

Department of Physics and Astronomy

University of Heidelberg

Master thesis

in Physics

submitted by

Pim van den Heuvel

born in Eindhoven

2021

**Arbitrary phase control
of X-ray pulses
with zeptosecond phase stability**

This Master thesis has been carried out by Pim van den Heuvel

at the

Max Planck Institute for Nuclear Physics

under the supervision of

apl. Prof. Dr. Jörg Evers

Beliebige Phasenkontrolle von Röntgenpulsen mit Zeptosekunden Phasenstabilität:

Bis vor Kurzem waren Fortschritte im Feld der kohärenten Quanten-Kontrolle auf optische und andere Bereiche mit niedrigen Frequenzen limitiert, da kohärente Lichtquellen bei höheren Photonen-Energien nicht existierten. In neuen Ergebnissen wurde zum ersten Mal kohärente Kontrolle im harten Röntgenbereich demonstriert, unter Verwendung einer Folie von Mössbauer Kernen zur Umwandlung eines Synchrotron-Pulses in einen Doppelpuls, dessen relative Phase durch mechanische Bewegung einstellbar ist. Dieses Kontrollschema weist die nötige Phasenstabilität auf dem Niveau weniger Zeptosekunden auf. In dieser These wird erstens die Datenanalyse-Prozedur, die dem Ergebnis der Phasenstabilität zugrunde liegt, untersucht, um sicherzustellen, dass diese kein limitierender Faktor bei der Erreichung der maximalen Phasenstabilität ist. Zweitens wird ein Grundsatzbeweis der beliebigen Kontrolle der relativen Phase experimentell demonstriert. Diese Untersuchung der beliebigen Kontrolle der Röntgenphase ermöglicht die allgemeine Manipulation der Quantendynamik von Mössbauer-Kernen und öffnet die Tür zum allgemeinen Formen von Röntgenpulsen.

Arbitrary phase control of X-ray pulses with zeptosecond phase stability:

Until recently, advances in coherent quantum control have been restricted to the optical and other low-frequency domains, because of the lack of coherent light sources at higher photon energies. Now, coherent control has been demonstrated in the hard X-ray region for the first time, using a target of Mössbauer nuclei to convert a synchrotron pulse into a double pulse whose relative phase is tunable by mechanical motion of the target. This control scheme has the necessary few-zeptosecond phase stability. In this thesis, first, the data analysis procedure underlying this phase stability result is investigated, to ensure it is not a limiting factor in reaching the maximum possible phase stability. Second, a proof of principle of arbitrary control of the relative double pulse phase is demonstrated experimentally. This demonstration of arbitrary X-ray phase control enables more general manipulation of nuclear quantum dynamics and opens the door to the arbitrary shaping of X-ray pulses.

Contents

1	Introduction	6
2	Theory	9
2.1	Mössbauer effect	10
2.2	Nuclear exciton	11
2.3	Nuclear Forward Scattering	13
2.4	Phase modulation	19
3	Methods	23
3.1	Experimental setup	23
3.2	Data analysis	24
3.2.1	Data preparation	25
3.2.2	Material characterisation	25
3.2.3	Energy calibration	28
3.2.4	Phase reconstruction	29
3.2.5	Phase stability	31
3.3	Arbitrary phase control	33
4	Results and discussion	35
4.1	Material characterisation	35
4.1.1	Alpha iron foil	35
4.1.2	Stainless steel foil	40
4.1.3	Summary	41
4.2	Energy calibration	42
4.2.1	Novel method	42
4.2.2	Material models	44
4.3	Phase reconstruction	45
4.4	Phase stability	47
4.4.1	Linear model	47
4.4.2	Scaling and step model	49
4.4.3	Without motion	50
4.4.4	Temporal deviations	51
4.4.5	Individual and combined phases	52
4.4.6	Discussion	53
4.5	Arbitrary phase control	54
4.5.1	Proof of principle	54
4.5.2	Phase stability	56

4.6 Summary and outlook	58
A Additional tables and figures	61
B Bibliography	66

1 Introduction

Quantum control is the active manipulation of quantum dynamics and, over the past three decades or so, this research field has delivered new ways of exploring physical processes on the atomic scale, as well as many practical applications [1]. The majority of these results have been achieved in the optical domain.

In the X-ray domain, however, coherent quantum control is experimentally difficult to achieve. Coherent light sources (lasers), on which most of the optical control techniques are based, are not available for X-rays. Extending such control techniques, and quantum optical techniques in general, into the X-ray domain would thus require creative new approaches [2].

Venturing into the X-ray domain is interesting for several reasons. For instance, X-rays can provide higher resolutions because of their smaller wavelength and their detection is highly efficient [3]. Usage of X-ray photons in the field of quantum information also holds potential advantages over optical photons. Namely, their smaller wavelength would allow smaller photonic circuits, they have a higher penetration depth and they could increase the capacity of information channels [4]. Finally, some of the novel X-ray quantum optics approaches, originally constructed to cope with the limitations in this frequency range, might also lead to a better understanding of, or to new methods for, traditional quantum optics [2].

Therefore, and because of exciting technological advances such as X-ray free electron lasers, the relatively new field of X-ray quantum optics has attracted significant attention [3]. In particular, hard X-rays allow the excitement of nuclear resonances in Mössbauer nuclei. These recoil-less transitions feature extremely narrow transitions and thus long lifetimes. Mössbauer nuclei served as the platform for many interesting quantum optics experiments, such as the measurement of a collective Lamb shift [5], electromagnetically induced transparency [6], group velocity control of X-ray pulses [7], photon echoes [8, 9, 10, 11], photon storage [12], coherent photon waveform control [4], nuclear quantum memory [13], sub-Ångström displacement measurements [14] and spectral enhancement of resonant photons [15].

The basic experimental technique behind many of these results [4, 8, 9, 10, 11, 13, 14, 15] is mechanical motion of either the emitter or the absorber of the Mössbauer radiation. This is also the main experimental method in this thesis, where it will be employed to control the relative phase of an X-ray double pulse. A short synchrotron pulse is converted into these two pulses by a target with resonant Mössbauer nuclei, with the motion of the target imprinting a phase shift on the second, delayed pulse. The mechanical motion is achieved through a piezoelectric foil, which is attached to the Mössbauer target. Coherent control of nuclear dynamics with this experimental setup, using a few specific motion-induced phase shifts, has been shown by Heeg

et al. [16].

Any coherent control scheme requires a certain stability of the light to successfully be able to manipulate quantum dynamical processes. In the extreme ultraviolet (EUV) region, stable coherent control schemes have been shown through high-order harmonic generation [17], with synchrotron radiation [18] and with a free-electron laser [19, 20]. The temporal stabilities for these experiments ranged from the few-attosecond [18, 19, 20] to even the 100 zeptosecond level [17]. For hard X-rays, specifically for the common 14.4 keV transition in iron-57, which is used here, even better timing or phase stability is required, on the order of a few zeptosecond. For our double pulse experimental scheme, phase stability on this level has recently been demonstrated by Heeg et al. [16].

In this thesis, firstly, the stability of our phase manipulation scheme is further investigated. On the one hand, the phase stability is simply monitored for a longer period of time. On the other hand, it is attempted to improve the accuracy of several steps in the data analysis chain shown in fig. 1.1, which ultimately yields the Allan deviation measure that quantifies the phase stability. This data analysis procedure was previously employed by Heeg et al. [16] to demonstrate the few-zeptosecond stability. If confidence about the accuracy of the data analysis is reached, it can be excluded as a potential source of systematic error in calculating the phase stability; only experimental systematic errors should remain. In other words, the phase stability we find should then be the fundamental limit of what is possible within our experimental setup. In chapter 3 the data analysis chain is explained in detail. It should already be clear from fig. 1.1, though, that modification of any step in the chain could have an impact on the eventual phase stability. In particular, in this work new models are constructed for characterising the targets with Mössbauer nuclei from which synchrotron light is scattered (step 2 in fig. 1.1). Additionally, a new method is introduced to calibrate data from experiments with two scattering targets against the corresponding theoretical calculation (step 3 in fig. 1.1). Such a comparison between theory and experiment is performed in steps 4 and 7 too, and the new material models from step 2 are used in steps 4 and 6 too (see fig. 1.1). So, the phase stability could indeed be influenced by these changes to the data analysis chain.

Secondly, arbitrary control of the phase imprinted by the target motion is studied. Phase control over the second, delayed pulse of the double pulse structure would allow the steering of nuclear quantum dynamics. As mentioned, such coherent control has already been realized for particular phase shapes [16]. Here, an approach is proposed to extend the phase control to arbitrary phase trajectories, which would enable more general manipulation of nuclear quantum states. This could also be an important step towards the arbitrary shaping of X-ray pulses. Our approach consists of deriving a relationship between the voltage pattern supplied to the piezo foil and the resulting phase shift. The latter can be reconstructed with our specific experimental setup and data analysis procedure. The inverse of this relationship can then be applied to any arbitrary phase trajectory, to construct the voltage pattern that should produce this phase through the piezo motion.

The remainder of this thesis is outlined as follows. Chapter 2 introduces the theory that is necessary to understand the experimental data. Chapter 3 describes our experimental setup and explains the data analysis chain shown in fig. 1.1 in detail. Chapter 4 discusses our results and finishes with a summary and outlook.

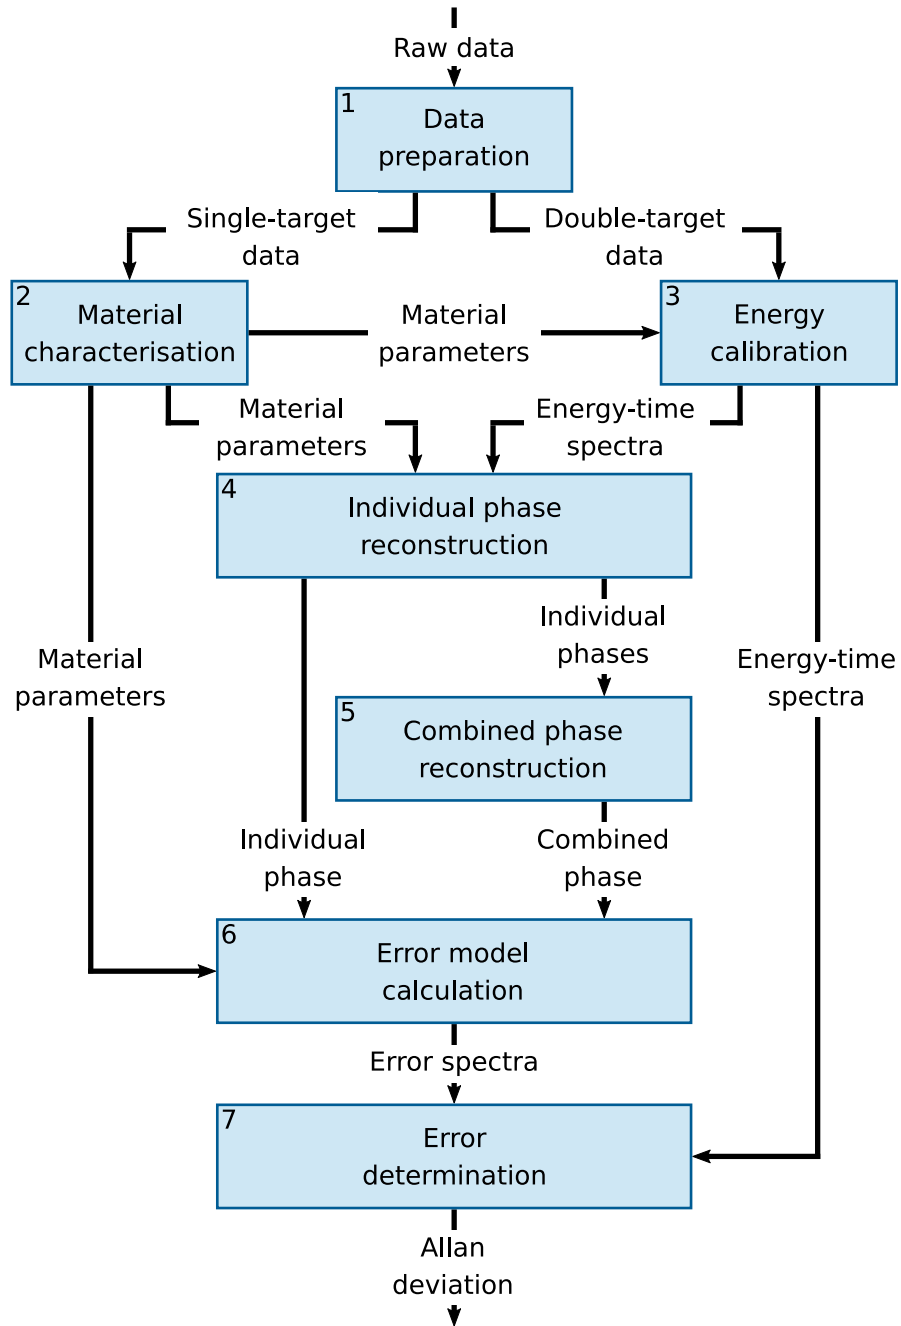


Figure 1.1: Schematic overview of the data analysis procedure to calculate the stability of our phase manipulation scheme. Explained in detail in section 3.2.

2 Theory

X-ray scattering and imaging techniques have been successfully used throughout the natural sciences for over a century. Typically, the X-rays are scattered off of electrons inside materials. In the discovery of the structure of DNA, for example, such an investigation proved to be a vital clue [21]. Instead, in this thesis, our interest lies with the interaction of X-ray light with nuclear resonances.

According to Röhlsberger [22], nuclear resonant scattering can be divided into four categories: it is either coherent or incoherent and either elastic or inelastic. In a coherent scattering process the phase correlations of the incoming light are (at least partially) conserved during the scattering process. Since this is necessary for any form of (quantum) phase control of X-ray pulses, only coherent scattering will be relevant for our discussion. Incoherent scattering, on the other hand, involves a loss of phase information and re-emission of radiation in all directions. Therefore, the experiments in this thesis are all performed in the forward scattering geometry (see fig. 2.1), where coherent scattering is the dominant scattering mechanism (see section 2.2).

Both elastic scattering, where the energy of the scattering system is not changed, and inelastic scattering, where energy is transferred to or from the system, are generally feasible scattering channels. Which of these is more likely to occur depends on the scattering system as well as experimental conditions. In scattering experiments with solid state samples containing some specific nuclear isotopes that can exhibit the Mössbauer effect, elastic scattering will usually be the dominant mechanism, as will be explained below. In this work, the scattering system will consist of one such resonant nuclear species: iron-57. The remainder of the theory section will therefore focus mostly on coherent elastic scattering.




Figure 2.1: Geometry of the Nuclear Forward Scattering (NFS) setup. The incident field E_i is transformed into the transmitted field E_{tr} by the nuclear resonant target.

2.1 Mössbauer effect

To determine the ratio between elastic and inelastic scattering, the conservation of energy and momentum of resonant photons must be considered. Any photon that is absorbed or emitted by a scattering system carries momentum proportional to its energy: $p = E/c$. Conservation of (linear) momentum requires that the scattering system gains the photon momentum in both emission and absorption (but in opposite directions), which leads to an increase in kinetic energy, $E = p^2/2m$, of the system. During emission, this energy can only arise from the resonant transition energy, so the energy of the photon will be less than the transition energy. During absorption, the photon energy has to provide both the transition energy and the kinetic recoil energy; it must be higher than the transition energy. Consequently, the absorption and emission peak are in general not located at the transition energy and do not overlap. Their separation is given by twice the kinetic recoil energy.

Now, if we assume the recoil energy to be small, and so the photon energy E_γ to be approximately equal to the transition energy E_0 , the recoil energy can be written as:

$$E_r = \frac{E_\gamma^2}{2mc^2} \approx \frac{E_0^2}{2mc^2}. \quad (2.1)$$

If some typical values for electronic transitions in atoms are plugged into this formula, a recoil energy is found that is considerably smaller than the natural linewidth of these transitions [23]. In this case the emission and absorption peak do actually overlap. Therefore, in most atomic resonance experiments, recoil energy does not have to be taken into account. Nuclear transitions, however, have much larger transition energies (14.4 keV for iron-57 [22], for instance). Their recoil energy is so large that emission and absorption peaks are indeed separated, which implies inelastic scattering; energy is transferred to the system in the form of recoil.

Nevertheless, it turns out that in practice, in nuclear resonance scattering with solid state samples, elastic scattering is usually more likely to occur. This is because the atoms in these materials are embedded in a crystal lattice structure. They have a fixed position due to their constant interaction with surrounding atoms. Any recoil can then be shared by the entire lattice, instead of it being absorbed by an individual nucleus [23]. Since the mass of the crystal will always be very large compared to the mass of one nucleus, the recoil energy, according to eq. (2.1), will be comparatively small. The emission and absorption peak overlap and elastic scattering indeed becomes the main scattering channel. This phenomenon is called the Mössbauer effect, after Rudolf Mössbauer, who first described it in his 1958 paper [24]. Background on his discovery, as well as a more detailed description, can also be found in his Noble lecture [25].

