universität wien

MASTERARBEIT

Titel der Masterarbeit

„Development and application of an analysis method for Schottky spectra“

verfasst von

Christoph Klaushofer, BSc

angestrebter akademischer Grad

Master of Science (MSc)

Wien, 2015

Studienkennzahl lt. Studienblatt:	A 066876
Studienrichtung lt. Studienblatt:	Masterstudium Physik
Betreut von:	Hon.-Prof. Dipl.-Phys. Dr. Eberhard Widmann

Contents

Abstract

This thesis deals with the data analysis of an experiment conducted in 2008 at the Experimental Storage Ring (ESR) at GSI in Darmstadt, Germany. Therein iodine ions with only one electron left in their shell were injected into the storage ring and their electron capture decay behavior was studied. Time-resolved Schottky mass spectrometry was used within the storage ring to monitor the stored ions and identify decays.
In the scope of this thesis a new computational analysis method was developed for these Schottky spectra. The emphasis was placed on an enhancement of the sensitivity for stored ions in the Schottky noise and the highest possible precision in the extraction of decay times. The analysis includes a visualization of the Schottky spectra, the identification of traces in subsequent Schottky spectra, corresponding to stored ions, and a maximum likelihood approach to identify decays in these traces. The frequency distribution of these decay times is fitted with an unbinned likelihood algorithm. Several simulations are supplied which address different aspects of the analysis procedure and the decay curves. They were used to optimize parameters and estimate uncertainties.
In previous experiments with similar settings an oscillatory modulation was found in the exponential decay power law. An attempt of this thesis was to confirm these findings with iodine. After the full analysis such an oscillation could not be unambiguously identified with the 2008 iodine data. However the main reason for this might be the lack of sufficient statistics.

Kurzfassung

In dieser Arbeit wird die Datenanalyse eines Experiments behandelt, welches 2008 am Speicherring ESR der GSI in Darmstadt, Deutschland, durchgeführt wurde. Darin wurden Elektroneneinfang-Zerfälle von Ionen mit nur einem einzigen Elektron im Speicherring untersucht. Zeitaufgelöste Schottky Massenspektrometrie wurde verwendet um Informationen über die Ionen zu erhalten und Zerfälle zu beobachten.
Im Rahmen der Masterarbeit wurde ein neues computergestütztes Analyseverfahren für diese Schottky Spektren entwickelt. Der Schwerpunkt lag darin, die Sensibilität für die Ionen im Schottky-Rauschen zu optimieren und Zerfallszeiten möglichst genau festzustellen. Die Analyse besteht aus der Extraktion der Schottky Spektren, Auffindung von Spuren in den zeitlich aufeinander folgenden Schottky Spektren, welche ein Indikator für das Vorhandensein von Ionen sind, und einer maximum likelihood Methode um daraus Zerfallszeiten zu extrahieren. Die Häufigkeitsverteilung dieser Zerfallszeiten wurde mit einem unbinned likelihood Algorithmus gefittet.
Für viele Aspekte der Analyse und dem Verständnis der Zerfallskurven wurden Simulationen durchgeführt, welche sich mit verschiedenen Aspekten der Analyse und der Zerfalls-Kurven beschäftigen. Das Hauptaugenmerk lag dabei auf der Optimierung von Parametern und der Fehlerabschätzung.
In früheren Experimenten mit ähnlichen Rahmenbedingungen fand man eine Schwingung im exponentiellen Zerfallsgesetz. In der Masterarbeit wurde versucht, diesen Effekt im Zusammenhang mit dem Jod-Experiment zu untersuchen. Nach der vollen Analyse konnte eine solche Oszillation nicht eindeutig in den Jod-Daten von 2008 festgestellt werden. Allerdings könnte der Grund dafür der geringe Umfang des Datensatzes sein.

Chapter 1

Introduction

Nowadays many experiments are performed to confirm some small part of a highly specialized theory. Thus the possible outcomes are already anticipated and once the experiment is concluded the theoretical part is confirmed or refuted. This holds especially for high-energy physics due to the high level of mechanization to measure quantities with maximum accuracy in a respectively small region of interest.
However cases occur when evidence arises from an experiment (looking for something completely different) that points towards something new, something that has not been observed before, something no theory has predicted so far. Although many of these findings can be attributed to instrumental errors or misinterpretations some experiments survive the critical assessment of the experts and the search for new physics starts.

1.1 The GSI experiments

In 2006 a group of scientists at the Gesellschaft für SchwerIonenforschung (GSI) in Darmstadt were investigating the decays of highly ionized ions in a storage ring. They observed the Electron-Capture (EC) decays of hydrogen- and helium-like ions (atoms stripped of all but one or two electrons respectively). Investigating the decay curve of these ions, they found that the decay constant does not only vary but is significantly larger for the hydrogen-like version of the ion. This is primarily unexpected since a higher abundance of electrons should lead to a higher probability of their capture. Even though this was not understood at first, it could later be explained by the introduction of new decay modes [1]. In addition to this effect they found something more in the data which came to be known as the "GSI puzzle": when observing the decay curve they found a deviation from the exponential power law manifesting in an oscillation. The experiment was performed with $^{140}Pr^{58+}$ and $^{142}Pm^{60+}$ and in both cases they found an oscillatory behavior with periods of $\sim$7 seconds and amplitudes of $\sim$20% [2].

The experiment was repeated in 2008 with $^{122}I^{52+}$ as an attempt to confirm these findings in a slightly different mass-to-charge region. The results of this experiment have not been published yet.
Up until then the analysis method, more specifically the identifying of the decay times, consisted of looking at the raw data in a case-by-case fashion and attribute a decay time to each case individually. In this thesis I developed a technique to analyze the data set of the iodine experiment automatically and to identify decay times by certain fixed criteria for all cases. The main motivation for this was to reduce possible effects of human bias and create an objective analysis tool for past and future experiments.

1.2 Overview

The thesis is organized in five larger parts, each with a dedicated chapter.
In the first part I provide an introduction into the experiment as well as the theoretical aspects. I start with the explanation of Schottky mass spectrometry and discuss it in detail. The attributes of the examined iodine isotope are presented and compared to the candidates of previous experiments. The results of these experiments are discussed in detail since they were the main motivation to carry out this thesis. A short overview of possible theoretical explanations is provided with references therein for interested readers. At the end of this part I give a short introduction to the visual data analysis which influenced the development of the automated data analysis significantly.
The second and third part contain the main work conducted during the time of my masters thesis: The development of the automated analysis technique. I explain in detail the several steps which are necessary to get from the raw data to a decay curve. These include the conversion of raw data into easily readable files, the identification of parent and daughter ions, the extraction of their information with minimal losses and the retrieval of a decay time. The latter is the most important and most difficult part of the analysis process. For this a new likelihood method has been developed. The unbinned likelihood fit used to retrieve fit parameters of an exponential and modulated exponential fit is described in the last section of this part.
The fourth part treats several simulations. The methods used for these simulations are briefly discussed. Several different simulations are explained on the one hand to optimize parameters of the method and on the other hand to argue the reliability of the results.
In the last part I present the results of the data analysis concerning the iodine experiment. The decay curves of different sets selected by different criteria are discussed. The results of the extracted decay times of the automated and visual data analysis are compared. At the end I provide a short

summary and discuss further possibilities to improve the method.

1.3 Acknowledgements

There are several people without whom this thesis would not have been possible.
I want to thank my supervising professor Eberhard Widmann for giving me the opportunity of realizing my masters thesis at the Stefan Meyer Institute. During the whole process I worked together with Paul Bühler who taught me many things about data analysis. I want to thank him for his continuous efforts during the development of the analysis technique as well as his input while writing this thesis.
Fritz Bosch gave me precious insights into the experiment at GSI and helped me understand the visual data analysis.
Furthermore I want to thank all of my colleagues and fellow students for various discussions and helpful motivation.

Chapter 2

Theory

This chapter is dedicated to some fundamentals of this thesis. The mechanisms of Schottky mass spectrometry are discussed and an introduction into the experimental setup at GSI is given. The decay properties of the examined iodine isotope are presented and compared with the other candidates examined in different experiments. These experiments are presented in detail with a focus on their results. A short overview of possible theoretical explanations is provided. At the end I explain the mechanisms used in the visual data analysis.

2.1 Schottky mass spectrometry

Schottky Mass Spectrometry (SMS) is a non-destructive method to determine the masses of single particles in a storage-ring. It is based on the measurement of the revolution frequency of charged ions in a magnetic field. The particle in the storage ring induces mirror charges on an electrostatic pick-up device installed inside the beam tube at every revolution. The more often the particle passes this pick-up and the higher its charge the more signals will be induced per time and thus the Signal-to-Noise (S/N) ratio will improve. Because of this the highest precision can be achieved with highly charged relativistic particles. The frequency spectrum of this beam noise (Schottky noise) shows peaks at every harmonic f of the mean revolution frequency f_{rev} of the particle.

$$f = h \cdot f_{rev} \tag{2.1}$$

$$f_{rev} = \frac{v}{C} \tag{2.2}$$

f ... h'th harmonic
h ... Integer
f_{rev} ... Revolution frequency
v ... Mean particle velocity
C ... Mean orbit circumference

By using electron-cooling in the storage ring, all of the beam components will be cooled to the same mean particle velocity v with a small velocity-spread. Due to different mass-to-charge ratios the particles will be bent differently by the dipole magnets and thus move in slightly different orbits with different circumferences C. Hence the revolution frequencies will be different for different masses and the particles can be distinguished. Up to first order the relative frequency difference of two different ion species is given by the relative difference of their mass-to-charge ratios:

$$\frac{\Delta f_{rev}}{f_{rev}} \simeq -\alpha_p \cdot \frac{\Delta(m/q)}{(m/q)} \tag{2.3}$$

α_p is the momentum compaction factor and can be considered constant in most cases (see [3] and [4]) if the velocity spread Δv is sufficiently small.
SMS is based on the measurement of frequency differences rather than absolute values. The relative widths of the peaks in a Schottky spectrum depend on the longitudinal momentum spread of the corresponding ion. The momentum spread is determined by the equilibrium between cooling and heating effects. The main cooling process used is electron-cooling whereas one of the main heating effects is intra-beam scattering which can be reduced to a minimum by reducing the number of induced particles. The contribution of this effect already vanishes at particle numbers of < 1000. With lower particle numbers the peak width mainly depends on the stability of the magnet power supplies [5].
The integrated noise power is proportional to $N \cdot q^2 \cdot f_{rev}^2$ where N is the number of particles and q the charge. From this correlation follows that the highest sensitivity can be achieved with highly charged ions at relativistic energies and the integrated noise power can be used as a measure of the number of stored ions.
In principle SMS is sensitive to single stored ions up to relatively high numbers. However, in the experiment considered here, due to large variances in the determination of the peak areas the actual counting of ions by taking the integrated peak area value is only feasible up to three ions. So for a maximum of three injected ions the number of particles in the ring can be determined unambiguously at any time.
The decay of a particle in the storage ring is associated with a change in its mass-to-charge-ratio. This change corresponds to a change in its revolution frequency. Hence when observing the disappearance of a particle at one frequency and the nearly simultaneous appearance of a particle at another

frequency one can conclude that a decay has occurred. If the disappearance of a particle can not be correlated with the reappearance of its decay product the cause can be either natural losses such as charge exchange with the residual gas or particles of the electron-cooling, or the daughter nucleus of the decay is not within the range of the observation window. The two effects can in principle not be distinguished.
The times of dis- and re-appearance of parent and daughter nucleus of a decay will not be exactly the same because due to the ejection of a decay particle the daughter nucleus will recoil and gain an additional momentum to its equilibrium velocity (determined by the electron-cooling). Hence the time difference of dis- and re-appearance will be determined by the duration of the cooling process of the daughter nucleus to its final frequency. Since only the velocity component in beam-direction plays a role in this process, the cooling times will be randomly distributed within the confines of the two extremes, namely an ejection in beam-direction or one perpendicular to it. To eliminate the systematic uncertainty of cooling times the most precise way of identifying a decay time would be to take the disappearance time of the parent nucleus. However in most experiments there are very many parent ions so the distinction of a single disappearance proves difficult while there are only very few decay products. Thus it is easier to extract the decay time by looking at the appearance of daughter nuclei although this introduces a small error due to the cooling time.
When examining EC-decays (the decay scheme is shown in eq. (2.4)) the charge state of parent and daughter ion is the same therefore the only relevant difference is their mass. The frequency difference Δf_{rev} is small ($\Delta f_{rev}/f_{rev} \sim 10^{-4}$) and it is sufficient to only look at a small frequency region improving the sensitivity. Also the mass difference of the two ions is very well known so Δf_{rev} can be predicted and used for decay identification.

$$p^+ + e^- \longrightarrow n + \nu_e \quad || \quad {}^{A}_{Z}I^{x+}_{1e^-} \longrightarrow {}^{A}_{Z-1}P^{x+}_{0e^-} \tag{2.4}$$

If a β^+ decay occurs, the atomic charge changes by one, shifting the revolution frequency by a large amount. This shift is usually much larger than the observation band and thus will only manifest in the disappearance of a parent ion. Because it can not be observed elsewise it can not be distinguished from other possible losses. However the latter effects can be estimated by examining the losses in the stable daughter ions.

2.2 The ion of interest

The atom examined in this thesis is an isotope of iodine, namely $^{122}_{53}$I. The neutral atom decays via EC and β^+ to the stable atom tellurium $^{122}_{52}$Te. The decay scheme is shown in Figure 2.1 lower part. The decay ratio of EC/β^+ is calculated in [6] and therein the half-life is shown to be approximately the

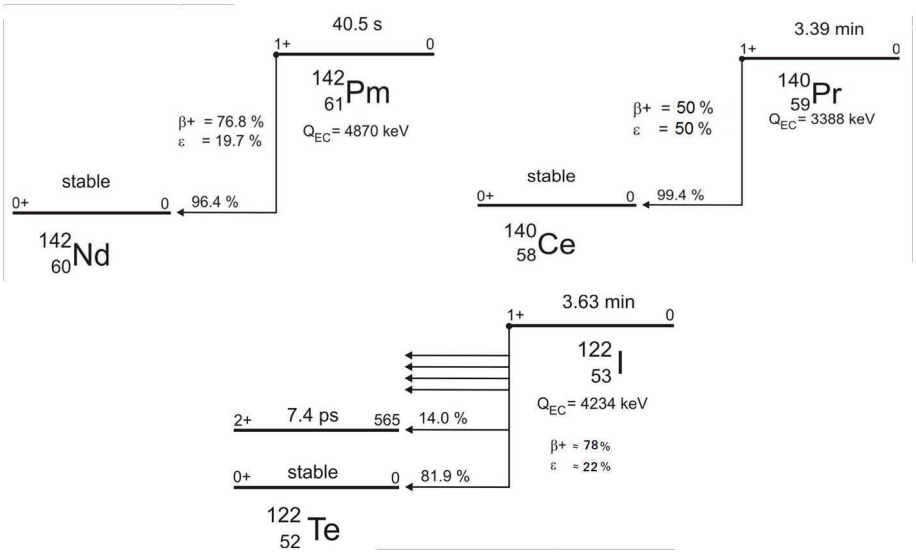

Figure 2.1: Decay schemes of the neutral atom $^{122}_{53}$I [7] and the neutral atoms $^{140}_{59}$Pr and $^{142}_{61}$Pm examined in [2] and [8].

same in the hydrogen-like $^{122}_{53}I^{52+}$ state.
The atoms which were examined in previous experiments were Praesodynium $^{140}_{59}$Pr and Promethium $^{142}_{61}$Pm (see Figure 2.1 upper part). Crucial criteria for picking these nuclei were the following:

- Occurrence of EC-decays
- Half-lifes large enough for the parent ions to make it into the storage ring but small enough to observe a decay in a reasonable measurement time
- Predominantly groundstate-to-groundstate decays
- Producibility
- Stable products

Meeting these criteria restricts the range of possible parent nuclei significantly. One of the main reasons to look at iodine rather than its well studied predecessors was to study a different mass-region (i.e. mass-to-charge region) and examine whether this puzzling effect of oscillatory decay curves remains, and whether its parameters change. However iodine is not a perfect candidate since it only decays to the ground-state of tellurium in 82% of EC decays. This decay mode is preferred because in ground-state-to-ground-state EC decays the neutrino carries in the rest frame of the parent ion the exact same amount of recoil momentum as does the daughter nucleus so by examining the daughter nucleus one can determine attributes of the neutrino.
Since $^{122}_{53}I^{52+}$ has a longer half-life and a smaller charge state than the other

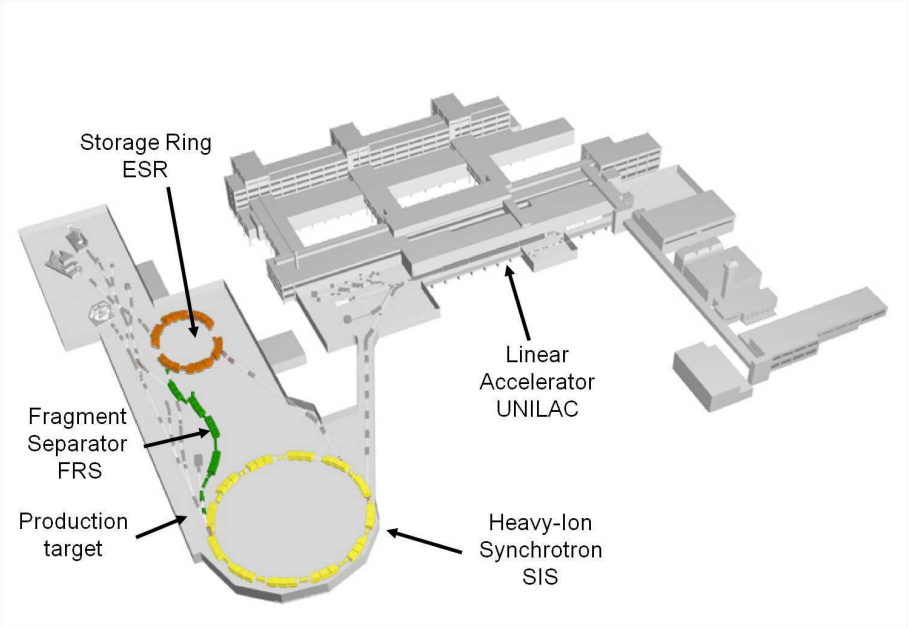

Figure 2.2: Experimental setup at GSI Darmstadt [10]. Samarium ions are accelerated first by the linear accelerator, then by the Synchrotron to energies of 602 MeV/u and hit a beryllium production target. The products are separated in-flight by the fragment separator and injected into the storage ring. Particles can be stored for long times while SMS is used to monitor them.

two ions it produces a weaker signal strength and thus the identification in Schottky spectra is more difficult. In fact the induced Schottky signals will be 1-(52/60) = 25 % smaller. However, due to the differences to the other two ions possibly more contributing effects could be observed in this ion.
A detailed discussion of the results of the ^{140}Pr and ^{142}Pm experiments will be given in Section 2.4.

2.3 Experimental setup

The experiment is set up at the GSI Darmstadt in Germany and was conducted in august 2008. A sketch of the full setup is shown in Figure 2.2. I will only give a short summary of the experiment since I was not involved myself. A more detailed description can be found in [9] and references therein.
The primary ions samarium ^{152}Sm are accelerated in the UNIversial Linear ACcelerator UNILAC and the ring accelerator SchwerIonenSynchrotron SIS18 to energies of 602 MeV per nucleon. They then hit a beryllium target of 2.5g/cm^2 thickness producing various species of ions in different charge

states. The extraction from the SIS18 takes less than 1 μs giving a very well defined time of creation for the ions. The FRagment-Seperator FRS [11] has the task to separate all these collision products in flight and can be adjusted very accurately to deliver any desired species into the Experimental Storage Ring ESR [12]. Since most of the desired ions are highly charged they allow an efficient electromagnetic separation. To further separate isobaric contaminants energy-degrader systems are used which are sensitive to Z^2. This is called the $B\rho - \Delta E - B\rho$ separation method [13]. The FRS can output either multiple elements in different charge states to get as much information out of single shots as possible or inject a single ion species to get the highest precision possible and suppress background effects. For this experiment mainly the latter functionality has been used. This possibility of producing and storing single ions is nearly unique in the particle physics sector.

After selection by the FRS the particles of interest are injected into the Experimental Storage Ring ESR at a high velocity spread. In the ESR Schottky mass spectrometry is performed as described in Section 2.1. In the first five seconds after the injection stochastic cooling is applied to cool down the particles to 400 MeV/u and reduce the velocity spread. During and after the stochastic cooling a continuous electron-cooling is used to reduce the velocity spread and keep it small such that the average particle velocity stays constant. This reduces the velocity spread to $\Delta v/v \sim 5 \times 10^{-7}$ [2].

The storage time can in principle be chosen arbitrarily long and was set to 64 seconds in this experiment, adjusted to the half-life of the observed ions. This way many decays could be observed while a reasonable number of measurement cycles were still possible. After these 64 seconds of measurement the remaining ions are extracted by closing a mechanical shutter in the Ring and after the reopening a new measurement cycle is started.

Since a nuclear decay is a statistical process there are measurement cycles with no observed decay at all. Taking the attributes of the neutral atom $^{122}_{53}$I every injected ion has an 18% chance to decay (via β^+ or EC) within the measurement window. The number of injected parent ions is subject to statistics itself and was set to an average of 10 injected ions. The two constraints leading to this number from below and above are that we want to observe as many decays as possible (hence as many injected ions as possible) but are very limited in our data analysis when it comes to multiple decays within one measurement cycle. This goes as far as that we can only unambiguously use the decay time of a measurement if we are sure that there was no second decay in it (this will be further discussed in Section 3.6). Because of these constraints there will be many measurement cycles without observed decays, so many thousand measurement cycles are necessary to observe enough decays to investigate a decay curve.

