.pdf>. [Online; accessed 07-May-2022].
- [33] Arri. Alexa lf credits. <https://www.arri.com/en/camera-systems/cameras/alexa-lf/alexa-lf-credits>. [Online; accessed 07-May-2022].
- [34] Max Roser Hannah Ritchie and Pablo Rosado. Energy. *Our World in Data*, 2020. <https://ourworldindata.org/energy>.

-
- [35] Division of Atomic Physics. Lund high-power laser facility. <https://atomic.physics.lu.se/research/laser-facility/>. [Online; accessed 25-April-2022].
- [36] Andor. ikon-l so. <https://andor.oxinst.com/products/high-energy-detection/ikon-l-so>. [Online; accessed 25-April-2022].
- [37] XCRO. X-ray interactions with matter. https://henke.lbl.gov/optical_constants/. [Online; accessed 04-May-2022].
- [38] H. Sinn, E.E. Alp, A. Alatas, J. Barraza, G. Bortel, E. Burkel, D. Shu, W. Sturhahn, J.P. Sutter, T.S. Toellner, and J. Zhao. An inelastic x-ray spectrometer with 2.2mev energy resolution. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 467-468:1545–1548, 2001. Proceedings of the 7th Int. Conf. on Synchrotron Radiation Instrumentation.