A more complete quantum-mechanical description of the Mössbauer effect must include phonons: the quantised vibrations of the crystal lattice. The recoil reaction can only excite integer numbers of phonons, because of their quantised nature. As the recoil energy grows larger, more phonons are likely to be created. Absorption and

emission spectra thus contain several peaks, one for each number of phonons that is created. The Mössbauer effect takes place when no phonons are exchanged with the scattering crystal, i.e., when the scattering is completely elastic. As no thermal motion can influence these zero phonon events, the zero phonon or recoil-free peak associated with the Mössbauer effect is narrow (4.66 neV for iron-57 [22]) and the corresponding lifetime of the excited state is relatively long (141 ns for iron-57 [22]). The ratio between elastic and inelastic scattering, finally, is determined by the Lamb-Mössbauer factor f_{LM} , which is the probability of a zero phonon event happening. For iron-57, at room temperature, the Lamb-Mössbauer factor is approximately 0.75.

So, the Mössbauer effect enables experimentally relatively easy, room-temperature, experiments of elastic nuclear resonant scattering. Mössbauer experiments can be implemented both with radioactive sources of Mössbauer nuclei and with synchrotron radiation (SR). The experimental results presented in this thesis were all obtained at the synchrotron facilities of ESRF (Grenoble) and DESY (Hamburg). Working with synchrotron radiation leads to some interesting particularities, as will be discussed in the next section.

2.2 Nuclear exciton

Since the zero-phonon peak is so narrow, synchrotron radiation is virtually 'white' compared to it, i.e., all frequency components are equally present [26]. In the time domain the SR pulses can then be seen as delta pulses; their duration is instantaneous compared to the lifetime of nuclear excitations. The absorption and emission processes can thus be considered independently [22]. Additionally, even though SR has a high brilliance, each pulse contains on average less than one resonant photon (due to the narrow resonant peak again). Thus, it suffices to consider only the single-photon or linear regime of the light-matter interaction [22].

If a single photon is incident on the scattering system, there is only enough energy to excite one nucleus. Yet, phenomena that require collective behaviour of several or many nuclei have been observed with SR. Nuclear Bragg scattering [27], for instance, is only possible if the one photon can somehow (self-)interfere, in a way reminiscent of the double-slit experiment. The explanation of such observations must be that many nuclei have a certain (equal and small) probability to be excited. The photon quantum is distributed over excited states of individual nuclei in a large superposition state, like such, for N nuclei that have been excited by a plane wave with wave vector k_0 :

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N e^{i\mathbf{k}_0 \cdot \mathbf{r}_i} |g\rangle |e_i\rangle, \quad (2.2)$$

where $|g\rangle |e_i\rangle$ is the state with the i th nucleus at position \mathbf{r}_i excited and all others in the ground state. The inclusion of phase factors indicates the coherence of the excited state, which is called a nuclear exciton in literature.

In a classical analysis, this collective excited state would instead consist of a simultaneous (small) excitation of all N nuclei; a dipole moment equal in energy to the incoming photon energy would be distributed equally over the nuclei. Eberly [28] shows explicitly that the conclusions drawn here will be the same for such a classical state as for the quantum state in eq. (2.2). A more detailed account of the nuclear exciton will therefore not be given here, but can be found in [29] and [30].

Following the explanation in Smirnov [26], the (classical) radiation emitted from N collectively excited nuclei contains contributions from each nucleus. These are added together to form, again for nuclei excited by a plane wave k_0 , a scattered field in the direction k_1 :

$$E \propto \sum_{i=1}^N e^{i\mathbf{s}\cdot\mathbf{r}_i}, \quad (2.3)$$

where $\mathbf{s} = \mathbf{k}_1 - \mathbf{k}_0$ is the scattering vector. The ‘wavelets’ from all nuclei interfere and the resulting intensity is

$$|E|^2 \propto N + \sum_i^N e^{-i\mathbf{s}\cdot\mathbf{r}_i} \cdot \sum_{i' \neq i}^N e^{i\mathbf{s}\cdot\mathbf{r}_{i'}}. \quad (2.4)$$

It can immediately be seen that for disordered samples, in which nuclei are randomly positioned, the double sum vanishes because the phases are uniformly distributed, except when $\mathbf{s} = \mathbf{0}$. In this case constructive interference occurs and this forward scattering direction is thus indeed where coherent elastic scattering can be observed, as mentioned at the beginning of this chapter. If order is present in a scattering sample, there might be certain other scattering vectors \mathbf{s} that allow constructive interference, which explains Nuclear Bragg scattering. Because the quantum state (eq. (2.2)) does not have a dipole moment, the preceding analysis should not immediately apply there. However, Eberly [28] argues that in the quantum case the excitation of nuclei is sequenced (according to the phase factors) too, and so the emission, even though it is a random process now, will still be sequenced, or phased, as well. One then arrives at the same conclusions for the directionality of emission as above.

Besides having a specific directionality, the scattered field also has an increased intensity, proportional to N^2 in the forward direction, as can be seen from eq. (2.4). This N -fold enhancement of the intensity, as compared to the intensity without any interference effects, is called super-radiance and was first described by Dicke [31]. The intensity increase can be deduced heuristically from the quantum state too [32]. There are N states, and thus N pathways for emission, in eq. (2.2), all with normalization factor $(\sqrt{N})^{-1}$. So, in total, the emission amplitude has a factor \sqrt{N} and the intensity then contains a factor of N , demonstrating the same N -fold intensity enhancement, as compared to the single incoming photon.

The nuclear exciton picture has a final, important consequence: the distributed nature of the excitation enables a macroscopic description of the scattering process [26]. This macroscopic description will be utilized in the following section.

2.3 Nuclear Forward Scattering

As motivated in the previous section, we are interested in the forward scattering geometry, which is also referred to as Nuclear Forward Scattering (NFS). In this section, we will find expressions for the field transmitted through a scattering material in the NFS setup (see fig. 2.1), using the macroscopic description. It will be assumed that the material contains only a single-line resonance, is homogeneous, isotropic and is not optically active. The latter two, respectively, mean that the polarization of the incoming field can be arbitrarily chosen and that it is not altered during propagation through the material. A more general explanation, including polarization effects, can be found in Röhlsberger [22].

In the macroscopic description all nuclei (that are in the beam path) contribute with their dipole moments, which are induced by the incoming electric field, to the overall (macroscopic) polarization density. As mentioned in the previous section, limiting ourselves to linear optics is adequate for our purposes. Linearity is often expressed by a linear relationship of the polarization density \mathbf{P} to the electric field \mathbf{E} inside the material:

$$\mathbf{P} = \epsilon_0 \eta \mathbf{E}, \quad (2.5)$$

where ϵ_0 is the electric permittivity and η is the susceptibility, which signifies the polarization response of the material. It is sensible that the polarization, via the susceptibility, is proportional to the number density of scattering nuclei N , as well as to their individual response functions, which are given by the scattering amplitude f . The susceptibility is a scalar due to the assumptions made about the material and is given by [26]:

$$\eta = \frac{4\pi}{K^2} N f, \quad (2.6)$$

with K the magnitude of the wave vector of the incoming field and the scattering amplitude

$$f(\omega) = -\frac{K}{4\pi} \sigma_0 f_{\text{LM}} \beta \frac{\Gamma/2\hbar}{\omega - \omega_0 - i\Gamma/2\hbar}, \quad (2.7)$$

where σ_0 is the cross-section, Γ is the natural width of the nuclear transition ($\gamma = \Gamma/\hbar$ denotes the width in frequency units), ω_0 its resonance frequency and β the fraction of resonant isotopes in the target. The Lamb-Mössbauer factor f_{LM} has been discussed in section 2.1. This nuclear response function f has a Lorentzian shape, as one might expect from a single-line resonance.

Now, if a plane wave with wave vector \mathbf{K} is incident on the material, it can be related to the wave vector inside the material, \mathbf{k} (which is in the same direction as \mathbf{K} in the NFS geometry), via the Maxwell equations in matter and the previous linearity assumption (see [26] for the complete derivation):

$$\mathbf{k} = \sqrt{1 + \eta} \mathbf{K} = n \mathbf{K}, \quad (2.8)$$

where n is the refractive index, which now solely characterises the response of the scattering material. Ignoring second order terms, and subsequently plugging in eqs. (2.6) and (2.7), the refractive index can be approximated as:

$$n \approx 1 + \frac{1}{2}\eta = 1 - \frac{1}{2K}\sigma_0 f_{\text{LM}}\beta N \frac{\Gamma/2\hbar}{\omega - \omega_0 - i\Gamma/2\hbar}. \quad (2.9)$$

Note that any electronic absorption is neglected in this equation. The complete refractive index can be found, e.g., in [2].

Finally, the transmitted field can be calculated. The incoming field is assumed to be a plane wave along the z -axis, for reasons of simplicity:

$$E_i = E_0 e^{i(\omega t - Kz)}. \quad (2.10)$$

The wave inside the material is then, with eqs. (2.8) and (2.9):

$$E_{sc} = E_0 e^{i(\omega t - kz)} = E_0 e^{i\omega t - i(1+\eta/2)Kz} = E_i e^{-i(\eta/2)Kz}. \quad (2.11)$$

The final term in this equation suggests that the exponential factor can also be seen as a transmission function, dependent on the thickness z of the material, that acts on the incoming wave:

$$E_{tr}(z) = E_i e^{-i(\eta/2)Kz} = T(z)E_i, \quad (2.12)$$

where the transmission function has been expressed as $T(z)$. In fig. 2.3a this transmission function is shown. Equation (2.12) is also valid for incoming wave packets, which consist of many plane waves. Indeed, for ‘white’ synchrotron radiation, all plane waves are equally present and the transmitted wave becomes equal to the transmission function. A quantum mechanical derivation [33] (with classical fields) finds the same transmission function in the linear regime (quantum mechanically this means that the nuclei must be very far from population inversion), if a similar assumption is made about the form of the susceptibility (eq. (2.6)).

Since the transmission function is an exponential, eq. (2.12) can always be broken down as follows:

$$E_{tr}(z) = E_i \sum_{n=0}^{\infty} \frac{(-i\eta Kz)^n}{2^n n!} = E_i \left[1 + \sum_{n=1}^{\infty} \frac{(-i\eta Kz)^n}{2^n n!} \right]. \quad (2.13)$$

Writing the transmitted wave this way implies that it consists of two separated, independent waves: the unaffected incoming one and the interacting scattered wave, like such:

$$E_{tr} = E_i + E_{fs}, \quad (2.14)$$

where the subscript fs stands for forward scattered. In a first approximation (valid for small z), only the first term of the (second) sum in eq. (2.13) is kept:

$$E_{tr}(z) \approx E_i (1 - i\eta Kz/2). \quad (2.15)$$




Figure 2.2: Transmission of a δ -peak through a single-line $0.5 \mu\text{m}$ thick nuclear resonant target, in the (a) frequency and (b) time domain. Electronic absorption is ignored.

At resonance, the susceptibility is purely imaginary (see eq. (2.7)) and the forward scattered wave destructively interferes with the incident wave (see fig. 2.2a). This explains the dip in transmitted intensity around resonance. The actuality of this picture with two independent waves can be proven if one can change the interference between them, as will be discussed in section 2.4. In fact, each next term in the sum of eq. (2.13) can be interpreted as an additional scattering event [34, 35], since it contains n times the single nucleus scattering amplitude. Equation (2.15), then, describes just a single scattering event, which corresponds to the small thickness z of this approximation.

So far, all formulas have (implicitly) been in the frequency domain. Corresponding time domain solutions can be found by performing Fourier transforms. For eq. (2.12), the Fourier transform is, according to the convolution theorem:

$$E_{tr}(t, z) = R(t, z) * E_i(t), \quad (2.16)$$

where $E_i(t)$ is the incident wave in the time domain (the Fourier transform of $E_i(\omega)$, which is arbitrary now), and $R(t, z)$ is the Fourier transform of the transmission function $T(z)$, called the response function. For SR, $E_i(t) = \delta(t)$ and the transmitted wave is actually equal to the response function, and so to the Fourier transform of $T(z)$. Equivalently, because SR carries all frequency components equally, the time representation of the transmitted wave can be found by simply integrating over all of these components (see eq. (2.11)):

$$E_{tr}(t, z) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} d\omega e^{i\omega t - i(1+\eta/2)Kz}. \quad (2.17)$$

Solving this integral involves contour integration (details are found in [33]), and

yields [26]:

$$E_{tr}(t, z) = \delta(t) - \frac{T}{2t_0} e^{i\omega_0 t - \tau/2} \frac{J_1(\sqrt{T\tau})}{\sqrt{T\tau}}, \quad t \geq 0, \quad (2.18)$$

where $t_0 = \Gamma/\hbar$, J_1 is the Bessel function of the first kind, and effective thickness $T = \sigma_0 f_{LM} \beta N z$ and reduced time $\tau = t/t_0$ are introduced as coordinates. The transmitted wave is a delta peak followed by an exponentially decaying resonant wave modified by the Bessel function J_1 (see fig. 2.3b). Indeed, the separation of the incident ($\delta(t)$) and forward scattered wave (eq. (2.13)) is immediately apparent here. Experiments with SR, therefore, are well-suited for time domain observation.

The final fraction in eq. (2.18),

$$\sigma(z, t) = \frac{J_1(\sqrt{T\tau})}{\sqrt{T\tau}}, \quad t \geq 0, \quad (2.19)$$

can be approximated as an exponential function for small values of the argument $T\tau$ of the Bessel function J_1 , that is, for initial times and thin materials [26]. Equation (2.18) then becomes:

$$E_{tr}(t, z) \approx \delta(t) - \frac{T}{4t_0} e^{i\omega_0 t - \tau/2 - T\tau/8}, \quad t \geq 0, \quad (2.20)$$

and it is plotted in fig. 2.2b. The extra exponential factor accelerates the initial decay of the delayed forward scattered radiation, which indicates super-radiant effects. Furthermore, the secondary radiation is proportional to T , so its intensity is proportional to T^2 . Since T^2 contains the factor $(nz)^2$, the intensity scales with the squared number of nuclei (per area). This is another manifestation of superradiance (see section 2.2).

For larger times, thicker materials, or both, the behaviour of the Bessel function J_1 must be taken into account. It has zeroes that are progressively spaced further apart (see fig. 2.3b), which are known as dynamical beats (DB) in the context of nuclear resonant scattering. Both longer times and thicker targets can in theory reveal more nodes in the dynamical beat pattern. In practice, however, the observation time is limited by the repetition rate of SR and only the thickness of targets can be varied.

The dynamical beats can be explained from different viewpoints. Van Bürck [36] suggests that the physical origin of the dynamical beats is best illustrated in the picture of a coupled system of a light field and (delocalized) nuclear excitation. In this view, energy is transferred back and forth between the nuclear system and the radiation, as it propagates through the material [26]. When the radiation reaches its maximum amplitude, the nuclei are least excited, and vice versa. Each additional antinode of the forward scattered radiation then represents an additional scattering event, since another cycle of transferring energy to and from the nuclear excitation has taken place. Each resonant scattering event, in turn, narrows the energy distribution of the propagating field (the Lorentz profile of eq. (2.7) is applied




Figure 2.3: Transmission of a δ -peak through a single-line $5 \mu m$ thick nuclear resonant target, in the (a) frequency and (b) time domain. Electronic absorption is ignored.

successively), causing its decay to slow, which explains the increasing distance between the dynamical beats. Also, if the target is thicker, or the observation time extended, more scattering events can occur and more dynamical beats are detected. So, as was the case in the frequency domain, the inclusion of multiple scattering events in this coupled-system picture is crucial for a complete description of the scattering process, and to explain all resulting features observed in time-domain NFS experiments (see fig. 2.3b).

Until this point, a single resonance line has been assumed. In many practical situations, two or several resonant peaks are present. This can, for example, be due to hyperfine splittings of the nuclear energy levels. In the case of a single splitting into two energy levels, two susceptibility functions will contribute to the transmitted field (compare with eq. (2.12)):

$$E_{tr}(z) = E_i e^{-i(\eta_1/2)Kz - i(\eta_2/2)Kz}. \quad (2.21)$$

Both susceptibility functions η_1 and η_2 look like eq. (2.9), but with different resonant frequencies for each. The more complicated exponential of eq. (2.21) prevents an analytical expression for the transmitted wave in the time domain. However, qualitative reasoning can be applied: the radiation originating from both resonant peaks will interfere in a beating pattern, whose frequency is equal to the difference in frequency between the peaks (see fig. 2.4a). The beating pattern is referred to as quantum beats (QB).

Level splitting is often more involved than the simple example above. In the case of magnetic hyperfine splitting, for example, several nuclear transitions can often be distinguished. This type of splitting can occur in materials with permanent magnetic fields at the locations of the nuclei. The magnetic field can have several origins. A common one is the internal magnetization in ferromagnetic materials,




Figure 2.4: Theoretical transmission of a δ -peak through a (a) $1\ \mu\text{m}$ and (b) $5\ \mu\text{m}$ thick α -iron foil. The magnetic hyperfine field is set to 33 T for both plots.

such as alpha iron (iron at room temperature), where the magnetic moments of individual atoms spontaneously align. The interaction of the nuclear magnetic dipole with this magnetic field causes the splitting of the nuclear energy levels, which is also referred to as the nuclear Zeeman effect. The magnitude of the splitting is determined by the precise value of the interaction energy. This depends on, among other factors, the total spin J of the nucleus, and specifically, on its z-component m_J (see, e.g., [22]). In alpha iron, as an example, the ground state has $J = 1/2$, while the excited state has $J = 3/2$. Respectively, two and four different values for m_J are allowed. Therefore, the ground state is split into two sublevels, while the excited state is split into four. Selection rules dictate that only six of the eight possible transitions between these levels are allowed. Ignoring any optical activity of the target material, the transmitted field is given by extending eq. (2.21) to include six nuclear susceptibilities, instead of just two. The corresponding beating pattern can be quite complicated in this case. However, when the internal hyperfine magnetic field and the magnetic polarization vector of a field incident on the target are aligned parallel to each other, four of the six transitions can be excluded [26], i.e., those with $\Delta m_J \neq 0$. With an external (weak) magnetic field the direction of the hyperfine field can be controlled, thereby allowing one to select the number of possible nuclear transitions (i.e., either six or two).