The signal from the capacitive pickup device in the storage ring is Fast-Fourrier-Transformed (FFT) online and only the frequency region of interest

is extracted and stored in an IQ-data-format. This format will be further explained in the following section.
In this step of frequency analysis a time-binning has to be introduced:
We start with a continuous signal, the Schottky noise. To gain frequency information one needs to select a certain time region and perform a Fourier transformation of it. The larger this time region is, the clearer the frequency signal, the smaller the time region, the better the time resolution. In this experiment the online binning and FFT could not be avoided although in later experiments (see [8]) the original Schottky noise could be stored for later offline analysis. The central frequency is aimed at the 42nd harmonic of the revolution frequency. The experimentalists chose a time binning of 64 ms length and a frequency span of 5 kHz around the central frequency of 60 MHz. This yields 1000 time bins of 64 ms and 641 frequency bins of ~ 8 Hz each. So a complete file consists of a 1000×641 matrix containing the full information of Fourier amplitudes of one measurement cycle. Every measurement cycle yields one file containing the full information which will be used for offline analysis as discussed in the next chapter.
The cooling processes can actually be monitored by SMS since the uncooled particles already induce signals on the capacitive pickup device directly after their injection into the ring. An example of the Fourier amplitudes of subsequent Schottky spectra is shown in Figure 2.3. A distinct signal of the parent ions appears after ~ 5 seconds of stochastic cooling and forms a trace in the subsequent Schottky spectra. A second, very faint trace appears after some time, which corresponds to one daughter ion. The traces are straight since the velocity is kept constant by the continuous electron-cooling. The frequency difference of the two traces Δf_{rev} is fixed by the Q-value of the decay.
The experiment lasted for two and a half weeks and a total of 16477 analyzable files were created for offline data analysis.

2.3.1 The IQ data format

The IQ data format is a means to store phase and amplitude information widely used in communication technology. Simply put it is a transformation of the angle and amplitude of a sine wave into Cartesian coordinates I and Q according to eq. (2.5-2.7):

$$I = A \cdot \cos(\phi) \tag{2.5}$$

$$Q = A \cdot \sin(\phi) \tag{2.6}$$

$$S = A \cdot \cos(2\pi f_c t + \phi) = I \cdot \cos(2\pi f_c t) - Q \cdot \sin(2\pi f_c t) \tag{2.7}$$

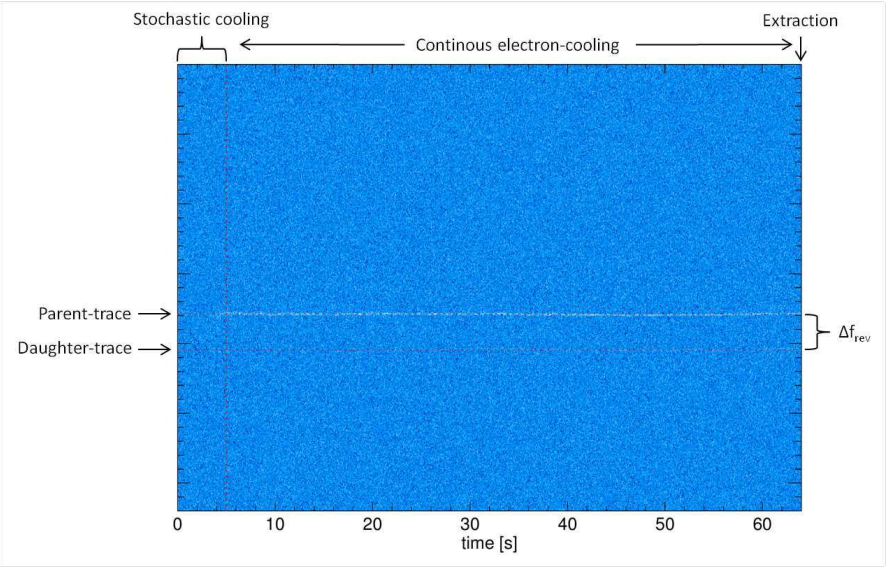

Figure 2.3: The different phases during one measurement cycle, as measured by SMS: Parent ions are injected into the storage ring and cooled by stochastic and electron-cooling to their final frequency. A decay can be identified by the appearance of a second trace at a fixed Δf_{rev}. After 64 seconds of measurement the remaining particles are extracted by closing a mechanical shutter.

A ... Signal amplitude
ϕ ... Signal phase
S ... Signal
f_c ... Signal frequency
t ... time

The right term of the signal shown in eq. (2.7) can be processed very well by hardware since the difference of a sine and a cosine is a simple phase shift by 90° which can be easily accomplished.
In our case we get a Fourier spectrum every 64 ms with 641 frequency bins. So f_c and t are known and the IQ-pair contain the full information. For the data analysis the signal amplitude has to be obtained by a retransformation of I and Q according to eq. (2.8). The resulting amplitudes are stored in [dBm] which is the unit for decibel milliwatt.

$$A = \sqrt{I^2 + Q^2} \tag{2.8}$$

2.4 Previous experiments

Most of the work in this thesis is based on or is a continuation of two experiments performed and analyzed in the past. The results of these are

published in [2] and [8] so most of the plots, results and conclusions are directly derived from these papers. This section can be skipped if the reader is familiar with the two papers.

2.4.1 Experiment in 2006

The first experiment to detect an oscillatory time modulation in exponential decay was performed in 2006 and published in 2008 [2]. Herein single orbital EC decays in hydrogen-like and helium-like (ions stripped of all electrons but two) ^{140}Pr ions were measured for the first time. It was observed in a previous experiment that the EC decay rate of ions with only one electron exceeds the one of ions with two electrons by $\sim 50\%$ [1]. Although it was surprising at first, this result could later be explained theoretically by standard weak decay theory [14].
The decay of hydrogen-like ^{142}Pm was studied more closely and examined together with the hydrogen-like ^{140}Pr results. The decay schemes of both neutral atoms are well known and have already been presented in Section 2.2. Both nuclei can either decay via the two-body EC-decay or the three-body β^+-decay. The product in the case of both ions is a stable daughter nucleus and the main decay branch is a ground state-to-ground state transition.
There were four independent runs in the experiment differing slightly in their parameters as shown in tab. 2.1.

Table 2.1: Different experimental runs of the experiment performed in 2006. The parameters given here are from left to right the run number, the ion of interest selected in the FRS, the energy of the primary samarium beam, the thickness of the beryllium production target and the number of injections in each run. This table was retrieved from [2]. In runs 3 and 4 an additional niobium foil was inserted in the FRS to separate the desired ions from isobaric contaminations.

Run	Ion	E (^{152}Sm) [MeV/u]	L (^{9}Be) [mg/cm^2]	#inj
1	^{140}Pr^{58+}	507.8	1032	453
2	^{140}Pr^{58+}	507.8	1032	842
3	^{140}Pr^{58+}	601.1	2513	5807
4	^{142}Pm^{60+}	607.4	2513	7011

Unlike in later experiments only two parent ions were on average injected into the ESR at each injection at an energy of 400 MeV per nucleon. The measurement window was set around 2 MHz corresponding to the 30th harmonic of the revolution frequency. The frequency difference of parent and daughter ion (corresponding to the mass difference i.e. the Q-value) amounted to $\sim$ 270 Hz and $\sim$ 310 Hz for ^{140}Pr and ^{142}Pm respectively at

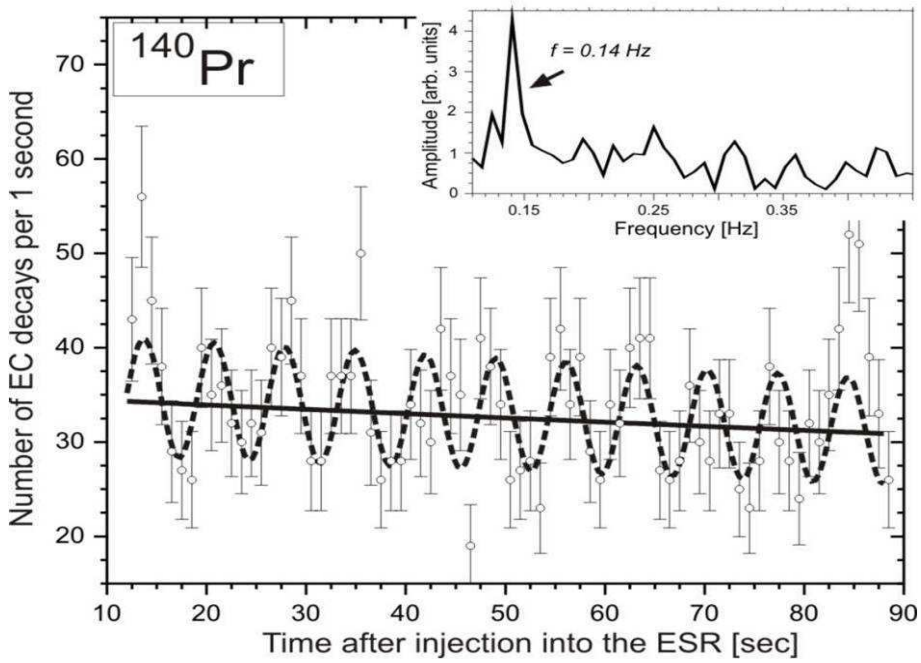

Figure 2.4: Results of the experiment performed in 2006 with $^{140}Pr^{58+}$. The data were fitted with an exponential decay function (eq. 2.9) and a modulated exponential decay function (eq. 2.10). Upper right corner: Fast Fourrier transform of the data, a peak at a frequency of 0.14 Hz is visible. Plot taken from [2].

the 30th harmonic. A decay was only acknowledged if a parent ion disappeared and a daughter ion appeared in a correlated way. Depending on the recoil energy and direction of the daughter ion its appearance is delayed with respect to the disappearance of the parent ion due to different cooling times.

The data analysis was performed visually as described in Section 2.6. Since the appearance times of the daughter nuclei could be determined more precisely than the disappearances of the parent ions only the former times were taken into account (although this includes the small error due to varying cooling times). The decay times of the three runs of ^{140}Pr were combined and the results are illustrated in Figure 2.4. The results of the ^{142}Pm run are shown in Figure 2.5.

The data were fitted with two functions: a standard exponential decay function (defined in eq. 2.9) and a modulated version of the exponential decay (defined in eq. 2.10):

$$\frac{dN_{EC}(t)}{dt} = N(0) \cdot e^{-\lambda t} \cdot \lambda_{EC} \tag{2.9}$$

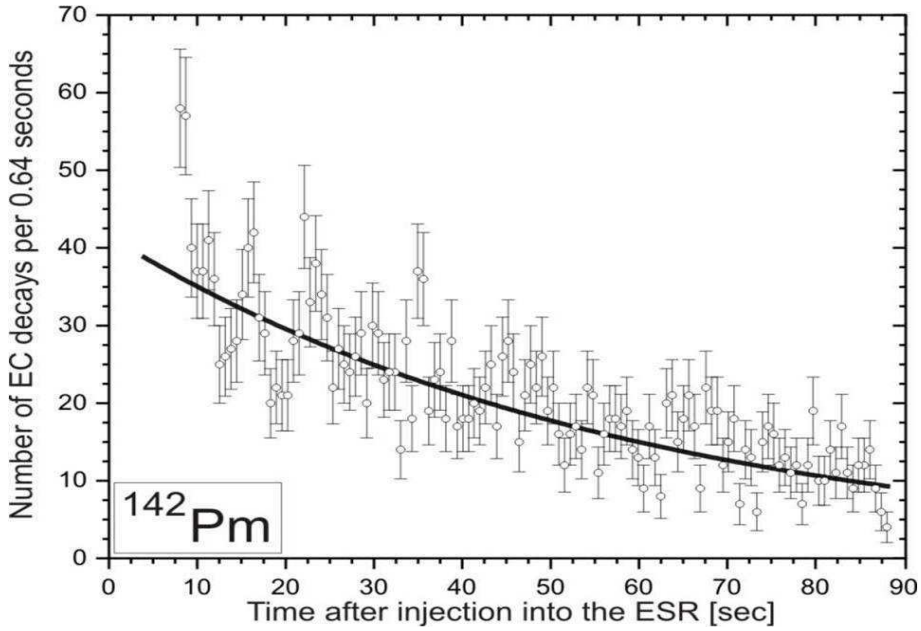

Figure 2.5: Results of the experiment performed in 2006 with $^{142}Pm^{60+}$. The data were fitted with an exponential decay function (eq. 2.9). No oscillation is visible in the full time window. Plot taken from [2].

$$\lambda = \lambda_{EC} + \lambda_{\beta} + \lambda_{loss}$$
$$\frac{dN_{EC}(t)}{dt} = N(0) \cdot e^{-\lambda t} \cdot \widetilde{\lambda_{EC}}(t) \tag{2.10}$$
$$\widetilde{\lambda_{EC}}(t) = \lambda_{EC} \cdot [1 + a \cdot \cos(\omega t + \phi)]$$

N_{EC}	...	Number of parent ions at time t which can possibly decay via EC
$N(0)$	...	Number of parent ions at $t = 0$, the time of injection
λ_*	...	Decay constant of respective decay channel
a	...	Amplitude of modulation
ω	...	Angular frequency of modulation
ϕ	...	Phase of modulation

Although the scattering of the data is relatively large, one can already observe without fitting some kind of periodic modulation in Figure 2.4. The fit of a modulated exponential decay (eq. 2.10) confirms this impression. In the Fast-Fourier Transformation(FFT), as shown in the upper right corner, a clear peak is visible at a distinct frequency of 0.14 Hz.
In the ^{142}Pm data shown in Figure 2.5 an oscillation can be observed in the first half of the time window but due to its shorter half-life it is damped very quickly. Hence only the first 33 seconds after the injection were fitted with eq. 2.10 and the results again point towards an oscillation as shown in Figure 2.6 with the same frequency as found in the Pr data.

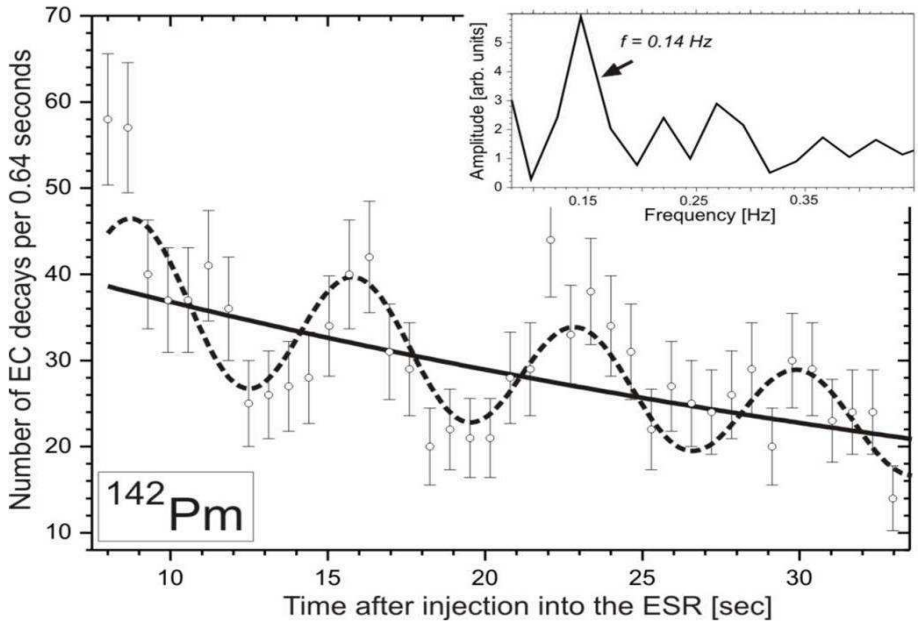

Figure 2.6: The first 33 seconds of the $^{142}Pm^{60+}$ data. A fit of a modulated exponential decay function (eq. 2.10) was added. Upper right corner: Fast Fourrier transform of the data, a peak is visible at a frequency of 0.14 Hz. Plot taken from [2].

All of the fit parameters are summarized in tab. 2.2. The ^{142}Pm data yield a significant difference whether the fit is performed on the whole data set or solely the first 33 seconds. The decay constants differ by 1.5 standard deviations and the oscillation disappears in the full time window. However this could be linked to a dampening effect due to the short half-life. The amplitudes a and angular frequencies ω of the two data sets of ^{140}Pr and ^{142}Pm agree astonishingly well. The phases ϕ don't overlap, however this fit parameter is strongly correlated with the angular frequency ω and has large error margins.

To summarize: in both hydrogen-like ^{140}Pr and ^{142}Pm a time modulation of the expected exponential decrease of the number of EC-decays was found. This modulation coincides from both experiments at a period of $\sim$ 7 seconds and an amplitude of $\sim$ 20%. Although a technical artifact could not be excluded at the time it was already anticipated back then that new physics might be involved. These experiments gave rise to follow-ups for further investigation of this puzzling effect of oscillatory behavior.

In a later offline analysis it was tried to analyse the data for β^+ decays. These manifest by the disappearance of a parent $^{142}Pm^{60+}$ ion without the appearance of a daughter $^{142}Nd^{60+}$ nucleus. However this was only possible for cases where at most two parent ions were injected into the ESR.

Table 2.2: Fit parameters of the plots shown in Figure 2.4, 2.5 and 2.6. The first two columsn refer to the fit equation and the figure in which the respective fit is illustrated. The first two rows correspond to the ^{140}Pr data and the last three rows to the ^{142}Pm data. Table taken from [2].

Eq.	Figure	$\lambda[s^{-1}]$	a	$\omega\,[s^{-1}]$	ϕ [rad]
2.9	2.4	0.0014(10)	-	-	-
2.10	2.4	0.0015(10)	0.18(3)	0.890(11)	0.4(4)
2.9	2.5	0.0170(9)	-	-	-
2.9	2.6	0.0240(42)	-	-	-
2.10	2.6	0.0224(42)	0.23(4)	0.885(31)	-1.6(5)

2.4.2 Experiment in 2010

Another experiment was performed in 2010. Although the data was collected later than in the iodine experiment (2008) the results have already been published in 2013 [8] due to less ambiguity in the data analysis. The experiment reinvestigated the β^+ and EC decays of ^{142}Pm^{60+}. For this purpose a new detector device had been developed, namely a 245 MHz resonater cavity. With this the S/N-ratio was improved by two orders of magnitude compared to the old capacitive pickup device. Due to this improvement it was even possible to observe the cooling traces of recoiling daughter nuclei. Thus fixing decay times with higher precision was possible compared to previous experiments where only the times could be determined when daughter nuclei arrived on their final frequency. An example of the drastic effects of this improved detector on the visibility of traces in the Schottky noise can be seen in Figure 2.7.
Also due to the better detection efficiency the time binning could be reduced to 32ms (compared to the 64ms binning in the previous experiments). The measurement window was set around the 124th harmonic of the parent ions at $\sim$ 1.98 MHz and has a frequency binning of 31.25 Hz. The frequency difference in an EC decay in this regime is 1560 Hz corresponding to a Q-value of 4830 keV.
From the observation of varying cooling times, distributions of cooling length and frequency change can be derived. With the according calculation one can derive the in flight component of the direction in which the neutrino was emitted in the two-body EC decay.
The production of the ^{142}Pm^{60+} ions was performed in the same way as in the 2006 experiment. The ions were injected by a kicker signal and extracted after 64 s of measurement time by another kicker signal which simultaneously injected the new ions of the next measurement. However the timing

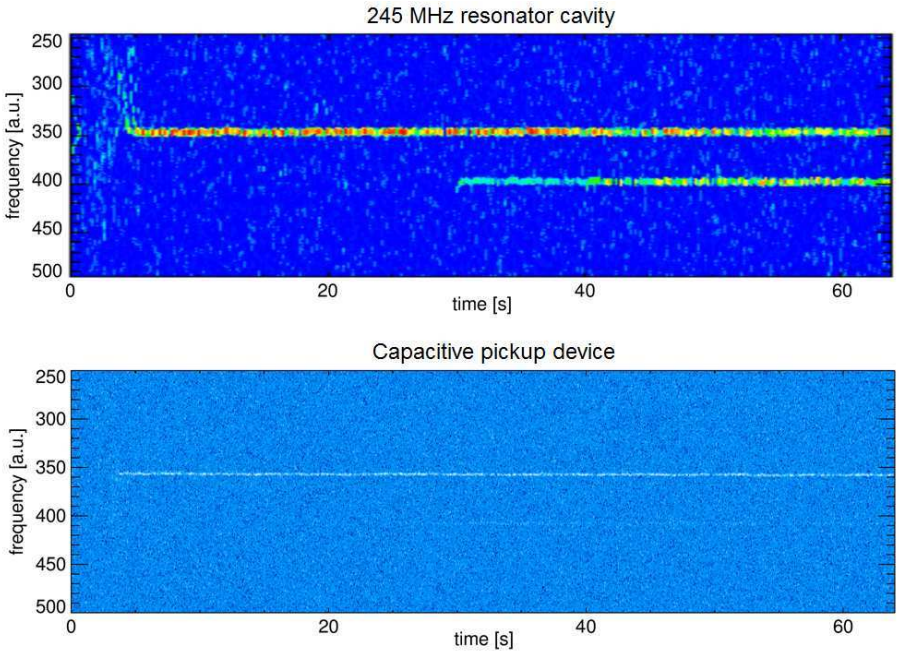

Figure 2.7: Comparison of the two detectors. Upper plot: Subsequent Schottky-spectra of the 245 resonator cavity show cooling traces and a very well defined daughter trace. Lower plot: Subsequent Schottky-spectra of the capacitive pickup device show no cooling traces and the daughter trace is barely visible. Both examples have a similar decay time around 30 seconds.

of this signal is crucial since the lack of an extraction would lead to ions remaining in the storage ring from previous measurements. This would obscure the starting time needed for a clear determination of comparable decay times. Unfortunately due to resets of this parameter after each failure in the machine hardware e.g. the ESR magnets the monitoring of this injection/extraction kicker signal was lost.
In total 17460 injection cycles have been performed with an average of four parent ions. 8665 EC decays were found in this sample. The full set of these decays does not show any significant oscillatory behavior in its decay curve. A subset of 7125 consecutive injections could be extracted where no hardware-failures (and thus kicker-signal resets) are documented. The 3594 EC decays that were found in these injections were used for the data analysis. However even in this subset it is not 100% clear whether the kicker-timing was set correctly so it could not be assured whether all stored ions were removed after each cycle.
Both detectors, the capacitive pickup device and the 245 MHz resonator were used in this experiment. The data of both devices were analyzed visually as in previous experiments. In addition an automated analysis was