Another common hyperfine effect is the quadrupole splitting. It results from the interaction of the nuclear electric quadrupole moment with an electric field gradient (EFG), which is caused by an asymmetric (local) charge distribution of the electrons in the environment around the nucleus [22]. Only when the nuclear charge is not distributed spherically symmetric, an electric quadrupole moment can exist and such an interaction can take place. Nuclei with $J = 1/2$ have a spherical charge distribution and hence no quadrupole splitting. Nuclei with $J = 3/2$, such as the

excited state in iron-57, have a quadrupole splitting, if a strong enough electric field gradient is present. For $J = 3/2$, the excited state is split into just two peaks.

The final hyperfine interaction usually encountered in nuclear resonant scattering experiments is the (chemical) isomer shift. Although the nucleus is small compared to the electron cloud, there is a non-negligible overlap (for s orbitals) of the electronic and nuclear charge distribution. The associated interaction energy leads to a shift (not a splitting) of the nuclear energy levels [22]. This isomer shift is always present, since there is always some overlap between the electron cloud and the nuclear electric monopole. The shift can only be measured relative to some reference material, which is usually taken to be alpha iron.

So, either a magnetic hyperfine splitting or a quadrupole splitting can generate quantum beats in nuclear resonant scattering with Mössbauer nuclei. For thick enough targets, the quantum beats are simply superimposed on the dynamical beats (see fig. 2.4b).

2.4 Phase modulation

Mechanical motion of the resonant target was already highlighted in chapter 1 as a possible way to achieve phase control in the nuclear regime. The motion imprints a certain phase shift on the field transmitted through the target. In the following section, the theory behind the phase modulation and some simple examples will be discussed.

In the general sense the effect of a mechanical displacement on the radiation emitted by a (single) target is best understood by considering (non-relativistic) transformations between the laboratory and target frame. Any incident field E_i can be transformed to the target rest frame, for arbitrary motion $x(t)$ (as seen in the laboratory frame) [15]:

$$\tilde{E}_i(t) = E_i(t)e^{-ikx(t)}. \quad (2.22)$$

In the target frame the transmitted wave can easily be calculated, either directly in the time domain (eq. (2.16)), or in the frequency domain (eq. (2.12)). In the latter case a Fourier transform must be performed before the (frequency) transmission function $T(z)$ is applied, followed by an inverse Fourier transform. In both cases, the transmitted wave \tilde{E}_{tr} is finally transformed back into the laboratory frame:

$$E_{tr}(t) = \tilde{E}_{tr}(t)e^{ikx(t)}. \quad (2.23)$$

Performing the transmission calculation in the frequency domain is in general computationally more straightforward, because of the multiplicative transmission function. This generic procedure is summarised in fig. 2.5.

In the case of SR, however, the response function in the time domain is well-known (eq. (2.18)). Since the incoming delta peak is not affected by the transformations (if

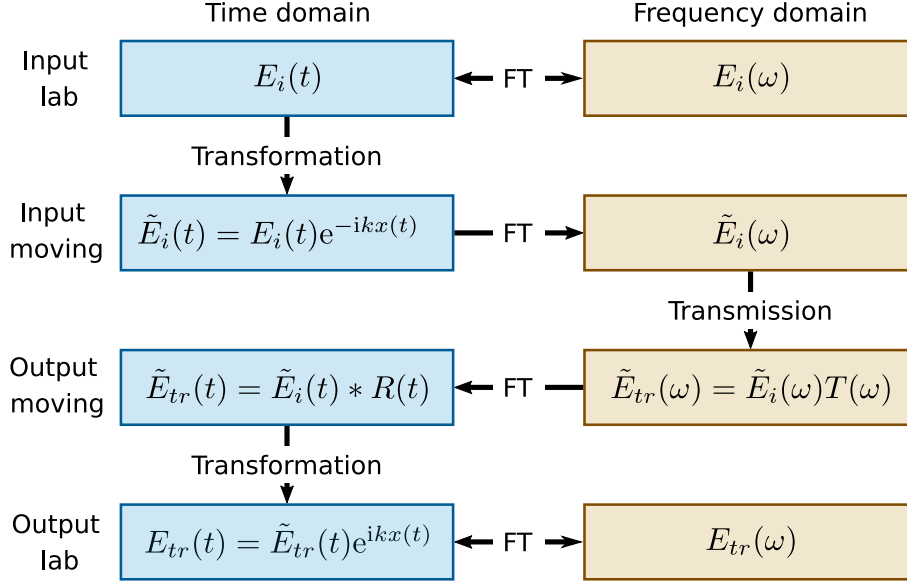


Figure 2.5: Computational strategy for calculating the field transmitted through a target that is (arbitrarily) moved. FT stands for Fourier Transform. The input and output field in the laboratory frame can be given in either the time or energy domain. The scheme is easily extended to several targets.

$x(0) = 0$ is chosen), the transmitted wave can be written directly in the time domain [15]:

$$E_{tr}(t) = \tilde{E}_{tr}(t)e^{ikx(t)} = \delta(t) - e^{ikx(t)} \frac{T}{2t_0} e^{i\omega_0 t - \tau/2} \frac{J_1(\sqrt{T\tau})}{\sqrt{T\tau}}, \quad t \geq 0. \quad (2.24)$$

The mechanical motion introduces a phase-shift $\phi(t) = kx(t)$ between the incident delta peak and the delayed, forward scattered radiation. However, this phase shift cannot be detected directly in the time domain, as only intensities are experimentally available.

One particular phase-shift of interest here is an instantaneous phase step. A phase step of ϕ_0 at $t = 0$ can be described as $\phi(t) = kx(t) = \phi_0\theta(t)$, with $\theta(t)$ the Heaviside step function. If $\phi_0 = \pi$ is chosen, the exponential in eq. (2.24) becomes -1 (except for at $t = 0$). The intensity over time still remains the same with such a phase step, though the spectrum changes dramatically. The Fourier transform of eq. (2.24) reads, namely, for SR ($E_i(\omega) = 1$):

$$E_{tr}(\omega) = 1 - \sum_{n=1}^{\infty} \frac{(-i\eta K z)^n}{2^n n!} = 2 - e^{-i(\eta/2)Kz} \quad (2.25)$$

The sign between the incident and forward scattered radiation has flipped (compare with eq. (2.13)). Indeed, in the small z approximation, eq. (2.25) becomes:

$$E_{tr}(\omega) \approx 1 + i\eta K z/2. \quad (2.26)$$

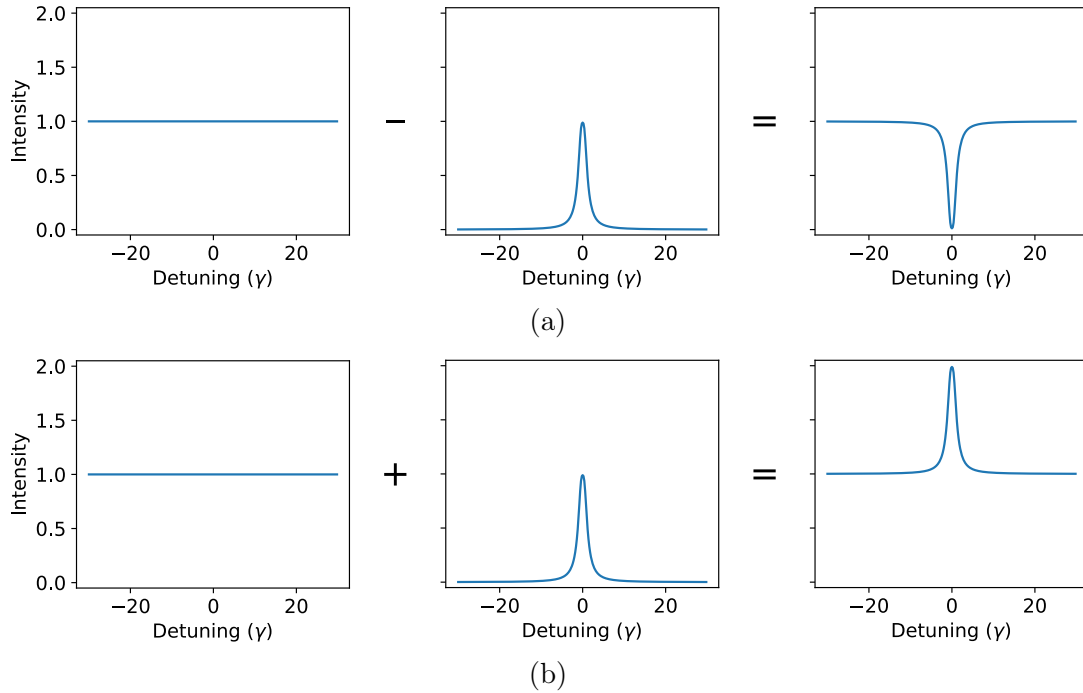


Figure 2.6: Illustration of resonant enhancement with an instantaneous phase shift of π (b), compared to the situation without phase shift (a). The phase shift causes the material spectrum (second plot) to be added to the white spectrum of the δ -peak (first plot), instead of being subtracted. This figure is just an idealised illustration of the concept; the theory actually predicts a larger resonant enhancement, which is also shaped differently.

At resonance the incoming and scattered wave are added together, they interfere constructively, as opposed to the regular situation in eq. (2.15). A peak around resonant frequencies now occurs, instead of a dip, which is illustrated in fig. 2.6. This resonant enhancement has also been verified experimentally by Heeg et al. [15]. In their paper, the resonant peak has a ‘double-hump’ profile (see also fig. 4.15b in section 4.5). Such a profile arises in thicker targets, where multiple scattering takes place, and the forward scattered wave is given by the entire sum in eq. (2.25). The double-hump profile provides an alternative picture to understand dynamical beats: the (intra-resonance) interference between the two peaks of the double hump, which are close in frequency, causes the beating pattern [26, 36].

Figure 2.6 at first sight suggests that energy conservation is violated; the increased spectral density around the resonance frequency is seemingly created out of nowhere. A closer look, however, reveals that the total area under the graph remains the same, with and without a phase shift. This ‘harvesting’ from other spectral regions onto the resonance is clearly seen in practice (see fig. 4.15b), where an instantaneous phase jump is impossible. Even in the theoretical limit of an instantaneous jump, energy is conserved, as the resonant peak is ‘harvested’ from an infinite spectral range [15, 33].

Another special case is a constantly increasing phase-shift, which corresponds

to linear motion of the target. It is a well-known fact, and an often employed experimental technique, that linear motion causes a Doppler shift of the resonance frequency of nuclei in a target. Even Mössbauer [24] applied this technique already, to scan over the (recoil-less) nuclear resonant peak. For a constant linear motion of $x(t) = vt$, with v the velocity of the target, one can deduce from eq. (2.18) that the Doppler frequency shift equals $\Delta\omega = kv = \omega_0 v/c$. In practice, an instantaneous phase-shift as discussed above is impossible and large parts of the motion pattern will typically resemble a constant linear motion. The resulting resonance shift because of the Doppler effect can be observed experimentally, which explains the ‘harvesting’ in at least part of the spectral range.

Theoretical computations of transmission through two or several targets, which may all move arbitrarily, is possible with the strategy outlined in fig. 2.5. The wave transmitted through any of the targets is then taken as the wave incident on the next downstream target. Again, it is usually more convenient to calculate the transmission through each target in frequency space. If (at least) two targets are present, the delayed part of the radiation transmitted through the first target can interfere with the delayed part of the radiation transmitted through the second target. Hence, the phase shift resulting from mechanical motion of the (first) target (see eq. (2.24)) can be revealed in the time domain. This concept is employed in our experimental setup too, as explained in the next chapter.

3 Methods

Both aims of this thesis, investigating the influence of the data analysis procedure on the fundamental stability of our phase manipulation scheme and arbitrary control of the phase, are reached through essentially the same experiment and succeeding data analysis, the outline of which was already shown in fig. 1.1. In this chapter, firstly our experimental setup and data acquisition method is described in section 3.1. Secondly, the data analysis is explained in more detail in section 3.2, including the attempts to improve its accuracy. Thirdly, section 3.3 considers our approach to achieve arbitrary phase control.

3.1 Experimental setup

It has been mentioned in chapter 2 that all experiments presented in this thesis were carried out at the synchrotron facilities of ESRF and DESY, specifically, at the Nuclear Resonance Beamline ID18 (ESRF) [37] and the PETRA III beamline (DESY) [38]. All experiments were in the NFS geometry. The simplest possible setup in the NFS geometry has been shown in fig. 2.1, where only one target is illuminated by the incident beam. Such a simple setup is sufficient to characterise the materials (step 2 in fig. 1.1) that are used in the following data analysis steps.

As is appropriate for SR, all observations are made in the time domain. The low spectral density of resonant photons in SR pulses has consequences both for the theory (see section 2.2) and in practice. It is highly unlikely that more than one photon is scattered per SR pulse; in fact, many pulses will not produce a detection event at all. Therefore, the photon detection events from many pulses must be accumulated to form a complete picture of the transmitted intensity. The sensitive (single-photon) detectors would be overloaded by the many non-resonant photons that pass through the target(s) unaffected (or are electronically scattered). In SR experiments, these photons are contained within the prompt delta pulse (eq. (2.18)) and are easily separated from the delayed resonant photons by shutting down the detectors for a short period of time (≈ 20 ns) after arrival of the pulse (which can be seen, e.g., in fig. 4.1). The observation time is limited, too, by the time between subsequent pulses (≈ 176 ns for ESRF and ≈ 192.5 ns for DESY).

As discussed at the end of section 2.4, measuring the phase-shift imprinted by our phase manipulation scheme directly in the time-domain requires at least two targets, so that an interference pattern is created. Here, the first target is chosen to be a $\approx 2 \mu\text{m}$ thick alpha iron slab, enriched in the resonant iron-57, while the second is a stainless steel ($\text{Fe}_{55}\text{Cr}_{25}\text{Ni}_{20}$) target, also enriched in iron-57, and $\approx 1 \mu\text{m}$ thick. The phase-shift of the delayed radiation is generated by mechanically moving

the α -iron target with a piezoelectric transducer that is attached to the target and driven by an arbitrary function generator. Additionally, the stainless steel target can be placed on a Doppler drive, allowing for a constant motion (along the path of the beam). Scanning over a range of velocities of the Doppler drive can be seen as scanning over the spectrum of the α -iron foil, if the second target peak is narrow compared to the first target peak(s). Each Doppler drive velocity corresponds to an energy detuning of the second target with respect to the first (see section 2.4). So, an interference pattern emerges in the transmitted field, with, for each Doppler velocity, a different detuning between the peaks of the first and second target. Combining these time-domain measurements, for all available velocities, yields a two-dimensional energy-time spectrum (see fig. 4.4a). Because of the interference between the two targets, this energy-time spectrum is then sensitive to the mechanically induced phase-shift in the first target [15, 16].

3.2 Data analysis

The analysis of the data, acquired as described in the previous section, is presented in this section and follows the data analysis chain in fig. 1.1. This section is broken down into subsections, with each subsection explaining either one or two steps from the chain. The attempts to improve the data analysis accuracy are explained in the relevant subsections: section 3.2.2 and section 3.2.3.

The first step of the data analysis chain is the preparation of both the single- and double-target data. It is described in section 3.2.1.

Step 2 of the chain, where the α -iron and stainless steel target are characterised by certain material parameters, is discussed in section 3.2.2. Previously, only two parameters described each material; here, combinations of many parameters are examined, to find the parameter set that best represents each material. These optimal material parameters are then used in theoretical calculations in the 3rd, 4th and 6th step of the data analysis chain (see the arrows originating from the second step in fig. 1.1). In this way, the optimal parameters could also influence the phase stability, which eventually comes out of the data analysis chain in the form of the Allan deviation. Section 3.2.2 also presents the evolutionary algorithm that is employed to fit the material parameters to the data. This same algorithm is used in other steps of the data analysis to perform fits against the data too.

In section 3.2.3, the experimental two-dimensional energy-time spectrum is compared against a theoretical calculation (performed with the optimal material parameters) to calibrate the experimental energy scale (step 3 of the analysis). Calibrated experimental spectra are later used in the 4th and 7th step of the analysis. In this subsection, a novel method of comparison between the theoretical and experimental energy-time spectra is introduced as well, which is also used in steps 4 and 7. This novel method itself, and the energy calibration that might be altered because of it, could thus also influence the final phase stability.

Step 4 and 5 of the analysis both concern the reconstruction of the piezo-generated

phase shift of the delayed radiation; they are both discussed in section 3.2.4. Step 4 reconstructs the phase for individual measurements, which are limited by the observation time (see section 3.1), by calculating theoretical energy-time spectra for different phases and optimizing the phase against the experimental spectrum. Step 5 chains the individual phases together into a combined and complete phase.

Finally, section 3.2.5 considers step 6 and 7 of the data analysis chain, which produce the stability of the reconstructed phase, in the form of the Allan deviation. Step 6 adds a range of error values to the reconstructed phase (which can be either an individual or combined phase) and calculates an error spectrum for each error value. In step 7 the experimental spectrum data is split into time intervals and compared against all error spectra, to find the best-fitting error for each interval. The errors, and thus the stability, can then be traced over the total measurement time.