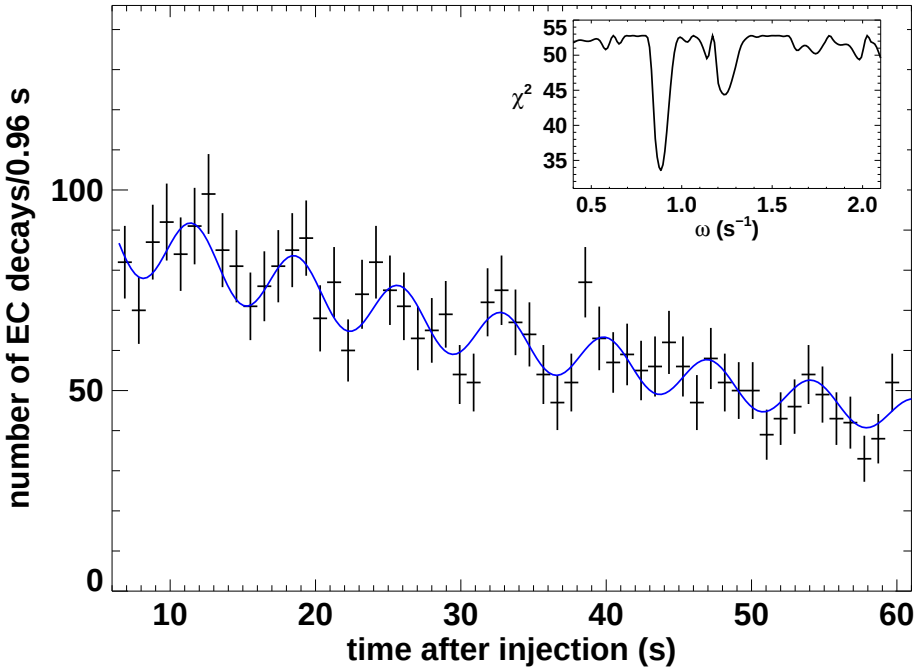

Figure 2.8: Results of the experiment performed in 2010 with $^{142}Pm^{60+}$, measured by the 245 MHz resonator. The data were fitted with an exponential decay function (eq. 2.9) and a modulated exponential decay function (eq. 2.10). Upper right corner: Distribution of the χ^2-values for varying ω-values in the modulated exponential decay fit. Plot taken from [8].

developed for the resonator spectra by an independent group.
Only the time range from 6 s to 60 s was used for the analysis because the first few seconds do not contain any usable data due to ongoing cooling processes and in the last few seconds it is very difficult to identify decays due to lack of on-trace information afterwards. This yields an effective analysis window of 54 seconds.
Again the data were fitted with the two models of a purely exponential decay and a modulated exponential decay as defined in the previous section in eq. 2.9 and 2.10. The results of these fits are shown in Figure 2.8 for the 245 MHz resonator and in Figure 2.9 for the old capacitive pickup device. For the latter only a smaller sample of 3098 EC decays could be used because measurements yielding more than one EC decay needed to be excluded. More than one decay time could not be identified unambiguously. A detection time difference arises between the two detectors corresponding to the cooling time of the recoiling daughter nuclei.
The fit results are shown in tab. 2.3.
Again a clear oscillation is visible. While the frequency is nearly the same as in the previous experiment with ^{142}Pm the amplitude is distinctly smaller

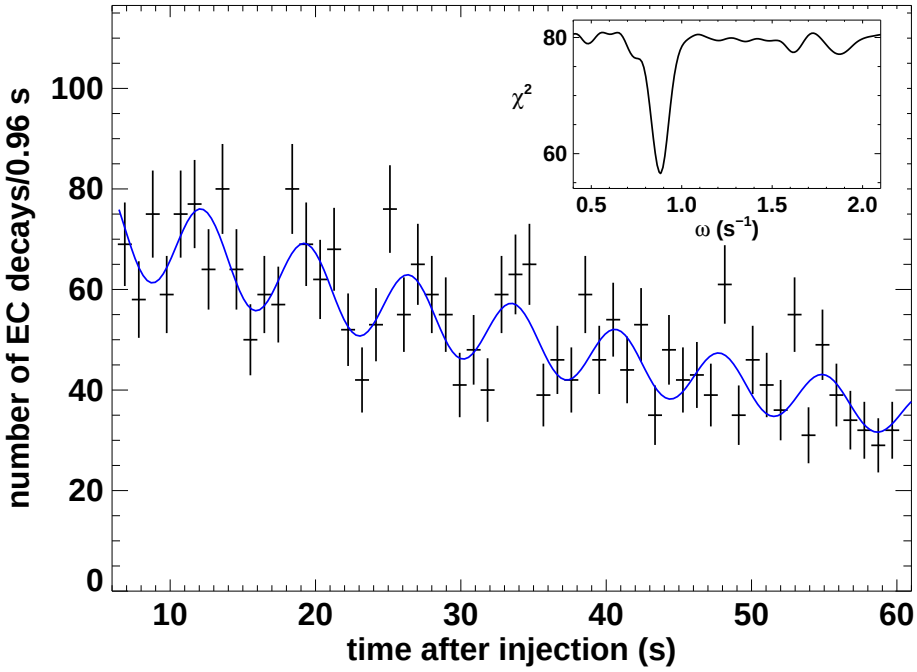

Figure 2.9: Results of the experiment performed in 2010 with $^{142}Pm^{60+}$, measured by the capacitive pickup device. The data were fitted with an exponential decay function (eq. 2.9) and a modulated exponential decay function (eq. 2.10). Upper right corner: Distribution of the χ^2-values for varying ω-values in the modulated exponential decay fit. Plot taken from [8].

Table 2.3: Fit parameters of the plots shown in Figure 2.8 and 2.9. The two rows correspond to the results of the 245 MHz **resonator** and the **capacitive** pickup. Table taken from [8].

$^{142}Pm^{60+}$	$\lambda[s^{-1}]$	a	$\omega\,[s^{-1}]$	ϕ [rad]
Resonator	0.0130(8)	0.107(24)	0.884(14)	2.35(48)
Capacitive	0.0126(7)	0.134(27)	0.882(14)	1.78(44)

at 10-13%. However with this experiment one has to keep in mind that an artificial selection of the data set has been made before the analysis because of the kicker signal problems. Without this selection, in the full data set no oscillation was visible.

2.5 Possible explanations

I will not try to give a thorough theoretical explanation about the origin of oscillations in this section but rather point out a few papers by theoreticians who have been concerned the topic.
The constancy of nuclear decay rates has already been discussed among physicists for a long time [15]. Small changes of the decay constant in the order of % could be observed by changing the environmental parameters such as pressure and temperature. It is clear that the probability of an EC decay will be influenced by the charge state of the parent ion i.e. by the electron density around the nucleus. For instance a fully ionized nucleus has no possibility to decay this way since there are no electrons to capture.
Since the discovery of an oscillatory modulation in the exponential EC decay of hydrogen-like ions in [2] many theoreticians have tried to explain this puzzling effect. However no consensus has been reached so far. The ground state-to-ground state EC decay is a two-body decay with a fixed known total energy (the Q-value). Hence information about the ejected neutrino can be obtained by observing the daughter nucleus. Energy- and momentum conservation can be used to characterize the neutrino without observing it. One proposed explanation for the oscillation is the mixing of several mass-eigenstates in a neutrino entangled with the daughter nucleus [16]. However this explanation has been criticized by others like [17] or [18]. Others favor quantum beats as the origin [19] or use different approaches [20].
So far the statistics of the experiments have been low and leave room for the assumption that the found effect may be due to an experimental error or statistical fluctuations. More theoretical work can be expected once the oscillation has been confirmed unambiguously and systematic experimental errors can be excluded.

2.6 The visual analysis method

Because of the low S/N-Ratio the data analysis of the spectra taken by the capacitive pick-up device (as used as sole detector in the 2006- and 2008-experiment) was performed visually.
In this approach a person looks at every measurement separately, creates a visualization of the subsequent Schottky spectra (hereafter called time-frequency plot/picture), and tries to determine a decay time visually. There are several tools which can be applied in this method. To receive a better S/N-ratio one can average over several subsequent Schottky-spectra (corresponding to a time-averaging in the frequency-time picture). Although this reduces the time resolution significantly it is a good approach to restrict the time region in which a decay may have occurred. After this restriction one can reverse the averaging and look into the designated area in more

detail. In this plot the colour-scale can be adjusted and variable thresholds introduced to highlight the traces of parent and daughter ions against the background. Applying these methods several times one can approach the actual decay time indicated by the start of the daughter trace. Once a sufficiently well-defined starting point has been found, it is written down and the next file is analyzed.
A problem with this technique is that an individual criterion has to be defined about what is identified as the true arrival time of an ion on a trace. This criterion has to be kept the same for the analysis of every file. For the identification of the decay time a significantly higher Fourier amplitude is necessary at the beginning of the trace which may not be present due to the high noise variations. Apart from this the method is very time-consuming since every file has to be analyzed separately and thus not easily repeatable.

Chapter 3

Development

The automated data analysis was developed as an alternative to the visual analysis. The goal was to produce an objective tool for the analysis of Schottky spectra from the iodine experiment. In this chapter I discuss in detail the various processing steps of this analysis.
The development of the analysis method was an iterative process. During each iteration of calculations and observations new information was gathered about the data which could be used to improve the method and optimize it. The goal of the full analysis is to extract a decay time as precisely as possible from every single file (corresponding to one measurement cycle). The following steps are necessary for this:

1. Extract the Fourier amplitudes from the raw data to create a time-frequency-plot.
2. Identity traces corresponding to stored ions in this plot.
3. Extract the maximum amount of information from a trace.
4. Identify the start of the daughter trace as precisely as possible.

3.1 Time-frequency-plot

The input files came in an IQ-Format which is a binary format optimized for frequency-time data (see Section 2.3.1). The first step is to read these binary files and extract the actual Fourier amplitudes which will be used in the analysis. As described in section 2.3 each file contains an array of values, namely an amplitude value for each pair of frequency and time. With these values one can create a 3D-plot to get a first impression of the data, an example can be seen in Figure 3.1(right-hand panel). The stored ions show up as straight horizontal lines in this picture[1]. The lines are straight

[1]Throughout the rest of this thesis the phrases “picture”, the “full picture” or the “time-frequency-plot” are used synonymously and represent the 3D-plot of Fourier amplitudes vs frequency vs time.

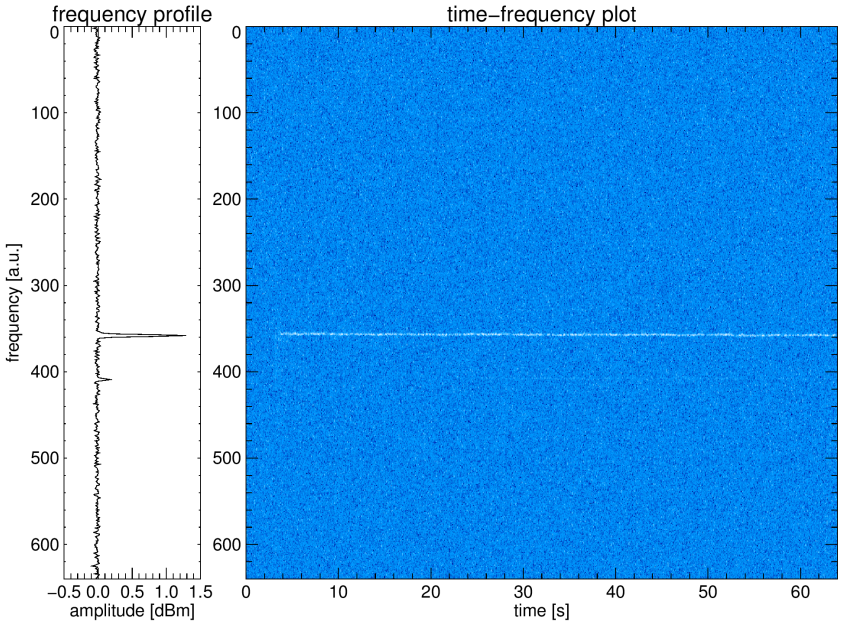

Figure 3.1: Example of a time-frequency plot. Left: The average over time. The two well defined peaks indicate the presence of a parent and a daughter trace. Right: A 3D plot of time, frequency and the Fourier amplitudes. The parent trace is clearly visible whereas the daughter trace is very faint.

because the electron-cooling keeps the velocity-spread very small and the average velocity constant. Hence the revolution frequency stays constant as observed. This “straightness” is one of the features used to detect the traces.

Most of the time-frequency plot only contains noise values with no information content. Only those amplitudes around a trace contain a signal. To center the noise level around zero an average noise value is subtracted from the whole picture. Since the noise level is subject to systematics of any time scales this average noise value is extracted for each file individually from an area of the picture where no traces are expected. To show that the noise actually does vary, the time-behavior of this average noise value is plotted in Figure 3.2. A clear drift is visible which is probably due to external effects and two distinct steps can be observed where some of the input parameters have been changed. There are also files which are completely off the scales. They are highlighted by crosses in the plot and show different kinds of extreme behavior which made us remove them from the data set. Some of these files for example show constant Schottky-spectra over the whole time indicating a problem of the spectrum analyzer. Others show many ion species traveling in an unpredictable manner indicating problems in the cool-

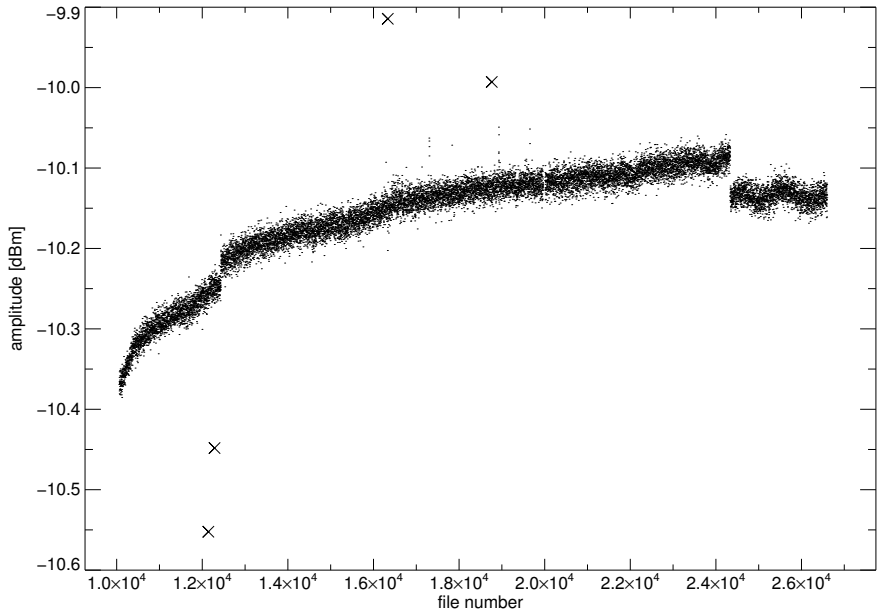

Figure 3.2: Evolution of the average noise level over the time of the experiment. The x-axis shows the file number which represents the time ordering of the measurements. The y-axis shows the average amplitude value of the noise in each file which is subtracted from the whole picture in order to center the noise level around zero. A clear systematic drift is visible and two sharp steps. The four crosses and 3 further files which are out of the amplitude scale could be identified as examples of technical errors.

ing system. Additionally to the four crosses in the plot there are 3 files with amplitudes completely off the scales which did not fit into the plot margins. All of these outliers have been examined in a case-by-case fashion and their removal from the data set could be approved.
To get an impression whether there is an ion present in the Schottky noise the whole picture can be averaged over time to create a frequency-amplitude profile. In this profile a peak corresponds to a trace as can be observed in Figure 3.1 (left). Two clear peaks can be discerned in the profile whereas only the line of the parent ions is clearly visible in the picture (the trace of the daughter nuclei is vaguely discernible). In this first step it is already apparent that it is hard to fix a precise starting point of a daughter trace by hand as was performed in the visual analysis although methods can be utilized to increase the contrast (see Section 2.6).
To improve the visibility of the traces subsequent time bins can be averaged as was done in previous experiments [2]. However this reduces the time resolution of the method significantly so we refrained from doing so.

3.2 Trace-identification

The next step in the analysis is to find out whether there are any traces in the picture and if so, how many and at which frequencies. Parent and daughter traces have different characteristics so they were treated individually.

3.2.1 Identifying parent traces

The traces of parent ions bear the following characteristics:

- The trace generally runs over the whole time-region (it is possible but very unlikely that all of the parent ions decay so the trace vanishes before the end of the time window).
- The amplitude is high because there are several ions stored in the trace.
- The amplitude is higher at the beginning than at the end (if a decay occurred).
- The trace starts as soon as the cooling process is finished.
- The frequency is always lower than that of a daughter trace.
- The frequency lies in a specific region.

Most of these points are purely physically motivated. The last point could also be confirmed through our experience: It is not at all secured that the frequency of the parent ions will stay constant over the whole beam-time. Variations can occur due to changes in the dipole magnets or the electron-cooling and external effects which can not be completely removed. A change of frequencies over time[2] has actually been observed as shown in Figure 3.3 upper plot. Although most of these fluctuations seem completely random we can observe systematic tendencies on small time scales. By identifying the parent traces of all measurements (which is rather easy because of the large amount of stored ions and hence a high signal) we could fix a frequency window containing all the found iodine parent traces. By restraining parent trace identifications to this window we prevented misidentifications in single cases e.g. stray ions of a different species. We fixed this allowed window to be from the 300th to the 500th frequency bin (out of 641 bins).
Using all this informations, the following procedure was used to find the parent traces:
The first eight seconds were cut out of the picture to make sure that the effects of stochastic cooling were absent. Then the picture was summed over time to get a similar profile as shown in Figure 3.1(left). In this plot we searched for the highest peak and checked whether it lies in the expected frequency domain and if its peak value is above a certain threshold (corresponding to single ions). Since the parent trace is very concise in most of

[2]The time mentioned here is always the beam-time, not the time of one observation. The frequency stays remarkably constant over single observation windows.

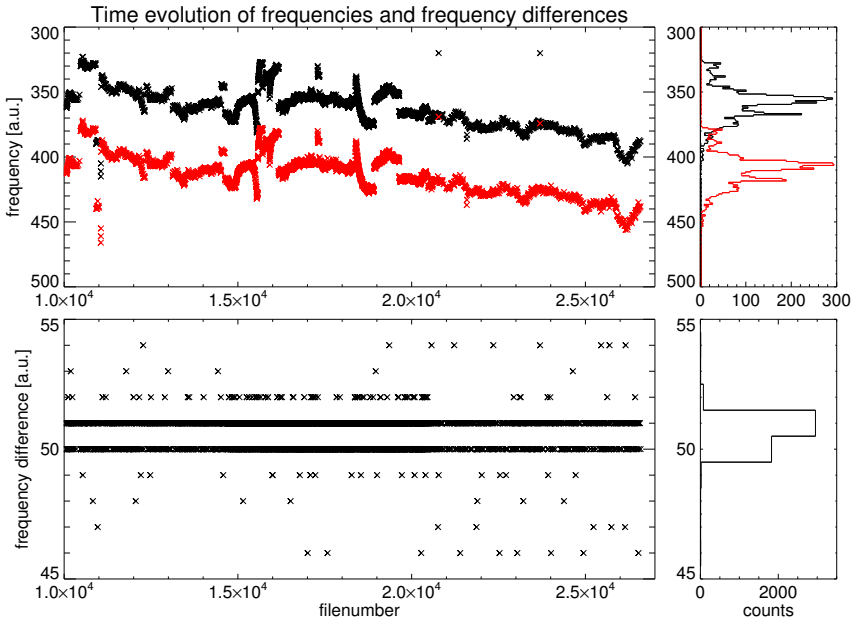

Figure 3.3: Upper plot: Evolution of the central frequencies of parent traces (black) and daughter traces (red) over the time of the experiment. Lower plot: Evolution of the frequency differences of parent and daughter traces $\Delta f = f_d - f_p$ over the time of the experiment. While the frequencies of parent and daughter trace vary largely in between 200 frequency bins the frequency differences stay remarkably constant with most cases at $\Delta f = 51 \pm 1$. The histograms to the right show the distributions of frequencies f and frequency differences Δf, respectively.

the cases (see e.g. Figure 3.1(right)) the identification of its frequency can be easily done .

3.2.2 Identifying daughter traces

The identification of a possible daughter trace is the first real challenge in the data analysis since it is generally much fainter than a parent trace. The following points characterize daughter traces:

- The trace does generally **not** cover the whole time region (rather it starts at some point when the first decay occurred).
- The amplitude is low because there is only one ion stored in the trace (or more in few cases).
- The amplitude is lower at the beginning than at the end.
- The trace starts at some arbitrary time (the decay time).

- The frequency is always higher than that of the parent nuclei.
- The frequency lies in a specific region with a well known Δf relative to the parent trace.
- Even if a parent trace was found there need not be a daughter trace.
- Decays that happen late will lead to very short and thus faint traces.

Again, most of the points stated here are physically motivated. Δf is expected to be constant because as discussed in Section 2.1 $\Delta f \propto \Delta(m/q)$. Even if the trace frequencies change, these changes should effect both the parent and the daughter trace in the same way. These considerations could also be confirmed with the data by looking at all the frequencies of parent and daughter nuclei we found and by checking that their frequency difference remains constant. This can be observed in Figure 3.3 lower plot. As can be checked in the histogram to the right of the plot in most cases $\Delta f \in [50, 51]$. The standard deviation of this distribution is lower than one frequency bin $\sigma(\Delta f) = 0.60$. This standard deviation is especially small compared to the standard deviation of f_p and f_d which is $\sigma(f_p) = \sigma(f_d) = 12.77$. The units in both cases are frequency bins.

The procedure of the daughter-trace-finding-algorithm is illustrated in Figure 3.4: As a first step the picture was split into three equally sized parts and averaged over time to receive frequency-amplitude-profiles. The splitting was done to account for the fact that the daughter trace does not cover the whole time region but rather occurs at some point. Averaging over the whole time would mean an averaging of the daughter trace with pure noise from before the decay. Looking at three parts separately reduces the averaging effect per interval although some of the information is disregarded for long traces (averaging the whole trace yields a better Signal-to-Noise ratio (S/N) than averaging only a part of it). I discovered that three intervals is the best compromise between loosing information and averaging with noise. An example of three profiles can be observed in Figure 3.4 right.