3.2.1 Data preparation

Before any analysis can take place, the raw data is compressed and binned (step 1 in fig. 1.1). Each raw data event marks the arrival of a photon at the detector and, for the double-target measurements, contains the velocity of the Doppler drive (at the time of photon detection) as well. All raw data events are binned into time or time-and-speed containers, respectively, for single-target or double-target experiments. The resulting histograms constitute an intensity measurement for the single-target case and a two-dimensional energy-time spectrum for the double-target case.

3.2.2 Material characterisation

The second step of the data analysis procedure is the characterisation of individual materials, according to specific parameters. The theory of single-target nuclear resonant scattering is well understood. For a single-line absorber, the analytical time-domain response has been given in eq. (2.18). If some kind of level splitting is present, the result can be obtained numerically, starting from eq. (2.21). These and all other numerical calculations are performed with the *pynuss* software package, written in Python by Killian P. Heeg.

Evolutionary algorithm

The optimal material parameters, both for the α -iron and stainless steel target, are determined through a fit of these theoretical calculations against the experimental results with an evolutionary algorithm. In analogy to the biological process, parameters are often called genes in an evolutionary algorithm. Any specific combination of parameter values, i.e., any possible solution to the fitting problem, is then referred to as a chromosome or as DNA [39]. Consecutive iterations (or generations) of the algorithm randomly ‘mutate’ and mix chromosomes. The best chromosomes are

selected to proceed to the next iteration (generation), according to some function that determines their ‘fitness’.

Since the data is collected in the single-photon regime, photon shot noise is likely to be the dominant type of random error. Shot noise is Poisson distributed, therefore, our fitness function is a Bayesian likelihood function with a Poisson distribution for each data point, i.e., for each time bin. If a certain theory (a certain chromosome) is given, the probability to detect N_{exp} photons experimentally, for a specific data point, is then:

$$P(N_{\text{exp}} | N_{\text{theo}}) \propto (N_{\text{theo}})^{N_{\text{exp}}} \frac{e^{-N_{\text{theo}}}}{N_{\text{exp}}!}, \quad (3.1)$$

where N_{theo} is the theoretically expected photon count. The total probability for the entire dataset is the product of these individual probabilities:

$$P(\text{exp} | \text{theo}) = \prod_i P(N_{\text{exp},i} | N_{\text{theo},i}). \quad (3.2)$$

According to Bayes’ theorem, and assuming uniform priors, the likelihood of a certain theory, given experimental results, is proportional to this total probability: $P(\text{theo} | \text{exp}) \propto P(\text{exp} | \text{theo})$. The optimal theoretical parameters can be found by maximizing the likelihood $P(\text{theo} | \text{exp})$, and thus by maximizing eq. (3.2). In numerical calculations, the logarithm of eq. (3.2) is often taken. The individual logarithmic probabilities for each data point are then simply added together. Because the logarithmic function is monotonically increasing, the optimal parameters remain equal.

Evolutionary algorithms incorporate a degree of randomness in their mutations and combinations. For this reason, they should cover most of the parameter space and thus be appropriate when the optimization landscape is irregular, with many local maxima. One issue with evolutionary algorithms is that they do not have clear termination conditions by themselves. They could be stopped when the fitness is not increasing significantly anymore or when several runs of the algorithm (with different, random starting points) seem to converge around the same fitness value. The latter strategy is employed here. Afterwards, conventional optimization techniques (e.g., gradient descent) are applied to further improve the fitness. The assumption is then that the convergence of the evolutionary algorithm runs indicates that the peak of the global maximum is found, and that the conventional optimization finds the absolute maximum value of this peak. The evolutionary algorithm, and the subsequent conventional optimization algorithm, are also applied to fits of the theory against the experiment in all later data analysis steps.

Original parameter combinations

The chromosomes in the evolutionary algorithm previously only contained two parameters for the material characterisation, both for the alpha iron and stainless

steel foil. The alpha iron foil has a hyperfine splitting, due to an internal magnetic field. An external magnetic field is employed to align this magnetization parallel to the magnetic polarization vector of the incoming field. As discussed in section 2.3, only two nuclear transitions are allowed in this case, whose interference leads to a (single) quantum beat in the time-domain. The strength of the hyperfine field determines the magnitude of the splitting and thus the frequency of the beating pattern. In any theoretical description of the transmission the strength of the hyperfine field must therefore be included. The thickness of the material is also essential, since it is related to the occurrence and frequency of the dynamical beats. Previously, fitting these two parameters against the experimental result was deemed to provide a sufficiently accurate model of the material.

For the characterisation of the stainless steel target, the thickness is of course a necessary parameter too. As stainless steel is not ferromagnetic, there is no hyperfine magnetic field that needs to be considered. An isomer shift, though, is usually present. Intensity measurements in the time-domain are not able to distinguish such a shift of the resonant peak. It is also possible that the isomer shift has a certain distribution, instead of a single value, for example because of impurities in the crystal structure that cause local changes in the electronic distribution around certain nuclei. This distribution of isomer shifts is assumed to be Gaussian and translates to a broadening of the resonant peak, which can be directly detected in the time-domain. So, for the stainless steel target, the two essential parameters were previously determined to be the thickness of the target and the width of the distribution of isomer shifts.

Improved parameter combinations

There is no principal limit to the number of parameters that can be included in a chromosome of the evolutionary algorithm. In this thesis, the optimal values are found for many combinations of parameters, following the evolutionary algorithm procedure. The parameter set that yields the highest fitness should explain the experimental results best. This set is found for both the α -iron and stainless steel target. For α -iron the parameter combinations all include the original parameters (the thickness and the strength of the magnetic hyperfine field) and additional ones are chosen from: a quadrupole splitting or a distribution thereof, a distribution of isomer shifts, a distribution of the hyperfine magnetic field strength, a distribution of thickness values, the direction of the internal hyperfine magnetic field or a distribution of this direction. The entire power set of this list of parameters is examined, i.e. all possible parameter sets, for any number of parameters. This amounts to over 100 different parameter combinations.

The direction of the internal magnetization field is parametrized by the polar angle ϕ and the azimuthal angle θ , with the z -axis in the direction of propagation of the incident wave. They are both set to $\pi/2$ if the angles are not included as parameters, which designates the desired direction parallel to the magnetic polarization vector of the incident field. The distribution of these angles and the thickness distribution are assumed to be uniform, and are described by their total width. All other distributions

are assumed to be Gaussian; they are described by their standard deviation σ .

Another distinction between the distributions is that some are coherent and some incoherent. Specifically, the distributions of the three hyperfine effects are coherent. They are caused by local, microscopic changes in the surroundings of nuclei, changes on the level of the nuclear Hamiltonian, which could be inserted into the coherent nuclear exciton state of eq. (2.2). The *pynuss* package contains built-in functions for these coherently averaged distributions. Because the thickness of the target is a macroscopically defined quantity, a coherent average is not possible. Instead, the transmission through a target has to be calculated with several thicknesses, and the resulting intensities averaged. Such an incoherent average is also implemented for the distributions of the magnetization angles. Note that incoherent averaging is computationally more demanding, as several transmission calculations have to be performed per chromosome.

In the case of the stainless steel target, the original parameters were the thickness and a distribution of isomer shifts. All tested parameter combinations for characterising the stainless steel include at least the thickness. As mentioned before, internal magnetization need not be considered. The possible additional parameters are then: a quadrupole splitting or a distribution thereof, a distribution of isomer shifts or a distribution of thickness values. Again, the optimal parameter values are found for the power set of this list of additional parameters, which are 15 different parameter sets.

The best material models found through this analysis, for both the α -iron and stainless steel foil, are used in theoretical calculations in succeeding steps of the data analysis too, specifically in steps 3, 4 and 6.

3.2.3 Energy calibration

In the third step of fig. 1.1, the energy scale of the experimental energy-time spectra is found by comparing them to corresponding theoretical calculations. Such a theoretical energy-time spectrum combines calculations of transmission through the double-target setup for a range of Doppler velocities of the second target, as described in section 3.1. All transmitted intensities are computed with the strategy from fig. 2.5. Experimentally, the velocities are binned, but the exact size of the bins is unknown. In the theoretical calculations, the velocities are precisely known, so, to find the best possible scaling of the experimental bins, many scalings are applied to them, and all scaled experimental spectra are compared with the (single) theoretical calculation. The best scaling is the one where theory and experiment best match, which is again determined by the Poisson metric of eq. (3.2). This metric can readily be applied to the two-dimensional energy-times spectra; there are merely more data points to consider now. In this third step of the data analysis, conventional optimization suffices to find the optimum scale, i.e., the evolutionary algorithm is not required.

The exact method of comparison between the theory and experiment is a point of consideration too, if one seeks to ensure the accuracy of our data analysis procedure. In principle, the energy-time spectrum can be seen as a continuum; with any Doppler

drive velocity (and at each point in time) it is possible to detect a photon. In practice, however, the experimental data points must be binned in bins of finite size, because of limited count rates, and calculating a continuum is not computationally feasible either. Therefore, some form of approximation is always required in matching the theoretical calculation with the experimental data. Previously, the theoretical grid calculation was linearly interpolated along the entire spectrum, to provide an approximation of the continuum. Each experimental bin was represented by its middle point, and the interpolated theory is evaluated at these same points. This results in two matching intensity values for each pair of points and, after normalization, these can be plugged into eq. (3.2) to assess the fitness.

Instead, it is possible to avoid the interpolation of the theory by simply comparing each experimental bin with the theoretical grid point that is closest to its middle point, i.e., its nearest neighbour. More preferably, an average is taken over all theoretical grid points that fall within the size of the experimental bin, and this average intensity is compared with the value for the experimental bin. Computationally, this is equivalent to applying a moving average to the theory, with a window size equal to the experimental bin size. This new method is expected to perform better especially when the data is non-linear. The counts within one experimental bin might then, for example, be concentrated more towards one of its edges, even though the theory is solely evaluated at its middle point in the method that was previously employed.

The performance of the novel method is evaluated according to the fitness of the calibrated energy-time spectra. Comparing these fitness values for theoretical calculations performed with different material parameter combinations can also serve as an additional assessment of the material characterisation. The energy-time spectra calibrated with the novel method of comparison, and the novel method itself, will also be used in steps 4 and 7 of the analysis.

3.2.4 Phase reconstruction

Once the material has been characterised and the energy range calibrated, the data from experiments where the first target is moved can be analyzed. This is step 4 and 5 of the data analysis chain in fig. 1.1, where the phase induced by the piezo motion is reconstructed. In these and following steps, theoretical and experimental energy-time spectra are compared with the novel method, as mentioned before. Since theoretical calculations are performed with the improved material models, it might not be clear which of these two changes is responsible for any possible improvements in the phase reconstruction and in later steps. This is not problematic in light of the overall aim of attempting to maximally increase the accuracy of the data analysis chain. The combination of these two changes, the improved material models and the novel method of comparison, will henceforth be referred to as the ‘enhanced model’. The previous version of the data analysis procedure, without these changes, will usually be called the ‘original model’.

In step 4 of the data analysis the motion of the piezo, and thus the phase imprinted on the delayed radiation, is reconstructed from a single experimental energy-time

spectrum. Energy-time spectra with piezo motion are calculated theoretically, according to the strategy in fig. 2.5, and compared to the experimental spectrum. The optimum phase is the one whose theoretically calculated spectrum matches the experiment best, which is again evaluated with the Poisson metric, and it is found through the evolutionary algorithm and the subsequent conventional optimization algorithm. So, the phase reconstruction essentially follows the same method as the energy calibration of section 3.2.3, with the difference that now only a single, calibrated, experimental energy-time spectrum is required, and many different theoretical spectra are calculated instead.

The phase is modeled as a (one-dimensional) B-spline, which can assume any arbitrary shape. In principle, the phase could be reconstructed for the entire observation time of each energy-time spectrum. In practice, however, it turns out that the phase reconstruction becomes insensitive after a period of roughly 70 ns ; any changes in the phase at later times only marginally influence its fitness value. Therefore, the phase is only computed for the first 70 ns , and it can be parametrized by eight slope points of the B-spline.

To reconstruct the entire phase, which in our experiment has a period of four times the synchrotron pulse period, several experimental energy-time spectra have to be acquired, all with a different delay of the phase with respect to the synchrotron pulse. This delay can be controlled through our arbitrary function generator. The individual reconstructed phases from each of these spectra can then be combined into the entire phase, if their respective delays differ by less than 70 ns (this difference is usually 25 ns in our experiments). This is done in step 5 of the data analysis chain.

During this step, individual phases can be arbitrarily removed from the combined phase, either if they are clearly anomalous or to ensure that the combined phase is indeed periodic. After this manual selection of individual phases, the combined phase can be optimized once again. It is then compared against all selected energy-time spectra at the same time; for each experimental spectrum, the section of the combined phase with the same delay as the experiment is taken to calculate the corresponding theoretical spectrum. The logarithmic Poisson score for all spectra can simply be added together to return a fitness value for the entire phase. The combined phase is parametrized by 50 or 60 B-spline slope points, for the ESRF and DESY data, respectively.

The optimization of the combined phase can be performed with the evolutionary algorithm and the succeeding conventional optimization once again. However, since the combined phase consists of individual phases that are already optimized, it is likely that the starting point for the combined phase is already close to the optimum. Thus, it is also attempted to use the conventional optimization algorithm only for optimizing the combined phase. The differences between these two optimization strategies will be considered.

It is also examined whether the ‘enhanced model’ returns different reconstructed phases than the ‘original model’, for both individual and combined phases. In the next data analysis steps explained in the following subsection, any of the reconstructed individual phases or the combined phase can be used (see fig. 1.1).

3.2.5 Phase stability

After the phases have been reconstructed, their stability over time can be examined in the final two steps of the data analysis chain. Long-term measurements are divided into equal time intervals and energy-time spectra are constructed for each interval. For each spectrum, in turn, the phase is reconstructed in a similar way as before. This allows one to observe changes in the reconstructed phase, which represent the noise and possible systematic errors in our experimental setup, over the course of the data acquisition period. Note that because of the continuous, pulsed operation mode of SR, indeed only these changes of the relative phase are relevant for the stability of the setup. The phase stability is not affected by absolute phase changes, which, for example, might be caused by drifts in the distance between experimental targets or components.

The phase stability is quantified here by the Allan deviation [40], which for a time interval τ is given by:

$$\sigma_y(\tau) = \left(\frac{1}{2(N-1)} \sum_{i=1}^{N-1} (y_{i+1} - y_i)^2 \right)^{\frac{1}{2}}, \quad (3.3)$$

where N is the number of time intervals and each y_i is the temporal deviation of the phase reconstructed with the data in the i -th interval, $\phi_i(t) = kx_i(t)$, as compared to the phase reconstructed with the complete energy-time spectrum, $\phi_0 = kx_0(t)$. Because for small τ the number of detected events per interval can be very small, it is not possible to reliably reconstruct the phases ϕ_i based on the data from their respective intervals alone. Instead, to determine these deviations y_i , a large library of theoretical energy-time spectra is created, where each spectrum is calculated with a phase that consists of the phase reconstructed with the complete dataset ϕ_0 and some additional noise. This is step 6 of the data analysis chain in fig. 1.1. The additional noise or error is constructed according to some error model, which depend on some parameter and are discussed below.

In step 7 of the data analysis, then, the experimental energy-time spectrum for each interval is compared to the library of theoretical error spectra to find the error parameter that best matches the experiment. Once more, the best fit is determined with the Poisson metric. Finally, this error parameter is converted to the temporal phase deviation y_i (the conversion depends on the error model), and the deviations can be plugged into eq. (3.3) to obtain the Allan deviation. This procedure is then repeated for a range of interval times τ , to create a complete picture of the phase stability over different time scales.

The Allan deviation finds the root mean square of the difference between subsequent deviations (or noise) y_i . For small interval times, the deviations will largely be random, and so will the differences between them. The Allan deviation is then large. Initially, increasing the interval time typically leads to a decrease in the Allan deviation, because the increasing statistics reduce the degree of randomness. The decrease is expected to be proportional to $1/\sqrt{\tau}$, since the interval time τ is itself proportional

to the number of data points in each interval. As the interval time grows larger, the drifts of the deviations on longer time scales, which are caused by systematic errors, exceed the effect of the increased statistics and the Allan deviation increases again.

As mentioned in section 3.2.4, the complete phase ϕ_0 upon which the error models are superimposed can either be an individual phase (step 4) or a combined phase (step 5). As the long-term measurements can only be performed for a single delay of the phase with respect to the synchrotron pulse, only the part of the combined phase with the same delay is required as ϕ_0 . For the individual phases, an extension of the reconstruction beyond the standard 70 ns is required, to include the entire observation window. These phases are now parametrized by 21 instead of 8 slope points. Note that these individual phases might show irregular behaviour for times larger than 70 ns , as the phase reconstruction becomes insensitive in this region (see section 3.2.4). Comparisons between the phase stability of an individual phase or a combined phase as ϕ_0 should be performed with the same experimental energy-time spectrum, i.e., with the same delay of the phase with respect to the pulse.

Three different error models were studied: a linear model, a scaling model and a step-like model. The linear model supposes that the motion reconstructed with the complete dataset x_0 is modified by a linear drift. For the i -th interval, the motion becomes:

$$x_i(t) = x_0(t) + A_i t. \quad (3.4)$$

This error model should cover all low-frequency noise for our observation times [16]. The maximum possible temporal deviation y_i with this model is $y_i = A_i t_2 / c$, where t_2 is the maximum observation time (170 ns for ESRF and 177 ns for DESY). The second error model is a scaling of the form:

$$x_i(t) = (1 + s_i)x_0(t), \quad (3.5)$$

with a scaling factor for each interval s_i . An upper bound on the temporal deviation is now $y_i = \max\{x_0(t)\}s_i/c$. The third error model is an additional step-like motion, as follows:

$$x_i(t) = x_0(t) + d_i \theta(t - 0^+), \quad (3.6)$$

where d_i is a displacement for the i -th interval and 0^+ is a time close to zero after the passing of the δ -peak, such that a phase offset between the δ -peak and the subsequent, delayed radiation is represented. For this model, the temporal deviation is $y_i = d_i/c$.

The Allan deviation is calculated for all three noise models, and for the situation when no motion is applied to the piezo. As for the phases in section 3.2.4, the potential differences between the enhanced and original model will be investigated as well. The difference between using an individual or combined phase as ϕ_0 is also considered.

3.3 Arbitrary phase control

The second aim of this thesis is to arbitrarily control the phase of the delayed radiation, which is induced by the movement of the piezo. As mentioned in section 3.1, the piezo, which is attached to the α -iron foil, is driven by an arbitrary function generator. Previously, the voltage supplied to the piezo by this function generator to obtain a desired phase shift was determined mostly through experience, with the basic assumption that the piezo motion would follow the applied voltage relatively closely. If, for example, a phase jump is required, a sudden jump would also be included in the voltage pattern. However, it is not likely that the piezo expansion can directly follow the voltage and the foil has a certain inertia as well.

Since our data analysis procedure allows us to reconstruct the motion and thus the phase shift generated by a certain voltage pattern, it should be possible to construct a relationship between them. Any desired phase could then be experimentally realised by reversing this relation, applying it to this phase, and driving the piezo with the resulting voltage pattern.

The relationship chosen here is a linear response function, which linearly relates the frequency components of the resulting phase $\phi = kx$ to those of the voltage U applied to the piezo:

$$\phi(\omega) = L(\omega)U(\omega), \quad (3.7)$$

where $L(\omega)$ is the linear response. It can be found from a known combination of phase and voltage as $L(\omega) = \phi(\omega)/U(\omega)$, because it is a linear function. The inverse $L^{-1}(\omega)$ of the response function is simply $1/L(\omega)$ and can be employed to find the required voltage pattern for a desired phase:

$$U(\omega) = \frac{1}{L(\omega)}\phi(\omega). \quad (3.8)$$

Since the phase is reconstructed in the time domain and the voltage is a function of time too, their Fourier transforms must first be taken to find $L(\omega)$. Only the first 15 components (excluding the constant, DC component) of the Fourier expansions of the phase $\phi(t)$ and voltage $U(t)$ are actually used to determine $L(\omega)$, as the piezo would likely not be able to follow the high-frequency components of any voltage applied to it anyway. So, the response function $L(\omega)$ (and its inverse) can computationally be represented as an array with 15 elements. Once the voltage for a specific phase has been determined with eq. (3.8), an inverse Fourier transform must be applied to it, back into the time domain, before it can drive the piezo.

The approach presented above is tested experimentally with an initial voltage pattern that was previously employed to generate a π step-like phase. This phase is reconstructed using our data analysis procedure and the linear response $L(\omega)$ is found. Then, as a proof of principle, it is attempted to improve the step-like nature of this phase, by plugging a perfect step function into eq. (3.8) and applying the resulting voltage pattern to the piezo. Finally, this improved, controlled phase is reconstructed as well, to examine if it indeed behaves more step-like.

The phase stability of the optimized phase is considered as well, with the data analysis procedure as described in section 3.2.5. For this calculation, and for the phase reconstructions, the ‘enhanced model’ (see section 3.2.4) was not implemented, because the response function $L(\omega)$ and the voltage pattern for the improved step-like phase were already found and applied during the experiment, using the ‘original model’ of the data analysis. Implementing the enhanced model only partially would lead to mixed results, which would be more difficult to interpret.

4 Results and discussion

The first part of this chapter, which comprises sections 4.1 to 4.4, discusses the accuracy improvements of the data analysis chain. These results will be presented in the same general order as in chapter 3, i.e., according to fig. 1.1. This means that the improved characterisation of the targets and the novel method for comparing theoretical and experimental energy-time spectra, together referred to as the ‘enhanced model’, will be evaluated at each step along the data analysis chain, starting with the single target data and culminating in the stability of the piezo-induced phase. The theoretical calculations are performed with a high resolution throughout, to ensure that this is not a factor limiting the eventual phase stability. In the second part of this chapter, in section 4.5, the arbitrary phase control approach is discussed.

4.1 Material characterisation

4.1.1 Alpha iron foil

Out of all the different parameter combinations that were tested, two parameters particularly stood out in the characterisation of the α -iron foil: the distribution of thickness values and the magnetization angles, which parametrize the direction of the hyperfine magnetic field. Even if these parameters are the only additional ones, differences with the original result, where only the thickness and the strength of the hyperfine magnetic field are considered, are noticeable. The optimal parameter values for these material models are presented in table 4.1 for the ESRF data. All thickness values are sensible and the magnetic field strength is close to the reported value of ≈ 33 T [41, 42] for all models. From now on the different models will be referred to according to the symbols introduced in this table. The transmitted intensities

Model	Thickness (μm)	Magnetic field strength (T)	Magnetic field angle φ (rad)	Magnetic field angle θ (rad)	Thickness distribution width (μm)	Fitness
<i>A</i>	2.52	32.7	-	-	-	-76700
<i>B</i>	2.52	32.7	-	-	0.30	-73000
<i>C</i>	2.66	32.6	1.70	1.37	-	-40200

Table 4.1: Optimal parameter values and fitness for the ESRF α -iron foil, for the original model (*A*), the model with an additional thickness distribution (*B*) and the model with an additional magnetization direction (*C*).




Figure 4.1: Optimized theoretical models of transmission through an α -iron foil at ESRF facilities, for different parameter combinations. Blue: experimental results, orange: model *A*, red: model *B*, green: model *C*. See table 4.1 for model descriptions.

achieved with these parameter combinations are plotted in fig. 4.1, together with the experimentally measured transmission.

Thickness distribution

The inset of fig. 4.1 displays a time interval where the experimental intensity appears irregular and featureless. This area is of particular interest, because model *A*, the original model, seems unable to correctly explain this irregularity: this theory shows a trough, whereas the experimental intensity remains largely constant. Comparing the experimental result against the expected curve for a combined dynamical and quantum beat (see fig. 2.4b), the first minimum of the dynamical beat should be located in this time interval. This explains the trough in this area for model *A*. The precise position of the minimum is determined by the thickness of the target. With a distribution of thickness values, the average is taken over many such minimum positions, each one slightly displaced with respect to the previous one. The structure of the dynamical beat minimum should then be blurred or vanish altogether. Model *B*, the model that includes a thickness distribution, indeed seems to exhibit this behaviour (fig. 4.1) and thereby (at least visually) approximates the experimental results more closely. A comparison of the fitness values, calculated as the logarithm of eq. (3.2), also indicates that model *B* is a better representation of the experiment than the original model *A*, as shown in table 4.1. Note that these fitness values are

Model	Thickness (μm)	Magnetic field strength (T)	Magnetic field angle φ (rad)	Magnetic field angle θ (rad)	Thickness distribution width (μm)	Fitness
<i>A</i>	2.51	32.5	-	-	-	-54600
<i>B</i>	2.52	32.5	-	-	0.25	-53400
<i>C</i>	2.60	32.5	1.54	1.74	-	-44900

Table 4.2: Optimal parameter values and fitness for the DESY α -iron foil, for the original model (*A*), the model with an additional thickness distribution (*B*) and the model with an additional magnetization direction (*C*).

the logarithm of eq. (3.2), which will always return negative values, as its terms are probabilities that lie between 0 and 1. The best fitness value is the maximum one, i.e., the one closest to zero.

From this outcome one could conclude that the α -iron target's thickness in fact has some distribution, i.e., that its surface is not smooth. The Gaussian thickness distribution's mean value is $2.52 \mu\text{m}$, its width (standard deviation) is $0.30 \mu\text{m}$, which appears reasonable. As well, model *B* leads to similar behaviour in the time interval around the first DB minimum for the DESY data too (see fig. 4.2). In fact, the same α -iron foil was mounted in both the DESY and ESRF experiments. The optimal parameter values for the DESY data, shown in table 4.2, indeed roughly correspond with those of the ESRF data. This reinforces the reliability of our fitting procedure with the evolutionary algorithm across different experiments. For the DESY data, the width of the thickness distribution was $0.25 \mu\text{m}$ instead. This difference with the ESRF result could be due to various (random) differences in the respective experimental setups; for example, the synchrotron beam could simply be hitting a different area of the target, which happens to be smoother.

Magnetization direction

The fitness improvement of model *B* for the ESRF data is slight, however, compared to that of model *C* (see table 4.1), which adds the direction of the internal magnetization field to the original parameters of model *A*. It can be difficult to quantitatively interpret these relative fitness differences. The fitness represents the likelihood of a theoretical model given the experimental results, and, for the ESRF model *C*, it is quantified as $\exp(-40200)$ (see eq. (3.2)). The fraction $\exp(-40200)/\exp(-76700)$ would then indicate how much more likely it is that model *C* correctly explains the data compared to model *A* (see table 4.1). However, because the product in eq. (3.2) contains many factors (and many of those are close to zero), both factors in this fraction are essentially zero and it is undefined. It can be helpful to consider the absolute maximum value the fitness can attain as a baseline to compare other results by. This maximum value is reached when the theory exactly matches the experiment. So, it can be calculated by substituting the experimental values into eq. (3.2) as if




Figure 4.2: Optimized theoretical models of transmission through an α -iron foil at DESY facilities, for different parameter combinations. Blue: experimental results, orange: model *A*, red: model *B*, green: model *C*. See table 4.2 for model descriptions.

they represented the theory. For both ESRF and DESY experiments this maximum fitness is -3300 .

In any case, model *C* is still the best few-parameter model in terms of its fitness. From an experimental point of view it is also sensible that including the direction of the internal magnetic field improves the theory model. The external magnet that controls the magnetization direction is aligned manually. Therefore, it is likely that the desired alignment, parallel to the magnetic polarization vector, is not exactly achieved. Model *C* is the best few-parameter model for the DESY data too, although the fitness difference with model *A* is not as dramatic there (see table 4.2).

Comparing the optimal magnetization vectors, parametrized by their spherical coordinates in table 4.1 and table 4.2, for different datasets directly can be misleading, as there are usually two vectors, mirrored around the axis of the desired alignment, that produce the same result. The optimization procedure cannot distinguish these degenerate vectors and randomly outputs one of them. The dot product between the preferential magnetization direction and the optimal direction found, though, can still serve as a good measure of the overlap between them. For the ESRF data, it is ≈ 0.97 , and for the DESY data ≈ 0.99 . The magnetization direction for the DESY data is closer to the preferred direction, which might explain the small improvement in fitness of model *C* against model *A*, relative to the same improvement for the ESRF data. The difference in the alignment along the preferential direction between

the two data sets is most likely due to the inherent inconsistency in the manual alignment of the external magnet.

Many-parameter combinations

The entire power set of all additional parameters mentioned in section 3.2.2 has been examined as possible material models for the ESRF α -iron foil. For parameter combinations with more than one additional parameter, the thickness distribution and the magnetization direction still emerge as the most crucial ones. A selection of interesting results with many-parameter combinations is summarised in table A.1 and plotted in fig. A.1. The combinations plotted there all include at least the magnetization direction, i.e., all are extensions of model *C*, since combinations with this parameter consistently demonstrate the largest improvements in fitness. The model that adds a thickness distribution to model *C* (model *D* in table A.1) seems to visually approach the experiment more closely, as was the case for model *B* compared against model *A*, in the irregular region around 100 ns. In some cases, inclusion of the quadrupole splitting as a parameter is favorable too. The highest possible fitness (-30900) is achieved by combining the following parameters (model *F* in table A.1): the magnetization direction, the thickness distribution, a quadrupole splitting and a distribution of isomer shifts. Some parameters that were not consequential on their own, such as the distribution of isomer shifts, thus improve the fitness if they are added to model *C*. This suggests a certain hierarchy of the parameters, with the magnetization direction being the most important one.

For the DESY data, the conclusions about many-parameter combinations are largely the same. The same selection of results as for the ESRF data is listed in table A.2 and plotted in fig. A.2. The model that adds a quadrupole splitting to model *C* performs well here too, although visually other parameter combinations seem to better capture the outline of the data in the irregular area around 100 ns. Therefore, and for reasons of computational complexity, only model *C* itself, with the single most important additional parameter, the magnetization direction, is chosen as the final characterisation of the α -iron target, which is used in the following steps

Model	Thickness (μm)	Isomer shift distribution width (mm/s)	Quadrupole splitting (mm/s)	Thickness distribution width (μm)	Quadrupole distribution width (mm/s)	Fitness
<i>G</i>	1.04	0.08	-	-	-	-10200
<i>H</i>	0.38	0.11	0.22	0.35	-	-7100
<i>I</i>	0.95	0.05	0.19	-	0.03	-8300

Table 4.3: Optimal parameter values and fitness for the ESRF stainless steel foil, for the original model (*G*), the best model obtained with the conventional optimization (*H*) and the best model obtained with the constrained optimization (*I*).

of the data analysis chain. This is motivated further in section 4.2.

4.1.2 Stainless steel foil

For the stainless steel foil, the material characterisation is somewhat less complicated, since fewer parameters need to be considered. Still, the entire power set of parameters in section 3.2.2 is examined again. For the ESRF data, the conventional optimization algorithm, which is applied to the result of the evolutionary algorithm, sometimes yielded parameter values that appeared unrealistic, such as too small thicknesses. Using a constrained optimization algorithm instead, where the parameter values are restricted within realistic bounds, prevents this. Table 4.3 lists the optimal parameter values for the best possible parameter combination for the conventional algorithm (model H) and the constrained algorithm (model I), together with the values found for the original model G . In fig. 4.3 the calculated intensities for these material models is shown. Both models H and I seem to approximate the experimental results better than the original model G , especially for late times, even though they contain different parameters.

Both of them, however, include the quadrupole splitting and the distribution of isomer shifts as parameters, although the latter is equivalent to the distribution of quadrupole shifts for many parameter combinations. It seems, thus, that several




Figure 4.3: Optimized theoretical models of transmission through a stainless steel foil at ESRF facilities, for different parameter combinations and optimization algorithms. Blue: experimental results, orange: model G , red: model H , green: model I . See table 4.3 for model descriptions.

combinations of parameters can all explain the experimental results well. Because of the limited observation time, it is also not clear if the somewhat irregular late time region (displayed in the inset of fig. 4.3), where the largest differences between models seem to occur, is the first minimum of the quantum beat, or if other effects are active.

The fitness values for both improved models H and I is higher than that of the original model G (see table 4.3). All fitnesses are significantly higher than for the α -iron foil, while the baseline of the maximum possible fitness is approximately the same: it is -3400 . It can then be concluded that the fits for the stainless steel foil are generally more accurate, which fig. 4.3 indeed seems to confirm. While model H yields the best fitness, the optimal parameter values it finds are physically improbable: the thickness distribution has a width of $0.35 \mu m$, around a mean thickness of only $0.38 \mu m$ (see table 4.3). The thickness for model I is $0.95 \mu m$, and its other parameter values appear physically feasible too. In fact, the values for the hyperfine interaction parameters of model I are close to values from literature for stainless steel foils, obtained with different experimental methods [43].

Therefore, model I is taken as the characterisation of the stainless steel foil for the next steps of the analysis. In the next section, further arguments in support of this decision are presented.

The conventional optimization algorithm does not return physically unlikely results for the DESY data. There, the model with just the original parameters is not improved significantly, neither visually (see fig. A.3) nor in terms of the parameter values and the fitness (see table A.3), by adding more parameters to it. Some many-parameter combinations even show non-physical results again. The original model (model G in table A.3) is thus kept to characterise the stainless steel foil for the rest of the analysis chain. This difference in the best model might be because a different foil, with possibly different properties, is characterised. It could also be because the DESY facilities allow a slightly longer observation time (the time between pulses is $\approx 192.5 ns$ instead of $\approx 176 ns$ for the ESRF facilities).

4.1.3 Summary

To summarise this section, the thickness distribution (model B) and internal magnetization direction (model C) were the most important parameters for the characterisation of the α -iron foil. Adding the thickness distribution improves results around the first minimum of the DB, while the magnetization direction is an overall better fit, which is experimentally reasonable. Model C is chosen as the best model for the next data analysis steps, both for the ESRF and DESY data.

The conventional optimization algorithm yielded non-physical parameter values for the ESRF stainless steel characterisation. Therefore, a constrained optimization algorithm was employed instead. The best parameter combination (model I) added a quadrupole splitting and a distribution thereof to the original model G . This model is used for the following steps of the ESRF data analysis, while for the DESY data analysis the original model suffices.