The three profiles of the parts are separately analyzed. Again we take the highest peak in the spectrum to be the parent trace[3] but this time go on looking for further peaks. We check for all the peaks to be in a specific frequency region, otherwise they will be ignored. In general more peaks will be found than there are traces. Because the daughter traces can be very faint, small thresholds are necessary to prevent missing real traces but this way noise peaks can be falsely identified as traces. To prevent this effect, further criteria are used to select true traces. Once all the peaks are found we look for correlations between the peaks, namely a specific frequency difference Δf corresponding to the known frequency difference of parent and daughter nuclei.

If no daughter trace has been found so far I specifically look for late decays

[3]There are usually more ions in the parent trace then in the daughter trace.

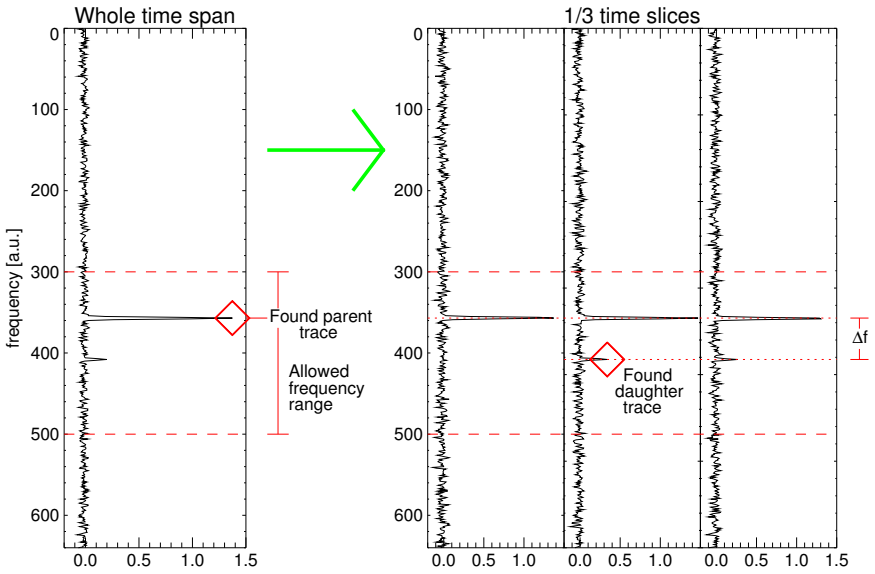

Figure 3.4: Explanation of the trace finding algorithm. The parent trace is identified by looking at an averaged profile of the whole time-span (without the first 8 seconds of cooling) as depicted in the left graph. To find the daughter trace, the time-frequency plot is split into three equal-sized parts and again averaged over time to produce three frequency-profiles as shown in the right graph. These graphs are analyses separately and the one with the highest second peak is believed to be the correct one. In this example the advantage of splitting the image is clear to see since there is no daughter peak in the first third whereas the peaks of the second and the third slice are enhanced in comparison to the whole time span.

corresponding to very short traces (thus faint if averaged with noise). For this the picture is cut into first 6 and then 8 parts and the above procedure is repeated. In this step one has to be careful because the difference between trace and noise peaks becomes very small and a precise fine-tuning of the thresholds used was necessary before I could achieve a considerable improvement. An example of the identification of such a faint trace is depicted in Figure 3.5. Herein a very small peak in the last part could be identified.
To account for special cases such as single-ion-decays, vanishing parent-traces and multiple decays we added some fail-saves. When everything is finished and two traces have been found matching the criteria, the central frequencies of these traces are handed to the next module of analysis.

It is important to note that traces that are not found during this step will remain untreated in the following process. So if a decay has been overlooked

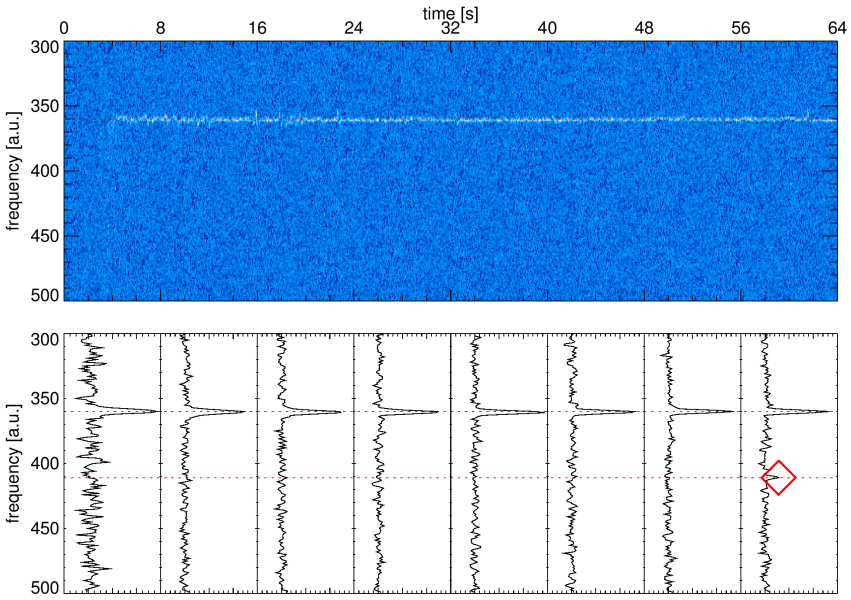

Figure 3.5: The identification of a very short and hence faint trace. The time axis of the profiles in the lower part coincide with the x-axis of the upper time-frequency plot. Every profile contains the time-information of exactly 8 seconds. It is possible only in the very last frequency-profile to make out a peak at the daughter-frequency. In the time-frequency plot it is not observable at all. The frequency-range has been restricted and dotted red lines indicate the parent and daughter frequencies.

in the trace-finding process it is lost for further analysis.

3.3 Trace-extraction

The appearance of an ion can be associated with an intensity increase at the respective trace frequency. Since the traces are very straight (due to the electron-cooling) it should be possible to extract a 2D-time-amplitude profile which contains all of the signal of a single ion. To do so, a first attempt could be to look at the central frequency of a trace as obtained in the previous section and take all the amplitude values over time from this specific frequency bin. This would assure that nearly all of the pixels contain trace information rather than just noise. However this would lead to a rather big information loss since trace pixels which do not lie directly on the central frequency[4] are disregarded.

[4]The existence of such pixels can already be seen in Figure 3.1 (left-hand panel) by noticing at the non-zero peak widths.

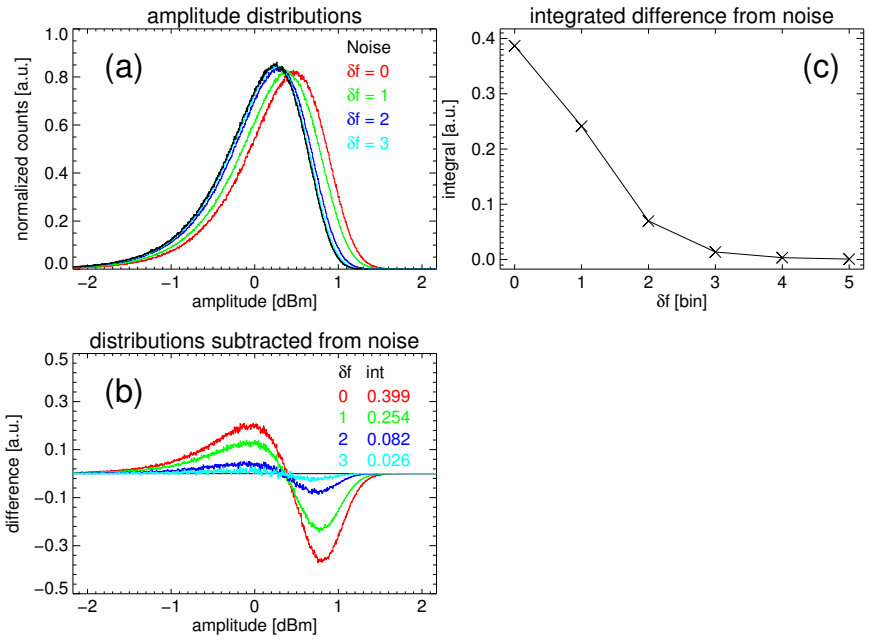

Figure 3.6: Amplitude distributions of trace and noise. (a) shows the different distributions at the frequency $f \pm \delta f$. $\delta f = 3$ already overlaps with most of the noise-distribution. (b) shows the difference of the distributions from the noise distribution and (c) shows the integrated differences as printed in (b) with further points added to make a long-range behavior visible.

To get an idea about the information-distribution among the frequency-bins around the central frequency f_c I analyzed the area around f_c more closely by looking at all the found daughter traces in the dataset and extracting their statistics:
In Figure 3.6 (a) several amplitude distributions are displayed. These were obtained by looking at one (or two) frequency bin(s), taking all the amplitude values at the frequency, and creating a histogram of their occurrence. The black curve corresponds to the average amplitude distribution of pure noise as can be obtained from any region where no ions stray. It is centered around zero because as described in Section 3.1 an average noise value has been subtracted from every amplitude in the time-frequency plot such that the centroid of the noise distribution is zero. For the other distributions only the time bins after the decay are considered (this analysis was of course only possible after a first set of decay times had been retrieved). The red curve (δf=0) is the amplitude distribution directly on f_c. Since most of the daughter traces contain exactly one ion, the difference between the red and the black curve contains the information input of one ion. The other coloured curves correspond to frequency bins around f_c. For this the

amplitude distributions in the two bins at $f_c \pm \delta f$ were considered. So these distributions contain the combined information of two frequencies. At this point I assumed that the amplitude distributions at $f_c + i$ and $f_c - i$ are the same so I combined them to improve statistics. It is already apparent in this plot that the greater the distance from f_c the more the distribution resembles the noise distribution.
To get a clearer distinction, the difference has been plotted in Figure 3.6 (b). A straight horizontal line at y=0 represents the noise distribution whereas deviations from it show the information content of the adjoining frequency bins which can be linked to the area between each individual distribution and zero. To get quantitative information these areas have been integrated and the resulting values are printed in the plot and also plotted versus δf in panel (c). It is obvious that the most information is stored in f_c. However here one can clearly see that the frequencies $f_c \pm 1$ must not be neglected since they still contain a considerable amount of information. Further out like $f_c \pm 2$ there is still an information content but its contribution is small compared to the amount on f_c and $f_c \pm 1$.
Figure 3.6 (c) shows the signal content vs δf. The information content approaches zero fast and frequencies further out than δf=2 do not need to be considered any more. Although the bins $f_c \pm 2$ still contain information, their contribution is small compared to their noise content. The inclusion of these bins may increase the noise content rather than the signal. There would not be any improvement compared to the contributions of f_c and $f_c \pm 1$ since their information content is significantly larger.
Therefore I conclude that the interesting information is contained within the three frequency bins $[f_c,\ f_c + 1,\ f_c - 1]$ This is already a reduction from 641 frequency bins to 3. Now the information content of these three frequency bins needs to be combined to get only one value per time bin with the maximum amount of information content stored in it.

3.3.1 Different parameter combinations

With the above information there are different approaches to extract the trace information. The goal is a 2D time-amplitude profile to pass on to the next step in the analysis. This profile should be the best representation of the number of ions at any time. Two different mechanisms have been considered:

- Averaging amplitudes over frequencies
- Looking in a certain range around the central frequency for the highest amplitude

Because there is a certain binning in the frequency domain it is possible that cases occur where the signal is distributed over neighboring frequency bins. Additionally the stored ion may jitter between two frequency bins. By

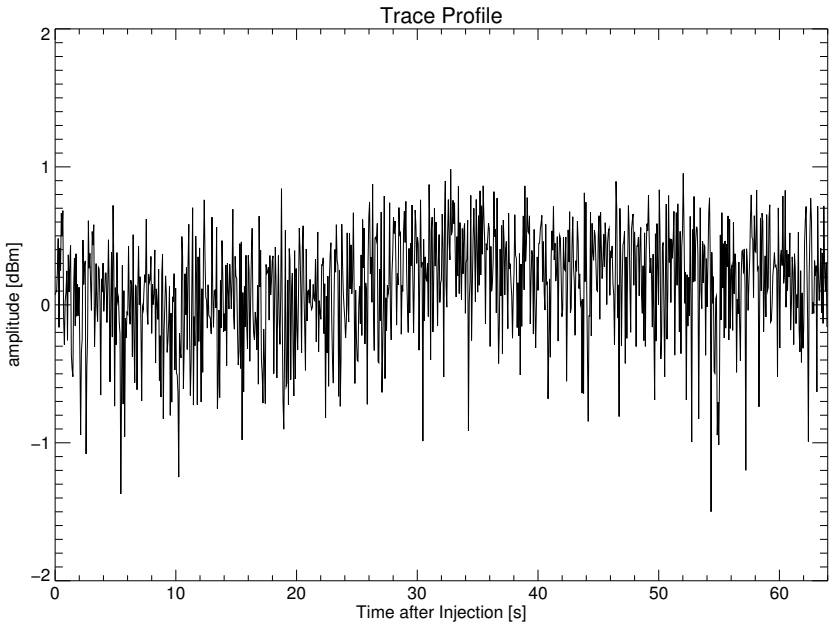

Figure 3.7: Example of a time-amplitude-profile. To obtain this profile two bins within $f_c \pm 1$ were averaged to get two values of which the higher one was taken for each time bin.

averaging over (at least) two frequency bins the information loss can be reduced. I have shown in Figure 3.6 that the signal is contained within up to 3 neighboring frequency bins. Since the intensity of each bin is subject to random noise, averaging several bins containing a signal contribution will reduce the noise influence. However this improvement is compromised once one or more of the averaged bins are dominated by noise, since averaging with noise reduces the signal strength. So the best amount of frequency bins to average needs to be identified empirically.
The second approach is to look at several frequency bins and consider the one with the highest amplitude to contain the full signal. This is not necessarily the case since the amplitude distribution of noise and trace overlap over a large region (as has been observed in Figure 3.6 (a)). However it is more probable for a signal bin to have a higher amplitude than it is for a noise bin. So when looking at a full trace and always taking the highest amplitude of the three bins $f_c \pm 1$ for each time bin the S/N-ratio within the trace will in general be increased.
I have experimented with these two approaches and applied them to the data to find out which yields the best result. Additionally I ran some simulations which are discussed in Section 5.2. The results show that the best way to extract a trace profile is not one of the above approaches but rather

a combination of the two. In the previous section we have already concluded that all we need to know is contained within the frequency bins $f_c \pm 1$. In this region we average twice, namely the amplitudes of the frequency bins $[f_c, f_c + 1]$ and $[f_c, f_c - 1]$ and from these two values we take the higher one. This procedure has proven to yield a good distinguishability of zero (noise) and one ion (trace). Once we have done this for every time bin we finally arrive at the time-amplitude profile. An example of such a profile is shown in Figure 3.7. Although noise is still dominating this picture one can already confirm visually that something happened somewhere in the middle of the profile which made the amplitudes shift to higher values. Although it is difficult to make out the exact time bin (which coincides with the delayed decay time) where this happened this is a further step of improving the signal which will in the end let us extract the decay time as precisely as possible.

3.4 Fit of a step-function

Looking at the case shown in Figure 3.7 it can be estimated pretty well in which time region the decay occurred (although there are cases where it is not at all obvious). However the precise decay time, the exact time bin, is difficult to deduce from this plot by eye. A first attempt to find the correct decay time is the following:
We expect the average noise amplitude to be zero, since we subtracted an average noise value from the whole picture. We also expect the average trace amplitude to be larger than zero. Concluding from this, if we already knew the decay time we could average the amplitudes before the decay and would expect a value around zero and average the values after the decay and would expect a non-zero value being significantly higher than the noise-value. Actually we would expect the latter value to be the same for all traces containing one ion. So we expect a function with a constant value (= zero) before the decay (= step) and a different constant value after the decay. This corresponds to a simple step function. So by fitting a step function to the profile one could determine the decay time (with the step position being one of the fit parameters).

3.4.1 Identifying the correct step height

The step height of the fit should be the same for all cases so excluding it as a fit-parameter increases the reliability of the fit considerably. However for this the correct step height needs to be found in the given data. This may sound pretty easy but contains some subtleties.
In principle the average value of a trace should mirror the number of particles stored in it. Since the number of particles can change during an observation

period (mainly due to decays but different losses are also possible) averaging over the entire observation time will always contain an error.
The most statistics is available for parent traces because in nearly every measurement cycle there are injected parent ions whereas there will not be a daughter trace in every run (in cases where no decay occurred). Taking the average value over the whole trace (starting after the stochastic cooling) is the first step to start ion counting in the traces. Of course there will be a lot of overlapping in the distributions since the number of ions can change in a trace over the duration of the measurement. But in general the average value should mirror the number of ions in the trace and averaging over the whole trace will maximize the statistics. Also we know exactly when the trace should start (after the stochastic cooling) and when it should end (At the end of the measurement window = time of extraction) which is not the case for daughter traces of which the starting time (the time of the decay) is unknown at this point. So we take the average amplitudes of all parent traces of the full sample of 16477 files and create a histogram of these amplitudes. The result can be observed in Figure 3.8.
The first three peaks are clearly distinguishable and two more can be discerned although they are not very sharp any more. The peaks have been fitted with gauss functions to increase their visibility and parametrize them. The peaks correspond to the number of ions in the trace i.e. the traces with an average value lying in the first peak house exactly one ion, the ones lying in the second peak two ions and so forth. The traces with amplitudes between peaks most likely have a decay event. So a trace starting with two ions of which one decays after 30 seconds (corresponding to half the measurement window) will have an amplitude value corresponding to 1.5 ions. Since the decay probability increases with the number of ions (each ion has a certain probability to decay) this effect will increase towards higher ion numbers and smear out the distribution as can be observed in Figure 3.8. To get an idea of the correlation between these average amplitudes and the ion numbers, the central values of the gauss fits have been plotted in Figure 3.9.
A simple polynomial correlation can be seen in the peak positions. The fit in this graph was performed according to eq. 3.1. This function seems to fit the data very well although from Section 2.1 a linear dependency would be expected. A possible reason for this could be a smearing effect: the more ions are injected the more decays will (on average) occur[5] so due to the averaging of the whole trace the amplitude value will shift to smaller values if a decay occurred somewhere in the meantime. This will shift the whole distribution to lower values as well as increase the peak width as can be observed in Figure 3.8.

[5]In addition to the observable EC decays unobserved β^+ decays occur and particles are lost due to charge exchange.

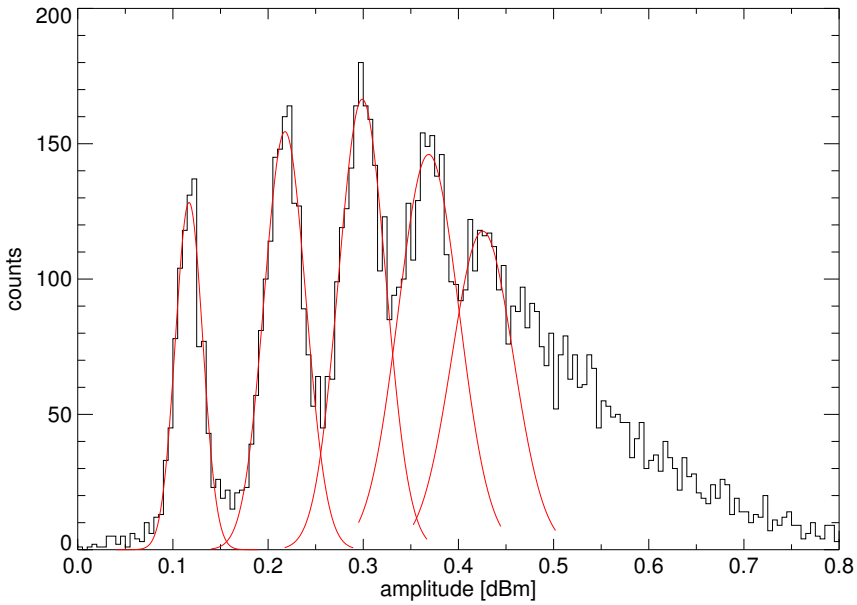

Figure 3.8: A histogram of the averaged amplitudes of the parent traces of the full set of 16477 measurements. Clear peaks are visible corresponding to the number of ions in the trace. These peaks have been fitted with gauss functions as indicated by the red curves. The peaks at higher amplitudes get broader and overlap and the peak distance becomes smaller. This is due to the fact that the decay probability increases with higher parent ion numbers which creates intermediate amplitudes.

$$A(n) = x_0 + x_1 \cdot n + x_2 \cdot n^2 \tag{3.1}$$

$A(n)$ ... Amplitude
n ... Number of peak or number of ions
x_i ... Fit parameters

The high statistics observable in Figure 3.8 could only be achieved by averaging over the whole trace. So when averaging over smaller parts of a parent trace statistical variations become large and it is more difficult to attribute a specific ion number to the trace. This also makes it difficult to identify decays using the parent trace which would have been a powerful additional source of information. These ideas are elaborated in Section 3.5.
However our first goal of doing this analysis has been achieved: The extraction of a step height value. Since the noise level has been set to zero the position of the first peak mirrors perfectly the step height of single decays.

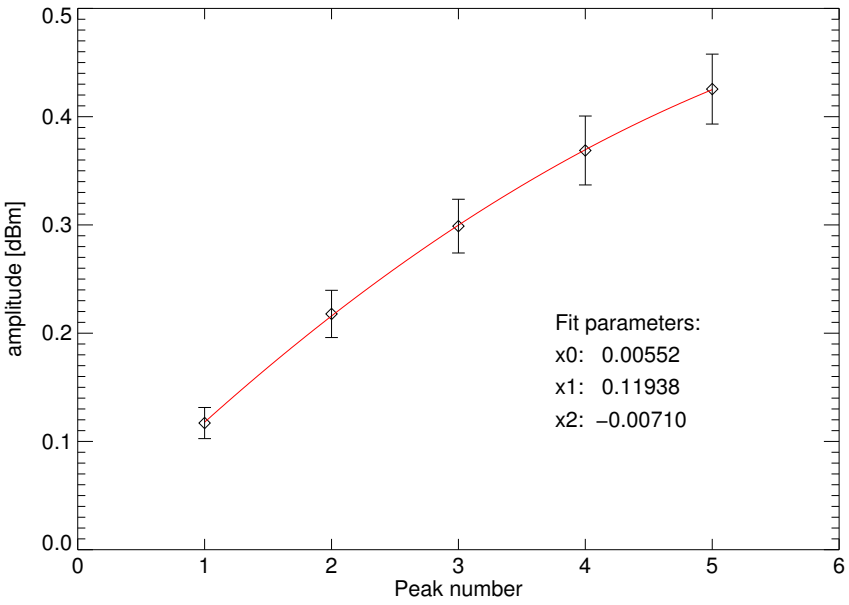

Figure 3.9: Positions of the peaks in Figure 3.8 versus peak number. The peak number is suspected to be equivalent with the number of ions stored in the trace. Error bars are derived from the peak widths. A second order polynomial fit has been applied to the data (red curve).