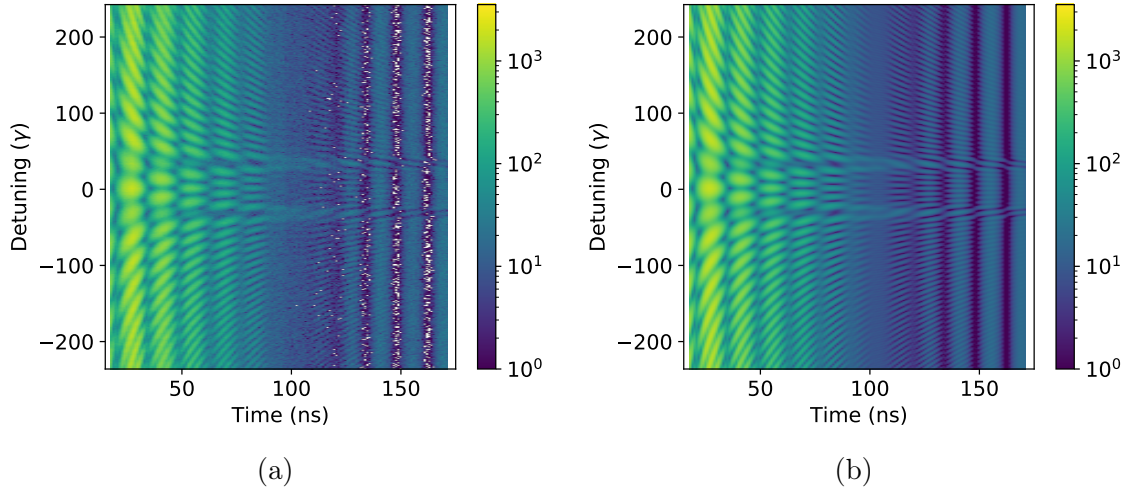


Figure 4.4: Energy-time spectra for the ESRF data, from (a) the experiment and (b) the corresponding theory, obtained with the novel method. White spaces in (a) are bins with zero photon counts.

4.2 Energy calibration

4.2.1 Novel method

The energy calibration obtained with the novel method of comparison between theoretical and experimental energy-time spectra is almost exactly the same as the one found with the previously employed (original) model. In fig. 4.4a a typical experimental energy-time spectrum is shown, which is calibrated with its corresponding theory shown in fig. 4.4b and using the novel method of comparison. The theory is calculated with the original models (see section 4.1) for both the α -iron and stainless steel target and seems to be in good agreement with the experiment. It is reasonable that the energy calibration for fig. 4.4 is virtually unchanged with the novel method, because it changes the comparison between theory and experiment only. The way in which the theory is calculated and the experimental data acquired remains unaltered.

The fitness associated with the comparison of theory and experiment in fig. 4.4 with the novel method was -468000 . With the previous method a fitness of -522000 could be achieved, while the maximum fitness possible was -170000 . For the DESY data too, the novel method improved the fitness values of energy-time spectra comparisons, albeit less significantly. Specifically, when the original material models (see section 4.1) are used, the novel method yields a fitness of -453000 , compared to -478000 with the previous method and a maximum possible fitness of -182000 .

The novel method performs better and so the theory that was manipulated with a moving average matches the experiment better than the interpolated theory that was previously employed. For this step of the data analysis, the novel method thus only improves the degree of confidence in the energy calibration found, not the energy calibration itself. The same conclusion is also drawn when other material models are




Figure 4.5: Difference in Poisson score of individual data points, for (a) novel comparison method versus method previously used and for (b) magnetization direction model versus original model. Blue indicates higher Poisson score and thus better agreement with experiment. In (a) the colorbar was clipped for better visibility, the original values ranged from -28 to 115.

considered, or if the DESY data is considered instead. Still, it is possible that for later steps in the data analysis chain the new method does alter results. If not, it would likely improve the confidence in our results. For that reason, it is also utilised in the comparison of different material models in section 4.2.2.

Note that the improved performance with the novel method is only possible when a high resolution is chosen, i.e., when a sufficient number of theory grid points fall within the experimental bin size. Otherwise, the moving average is not possible and the novel method actually performs worse, also because nearest neighbours are then likely remote.

The novel and original method of comparison between theory and experiment can also be visually compared by their Poisson metric scores. Figure 4.5a shows this graph, where for each individual data point the difference in the logarithmic Poisson metric between the two methods is plotted. The novel method seems to have an equal or better Poisson score in most areas of the energy-time spectrum, but the difference is especially large for early times and large detunings. If the stainless steel foil is far detuned by means of the Doppler shift, the beating pattern between its resonance and the α -iron resonance has a high frequency. During early times the count rates are still high (see fig. 4.4), so the high-frequency oscillation will also have a large amplitude there. The areas with the largest differences in fig. 4.5a are thus relatively non-linear. These are exactly the areas where the novel method was expected to perform better, because it averages over these fast-changing areas in the theory. Experimentally, a similar averaging is always present in the form of binning.

4.2.2 Material models

Both notable additional parameters for the α -iron material model that were listed in table 4.1 for the ESRF data, the thickness distribution (model *B*) and the magnetization direction (model *C*), improve the fitness for the comparison of energy-time spectra too. With model *B* a fitness of -463000 can be reached (compared to -468000 with the original model *A*) and model *C* returns a value of -440000 . The original stainless steel model was used for these fitness calculations. So, in terms of fitness, model *C* performs better than model *B*, corroborating the outcome of the single-target data analysis in section 4.1.1.

The different material models can also be compared against the original model *A* in graphical form through their metric score. In fig. 4.5b this is done for model *C*. This figure shows that this model actually performs worse than the original model in some areas of the spectrum, though this is compensated for mostly during very early times. This early times improvement of model *C* is not easily recognized from the α -iron foil data alone (see fig. 4.1). The same graph for model *B* is plotted in fig. A.4a. This figure is less clear because the total fitness difference between model *B* and the original model *A* is relatively small. However, it is noteworthy that the area around 100 ns in this figure, where model *B* visually seemed to achieve better results before (see fig. 4.1), provides an insignificant contribution to the total fitness difference between the models. So, the Poisson metric comparisons for both improved models show that visual inspection of the single-target data as presented in fig. 4.1 is perhaps not the best indication of the actual performance of the models.

Adding more parameters to model *C*, for example those in table A.1 and table A.2, only slightly increases or even decreases the fitness for the energy-time spectra. In fact, some of these models have a better fitness than model *C* for the single-target experiments, but perform worse with the double-target data. For example, for the DESY data (see fig. A.4b), the fitness of model *E*, which includes quadrupole splitting as a parameter (see table A.2), is -491000 , compared to a value of -440000 for model *C* (see table 4.2), while model *E* performed significantly better for the single-target data (see table A.2). Since α -iron has a regular crystal structure, the local charge distributions of its nuclei should be symmetric, and no quadrupole splitting is actually expected to exist [41, 42]. Even if it is physically not present, it could still express itself in the theoretical calculations in an advantageous way. However, the preceding result seems to exclude this possibility.

For the stainless steel models, the discrepancy between single-target and double-target experiments is more striking. For the ESRF models *H* and *I* (see table 4.3), the fitness values were respectively -719000 and -457000 , compared to -468000 for the original model *G*. These calculations were performed with the original α -iron foil model *A* (see table 4.1) and the metric comparisons for both models with the original stainless steel model *G* are shown in fig. A.5. As discussed in the previous section, the model obtained via the conventional optimization algorithm contained some unrealistic material parameters. Whereas for the single-target data the two improved models yielded similar results (see table 4.3), the comparison of their energy-time




Figure 4.6: Comparison of enhanced model (orange) versus original model (blue) for ESRF data, for (a) an individual reconstructed phase and (b) the combined phase. The latter is obtained with the conventional optimization algorithm only. Differences between models can barely be distinguished.

spectra seems to, again, rule out the model with non-physical parameter values. It is reasonable to conclude that the energy-time spectrum provides a better assessment of which material model is the best, since it is simply a more complete picture.

It is then also justified to indeed consider only model *C* for the α -iron foil in the next steps of the analysis. For the stainless steel foil, model *I*, which includes the quadrupole splitting and its distribution, is taken for the ESRF data, while for the DESY data the original stainless steel model *G* suffices. The novel method for comparing theoretical and experiment energy-time spectra is also employed in all following steps of the data analysis, if applicable.

4.3 Phase reconstruction

The ‘enhanced model’ (see section 3.2.4), which combines the optimum material models, summarised in section 4.1.3, and the novel method of comparing theory and experiment from the previous section, is tested by comparing the phases reconstructed with it against the phases obtained with the original data analysis. For a particular reconstructed phase of the ESRF data, which approximates a π phase-step, this comparison is shown in fig. 4.6a. It can be seen that the phase obtained with the enhanced model is hardly different. Its fitness is -197000 , while the original model yields a value of -203000 , compared to a maximum possible fitness baseline of -44000 . The fitness is slightly increased with the enhanced model, although, based on the individual improvements of the material model (section 4.1) and the novel method of comparison (section 4.2), one would perhaps expect a larger difference. For smoother reconstructed phase functions, the improvement might be even smaller.




Figure 4.7: Comparison of enhanced model (orange) versus original model (blue) for DESY data, for (a) an individual reconstructed phase and (b) the combined phase. The latter is obtained with the conventional optimization algorithm only. Differences between models can barely be distinguished.

An example of such a phase function for the DESY data, which is more or less linear, is shown in fig. 4.7a. There, the fitness of the enhanced model is actually worse than with the original model. However, comparing the fitness of these models directly is ambiguous. For the DESY data, the experimental energy resolution is higher; the bin sizes are smaller in the energy direction. Previously, for the original model, it was customary to merge a few bins together along the energy axis. For the enhanced model this merging was avoided, to ensure that the resolution would not limit the accuracy of this and further data analysis steps. Therefore, the experimental data contains more bins and the sum of the logarithmic Poisson metric (eq. (3.2)) runs over more terms. The fitness will then always be worse.

It is unsure why the enhanced model does not demonstrate similar increases in the Poisson score as were seen in the previous sections. For the ESRF data, the range of detunings considered for the phase reconstruction was customarily reduced to $(-120\gamma, 120\gamma)$. It might be that the areas where the novel comparison method performed well, namely for large detunings (see fig. 4.5a), were cut off and that the enhanced model consequently did not perform well. It is also possible that the piezo motion simply has a larger influence on the energy-time spectra comparison than changes in the material models or in the method of comparison. To understand the enhanced model better, it would be helpful to investigate the impact of the material models and the novel comparison method individually.

The enhanced model also hardly alters the full-period phases, as is expected, since they are constructed by combining the largely unchanged individual phases. This result is shown in fig. 4.6b for the ESRF data and in fig. 4.7b for the DESY data. Much like the individual phases, the enhanced model only slightly improves the fitness of the combined ESRF phase, whereas for the DESY data the fitness is worse

again, because of its different experimental resolution.

For both datasets, the simple concatenation of individual phases was already close to the optimum combined phase, as was expected. This is shown in fig. A.6, where both the evolutionary algorithm and the conventional optimization algorithm do not alter the combined phase much. It is interesting to note that there might be a slight difference between enhanced and original model for the unoptimized combined phase, but applying the conventional optimization algorithm allows them to converge (see fig. 4.6b). In the next section, which covers the final two steps of the data analysis chain, the combined phases obtained with the conventional optimization algorithm will be used.

4.4 Phase stability

The phase stability results are divided over subsections. In section 4.4.1 the main stability results with the linear noise model (eq. (3.4)) are presented. Section 4.4.2 considers the stability with the scaling (eq. (3.5)) and step-like (eq. (3.6)) noise models. In section 4.4.3 the stability without any piezo motion is discussed. The analysis of the temporal deviations in section 4.4.4 complements the Allan deviation results in the previous subsections. It is also beneficial for the comparison between the stability of individual and combined phases in section 4.4.5. Section 4.4.6, finally, discusses and summarises some results from previous subsections.

4.4.1 Linear model

The phase stability as expressed by the Allan deviation is shown in fig. 4.8a for the ESRF data. For this calculation, the combined phase in fig. 4.6b, which is reconstructed with the entire dataset, is taken as ϕ_0 (see section 3.2.5), and the linear noise model is used. The experimental delay of the combined phase with respect to the synchrotron pulse was the same as that for the individually reconstructed phase in fig. 4.6a, so the phase is a step-like phase jump of π . The few-zeptosecond level of stability is reached, as has already been shown [16]. Here, though, the phase stability achieved with the enhanced model is also shown in fig. 4.8a; it is almost the same as with the previous model.

For the DESY data, the phase stability is shown in fig. 4.8b: a similar stability range is obtained. For this calculation, the extension of the individual phase in fig. 4.7a, which is approximately a linear function, to the maximum observation time of 177 ns is taken as ϕ_0 , and the linear noise model is used again. The extended measurement time for the DESY data allows the determination of the Allan deviation for larger interval times, where the phase stability remains at the few-zeptosecond level. At select values of the interval time τ , it even drops below the zeptosecond barrier. No distinct rise in the Allan deviation for the larger interval times, which would indicate systematic errors at those time scales, is discerned.

It seems that for these larger intervals, the counteracting effects of systematic




Figure 4.8: Phase stability for (a) the ESRF data and (b) the DESY data. Blue: original model, orange: enhanced model. The shaded areas represent error regions and the diagonal lines denote $1/\sqrt{\tau}$ scaling.

errors and of increased statistics balance each other. The DESY stability does, however, show good agreement with the expected $1/\sqrt{\tau}$ scaling for small interval times. Comparing the phase stability results from both datasets in fig. 4.8, the steep drop in the ESRF stability for the largest interval times (around $\tau = 10^3$) is not reproduced in the DESY data. For these interval times the statistics for the Allan deviation measure are low, so it seems possible that this drop is an expression thereof, instead of a physically meaningful result.

The Allan deviation obtained with the enhanced model is actually slightly higher for the DESY data, which signifies a slightly worse stability. For the ESRF data, the enhanced model attains similar, or perhaps, in some places, a slightly improved phase stability. Thus, it can not be concluded that the enhanced model is clearly better or worse than the original model. However, for both datasets, the differences between the enhanced and original model are relatively small. It seems fair to conclude that the final phase stability is robust against the changes made in the data analysis chain, and so possible inaccuracies in the data analysis will not be a major factor of systematic error in determining the (maximum) phase stability.

Whether the fundamental limit of the phase stability of our setup has indeed been reached, can not be inferred from the DESY data, as the expected increase in the Allan deviation for large intervals is not yet seen. It is possible that measurements over even longer periods of time uncover a lower Allan deviation still. If experimental errors that manifest themselves on medium time-scales could somehow be reduced, the Allan deviation could be improved further, as the clear $1/\sqrt{\tau}$ scaling seen in fig. 4.8b could continue for longer. In any case, the few-zeptosecond phase stability that is found for both datasets is sufficient to implement reliable and stable quantum control of nuclei, as has been shown by Heeg et al. [16].




Figure 4.9: Phase stability for (a) the ESRF data and (b) the DESY data, with a linear error model (blue), a scaling model (green) and a step-like model (red). The green and red curve overlap for (a). All results were obtained with the original model. The shaded areas represent error regions and the diagonal lines denote $1/\sqrt{\tau}$ scaling.

4.4.2 Scaling and step model

The results presented in fig. 4.8 were all obtained with the main linear noise model. The Allan deviations for the other noise models, the scaling and step-like model, are shown in fig. 4.9. Since the enhanced data analysis model did not perform particularly better or worse than the original model, these results are all obtained with the original model. For the ESRF data, the enhanced model performance was tested explicitly for the scaling and step-like models too, and the resulting Allan deviations overlapped those in fig. 4.9a. The phase stability of these other two noise models reaches well into the sub-zeptosecond (yoctosecond) region; they are all a factor of two to five more stable than the linear model, which is thus indeed the largest factor of noise.

Whereas the stabilities for the scaling and step-like model overlap for the ESRF data, the scaling model has a significantly lower stability (higher Allan deviation) for the DESY data. This difference in the prominence of the different noise models is likely due to the different shapes of the phases for which the Allan deviation is calculated. For a perfect step-function, magnifying the step with the scaling model has the same effect as adding a small additional step with the step-like model. It is thus likely that the results of these two models are close for a step-like phase function. This is indeed the case for the ESRF stability (see fig. 4.9a), which is based on a step-like phase (see fig. 4.6a).

For a perfectly linear function, a scaling of the linear function has the same effect as adding an additional linear drift through the linear model. Since the DESY phase looks nearly linear (see fig. 4.7a), the scaling and linear model would likely overlap or be very close. However, the additional displacement $s \max\{x_0(t)\}$ to the base motion $x_0(t)$ caused by the scaling model (see eq. (3.5)) must then equal the additional




Figure 4.10: Phase stability without piezo motion for (a) the ESRF data and (b) the DESY data, with a linear error model. Results for both plots were obtained with the original model. The shaded areas represent error regions and the diagonal lines denote $1/\sqrt{\tau}$ scaling.

drift At_2 caused by the linear model (see eq. (3.4)). Here a value of 4 is chosen as $\max\{x_0(t)\}$, which is approximately the phase value at 70 ns (see fig. 4.7a), not at the maximum observation time of $t_2 = 177\text{ ns}$, since the individual phase becomes irregular for larger times (see section 3.2.4). This explains why the scaling model has a higher stability than the linear model: it is not a true upper bound. If instead the linearly interpolated phase value at $t_2 = 177\text{ ns}$ is chosen as $\max\{x_0(t)\}$ (namely, $(177/70) \cdot 4$), the phase stability achieved with the scaling and linear model is close.