3.4.2 Performing the fit

Once we have a fixed step height we can apply a simple χ^2 fit of a step-function to the time-amplitude profiles of the daughter traces. Since the height of the step is a fixed value, the only variable fit parameter is the step position corresponding to the decay time. The reduced χ^2 is defined as indicated in eq. 3.2. A minimization of this value will correspond to the best fit of the data.

$$\chi^2(p) = \sum_i \frac{(M(p,t_i) - O(t_i))^2}{\sigma_i^2} \quad || \quad \chi^2_{red}(p) = \frac{\chi^2(p)}{\nu} \tag{3.2}$$

$M(p,t_i)$	...	Model value of parameters p at time t_i
$O(t_i)$	...	Observed value at time t_i
σ_i^2	...	Variance of $O(t_i)$
ν	...	Number of degrees of freedom

The variance σ_i is the same for every data point i since each point is subject to the same statistical variations. Thus $\sigma_i = \sigma \; \forall \, i$ and it can be dropped from the fit. Since we only have one fit parameter, namely the decay time t_{dec} the definition in eq. (3.2) simplifies to eq. (3.3). The number of degrees

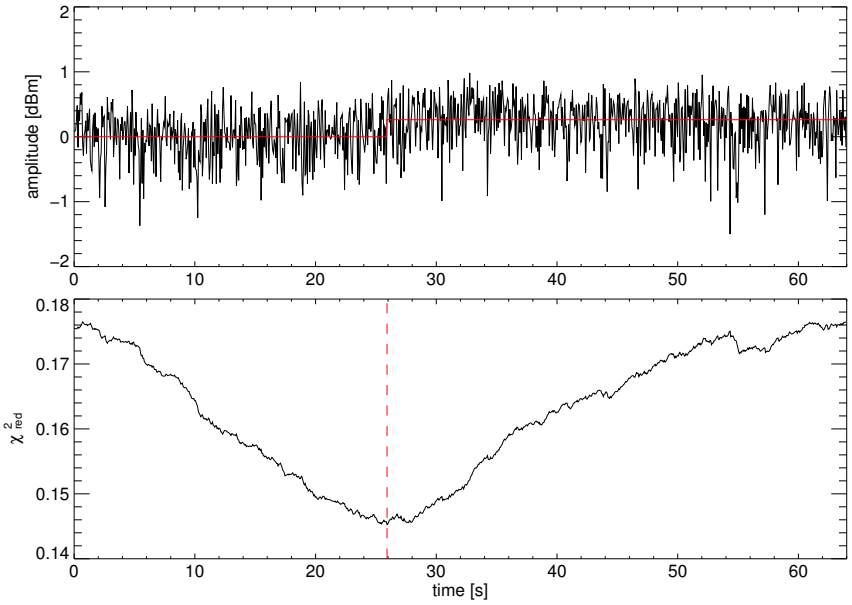

Figure 3.10: Reduced χ^2-distribution of the fit of a step-function with fixed step-height to a trace-profile. Minimum of the distribution corresponding to the best decay-time is indicated by a red dashed line.

of freedom ν in our case is 1000-1 corresponding to the number of time-bins subtracted by the number of fit parameters.

$$\chi^2_{red}(t_{dec}) \sim \frac{1}{\nu}\sum_i (M(t_{dec}, t_i) - O(t_i))^2 \tag{3.3}$$

An example of such a χ^2_{red}-distribution and the corresponding trace profile are shown in Figure 3.10. The minimum in this distribution corresponds to the best fit value of t_{dec} so the best estimate of the decay time for this data. This has just given us the means of finding a clear decay time in otherwise very noisy data.
However already in this example it is observable that there is a broad valley in the χ^2_{red}-distribution. So the value of t_{dec} can not be fixed unambiguously. This was the motivation for the development of the likelihood method discussed in Chapter 4.

3.5 Use of the parent trace

In every file there are two sources of information, namely the two traces: the parent trace and the daughter trace. So far the combined statistics of all the

parent traces were used to define a step height and the daughter trace was used to find a decay time. In principle the decay should also be observable in the parent trace manifesting in a decrease of the trace amplitude. However when I investigated this possibility I encountered several problems which made me discard the parent trace as a source of decay times in the end. In this section I provide a short discussion of my approach and the difficulties encountered.
When trying to extract decay times from the parent trace there are several fundamental differences compared to the daughter trace:

- There are many ions stored in the trace (10-20) before a decay as well as afterwards.
- The number of injected ions varies with every file.
- A decay manifests by the loss of a particle rather than the gain so the trace amplitude decreases.
- An EC decay can not be distinguished from a β^+ decay or losses due to charge exchange.
- Multiple particle losses are possible and probable (depending on the number of injected ions).
- The step height of a particle loss varies with respect to the number of stored ions.

These points also represent the main difficulties. As could be observed in the previous section in Figure 3.8 it becomes difficult to clearly define the number of ions in a trace once this number exceeds 5. This is the case for most files, especially those wherein decays occur. Even if the average stored number of ions in one file could be determined (e.g. by the polynomial shown in Figure 3.9) the amplitudes at the start and the end of a trace vary considerably depending on the amount of particle losses. The difference of these amplitudes can be used as an indicator for the amount of lost particles. However due to the high noise level averaging of several time bins is necessary which smears out the particle counting if too many bins are used and increases its variance if too few are used. Since the number of injected ions is uncertain, the step height can not be attributed a distinct value. If the number of lost particles is large the time intervals with an (expected) constant amplitude, corresponding to a specific number of stored ions, become small. As these time intervals become smaller the variation of their averaged values increases. This makes it difficult to find steps in traces with many ions/losses.
Rather than investigating the parent trace independently one could use it to check the results of a daughter trace analysis. However since the parent trace does not only experience particle losses due to EC decays but is actually dominated by β^+ decays ($\sim 78\%$ see Section 2.2) it is difficult to decide

whether an (uncertain) step in the parent trace is due to the same decay observed in a daughter trace or due to some other effect.

Because of these reasons the daughter trace was used as the sole information source for EC decay times.

3.6 Selection of single-decay cases

So far the decay time detection mechanisms are only sensitive to single decays. Due to the high noise level and the inability to observe cooling curves it is very difficult to reliably detect more than one decay. Similar problems as discussed in the previous section arise such as the variability of step heights and difficulties in the ion counting. Rather than trying to find multiple decay times, which would lead to misidentifications and increased uncertainties, we only strive to find to find one per file. This may not seem like a big problem, however if decays following the first one are ignored this will affect the resulting decay curve. A simulation about the behavior of the decay constant in such a case is provided in Section 5.5. Therefore we would like to exclude files containing multiple decays from the dataset.
In this section I explain how this selection of single-decay files is performed i.e. how it can be concluded that a file does not contain multiple decays.
The best place to look for multiple decays is at the very end of the daughter trace. If more than one decay occurs within the measurement window then at the end of it there are more than one ion in the daughter trace. So a high amplitude at the end of the trace will be an indicator for multiple decays. Remembering the view of a trace profile (e.g. Figure 3.7) it is not a good idea to look at single amplitude values. Rather it is necessary to average over several time bins to obtain a significant difference from the noise level. I took the last 100 timebins of every trace found and averaged them and did the same for the first 100 in every frame as a comparison. The results can be observed in Figure 3.11. The lower cloud of dots corresponds to the average of the first 100 bins whereas the cloud at higher amplitudes corresponds to the last 100 bins.
It can be observed that the experiment can be segmented into three time regions. The first and third region contains rather confined amplitudes while the amplitudes of the second region reach up to higher values. This can be explained by the fact that the number of injected ions was regulated to different values in the course of the experiment. While in the first and third region the average number of injected parent ions was ~ 10 it was increased to ~ 20 in the second region. This is why files with more than one decay occur more often during the second period. However the probability of multiple decays also increases with the number of injected ions . Hence the three regions will be treated separately in the following analysis.

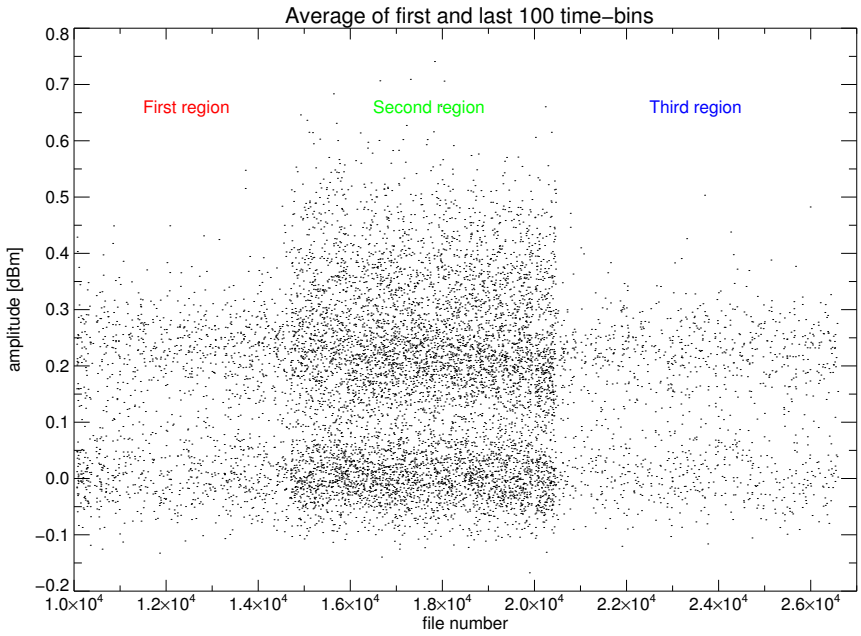

Figure 3.11: The averaged amplitudes of the first and last 100 time-bins of all files with detected daughter traces. Three distinct regions are observable with the middle region having the highest decay density and the files with the highest averaged amplitudes.

Especially in the second region many amplitudes clearly exceed the region which could be attributed to single ions (well distinguished in region one and three). These are obviously cases with multiple decays which we want to eliminate from the data set.
The next step is to create histograms of the amplitudes observed in Figure 3.11, separated by their regions to be able to observe possible characteristics which allow to distinguish files with more than one decay. The results are plotted in Figure 3.12 in the upper right graph. Here we see more or less well separated peaks, corresponding to the number of stored ions in the trace. It is also obvious that the main set of decays (single as well as multiple) comes from the second region rather than the other two. We expect the structure corresponding to the average of the first 100 bins to be dominated by the noise. It may also have a small contribution of a trace-component coming from early decays. This is e.g. visible in a second green peak reaching into the area of the last 100 bins. However by the shape of this structure we can deduce that it could probably be well described by a smooth distribution if the statistics were high enough.
Such a distribution is mainly characterized by the noise statistics, although its peak is shifted to a higher value due to its stored ion component. Since

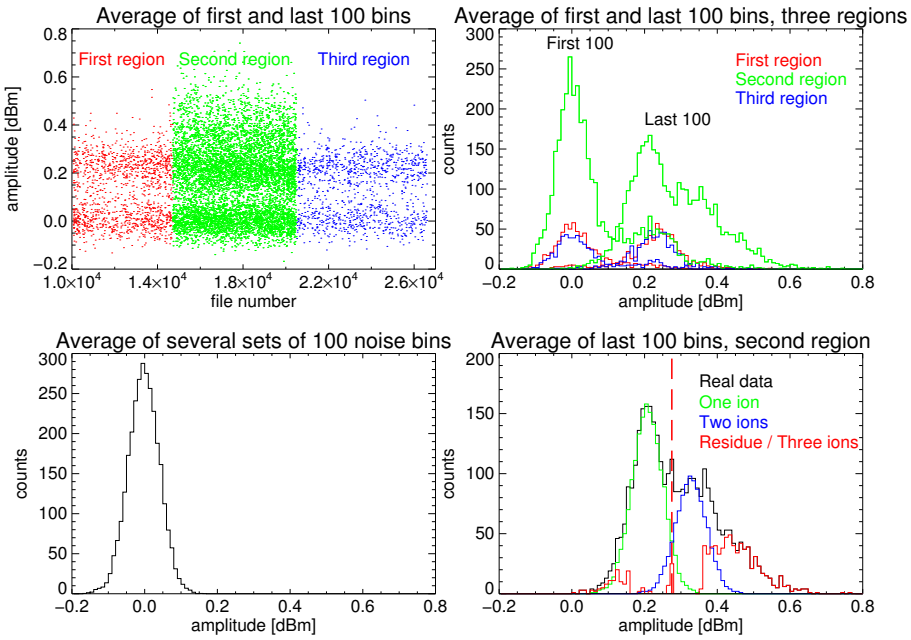

Figure 3.12: The first and last 100 time bins on every found trace were averaged and are plotted in several different configurations. A detailed explanation is given in the text.

there are enough trace-free areas in the time-frequency plots we can easily derive these noise statistics. To achieve this I took several areas within each file which were guaranteed trace-free and created sets of averaged 100 noise bins (modulated by the parameter-selection of average two and take the higher value). The corresponding plot is provided in Figure 3.12 in the lower left graph.

Because most of the decays lie in the second region and even more so most of the multiple decays lie within that region the other two regions were disregarded in the following analysis. If a decay happened in the last 100 time bins it would shift the distribution towards lower values and distort the structure. Because of this I removed these files from the sample. What is left is shown in Figure 3.12 in the lower right graph.

Now I tried to fit the noise distribution into this structure to account for the peak corresponding to exactly one ion. To do this I only varied the height of the distribution and its central position. The result is the green curve in the graph which seems to fit the first peak very well. After this I fitted a second noise distribution into the data, shown as a blue curve. The red curve is the sum of the two peaks subtracted from the real data. This red curve shows that the two approximations are very good and hints towards a third peak which can be identified as three or more ions. So the distribution

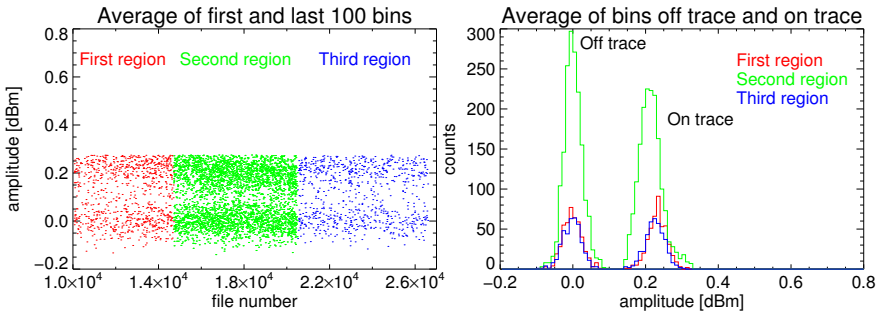

Figure 3.13: Data set of all decays after the selections explained in the text have been applied. The left plot shows again the amplitude distributions of averaged first and last 100 time bins. The right plot shows the average of all bins before a decay and after a decay.

of the averaged last 100 time bins of the second region can be characterized very well.
With the two fits one can attribute to every file the number of final ions i.e. a number of decays: Files included in the green curve can be assumed to contain one decay whereas files of the blue curve more probably contain two decays and so on. The problem is the area where the two distributions overlap. Those cases can not be clearly assigned to one or the other set of cases.
I set the threshold, indicated by the red dashed line, at the amplitude, where the two distributions cross. This way as many single-decay-files are kept in the data set while as few as possible two-decay-files are absorbed. That it is not possible to completely rule out multiple decays without significantly loosing single-decay-files is obvious in the graph. So all files which have an averaged amplitude of their last 100 time bins on the daughter trace exceeding the threshold were removed from the data set. Since the distribution could be well approximated by the well defined noise distribution it can be calculated that we loose 6% of single-decay files and discard 82% of multi-decay files this way.
An additional selection has been applied which has nothing to do with multiple decays. I looked at the amplitude distributions of averages from the decay time to the end of the frame (= amplitude values on trace) as well as the averages of the first time bin to the decay time (= amplitude values off trace). These should show two smooth peaks respectively after the above selection of single decay files. To remove difficult cases such as very faint traces which might have been mistakingly identified I restricted these amplitudes to lie within 40% of the maximum of their respective peaks. This selection only reduces the data set by a very small fraction (especially compared to the above selection) and makes the results of the analysis more secure.

With these two selection rules the new data set has an amplitude distribution as plotted in Figure 3.13. The panel on the right-hand side shows the amplitude distribution of the averaged values before a trace (first time bin until decay time) and on the trace (decay time until last time bin).

Chapter 4

Likelihood method

The main goal of this thesis was to develop a technique to retrieve the decay times from Schottky spectra as precisely as possible. A relatively simple mechanism has already been presented in Section 3.4 where the trace profile was fitted with a one-parametric step function via a χ^2-minimization. While this is a valid first approach there is still room for improvement.
The more information is used for the analysis is in general the better. For the step fit a well-defined step height was used, retrieved from the data of the parent-traces. This and the fact, that the average noise level is reduced to zero is the only input into this function. However we have much more information which we can and will use in the so-called likelihood (lh) method. The idea of this method is to attribute to each amplitude value on the trace a probability value of whether it lies on the trace or in the noise. By using this information one can deduce the most probable time of a decay occurrence. A general introduction into the usage of lh calculations can be found in [21].

4.1 The noise distribution

First of all a clear description of the involved statistics is necessary. This means getting a precise distribution of the Fourier amplitudes of pure noise. Since most of our data consists of noise (i.e. frequencies that do not coincide with parent- or daughter traces) and this noise is not different in any way from the noise before daughter traces we can use these regions to characterize the noise very well.
For this we look at the amplitude distribution of the noise. We create histograms of the amplitudes of noise which show the frequency of occurrence of each amplitude. Since in the end we want to compare these distributions to the on-trace data, we need to take the parametrization into account which was used during the process of trace-extraction (see Section 3.3). There we took the central frequency $f_c \pm 1$, averaged over two bins and took the higher one of these values. In comparison with the distribution of single bin ampli-

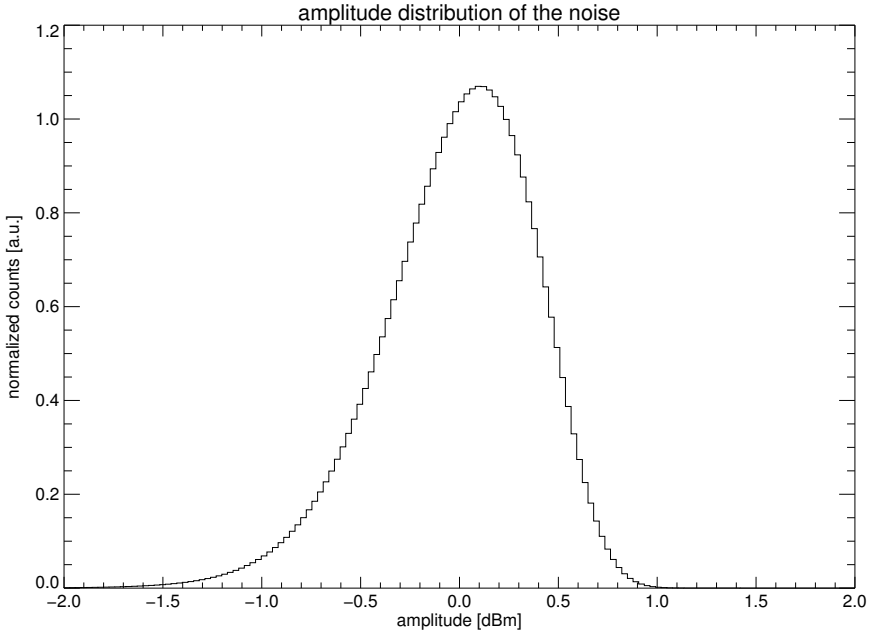

Figure 4.1: Amplitude distribution of the noise. For every element three neighboring frequency bins were taken, two of them averaged and the higher value retrieved. The integral of the histogram was normalized to 1.

tudes this will of course change our amplitude distributions (in general to higher values).
So we look at frequency-regions which have proven to be empty (of ions) and extract 'fake'-traces at fixed frequencies. Then we study the amplitude distribution of these traces. Since we have a data set of 16477 files and every fake trace only takes up 3 frequency-bins ($f_c \pm 1$) we can do the characterization arbitrarily well and soon reach a very smooth distribution. The result is plotted in Figure 4.1.
We can do the same for the on-trace data. This is a little more tricky and there are natural limits to the statistics achievable. For the set of trace data (I will always refer to single-ion traces as "the trace") we can use the amplitudes of every file with one decay. The time bin where we start extracting amplitudes is the decay time. So to get these distributions we first need to already have decay times. We take these decay times from the step fit performed earlier and start a few time bins later with extracting the amplitude values to compensate for possible errors in the decay time identification. The results are shown in Figure 4.2. I realized during this analysis that the shapes of the noise and one-ion amplitude distributions are very similar. I observed this when I compared the trace distribution with a noise distribution shifted by a fixed value, the step height. The two distributions overlap

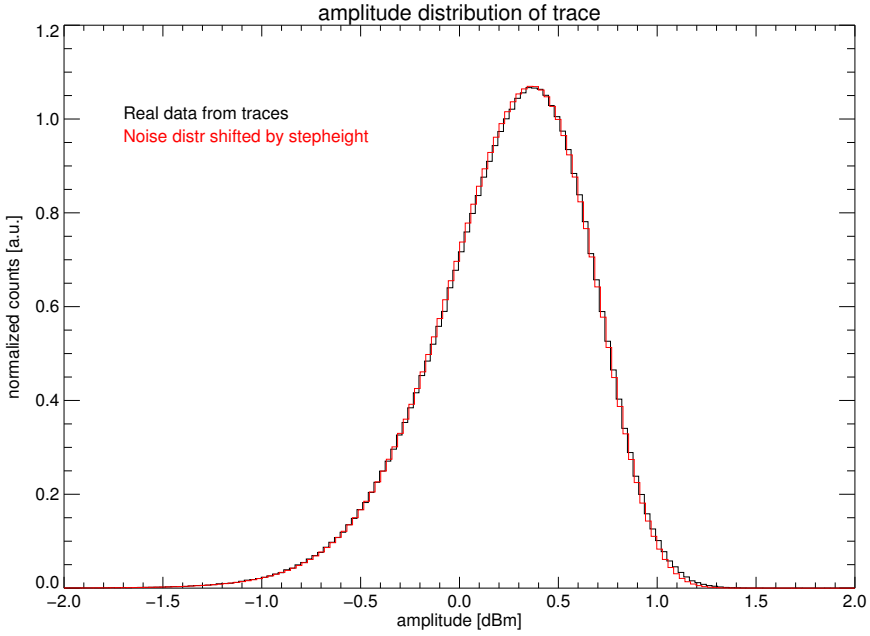

Figure 4.2: Amplitude distribution of the trace. For every element the three neighboring frequency bins $f_c \pm 1$ were taken, two of them averaged and the higher value retrieved. The integral of the histogram was normalized to 1. The data of the black curve is taken from every found trace starting at the decay time found by the step fit + 100 bins. The red curve is the noise distribution shifted by the step height value. The two curves clearly coincide.

perfectly once the step height is slightly adjusted as shown in the plot.
Not only does this give us an insight into the correlation between noise and trace data, we can now also use the noise distribution shifted by the step height instead of an individual distribution. This allows to calculate the step height more precisely than performed in Section 3.4.1.