4.4.3 Without motion

Apart from the other noise models in the previous section, a higher phase stability is also reached if no piezo motion is applied to the α -iron foil, as shown in fig. 4.10. For both datasets, the phase stabilities were obtained with the original model, and using the linear noise model. Again, for the ESRF data the enhanced model was also explicitly checked and it overlapped the result for the original model in fig. 4.10a. The behaviour of the phase stability is similar for both datasets. Both go well into the sub-zeptosecond region, with the DESY phase stability perhaps even reaching 0.1 zs for large interval times. They also seem to follow roughly the same trajectory, close to the $1/\sqrt{\tau}$ scaling, and systematic errors do not seem to introduce themselves yet. The latter is sensible, as the largest errors are expected to be related to the dynamics of the piezo motion.

For our phase manipulation scheme, the stability without motion is of course not too relevant. However, these results do demonstrate that there are no large issues with our basic setup that could negatively influence the actual phase stability, i.e., the stability when motion is included. The fact that the stability trajectories for both datasets are similar also suggests that our basic experimental setup is well repeatable (which the phase stability results with motion in fig. 4.8 already suggested too).




Figure 4.11: Temporal deviations over the total measurement time for the stability results presented in fig. 4.8. For the ESRF data (a) the interval time τ was 200 s and subsequent intervals were offset by 2 s, to allow a better view of early times. For the the DESY data (b) the interval time τ was 100 s and there is no overlap between intervals. Blue: original model, orange: enhanced model. The shaded areas represent error regions.

4.4.4 Temporal deviations

To attain the phase stabilities of sections 4.4.1 and 4.4.2, which are plotted in figs. 4.8 and 4.9, specific parts of the measurement period were excluded. Plots of the temporal deviations y_i (see eq. (3.3)) over the total measurement time, as shown in fig. 4.11, determined which regions were discarded. Specifically, fig. 4.11 shows the deviations as obtained for the linear noise model, which correspond to the phase stability results in fig. 4.8. For the ESRF data, the deviations are significantly higher and less stable during the first ≈ 400 s, which is identified as the initialisation period of the piezo. This area was thus removed to obtain the phase stability in fig. 4.8a. A comparison of the stability with and without this initialisation period is shown in fig. 4.12a.

The deviations for the DESY data seem relatively stable for most of the measurement period, compared to the large peak at around 35000 s. The origin of this seemingly sudden peak is unknown. As the peak is ≈ 80 zs high, which corresponds to a difference in phase of ≈ 1.75 rad, this is not a complete breakdown of the motion (see fig. 4.7a). The deviations do seem more unstable after the peak than before it, indicating something might have permanently changed in the setup with the emergence of the peak. The phase stability can indeed be improved when the area from the start of the peak onwards is removed, as is shown in fig. 4.12b. If additionally the area up to and including the first smaller peak, which appears at around 3500 s, is removed, the stability of fig. 4.8b is obtained. With the elimination of this peak, fluctuations in the phase stability for interval times around $\tau = 100$ s also disappear (see fig. 4.12b). These interval times in fact correspond to the width of this first smaller peak. It thus seems that these fluctuations at intermediate times

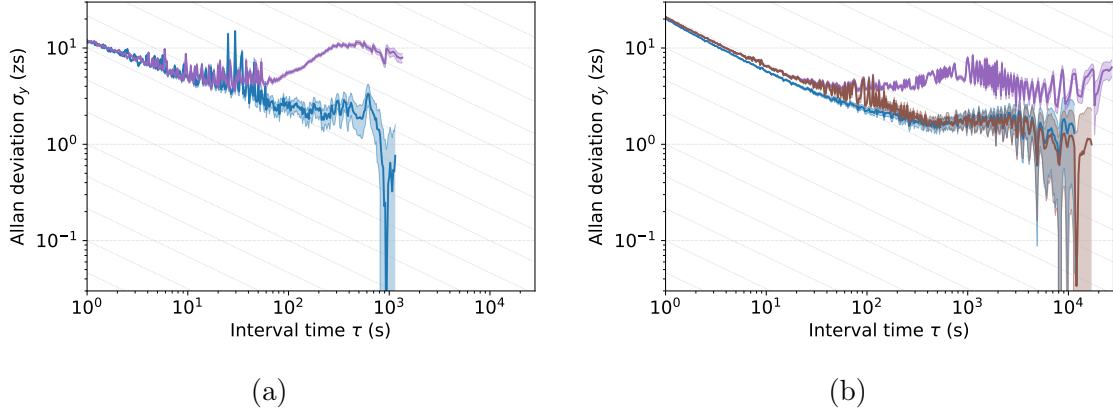


Figure 4.12: Phase stability for different periods of the measurement time, for (a) the ESRF data and (b) the DESY data, with a linear error model. For both plots, purple represents the stability for the entire measurement time. In (a), the first 400 s were removed for the blue curve. In (b), the deviations from 34500 s onwards were removed for the brown curve. For the blue curve, the first 4000 s were additionally removed. The blue curves in both plots are identical to the ones in fig. 4.8. Results for both plots were obtained with the original model. The shaded areas represent error regions and the diagonal lines denote $1/\sqrt{\tau}$ scaling.

are not related to detector dead times, which was the case for the ESRF data [16].

Analyzing the temporal deviations for the enhanced and original model (see fig. 4.11) could clarify their influence on the phase stability results. For both datasets, the enhanced model deviations follow the same general trajectory as the deviations for the original model, which leads to the similar stability results observed in fig. 4.8. The offset of the deviations for the enhanced model with respect to those of the original model for the ESRF data (see fig. 4.11a) moves the deviations closer to zero, which indicates that smaller errors are found with the enhanced model. This offset, however, does not change the conclusions about the phase stability, as the trends in the deviations over time are the same as for the original model.

For the DESY data, no such offset larger than the error regions is seen. From fig. 4.11b also no clear reason can be discerned for the slightly worse stability results of the enhanced model (see fig. 4.8b).

4.4.5 Individual and combined phases

The DESY phase stabilities from sections 4.4.1 and 4.4.2, plotted in fig. 4.8b and fig. 4.9b, were computed for the extended individual phase of fig. 4.7a. The stability for the combined phase (see fig. 4.7b) for the same experimental energy-time spectrum, and thus with the same delay, as the individual phase, is shown in fig. 4.13a, for the linear noise model. It is slightly worse than the stability of the individually reconstructed phase. This is unexpected, as the combined phase should improve




Figure 4.13: Phase stability (a) and the corresponding temporal deviations (b) over the total measurement time for the DESY data, obtained with the linear noise model, for the individual phase (blue) and the combined phase (grey). The individual phase was also plotted in fig. 4.8b. The shaded errors represent error regions. For (a), the diagonal lines indicate $1/\sqrt{\tau}$ scaling.

the accuracy of the total phase reconstruction by overlapping the individual phases and retaining only the early time periods of each phase, which should be the most well-defined periods due to higher count rates.

The deviations y_i for the individual and combined phase are compared in fig. 4.13b. For the individual phase the deviations are remarkably closer to zero than for the combined phase, which has an offset of ≈ 65 zs. So, the individually reconstructed phase seems to represent the phase better throughout the measurement time. Apparently, the fact that the (extended) individual phase is not well-defined for times larger than 70 ns (see section 3.2.4) is no issue. This is actually reasonable, since these late times will not be relevant for the spectra calculated with errors added to the phase either. This result also implies that the DESY combined phase is somehow inaccurate. This could for example be because the individual phases do not seamlessly overlap. It would be interesting to examine an individually constructed phase for the ESRF data too.

4.4.6 Discussion

The enhanced model performs differently for the ESRF and DESY data, both for the phase stability (see section 4.4.1 and fig. 4.8) and the underlying deviations (see section 4.4.4 and fig. 4.11). Several causes could explain this difference in the behaviour of the enhanced model. Firstly, the elements of the enhanced model, the improved material characterisation (see section 4.1) and the novel method of comparing theoretical and experimental energy-time spectra (see section 4.2), yielded relatively smaller improvements for the DESY data than for the ESRF data. This could be one reason why the difference in deviations between the two models in

fig. 4.11 is smaller for the DESY data. Secondly, the different energy resolution (discussed in section 4.3) of the experiments could also cause differences in the performance of the enhanced model. Thirdly, there are noticeable differences between the (experimental) conditions for which the phase stability was calculated, which the enhanced model might respond to differently as well. The shape of the investigated phase, for example, was a π step-like function for the ESRF data (see fig. 4.6a) and a linear function for the DESY data (see fig. 4.7a). For the ESRF stability calculation a combined phase was used, while an (extended) individual phase was chosen for the DESY data (see section 4.4.5). Lastly, even though the same voltage pattern was applied to the piezo for both datasets, albeit with a different polarity and scale, their respective combined phases (compare figs. 4.6b and 4.7b) look completely different. This suggests that the voltage response of the piezo was very different for both experiments. Note that even under these different conditions, the phase stability as shown in fig. 4.8 looks similar for both ESRF and DESY experiments, which reflects well on the repeatability of the phase stability with our experimental setup.

In summary, the phase stability of our phase manipulation scheme was confirmed to be in the few-zeptosecond range. It could not be determined whether this constitutes a fundamental limit to the phase stability of our setup. However, inaccuracies of the data analysis could likely be excluded as a factor that negatively influences the maximum phase stability, as the enhanced data analysis model did not alter results by much. Below-zeptosecond phase stability was achieved for the scaling and step-like noise model, as well as for the case when no motion is applied to the piezo. For the DESY data, the individually reconstructed phase surprisingly performed better than the combined phase.

4.5 Arbitrary phase control

4.5.1 Proof of principle

The results for the proof of principle of arbitrary phase control are shown in fig. 4.14. The phase induced by the voltage pattern that was generated with our linear response approach (bottom right in fig. 4.14) closely resembles the desired phase (middle right in fig. 4.14). Its central features, the phase steps, are clearly seen and occur at the correct times. The magnitude of the phase steps is also approximately the desired 2π . A first example of arbitrary phase control of an X-ray pulse with our linear response method has thus been demonstrated.

The initial voltage pattern in the top panel of fig. 4.14 was arrived at through experience in earlier experiments. It contains a sudden jump followed by a gentler ramp, which is intended to counteract any recoil of the piezo after the sudden jump. The initial phase actually seems to follow the applied initial voltage rather closely, although its sign is flipped. That is, the general shape of the voltage can still be discerned in the phase. For the voltage pattern created for our desired phase (middle panel of fig. 4.14) this is decidedly not the case: it is seemingly random at first sight.

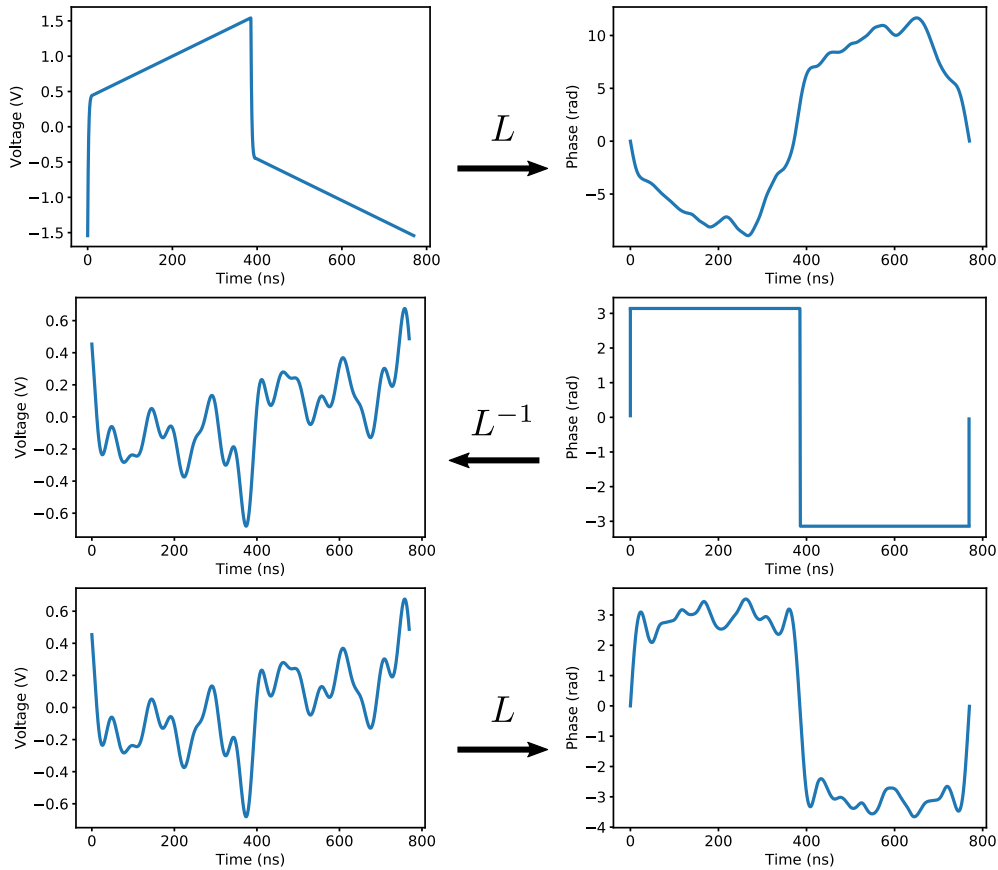


Figure 4.14: Phase control with linear response method. Top panel shows initial voltage and the resulting reconstructed phase. The linear response L thus found is inverted and used in the middle panel to create a voltage pattern corresponding to the desired phase. Bottom panel then shows the phase reconstruction with this voltage pattern.

However, it also contains relatively large jumps at the same positions as the jumps in the desired phase, although the jumps have the opposite polarity again. Additionally, this voltage pattern shares a certain symmetry with both the desired phase and the initial voltage pattern, where the second half of the pattern is (approximately) equal to the negative of the first half. In any case, this voltage pattern manages to produce a phase that looks similar to the desired phase.

If the correct delay of the combined controlled phase in fig. 4.14 (bottom right) is chosen (zero delay in this case), it should approximate a π phase-step and a resonant peak instead of a resonant dip should appear in the spectrum of the transmitted radiation. The individually reconstructed phase for delay zero is shown in fig. 4.15a. It displays a phase-step slightly larger than π and the spectrum calculated with this phase indeed shows resonant enhancement, as shown in fig. 4.15b. Around the resonant frequency the intensity of the incoming field is multiplied by a factor of 3 or more, which is a similar result to the one published by Heeg et al. [15]. With the




Figure 4.15: Transmission spectrum of the α -iron foil (b), calculated with the phase in (a). For (b), the black curve is the spectrum without piezo motion and the blue curve the spectrum with motion. Electronic absorption is included for both spectra.

original phase (top right in fig. 4.14), these results could not be achieved.

Compared to the π step-like phase achieved previously for the ERSF data (see fig. 4.6a), however, the phase in fig. 4.15a has a larger correction in the opposite direction following the phase jump. This might be caused by the more jittery voltage pattern. To prevent this, it could be attempted to create a voltage pattern for a phase-step with a larger magnitude, in case this would scale the jump itself but not the corrections (as seemed to be the case for the ESRF data, see fig. 4.6b). Moreover, a larger magnitude could increase the slope of the phase jump as well, thereby more closely approaching an ideal phase-step, which might increase the resonant intensity even more.

4.5.2 Phase stability

The stability of the controlled phase has also been analysed and is shown in fig. 4.16a. The Allan deviation again reaches the few-zeptosecond scale with the linear noise model, although the minimum attainable stability is around 7 zs , which is higher than for previous results (see fig. 4.8). For small interval times, the stability follows the $1/\sqrt{\tau}$ scaling well. Already for tens of seconds, however, systematic errors appear and the Allan deviation increases. The temporal deviations in fig. 4.16b seem to confirm this; large fluctuations on relatively short time scales are seen, as well as a long-term drift.

The phase stability is obtained for the extended version of the individual phase shown in fig. 4.17. The delay of the combined phase (bottom right of fig. 4.14) that was chosen to acquire this phase means its starting point is close to, but not exactly at, the zero point of the combined phase. Therefore, it is a step-like phase, but its




Figure 4.16: Phase stability (a) for the motion shown in fig. 4.17, obtained with the linear noise model, and the corresponding temporal deviations (b) over the total measurement time. The shaded errors represent error regions. For (a), the diagonal lines indicate $1/\sqrt{\tau}$ scaling.

magnitude is smaller than π . The phase stability in fig. 4.16a is thus not exactly representative of the resonant enhancement scheme. However, it is expected that its stability should be close.

The worse stability for the controlled phase jump, as compared to the results presented in section 4.4.1, could be caused by the different phase investigated here. However, the stability of the similar step-like phase function (fig. 4.6a) for the ESRF data was also found to be better (fig. 4.8a). These two phases were generated by different voltage patterns, though. The ESRF voltage is the same as the initial voltage in the top left panel of fig. 4.14; it is smoother than the voltage for the controlled phase (middle and bottom left of fig. 4.14). It is sensible that the latter, more jittery voltage pattern is more difficult for the piezo to consistently follow. This would lead to less stable motion of the piezo, and thus a less stable phase.

In any case, the phase stability is still in the few-zeptosecond range, which is clearly enough to distinguish the effects of different phases in the experimental data, also over longer periods of time, and to reliably reconstruct these phases. The steering of nuclear quantum dynamics, as was demonstrated in [16], is also possible with our specific controlled phase.

The enhanced model for the data analysis procedure was not applied to the controlled phase and its stability (see section 3.3). Given the results of section 4.4.1, one would not expect that implementing the enhanced model would drastically alter the results and conclusions presented in this section. Another improvement could be extending the linear response function that models the response of the piezo to include higher frequencies, or to replace it by a function that includes non-linear effects. However, it is questionable whether the piezo could respond quickly enough to the resulting voltage patterns. If the voltage patterns become too irregular, this might also negatively influence the phase stability.




Figure 4.17: Individually reconstructed phase with which the phase stability and temporal deviations in fig. 4.16 are calculated.

In future experiments, it would obviously be interesting to attempt the generation of other phases than the π phase jump. This would also allow much more general steering of the nuclear quantum state than has previously been achieved. The results presented here indicate that this is an achievable goal. Note that the generation of phases will always be limited by the dynamics of the piezo; perfectly instantaneous phase-steps will never be possible in practice. In that sense, no completely arbitrary phase control can be achieved.

It would also be valuable to prove that our approach is repeatable. As already mentioned in section 4.4.6, it is possible that for different experimental runs, the same voltage applied to the piezo can result in different phases (compare figs. 4.6b and 4.7b). Therefore, it is necessary to repeat the same steps as presented in fig. 4.14 for each experimental run; it can not be expected that the voltage pattern that produces a desired phase for one run would produce the same phase for a second run. A second limitation of our experimental scheme is that it is event-based. The phase reconstruction and thus the construction of a voltage for a desired phase requires sufficient statistics; this construction can only occur after a certain period of time. Any real-time feedback systems would therefore be difficult to implement. However, successive iterations of the linear response approach could perhaps refine the obtained phase even further.

4.6 Summary and outlook

The modifications made to the data analysis of our phase manipulation scheme were investigated at all steps along the data analysis chain introduced in fig. 1.1. Significantly improved material models were found for the α -iron foil, for which the most important parameter was the direction of the internal magnetization field. For the stainless steel foil, the improvements were smaller and no parameters particularly

stood out. The novel method of comparing theoretical and experimental energy-time spectra also increased the confidence of the fit between them, even though it did not alter the energy calibration.

However, the enhanced model, which combines all changes to the data analysis procedure, did not consistently improve nor worsen the phase reconstruction and phase stability calculations. Thus, it was concluded that the data analysis is not a significant source of systematic error in determining the phase stability. The measurement time of the phase stability experiments did not allow a definite conclusion about the fundamental phase stability limit. Nonetheless, the few-zeptosecond phase stability of our experimental setup was confirmed here and is sufficient for (future) coherent control experiments. For certain noise models, or without piezo motion, the phase stability even reaches below the zeptosecond barrier.

A proof of principle experiment has demonstrated arbitrary phase control. In particular, our linear response approach produced a phase that closely resembled the desired 2π phase step. Spectral enhancement of resonant photons has been shown with this specific controlled phase. Its phase stability was in the few-zeptosecond range, which again meets the requirements for coherent control experiments.

It is expected that our method can indeed be successfully applied to any arbitrary phase. Our phase-controlled X-ray double pulse would then allow more general manipulation of nuclear dynamics than has previously been achieved. For example, nuclear Ramsey spectroscopy and X-ray pump-probe experiments could be made possible. Phase control is also a large step towards the arbitrary shaping of X-ray pulses, which would enable more exotic nuclear quantum optics experiments in the future.

Appendix

A Additional tables and figures

Model	Thickness (μm)	Magnetic field strength (T)	Magnetic field angle φ (rad)	Magnetic field angle θ (rad)	Thickness distribution width (μm)
<i>D</i>	2.66	32.6	1.70	1.37	0.14
<i>E</i>	2.69	32.5	1.70	1.36	-
<i>F</i>	2.60	32.6	1.70	1.35	0.30

	Quadrupole splitting (mm/s)	Isomer shift distribution width (mm/s)	Fitness
<i>D</i>	-	-	-40100
<i>E</i>	0.51	-	-35500
<i>F</i>	0.55	0.05	-30900

Table A.1: Optimal parameter values and fitness for the ESRF data, for several interesting many-parameter models. The models are defined by the parameter values that are included for them.




Figure A.1: Optimized theoretical models of transmission through an α -iron foil at ESRF facilities, for selected many-parameter combinations. Blue: experimental results, orange: model *C*, red: model *D*, green: model *E*, purple: model *F*. See table 4.1 and table A.1 for model descriptions.

Model	Thickness (μm)	Magnetic field strength (T)	Magnetic field angle φ (rad)	Magnetic field angle θ (rad)	Thickness distribution width (μm)
<i>D</i>	2.60	32.4	1.61	1.40	0.18
<i>E</i>	2.70	32.0	1.63	1.76	-
<i>F</i>	2.68	32.0	1.63	1.76	≈ 0
	Quadrupole splitting (mm/s)	Isomer shift distribution width (mm/s)	Fitness		
<i>D</i>	-	-	-44600		
<i>E</i>	1.18	-	-23500		
<i>F</i>	1.18	0.02	-23500		

Table A.2: Optimal parameter values and fitness for the DESY data, for several interesting many-parameter models. The models are defined by the parameter values that are included for them.




Figure A.2: Optimized theoretical models of transmission through an α -iron foil at DESY facilities, for selected many-parameter combinations. Blue: experimental results, orange: model *C*, red: model *D*, green: model *E*, purple: model *F*. See table 4.2 and table A.2 for model descriptions.

Model	Thickness (μm)	Isomer shift distribution width (mm/s)	Quadrupole splitting (mm/s)	Thickness distribution width (μm)	Quadrupole distribution width (mm/s)	Fitness
<i>G</i>	1.26	0.04	-	-	-	-8100
<i>H</i>	1.26	0.03	0.03	≈ 0	-	-8100
<i>I</i>	1.26	0.01	0.03	-	0.07	-8100

Table A.3: Optimal parameter values and fitness for the DESY data, for the same models as in table 4.3. Here, all models are obtained with the conventional optimization algorithm.




Figure A.3: Optimized theoretical models of transmission through a stainless steel foil at DESY facilities, for different parameter combinations and optimization algorithms. Blue: experimental results, orange: model *G*, red: model *H*, green: model *I*. See table A.3 for model descriptions.




Figure A.4: Difference in Poisson score of individual data points, for (a) model *B* versus model *A*, for ESRF data (see table 4.1) and (b) model *E* versus model *C*, for DESY data (see tables 4.2 and A.2). Blue indicates higher Poisson score and thus better agreement with experiment.




Figure A.5: Difference in Poisson score of individual data points, for (a) model H versus model G and (b) model I versus model G . See table 4.3 for model descriptions. Blue indicates higher Poisson score and thus better agreement with experiment.




Figure A.6: Combined phase reconstruction with different optimization strategies, for (a) ESRF data and enhanced model and (b) DESY data and original model. Blue: unoptimized combined phase, orange: optimized with conventional optimization algorithm, green: optimized with evolutionary algorithm. Differences between optimization strategies are difficult to distinguish.

B Bibliography

- [1] C. Brif, R. Chakrabarti, and H. Rabitz. “Control of quantum phenomena: past, present and future”. In: *New Journal of Physics* 12.7 (2010), p. 075008. DOI: [10.1088/1367-2630/12/7/075008](https://doi.org/10.1088/1367-2630/12/7/075008).
- [2] K. P. Heeg. “X-ray quantum optics with Moessbauer nuclei in thin-film cavities”. PhD thesis. 2014.
- [3] B. W. Adams et al. “X-ray quantum optics”. In: *Journal of Modern Optics* 60.1 (2013), pp. 2–21. DOI: [10.1080/09500340.2012.752113](https://doi.org/10.1080/09500340.2012.752113).
- [4] F. Vagizov et al. “Coherent control of the waveforms of recoilless γ -ray photons”. In: *Nature* 508.1 (2014), pp. 80–83. DOI: [10.1038/nature13018](https://doi.org/10.1038/nature13018).
- [5] R. Röhlsberger et al. “Collective Lamb Shift in Single-Photon Superradiance”. In: *Science* 328.5983 (2010), pp. 1248–1251. DOI: [10.1126/science.1187770](https://doi.org/10.1126/science.1187770).
- [6] R. Röhlsberger et al. “Electromagnetically induced transparency with resonant nuclei in a cavity”. In: *Nature* 482.7384 (2012), pp. 199–203.
- [7] K. P. Heeg et al. “Tunable Subluminal Propagation of Narrow-band X-Ray Pulses”. In: *Physical Review Letters* 114.20 (2015), p. 203601. DOI: [10.1103/PhysRevLett.114.203601](https://doi.org/10.1103/PhysRevLett.114.203601).
- [8] P. Heliö et al. “Gamma echo”. In: *Physical Review Letters* 66.15 (1991), pp. 2037–2040. DOI: [10.1103/PhysRevLett.66.2037](https://doi.org/10.1103/PhysRevLett.66.2037).
- [9] G. V. Smirnov et al. “Nuclear Exciton Echo Produced by Ultrasound in Forward Scattering of Synchrotron Radiation”. In: *Physical Review Letters* 77.1 (1996), pp. 183–186.
- [10] R. N. Shakhmuratov, F. Vagizov, and O. Kocharovskaya. “Radiation burst from a single γ -photon field”. In: *Phys. Rev. A* 84.4 (2011), p. 043820. DOI: [10.1103/PhysRevA.84.043820](https://doi.org/10.1103/PhysRevA.84.043820).
- [11] R. N. Shakhmuratov, F. Vagizov, and O. Kocharovskaya. “Single gamma-photon revival from sandwich absorbers”. In: *Phys. Rev. A* 87.1 (2013), p. 013807. DOI: [10.1103/PhysRevA.87.013807](https://doi.org/10.1103/PhysRevA.87.013807).
- [12] Y. V. Shvyd’ko et al. “Storage of nuclear excitation energy through magnetic switching”. In: *Physical Review Letters* 77.15 (1996), pp. 3232–3235. DOI: [10.1103/PhysRevLett.77.3232](https://doi.org/10.1103/PhysRevLett.77.3232).
- [13] X. Zhang et al. “Nuclear quantum memory and time sequencing of a single γ photon”. In: *Physical Review Letters* 123.25 (2019), pp. 1–5. DOI: [10.1103/PhysRevLett.123.250504](https://doi.org/10.1103/PhysRevLett.123.250504).

- [14] S. Goerttler et al. “Time-Resolved sub-Ångström Metrology by Temporal Phase Interferometry near X-Ray Resonances of Nuclei”. In: *Physical Review Letters* 123.15 (2019), p. 153902. DOI: [10.1103/PhysRevLett.123.153902](https://doi.org/10.1103/PhysRevLett.123.153902).
- [15] K. P. Heeg et al. “Spectral narrowing of x-ray pulses for precision spectroscopy with nuclear resonances”. In: *Science* 357.6349 (2017), pp. 375–378. DOI: [10.1126/science.aan3512](https://doi.org/10.1126/science.aan3512).
- [16] K. P. Heeg et al. *Coherent x-ray-optical control of nuclear excitons with zeptosecond phase-stability*. 2020. arXiv: [2003.03755](https://arxiv.org/abs/2003.03755) [quant-ph].
- [17] D. E. Laban et al. “Extreme ultraviolet interferometer using high-order harmonic generation from successive sources”. In: *Physical Review Letters* 109.26 (2012), pp. 1–5. DOI: [10.1103/PhysRevLett.109.263902](https://doi.org/10.1103/PhysRevLett.109.263902).
- [18] Y. Hikosaka et al. “Coherent control in the extreme ultraviolet and attosecond regime by synchrotron radiation”. In: *Nature Communications* 10.1 (2019). DOI: [10.1038/s41467-019-12978-w](https://doi.org/10.1038/s41467-019-12978-w).
- [19] K. C. Prince et al. “Coherent control with a short-wavelength free-electron laser”. In: *Nature Photonics* 10.3 (2016), pp. 176–179. DOI: [10.1038/nphoton.2016.13](https://doi.org/10.1038/nphoton.2016.13).
- [20] D. Iablonskyi et al. “Observation and Control of Laser-Enabled Auger Decay”. In: *Physical Review Letters* 119.7 (2017), p. 073203. DOI: [10.1103/PhysRevLett.119.073203](https://doi.org/10.1103/PhysRevLett.119.073203).
- [21] J. D. Watson and F. H. C. Crick. “Molecular Structure of Nucleic Acids: A Structure for Deoxyribose Nucleic Acid”. In: *Nature* 171.4356 (1953), pp. 737–738.
- [22] R. Röhlsberger. *Nuclear Condensed Matter Physics with Synchrotron Radiation*. Springer, Berlin, Heidelberg, 2004.
- [23] R. A. Dunlap. *The Mössbauer Effect*. Morgan & Claypool Publishers, 2019. DOI: [10.1088/2053-2571/ab0add](https://doi.org/10.1088/2053-2571/ab0add).
- [24] R. L. Mössbauer. “Kernresonanzabsorption von Gammastrahlung in Ir191”. In: *Naturwissenschaften* 45.22 (1958), pp. 538–539. DOI: [10.1007/BF00632050](https://doi.org/10.1007/BF00632050).
- [25] R. L. Mössbauer. “Recoilless Nuclear Resonance Absorption of Gamma Radiation”. In: *Science* 137.3532 (1962), pp. 731–738. DOI: [10.1126/science.137.3532.731](https://doi.org/10.1126/science.137.3532.731).
- [26] G. V. Smirnov. “General properties of nuclear resonant scattering”. In: *Hyperfine Interactions* 123-124.1-4 (1999), pp. 31–77. DOI: [10.1023/A:1017007520099](https://doi.org/10.1023/A:1017007520099).
- [27] U. Van Burck et al. “Enhanced nuclear resonance scattering in dynamical diffraction of gamma rays”. In: *Journal of Physics C: Solid State Physics* 13.24 (1980), pp. 4511–4529. DOI: [10.1088/0022-3719/13/24/014](https://doi.org/10.1088/0022-3719/13/24/014).
- [28] J. H. Eberly. “Emission of one photon in an electric dipole transition of one among N atoms”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 39.15 (2006). DOI: [10.1088/0953-4075/39/15/S07](https://doi.org/10.1088/0953-4075/39/15/S07).

- [29] Y. Kagan. “Theory of coherent phenomena and fundamentals in nuclear resonant scattering”. In: *Hyperfine Interactions* 123-124.1-4 (1999), pp. 83–126. DOI: [10.1023/A:1017059504169](https://doi.org/10.1023/A:1017059504169).
- [30] J. P. Hannon and G. T. Trammell. “Coherent γ -ray optics”. In: *Hyperfine Interactions* 123-124.1-4 (1999), pp. 127–274.
- [31] R. H. Dicke. “Coherence in spontaneous radiation processes”. In: *Physical Review* 93.1 (1954), pp. 99–110. DOI: [10.1103/PhysRev.93.99](https://doi.org/10.1103/PhysRev.93.99).
- [32] B. W. Adams. “Nuclear γ -ray superradiance”. In: *Journal of Modern Optics* 56.18-19 (2009), pp. 1974–1984. DOI: [10.1080/09500340903199921](https://doi.org/10.1080/09500340903199921).
- [33] P. Reiser. “Time domain control of x-ray quantum dynamics”. MA thesis. 2014.
- [34] R. N. Shakhmuratov et al. “Slow γ photon with a doublet structure: Time delay via a transition from destructive to constructive interference of collectively scattered radiation with the incoming photon”. In: *Physical Review A - Atomic, Molecular, and Optical Physics* 80.6 (2009). DOI: [10.1103/PhysRevA.80.063805](https://doi.org/10.1103/PhysRevA.80.063805).
- [35] R. Shakhmuratov et al. “Coherent forward scattering of a single photon wave packet in a resonant medium”. In: *Bulletin of the Russian Academy of Sciences: Physics* 74 (2010). DOI: [10.3103/S1062873810070142](https://doi.org/10.3103/S1062873810070142).
- [36] U. Van Bürck. “Coherent pulse propagation through resonant media”. In: *Hyperfine Interactions* 123-124.1-4 (1999), pp. 483–509. DOI: [10.1023/A:1017080008712](https://doi.org/10.1023/A:1017080008712).
- [37] R. Ruffer and A. I. Chumakov. “Nuclear Resonance Beamline at ESRF”. In: *Hyperfine Interactions* 97.1 (1996), pp. 589–604.
- [38] *PETRA III*. https://www.desy.de/research/facilities__projects/petra_iii/index_eng.html. Accessed: 2021-01-15.
- [39] C. A. Coello Coello. “An Introduction to Evolutionary Algorithms and Their Applications”. In: *Advanced Distributed Systems*. Ed. by F. F. Ramos, V. Larios Rosillo, and H. Unger. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 425–442.
- [40] D. W. Allan. “Statistics of Atomic Frequency Standards”. In: *Proceedings of the IEEE* 54.2 (1966), pp. 221–230.
- [41] B. Fultz. “Mössbauer Spectrometry”. In: *Characterization of Materials*. Ed. by E. Kaufmann. John Wiley, New York, 2011.
- [42] C. E. Violet and D. N. Pipkorn. “Mössbauer line positions and hyperfine interactions in α iron”. In: *Journal of Applied Physics* 42.11 (1971), pp. 4339–4342. DOI: [10.1063/1.1659777](https://doi.org/10.1063/1.1659777).
- [43] R. C. Renot and L. J. Swartzendruber. “Origin of Mössbauer Linewidth in Stainless Steel”. In: *AIP Conference Proceedings* 10.1 (1973), pp. 1350–1354. DOI: [10.1063/1.2946790](https://doi.org/10.1063/1.2946790).

Erklärung:

Ich versichere, dass ich diese Arbeit selbstständig verfasst habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 15.01.2021

PHvdKeeuw
.....