4.2 Calculating the likelihood

Once we have the noise and trace distribution we can compare them as done in Figure 4.3. The two distributions overlap in most of their range therefore using amplitudes one can not decide with absolute certainty whether a time-frequency bin lies in the noise or on the trace. However what we can deduce from this plot is the probability of each amplitude to be part of the noise or the trace and this probability is related directly to the counts in the histogram.
If we look at a trace-profile we can attribute to each amplitude two probability values: to lie in the noise or on the trace. An illustrative example

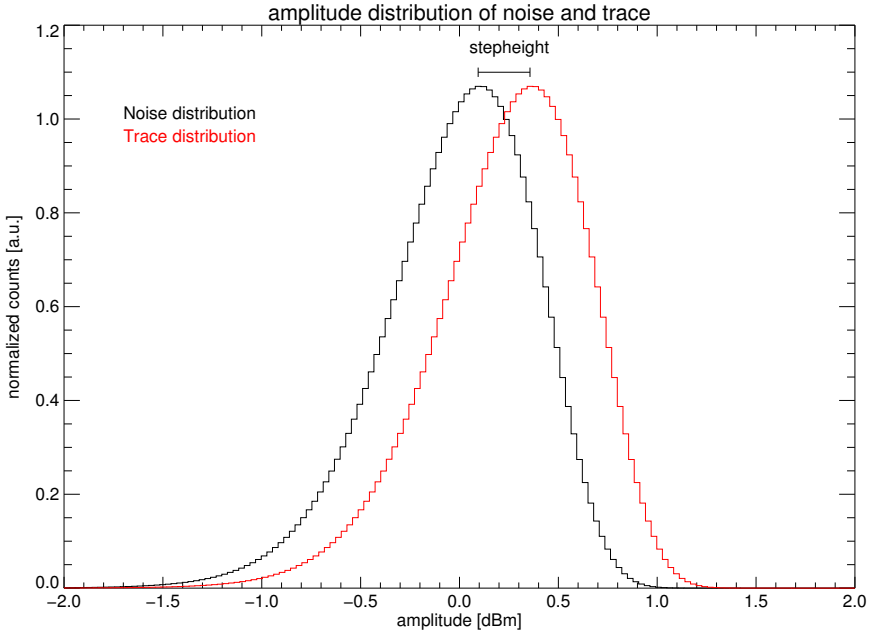

Figure 4.3: Amplitude distribution of noise and trace. Both histogram integrals were normalized to 1. The distributions overlap over most of their range.

is shown in Figure 4.4. In the first picture we see a simplified trace-profile with 7 amplitudes corresponding to 7 time bins. The second plot shows the amplitude distributions of noise(black) and trace(red) with the x-axis being rescaled to represent the probabilities. The dotted green lines connect each amplitude value with two probability values, one for being part of the noise, the other for the trace. The probabilities can be very different or be nearly the same, depending on the amplitude value. For example the fourth point in the trace-profile lies at a very low amplitude where the probability-difference between noise- and trace-profile is also very small. The last point is at the amplitude where the two distributions overlap so it is actually just as probable for this point to be part of the trace as of the noise. The truth about this will only be revealed once the data point is set into relation with its neighboring amplitudes. Other data points such as the first and the fifth point show very large differences in their probabilities so they may have a larger impact on the later calculation.

Now we calculate the likelihood of decay occurrence for every time bin t_i. To do this we multiply the noise-probabilities of the timebins before t_i with the trace-probabilities of the time bins after (and including) t_i (see eq. (4.1)). This mirrors the likelihood of a decay happening at time t_i. We calculate

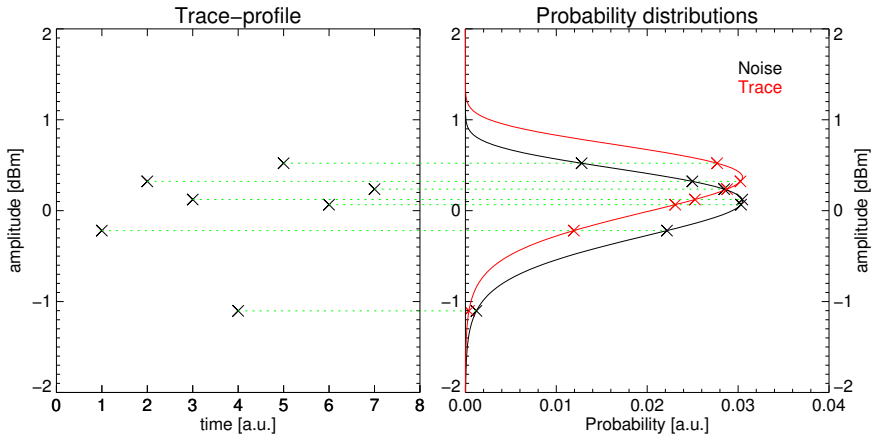

Figure 4.4: Explanatory example of the likelihood-method. To every amplitude in the trace-profile two probability values are attributed, namely one for lying in the noise and one for lying on the trace.

this likelihood for every time bin except the first and the last one. We skip these two because they do not have time bins before or after them respectively to do the calculation. In general decays happening in the first or last time bin can not be distinguished from decays happening earlier or later, respectively.
Multiplying small numbers we run into machine precision problems very soon (when calculating on a digital device). So in practice it is much easier to use the logarithm of the probabilities for calculations. This reduces the multiplications to simple sums and yields the log-likelihood (llh). It is defined in eq. (4.2). The llh is used in most calculations because it is much easier to handle and also in plots it is more illustrative than its exponentiated version, the likelihood. Since it contains the same information as the likelihood the two can be interchanged (although the scaling has to be kept in mind).

$$lh(t_i) = \prod_{j=0}^{j=i-1} P_{noise}(t_j) \cdot \prod_{j=i}^{j=N} P_{trace}(t_j) \tag{4.1}$$

$$llh(t_i) = \sum_{j=0}^{j=i-1} \log(P_{noise}(t_j)) + \sum_{j=i}^{j=N} \log(P_{trace}(t_j)) \tag{4.2}$$

Every llh-value of the trace-profile uses the information of all amplitudes as well as the information about the amplitude-distributions of noise and trace. If we look at our previous example, the llh-values are plotted in Figure 4.5. While it is not apparent at all in the trace profile at which time the decay might have occurred, there is a clear peak in the llh-distribution pointing

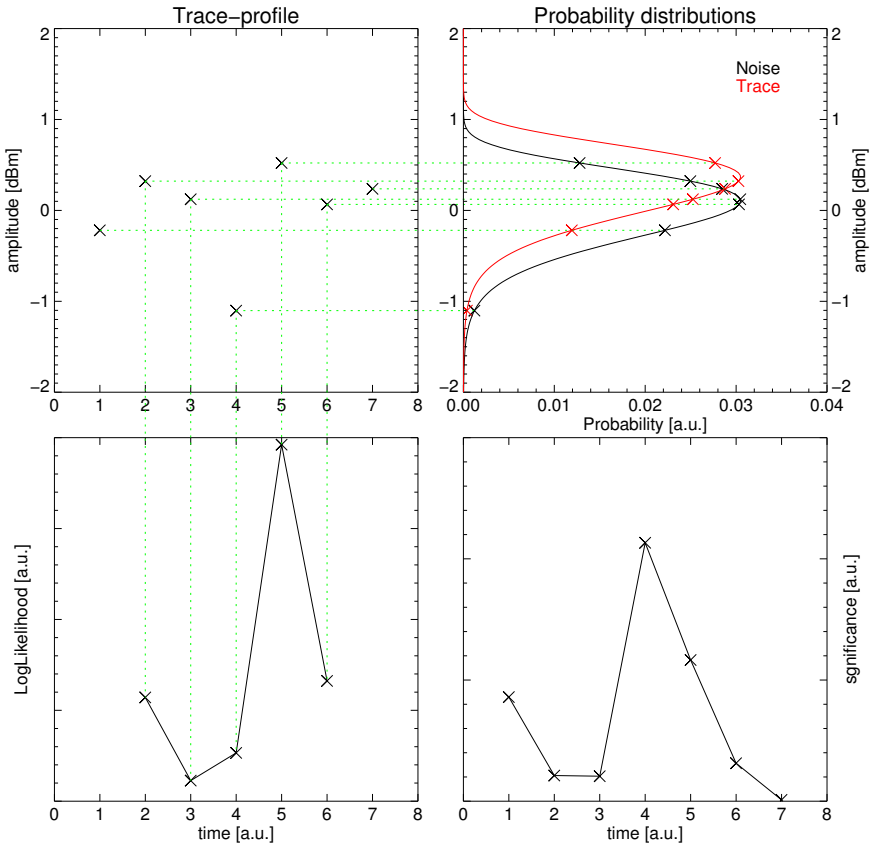

Figure 4.5: Explanatory example of the likelihood-method. To every time-bin in the trace-profile a log-likelihood-value of decay-occurrence is attributed. The maximum in this distribution corresponds to the most likely time of decay. The significance of each point depends on where its amplitude lies in the probability distribution.

towards the most likely time of decay.
The position of the peak and its height in comparison to the other llh-values depends on all the amplitude values of the trace rather than the single relatively high amplitude of the decay time bin. The very low amplitude of time bin 4 and the high one of time bin 5 have played an important role. However it is just as important for the first time bin (having a much larger probability to be in the noise than the trace) to be before the decay as it is for the fifth time bin (having the probabilities more or less reversed in comparison to the first one) to be on or after the decay. Since the probabilities are multiplied (i.e. the log probabilities summed) the ratios of the probabilities are of importance, rather than their difference. So the largest noise

and on-trace influence will lie in the bins with amplitudes corresponding to the largest probability ratios of belonging to trace or noise. E.g. this ratio is rather large for time bins 1 and 5 and relatively small for time bins 2, 3 and 6. At time bin 7 it is ~ 1 so this time bin has nearly the same probability to belong to the trace or the noise. The influence of time bin 4 is also rather big because even though the difference of probability is small, their ratio is large. The significance of each point for the likelihood analysis, defined as the ratio of the probabilities, has been plotted in the last graph (lower right) in Figure 4.5.
At this point we don't need to fix a decay time but have actually much more information available: the full likelihood information over time of where the decay may have occurred. So instead of stating that due to the data we conclude the decay has occurred at time bin 5 we can make a quantitative statement about how much more likely the decay happened at time bin 5 compared to all other times.
This is the main advantage of the likelihood method. Because of the high noise level in the data there is no way to deduce an exact decay time for each measurement with 100% certainty. But with the likelihood method we can state probabilities at which time the decay may have occurred containing the full information of the specific measurement.

4.2.1 Extraction of a decay time

In most cases it is not a good idea to extract a fixed decay time from a llh distribution. Even though this can be done by taking the time of the maximum (every thinkable confined distribution has a maximum although it may not be unique) we loose a lot of information by doing so.
This being said, for various purposes e.g. comparison to other analyses it is necessary to be able to attribute a fixed decay time to every measurement so this is done by taking the time-value of the most likely decay. For all purposes where this can be omitted the full llh distribution is used.
One of these cases where a reduction of the llh to a decay time value can not be omitted is when defining the time window in which to perform a decay curve fit. This constraint is necessary because of the cooling effects at the start of the measurement and the low detection efficiency at the end. This window was set to 6.4 seconds to 60 seconds. Files in which the position of the maximum likelihood was found to lie outside of this window will not be considered in the fit.

4.3 Fit of a decay curve

Once we have figured out the decay times (or decay information) of all measurement cycles we want to deduce a decay curve. In the end we want

to find out, whether we see an oscillatory behavior as observed in other experiments and if so with which parameters. The decay curve corresponds more or less to a histogram of the decay times. However we don't really want to create a histogram because we are already very limited by statistics and the binning necessary for a histogram would reduce our accuracy even further.
So what we use is a so-called unbinned likelihood fit. It allows usage of the initial decay times and fits a model function to this 1D-data rather then the usual method of fitting a function into a plot of 2D-data. An elaborate explanation will be given in the following section.

4.3.1 The unbinned likelihood fit

The approach in this method is somehow similar to the likelihood method described in Section 4.2 although the calculations are performed differently. Let us consider n decay-times within the time window t_{min} to t_{max} and a Model M which we want to fit to the data. M could e.g. be the pure exponential decay curve as in eq. (4.3) or the modulated one as in eq. (4.4).

$$M_e(t,p) = p_0 \cdot e^{-p_1 \cdot t} \tag{4.3}$$

$$M_{mod}(t,p) = p_0 \cdot e^{-p_1 \cdot t}(1 + p_2 \cdot \cos(p_3 \cdot t + p_4)) \tag{4.4}$$

M ... Model
t ... Time
p_* ... Fit parameters

The values of M correspond to the frequency of occurrence of a decay time. Since these values are only interesting to us in the confined region between t_{min} and t_{max} we can normalize the function between these values to one. The resulting function will give us the probability of decay occurrence at any given time between t_{min} and t_{max}. N (defined in eq. (4.5)) is kept in the equation.

$$N(p) := \int_{t_{min}}^{t_{max}} M(t,p)dt \tag{4.5}$$

Note that p will in general be a vector of several fit parameters. Let us go back to the n decay times. We now attribute to each of these times a probability value of occurrence. When we multiply all of these values we will obtain a likelihood-value of the fit-parameters. The higher this value, the better the model fits the data. So by maximizing this product (eq. (4.6)) with respect to the fit parameters we get those parameters which best fit the data.

$$\prod_i^n \frac{M(t_i,p)}{N(p)} = N^{-n}(p) \cdot \prod_i^n M(t_i,p) \tag{4.6}$$

In general we would expect the normalization constant N to equal the number of decays n. However this would only be the case if M fit the data

perfectly which in general is not the case due to statistical variations. To correct for this fact we multiply the above expression with a poisson factor. This factor corresponds to the probability that n decays occur if N corresponded to the integral of the true distribution. Due to this multiplication the expression in eq. (4.6) transforms to eq. (4.7).

$$\frac{N^n}{n!}e^{-N}\cdot\prod_i^n\frac{M}{N}=\frac{e^{-N}}{n!}\cdot\prod_i^n M \tag{4.7}$$

As in the previous likelihood method we would rather work with sums than products so we take the logarithm of the above expression. Since it is standard to minimize a function rather than to maximize it we also multiply it with -1. This will give us the expression $\mathcal{L}$ as defined in eq. (4.8) to minimize.

$$\mathcal{L}(p)=-\log(1/n!)+N(p)-\sum_i^n\log(M(t_i,p)) \tag{4.8}$$

The first term is a constant and can be dropped for the fit process.
This is the general idea of an unbinned likelihood fit. We can apply it to any set of decay times to get a decay curve fit without loosing time resolution. However from the decay analysis we do not have one precise decay time but rather a likelihood distribution of decay times for every case and we want to include this excess of information into the fit mechanism. We can write $M(t_i)$ as:

$$M(t_i)=\int M(t)\cdot\delta(t_i)dt\longrightarrow\int M(t)\cdot lh_i(t)dt \tag{4.9}$$

$\delta(t_i)$ is the δ-distribution. With the likelihood method we do not have precise t_is (= decay times) but llh-distributions for every case i. After normalizing the likelihoods we can replace the δ-distributions with the likelihoods as indicated in eq. (4.9). This changes the expression in eq. (4.8) to eq. (4.10).

$$\mathcal{L}=N(p)-\sum_i\log\left(\int_{t_{min}}^{t_{max}}M(t,p)\cdot lh_i(t)dt\right) \tag{4.10}$$

This is the final expression which has to be minimized using e.g. the Minuit package [22]. Since we have a time binning of 64ms the integral in eq. (4.10) becomes a sum and the delta-distribution a delta-function. When we want to use fixed decay times e.g. for comparison we introduce them as delta-functions with its peak at the corresponding time.
Additionally to the full-parameter fit we also did a frequency-scan. In this scan we performed the fit with several fixed values for the parameter ω (denoted p_3 in this section) and compared their likelihood values. Since the likelihood value does not allow to judge the goodness of fit (as does e.g. a χ^2-value) only a comparison of likelihoods allows statements about the quality of the fit. By performing a frequency-scan we can investigate whether there is one distinguished frequency with a much higher likelihood than all others.

4.3.2 Plot of the decay curve

It is not trivial to visualize the real data. If one plots a histogram of the decay times (as taken from the maximum likelihood) this will not mirror exactly what is used during the fit process. So artificial deviations between the fit function and the data plot will occur. The best way to plot the data is to sum up all the likelihoods. Additionally I averaged 16 time bins for every data point to increase visibility. This way the fit results can be properly plotted over the data which is actually fitted. The error-bars can be estimated from counting-statistics.
The results will be displayed and discussed in Chapter 6.

Chapter 5

Simulations

Various aspects of the analysis have been simulated. The two main goals were parameter optimization and an estimation of the detection and algorithm efficiency.

5.1 Creating a trace profile

It is currently too far reaching and and not feasible to simulate the whole measuring process because several aspects are not known or not well understood. This is why the simulations start at the point described in section 3.3, the extraction of a trace profile. As described in the section the trace profile can be extracted with different parameter sets of "range" and "average". Range means we look at the frequency bins $f_c \pm$ range and take the highest amplitude value. Average means, we average over a certain number of frequency-bins. The two parameters combined yield a certain range in which as many sets of neighboring frequency bins as possible are averaged and then the highest value of these averages is taken. The different parameter pairs lead to different amplitude distributions.
These amplitude distributions for different parameter pairs can be extracted from the real data (an example was shown e.g. in Figure 4.1). They correspond to the probability for a time bin with a given amplitude to be part of the noise or a trace. With these distributions one can create artificial trace profiles. With a fixed decay time one attributes to the time bins before the decay amplitude values randomly selected from the noise distribution and amplitude values from the trace distribution to the bins after the decay. This will create profiles like the real profiles[1] with the difference that the decay time is known in advance since we set it ourselves. An example of such a simulated trace is shown in Figure 5.1 with the set decay time indicated by

[1] This is actually only possible if there are no systematic changes during the experiment. The data has already been corrected for small changes such as the change of the average noise level, so they do not influence the simulations.

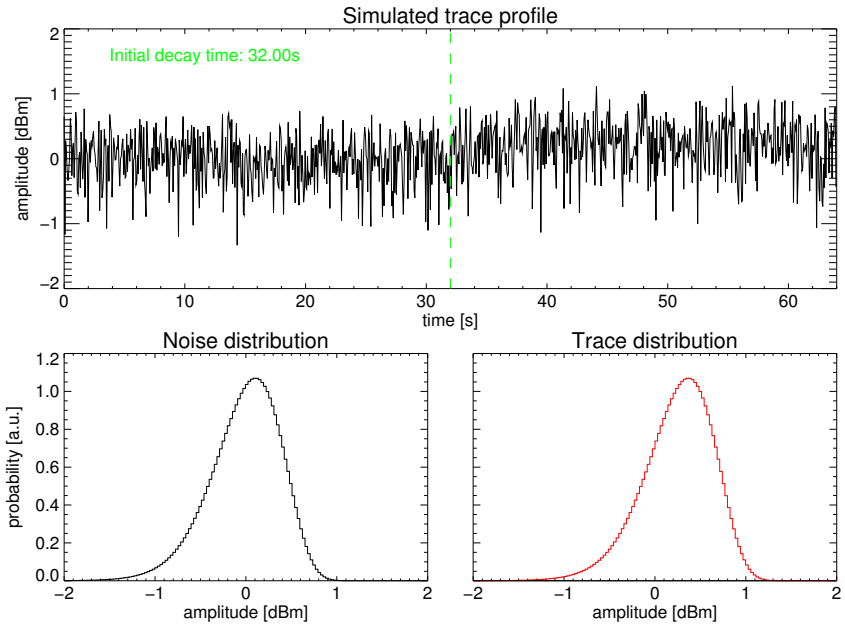

Figure 5.1: Example of a simulated trace profile. The amplitudes before the set decay time are distributed according to the noise distribution, the amplitudes afterwards according to the trace distribution. These are shown beneath the trace profile for the parameters range=1 and average=2.

a dashed green line.
So creating such a simulated trace profile is relatively easy once we have the necessary amplitude distributions for the parameters we want to investigate. Since they are derived from the real data we can assume that the simulated traces mirror the fluctuations of the real traces pretty well.

5.2 Simulation of different parameter combinations

Above I have already recapped the possibilities to extract the trace information out of the time-frequency plot with different parameter settings. In this section I provide some simulations aimed to find out which is the best parameter combination.
I simulated several traces with a fixed decay time (namely at time bin 500 corresponding to 32 seconds after injection) and ran the lh-analysis over them to check whether the found decay time coincides with the initial one or how far it deviates. The amplitude distributions depend on the parameters so also the outcome of these simulations will be different for each set. The first of the two parameters which can be adjusted is the range, "rge". A range of 1 means that the highest value of $f_c \pm 1$ will be extracted. The

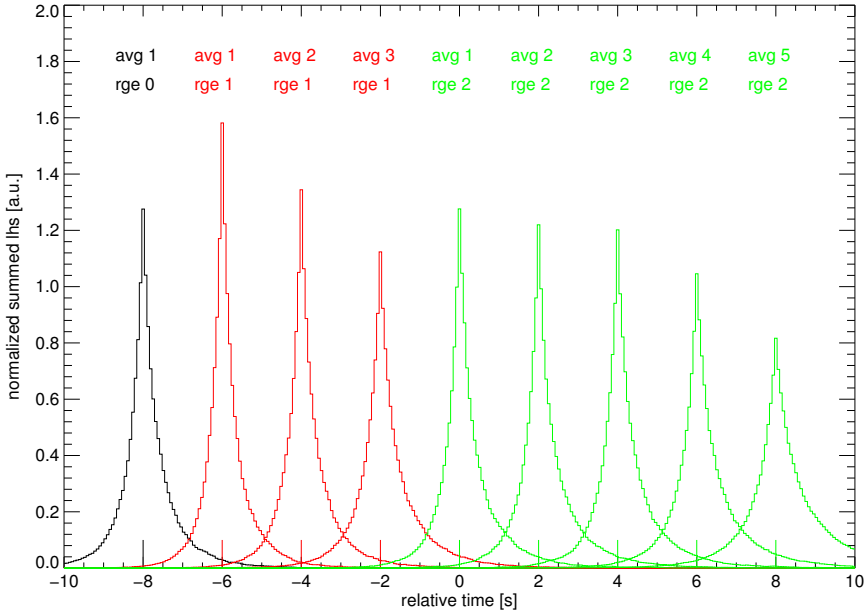

Figure 5.2: Found decay times vs initial set decay time of several simulated parameter-combinations. These distributions were created by summing up the corresponding likelihoods and normalizing the distributions to 1. The distributions are shifted such that the center of every distribution is at a distance of two seconds to the previous parameter set to make multiple distributions visible and allow their comparison. All distributions are centered around zero, their individual zero-mark indicated by a thick tick in their respective colour.

second parameter is how many bins are averaged, abbreviated with "avg". For instance rge=1,avg=2 corresponds to looking at $f_c \pm 1$, averaging over two bins and taking the higher value. Of course the value of rge will limit the possible values of avg. There are exactly rge*2+1 possible ways to average and rge*2-avg+2 values to choose the highest from.
To compare these parameter combinations I simulated 5000 traces for every combination up to rge=3. This yields exactly 16 parameter combinations which I indexed from 0 to 15. The amplitude distributions were extracted for each parameter-combination separately from the real data. I ran the lh-analysis over the 5000 traces which yields 5000 likelihood-profiles. To keep the full information, rather than looking at the decay times (as taken from the maximum likelihood) I summed up the 5000 likelihoods and normalized them. Then I subtracted the initial decay time from the set to get a distribution around zero of the differences between found and initial decay times. The results for the first nine parameters can be observed in Figure 5.2.
The distributions are all centered around zero but have been shifted in the

plot to make multiple distributions visible and to be able to compare them. A colour-coding has been introduced to highlight distribution of the same range. A high peak and a small width correspond to the decay times being found very closely to the real one. It is apparent that the best results are achieved at rge=1 which supports the arguments brought in Section 3.3. Even though this plot might suggest that the parameter combination avg=1,rge=1 is the optimum which corresponds to no averaging at all, I decided to go with the combination of avg=2, rge=1, which is the second-best choice, according to Figure 5.2. Although this might lead to worse results in some cases where the daughter trace is extremely straight and does not spread its intensity over more than one frequency-bin it promises to cope better with jittering traces, an effect which is not included in the simulation. In the end the simulations showed that the choice of parameters only has a small influence on the outcome. Even in an obvious bad choice of parameters e.g. rge=2,avg=5 in most cases the decay time is still found within ± 2 seconds of the original one, as can be observed in the last distribution of Figure 5.2. Most of the distributions look rather similar.
These simulations strengthened the decision to fix our parameters to rge=1, avg=2.

5.3 Detection efficiency

The simulation of trace profiles is a perfect tool to test the detection efficiency over the whole time-window. In principle it is imaginable that the efficiency should be best for decay times somewhere in the middle of the time frame because that would contain the same amount of bins in the noise as on the trace. Simulations are suited to test these assumptions.
Similarly to the previous section I simulated 5000 traces with the same decay time and ran the lh-analysis over them. Rather than varying the parameter settings (as in the previous section) I varied the decay times and compared the decay time distributions afterwards. Again I used the summed likelihoods rather than the fixed decay times from the likelihood peaks. The results can be observed in Figure 5.3.
From a first glance the distributions do not seem to change very much over the whole time frame. To get a closer look I plotted three of the distributions again in Figure 5.4. To get the best systematic impression I chose for this one very early decay time, one in the middle and one at the very end.
Here it is obvious that there is no significant difference between the distributions. The detection efficiency of our method is the same for decays at the very beginning, in the middle or at the very end of the time spectrum. How can this be explained?
The answer lies in the very mechanism of the likelihood method. As explained in the small example in Section 4.2 every single time bin is of the

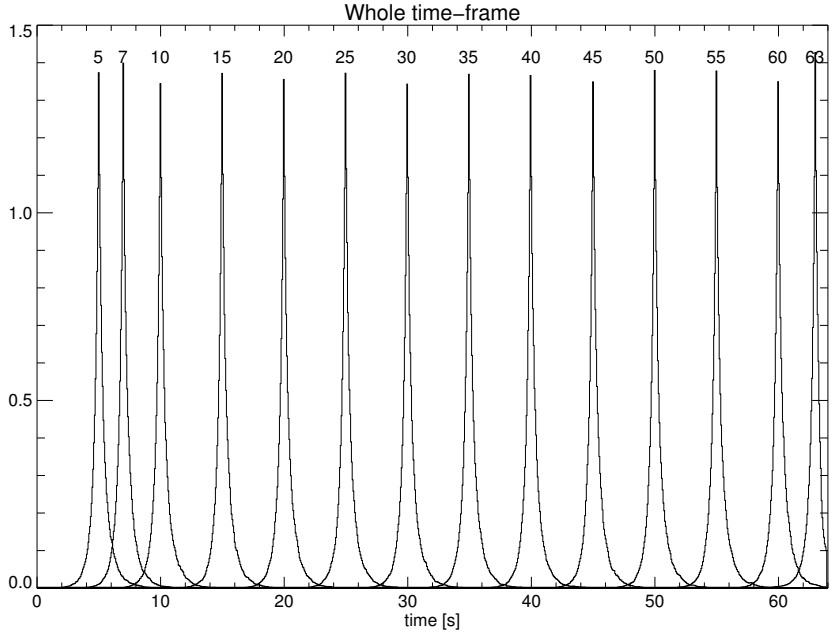

Figure 5.3: Simulated detection efficiency over the whole measurement window. The distributions consist of the sum of 5000 likelihood profiles each, normalized to one.

same importance in the calculation process. The time bins after the decay will show an amplitude distribution as the trace distribution while the time bins before the decay will show an amplitude distributions as the noise distribution. So if the decay happened very early most of the amplitudes will more likely belong to the trace than to the noise, making an early decay time most likely. The contrary is the case for a very late decay time. If the decay occurred in the middle of the time frame roughly half of the amplitudes will more likely belong to the trace and half to the noise. However the efficiency of determining the time bin in which the probability of belonging to the noise or the trace changes will not be dependent on the position of this time bin since all time bins make the same contribution to the likelihood-profile independent on whether they are near or far away from the actual decay time.

The point where a time-dependence may enter the detection efficiency is the trace-identification from the time-frequency-plot. Here it is obviously more difficult to find a very short trace (corresponding to a late decay time) than finding a very long trace which will correspond to a large peak in the averaged spectrum (see Section 3.2). However this part of the analysis we are currently not able to simulate.

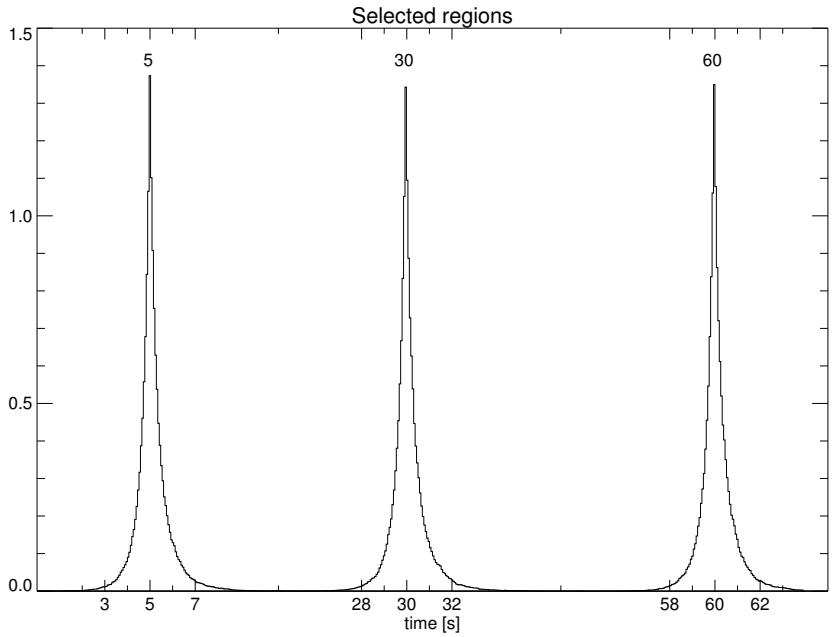

Figure 5.4: A closer look at the detection efficiency at decay times 5s, 30s and 60s. No clear differences between the distributions are visible.

5.4 Simulate decay times

The next step is to check with simulations whether the analysis method is sensitive enough to detect oscillations in the decay curve if they are present. One of the main experimental limiting factors has been the lack of statistics so I wanted to check whether this modulation is at all detectable with the present data set.
To do this I created artificial decay times distributed as if they belonged to an oscillatory modulated decay curve. The five input parameters defined in eq. 5.1 can be varied according to realistic assumptions.

$$N(t) = N_0 \cdot e^{-\lambda t} \cdot [1 + a \cdot \cos(\omega t + \phi)] \tag{5.1}$$

5.4.1 Decay curve results

Already after this step one can perform an unbinned likelihood fit with the decay times as input delta functions to see whether the statistical variations of the decay times have smeared out the oscillation and rendered it undetectable or have left it intact. An example of such a fit of 3000 initial decay times is plotted in Figure 5.5. It is obvious that the fit parameters and the initial parameters fed into the decay time-creation-algorithm are not exactly

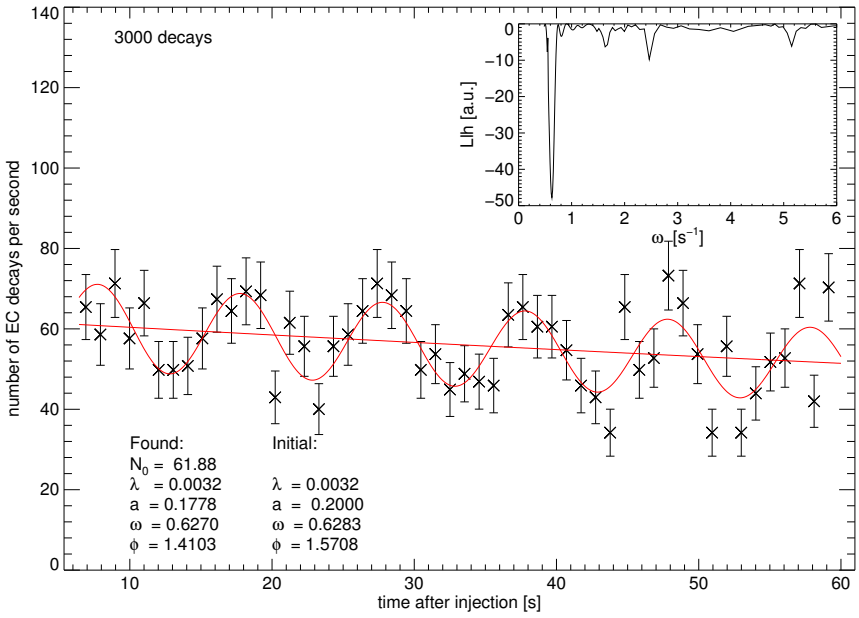

Figure 5.5: Fit of initial decay times of a simulations, distributed according to a modulated decay curve with the parameters indicated in the plot. N_0 is not an input-parameter in here but rather fixed by the number of decays (here 3000). The unbinned likelihood fit was performed with the decay times as delta functions.

the same. It has to be emphasized that this is not due to any kind of error in the analysis but rather has to be attributed to statistics. The only way to reduce this difference is to increase the sample size. This effect can't be removed by improving the analysis technique. In a general case if the decay times are distributed according to a modulated decay curve the result may not show the exact 'real' outcome (as set by the initial parameters here). The magnitude of these deviations will be discussed in Section 5.4.3.
After this first check, a simulated trace profile is created for every single decay as described in Section 5.1. Then the standard lh-analysis is run over these profiles and a likelihood distribution is extracted for every decay. When performing the unbinned likelihood fit with these likelihood distributions generally a few cases will be dropped because their maximum likelihood is not in the allowed time region (see Section 4.2.1). Hence the number of decays analyzed will not exactly be the same as in the initial data set. However this will also happen to the real data when the maximum likelihood is shifted compared to the real decay time due to statistical effects. So the simulations can also give us a rough estimate of how many cases are lost. In our above example of 3000 initial decays 20 decays are lost. The resulting

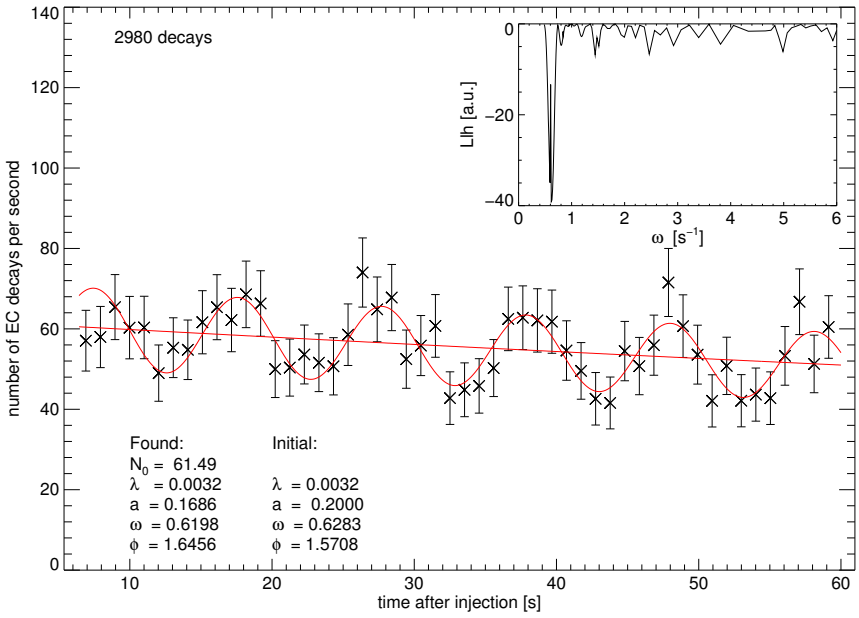

Figure 5.6: Unbinned likelihood fit of the analyzed simulation. The found fit parameters coincide with the ones of the fitted initial decay times shown in Figure 5.5.

decay curve is plotted in Figure 5.6. Apart from the loss of decays the fit results of the initial and statistically distorted decay curve are nearly the same and agree within their error margins with the initially set parameters. In the plots the frequency scan mentioned in Section 4.3.1 is provided as an inset. It is observable that in both cases there is one clearly preferred modulation frequency.

5.4.2 Varying input parameters

The outcome of such simulation tests heavily depends upon the input parameters. They will also determine the detectability of an oscillatory modulation in the decay curve. The input parameters in the simulations are the number of decays N, the decay constant λ, the amplitude of the oscillation with respect to the rest of the decay curve a, the angular frequency of the oscillation ω and its phase ϕ. The phase, as already mentioned in the discussion of the previous experiments (Section 2.4) is not a very reliable fit parameter since it is strongly correlated with the frequency.
One could vary these input parameters in an infinite number of combinations so I picked out a few interesting cases. The parameters of these cases

are printed in table 5.1. Simulations 1 and 2 simulate the results of the 2006-experiment mentioned in Section 2.4.1 of Pr and Pm respectively. The exact number of found decays is not mentioned in [2] so I assumed 3000 decays. Simulations 3 and 4 simulate the results of the 2010-experiment discussed in Section 2.4.2 of the two detectors, the resonator and the capacitive pickup device, respectively. Simulation 5 simulates the decay constant of the neutral atom ^{122}I with some added realistic parameters.

Table 5.1: Different sets of input parameters for the simulation of an oscillatory decay curve.

Simulation	N	λ $[s^{-1}]$	a	ω $[s^{-1}]$	ϕ [rad]
1	3000	0.0015	0.18	0.890	0.4
2	3000	0.0224	0.23	0.885	-1.6
3	3600	0.0130	0.107	0.884	2.35
4	3600	0.0126	0.134	0.882	1.78
5	2500	0.0032	0.150	1.000	0.00

5.4.3 Increase statistics

It has been shown above that the results of a fit through the initial decay times will not yield the exact same values as the input parameters. This is because the decay times are created anew with every simulation so every simulation yields slightly different results even if the input parameters are the same. Once we include the likelihood analysis this effect will be enhanced because the simulation of the single traces introduces an additional statistical randomness.
So if no oscillation is visible in the decay curve of a single simulation, the only deduction possible from this observation is the following: There is a set of decay times which initially had an oscillatory behavior. However due to statistical variations this behavior will not be visible or detectable with our analysis tools.
To account for this fact I repeated every simulation 100 times with the same input parameters. Because the computing time for such a sample of 100 simulations takes around two days I did not increase the sample size any further. This was sufficient to make qualitative statements about the dependencies on the input parameters.
Each simulation creates an output of the four fit parameters of the decay curve. I created histograms to see how these outcomes are distributed around the input value. To get an idea about whether the observed effects originated from the decay time statistics or the analysis I also considered the fit pa-

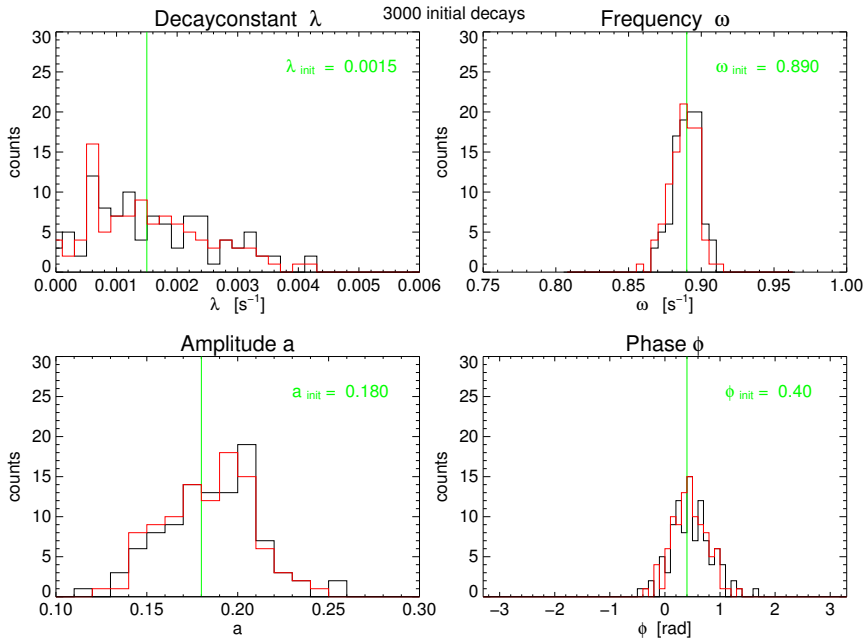

Figure 5.7: Histograms of the fit results of 100 simulations. Black histograms correspond to the full analysis, red histograms to a fit of the initial decay times. The initially set input parameters are indicated by green lines and printed as insets to their respective plots.

rameters resulting from the fit of the initial decay times. These were input into the unbinned likelihood fit as delta-functions. A delta function is the best possible profile of a likelihood indicating an absolutely unambiguous decay time.

An example of such a set of histograms for the parameters of simulation 1 from tab. 5.1 is shown in Figure 5.7.

In these plots the black histograms show the fully analyzed simulation and the red histograms the fit parameters of the initial decay times. The initially set parameters of the simulations are printed in their corresponding plots and indicated by green lines. In all four fit parameters the two histograms coincide very well. This means that our analysis technique finds the correct decay times relatively close to delta-functions (which would be the best case) such that nearly the same decay curve fit results are obtained. Let us take a closer look at the single plots and try to understand them:

The distribution of the fitted decay constants λ looks rather broad. However this is a purely statistical effect. If the sample size was increased to several thousand simulations it would be clear to see that what is shown in Figure 5.7 is merely the peak of a distribution with lacking statistics. This will be proven in the next section, Section 5.5.

The distribution of the modulation frequencies ω looks rather sharp and well centered around the real (initial) set ω_{init}. This shows that with these initial parameters the oscillation should still be visible after the full analysis.
The distribution of the modulation amplitudes a is rather broad but well centered around the initial set value. Hence the uncertainty of a fitted a will be rather large in this setting but the value itself still lie in the correct region.
Even the distribution of the phase ϕ is well centered around its original value. Since by definition ϕ is confined to the region $[-\pi,+\pi]$ the fit results span a rather large area $\sim 20\%$.
So for this parameter set it can be acknowledged that if such an oscillatory behavior was present in the initial decay times it could also be found with the presented data analysis technique.
However this depends on all of the five input parameters (except maybe the phase). To get a feeling for these dependencies the average values and standard deviations of the found fit parameters of the five simulations introduced in tab. 5.1 are presented in tab. 5.2.

Table 5.2: Results of 100 sets of simulations for different input parameters. I and F denote the initial settings and the found fit parameters respectively. The found values derive from the average and standard deviation of 100 simulations. S denotes the simulation number and N the number of decays.

	S	N	λ $[s^{-1}]$	a	ω $[s^{-1}]$	ϕ [rad]
I	1	3000	0.0015	0.180	0.890	0.40
F	1	3000	0.0015 ± 0.0010	0.186 ± 0.026	0.888 ± 0.010	0.48 ± 0.38
I	2	3000	0.0224	0.230	0.885	-1.60
F	2	3000	0.0225 ± 0.0013	0.231 ± 0.028	0.892 ± 0.008	-1.83 ± 0.25
I	3	3600	0.0130	0.107	0.884	2.35
F	3	3600	0.0128 ± 0.0013	0.121 ± 0.023	1.006 ± 0.651	1.38 ± 2.20
I	4	3600	0.0126	0.134	0.882	1.78
F	4	3600	0.0124 ± 0.0011	0.140 ± 0.025	0.922 ± 0.529	1.99 ± 0.58
I	5	2500	0.0032	0.150	1.000	0.00
F	5	3000	0.0030 ± 0.0013	0.165 ± 0.031	1.098 ± 0.618	0.06 ± 0.51

In most cases the input parameters are reproduced with error behaviors for each parameter as discussed above. However in the cases of simulations 3, 4, and 5 the large errors and increased value of ω indicate that something happened. A closer investigation shows that in most simulations a completely normal behavior is found which mirrors the distribution behavior shown and discussed in Figure 5.7. However there are cases where the best ω is found at a very high value. These frequencies approach periods in order of the time binning of 64 ms so they clearly have nothing to do with the original oscillation. I defined the region in which fits are still accepted as showing

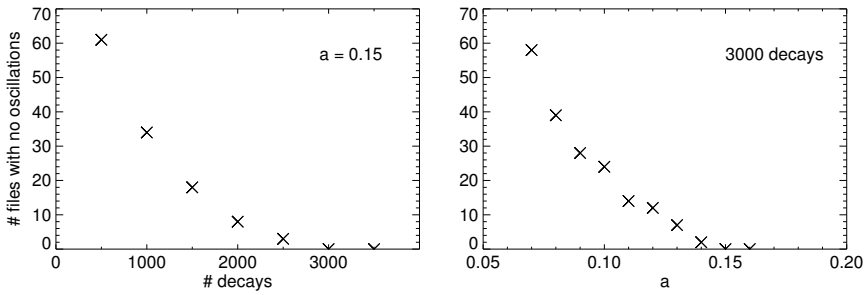

Figure 5.8: Different sets of simulations with varying input parameters and the number of simulations which show no oscillation after the analysis.

the initial frequency as:

$$\omega_{allowed} = \omega_{init} \pm_{0.1}^{0.15} [\mathrm{s}^{-1}] \tag{5.2}$$

So in the cases of simulations 3, 4, and 5 it is possible to create statistically distributed decay times which will not show an oscillatory behavior even though they initially possess it. What is special about these three cases? They have either small modulation amplitudes or small decay numbers so we can assume that these are the two main parameters which can make an intrinsic oscillation remain undetected.
To investigate this effect further I took the input parameters of simulation 5 and varied a and N to check the amount of simulations in which the oscillation is not found correctly (according to the criterion defined in eq. 5.2).
The results are shown in Figure 5.8. It can be observed that if the oscillation amplitude is small (< 0.12) in more than 15 % of the cases an oscillation would not be observable. Additionally if the decay statistics are low (< 2000 decays) the oscillation with an amplitude of 0.15 will vanish in more than 10 % of the cases. A combination of the two effects will of course increase the non-observability even further. While a is a constant which can not be influenced by the experimentalist, N could be increased by extending the beam time.

It has been shown in this section that for a wide range of cases, if an oscillation was present in the initial decay times, it would also be found with the analysis method. However depending on the size of the data set and the oscillation amplitude it is also possible that an oscillation can remain undetected by the analysis method.

5.5 Multiple decays

In Section 3.6 I discussed the necessity to select those files out of the data set which only show a single decay. While this reduces our data set considerably I will argue the necessity for this selection with some simulations in this section.
I simulated files with 10 injected ions and thus 10 possible decays. The decay times of these 10 ions were distributed according to a purely exponential power-law with a decay constant[2] of $\lambda = 0.012\ \mathrm{s}^{-1}$. The decay constant was chosen such that in sets of 16000 files there would be ~ 3000 files with decays within a time window of 64 seconds. These numbers roughly mirror our data set.
Since the decay times in each file are exactly known I could choose different sets of files as input for a decay curve fit. To check the statistical variations of the decay times I first performed a fit through all of the decay times. Then I looked at all the files with decays but extracted only the first decay time. In most cases this corresponds to single decays. However in about 300 files of the set (~ 10 % of cases), later decay times were ignored. A third fit was performed with the files that only housed exactly one decay.
To account for statistical variations and to improve statistics I performed this simulation 10000 times. For each simulation I performed the three fits and extracted the fit values of the decay constant. The results are plotted in Figure 5.9. The black line shows the initially set decay constant of $\lambda = 0.012\ \mathrm{s}^{-1}$.
The red plot is the histogram of the fit of all decay times. It is centered around the real value and its width corresponds to the statistical variation of the decay times. This width also explains the distribution of the values of λ in the previous section over a seemingly large region. Here it is observable that nearly all of the values lie within the region of $\lambda_{init} \pm 0.03$ as was also observed in the previous section.
The blue histogram is the histogram of decay constants of only those files in a set which have exactly one decay within the measurement window of 64 seconds. Here it is clear to see that it coincides very well with the real distribution, the red curve. This proves that the selection of these files is valid and will yield the correct decay constant (within certain error margins).
The green histogram contains the decay constants of those sets of files which contain at least one decay but only the earliest decay time was extracted. It can be observed that the decay constant in these cases is significantly shifted to higher values. Therefore even though the fraction of files contained in the set which contain multiple decays is only around 10 %, these files, if not eliminated from the set, greatly influence the fit of an exponential decay. This can be understood as follows:

[2] This corresponds to a half-life of $T_{1/2} \sim 1$ min.

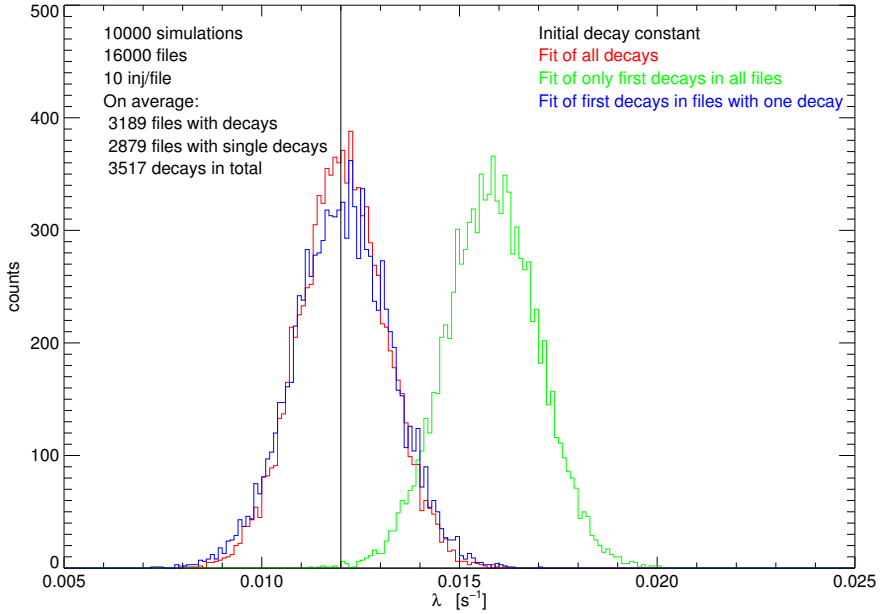

Figure 5.9: Histograms of decay constants retrieved from simulated data-sets in three different ways. Several parameters of the simulation are printed on the left side of the plot. The black line shows the initial set decay constant and marks the center of the blue and red distributions. The red curve is the result of decay curve fits through all of the decay times (including multiple decays) of the data sets. The blue curve only includes files with single decays. The green curve includes all files with decays but only uses the first decay time of each file.

Second or third decays in a file will generally happen rather late. So if these decay times are ignored there will be a lack of late decay times so earlier decay times will receive a higher significance in the fit than they should. This will increase the steepness of the decay curve which results in a higher value of λ.

I have shown in this section that the selection of single-decay files is valid and necessary. If files with multiple decays are left in the data set this will lead to wrong results including an artificially increased decay constant.

Chapter 6

Results and conclusions

In this chapter I present the results of my work. I compare them with results from the visual data analysis, provide a short summary of the thesis and draw some conclusions.

6.1 Results of the automated analysis

In the full data set of 16477 files there are 5150 files with found daughter traces. The first 6.4 seconds are ignored due to the stochastic cooling. During this time traces can not be well defined and thus no decay times assigned. In addition to this the last 4 seconds of the 64 seconds-measurement frame are ignored since much smaller detection efficiency is expected for such late starting traces. Daughter-traces may not be able to be identified since the length of their trace (i.e. the strength of an averaged signal) may be too small.
With this selection of a time window and the selection of files with exactly one decay (as described in Section 3.6) a set of 2664 files is left. The resulting decay curve is plotted in Figure 6.1.
The main aim of the data fits is to check whether the pure exponential fit M_0 is a good description of the decay behavior or if a modulated exponential fit M_1 significantly better describes the data. For this purpose both fitted curves have been plotted in red in all plots showing decay frequencies. It is difficult to see an oscillatory behavior with low statistics and a limited time window. To get an impression about the significance of a fitted frequency a frequency scan is supplied in every plot. To create this the data was fitted with M_1 with the angular frequency ω fixed at different values in the range of $[0.30, 6.28]\,\mathrm{s}^{-1}$. This corresponds to a period range of $[21, 1]\,\mathrm{s}$. The resulting negative llhs are plotted against ω (a low value in this plot corresponds to a high likelihood).
The frequency scan in Figure 6.1 indicates one designated frequency with a significantly higher llh-value than all the other frequencies at $\omega \sim 0.95$

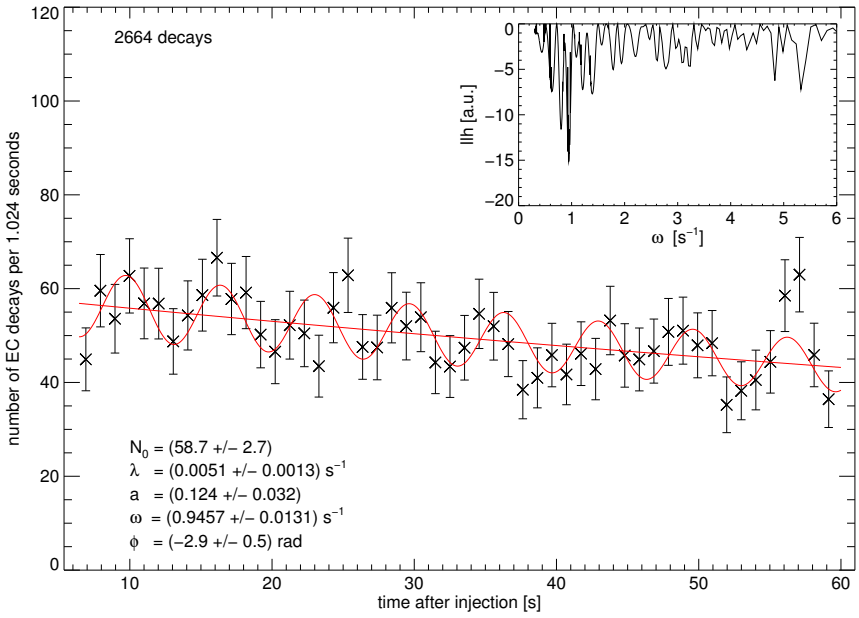

Figure 6.1: Final result of the automated data analysis: The binned sum of likelihoods fitted by a pure exponential M_0 and a modulated exponential M_1. Only files with a maximum lh-value within the time window of 6.4 to 60 seconds are taken into account. The frequency-scan in the upper right corner shows two peaks significantly higher Llh-values than all others. The fit parameters of the best fitting M_1 are printed in the lower left corner.

s^{-1}. The corresponding fit curve is the one shown in the figure. A second neighboring peak at $\omega \sim 0.80$ s^{-1} is also visible with a high llh-value. The difference of the likelihood values between those two peaks and between the second peak and the average level of the distribution is five orders of magnitude. This is a clear indication of an oscillation with $\omega \sim 0.95$ s^{-1} which corresponds to a period of $T \sim 6.6$.

To check the stability of the fit with respect to the chosen time window and rule out effects caused by the very beginning and the end of the time frame (these are the regions with the highest uncertainties with respect to finding traces and identifying them correctly) I repeated the fit with varied time windows. Two of these fits are plotted in Figure 6.2. The collective fit results are presented in tab. 6.1.

In these two examples a preferred frequency can not be easily distinguished any more. The peak around $\omega \sim 1$ observed in Figure 6.1 has become smaller and does not significantly exceed the llh-values of other frequencies. Since there is no unambiguously preferred ω in this plots a substitution of M_0 with M_1 can not be supported. These three fits show that the fit of M_1 is

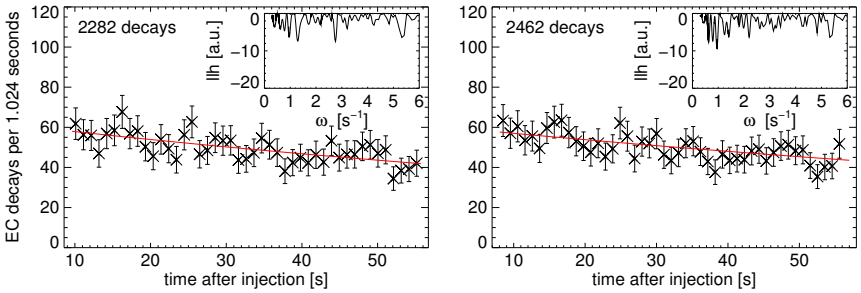

Figure 6.2: Decay curve fits of two different time windows. Left: time window from 9.6 to 55.7 seconds. Right: time window from 8. to 57. seconds. The new confines reduce the data set to 2282 and 2462 decays for left and right case, respectively. The fit results are displayed in tab. 6.1.

Table 6.1: Fit results of the three decay curves shown in Figure 6.1 and Figure 6.2. All three fits yield different parameters although most of the data are the same.

Figure	6.1	6.2 left	6.2 right
t_{min} [s]	6.4	9.6	8
$t_{max}[s]$	60	55.7	57
N_{dec}	2664	2282	2462
N_0	58.7 ± 2.7	62.1 ± 4.4	59.9 ± 2.7
λ [s^{-1}]	0.0051 ± 0.0013	0.0071 ± 0.0020	0.0056 ± 0.0013
a	0.124 ± 0.032	0.133 ± 0.047	0.101 ± 0.032
ω [s^{-1}]	0.9457 ± 0.0131	2.7497 ± 0.0258	0.9497 ± 0.0186
ϕ [rad]	-2.9 ± 0.5	-1.6 ± 0.9	-2.9 ± 0.7

not stable with respect to the chosen time window.
One part or this instability arises from the reduction of the data set. As has been shown in Section 5.4.3 a lower number of decay times reduces the detectability of an oscillation significantly as does a reduction in the oscillation amplitude. This may provide a reason for not detecting the same oscillation frequency.
To analyze the detectability of an oscillation with parameters as in the shown fits I ran simulations with the three parameter sets of tab. 6.1 as input. Out of 100 simulations 94 were correctly identified in the first set (2664 decays), only 25 in the second set (2282 decays) and 81 in the third set (2462 decays). This means that even if, by chance, the true decay distribution corresponded to one of the above parameter set, it would not necessaryly be detectable. The large effect in the second set can not be singularly attributed to the amplitude and the number of decays. In this case the large oscillation fre-

quency seems to play a role.

6.1.1 The Akaike Information Criterion

A more quantitative approach to decide whether the modulated exponential fit M_1 is a significantly better fit of the data than the pure exponential fit M_0 is the Akaike Information Criterion (AIC). It is an information-theoretic approach developed by Hirotugu Akaike [23]. The AIC value is a measure of how much information is lost if a model is used for the data rather than 'the truth'. Akaike starts with the Kullback-Leibler (K-L) distance [24] as a measure of information content, which is approximated (up to a constant) by eq. (6.1) with the maximized likelihood value L.

$$AIC = -2\log(L) + 2K \tag{6.1}$$

K is the number of parameters of a model. The model with the smallest AIC value is the best approximation for the information in the data. Since AIC is on a relative scale usually the AIC differences ΔAIC of models are considered as defined in eq. (6.2).

$$\Delta_i AIC = AIC_i - \min(AIC) \tag{6.2}$$

$\min(AIC)$ denotes the lowest AIC value of the set of models i considered. Since we only consider two models M_0 and M_1 there are only two AIC values. The model with the lower AIC value in every case was the modulated one. The ΔAIC values for the three cases discussed above are provided in tab. 6.2.

Table 6.2: ΔAIC results for three different time windows.

Figure	# decays	ΔAIC
6.1	2664	24.3
6.2 l	2282	9.1
6.2 r	2462	3.4

These values point towards the modulated decay model containing more information from the data set than the purely exponential one. However they extend over a large range considering that all of these sets describe the same data. It has to be remarked that a fit with more parameters will always fit a given set of data better than a fit with fewer parameters. Although this fact is included in the AIC definition the extent to which this is done may not be sufficient.
To check this I ran 100 simulations with a pure exponential as input and

fitted the resulting decay curves with M_0 and M_1. I found the corresponding ΔAIC values to be Gaussian distributed around $\Delta AIC = 14.4 \pm 5.1$ favoring the modulated decay in every case. This suggests that the minimum AIC value will always be from M_1 and a $\Delta AIC \gg 14.4$ is necessary to clearly favor M_1 above M_0. This is only marginally true for the first case but not for cases 2 and 3. Therefore the Akaike information criterion does not unambiguously support the rejection of M_0 in favor of M_1.

I have shown in this section that the fit of a modulated exponential decay to the data is instable with respect to the range of considered decay times. Additionally the application of the Akaike information criterion suggests that there is not enough evidence to reject the pure exponential decay in favor of an oscillatory modulated exponential decay.

6.2 Comparison with the visual analysis

Apart from the automated data analysis presented in this thesis an independent visual data analysis has been performed by another group. This was the sole analysis tool in the 2006 experiment [2]. The preliminary results allow a comparison of the analyses regarding the identification of daughter traces and decay times. As in previous instances a small error will be introduced by identifying the maximum of the likelihood-distribution as the actual decay time rather than using the full information of the likelihood-distribution.
The visual data analysis finds a total of 2570 files with single decays within the time window of 6.4 to 60 seconds compared to 2664 files in the automated analysis. These two data sets only have 2190 files in common. Therefore there are around 400 files in both sets in which only one of the analysis methods found a single EC decay. I took a closer look and compared the visual files with the full automated data set of 5150 files which includes multiple decays and decays outside the time window. In this comparison there are now 2558 files in common which leaves only 12 files where the visual data analysis finds a decay and the automated one does not. This points towards the multiple-decay-selection-criterion being the main reason for this large difference between the visual and automated data sets. This means that files with multiple decays are rejected differently in the two analyses. While the automated analysis selects mainly on the basis of the amplitude of the last 100 timebins of a frame, the visual analysis selects more intuitively by checking whether there is another sudden rise of the amplitude somewhere on the trace.
These differences can not at the moment be resolved, however they may lead to slightly different results of both analyses.
Considering the files in common, where both analyses find single decays within the time window, one can check whether the found decay times of

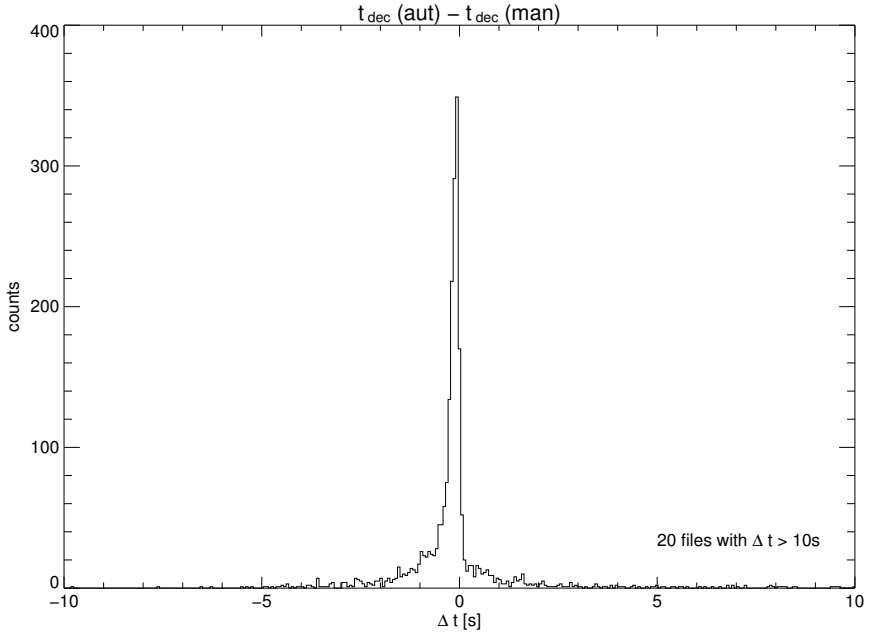

Figure 6.3: Differences of detected decay times in 2190 files found with both the visual and the automated data analysis. In most cases the difference is small and in the range of few time bins. 20 files have Δt values which did not fit into the frame.

these cases coincide. Hence the distribution of the differences in decay time of the common files are plotted in Figure 6.3.
The distribution is sharply peaked at close to $\Delta t = 0$ s. 80 % agree within ± 1 second, already 97% agree within ± 5 seconds. Thus with respect to periods of several seconds of possible oscillations this distribution is narrow.
The automated analysis seems to find the decay times on average a few time bins earlier than the visual analysis. This can be attributed to the functionality of the likelihood method: it may attribute a decay time to an earlier time bin even if this bin does not show an amplitude well above the average. However such a value is necessary for the visual analysis to identify the start of a trace.

6.3 Conclusions and outlook

6.3.1 Summary

The aim of this thesis was to create an analysis tool for subsequent Schottky spectra. An objective method was desired to retrieve most precise and

reproducible data about decay times which could be used to characterize the EC decay behavior of hydrogen-like ions. The motivation for this was to solve the "GSI puzzle" and remove a possible human bias from the visual data analysis.
A method was developed to find traces within Schottky spectra with a high noise level. The information of these traces was extracted in a way that minimizes the noise contribution and maximizes the information gain. A new likelihood method was employed to find the most likely time of a decay. Criteria were applied to rid the data set of unwanted properties. Finally unbinned likelihood fits were performed with the collective data sets for two models: a purely exponential decay M_0 and an oscillatory modulated exponential decay M_1. Several simulations were performed to characterize and optimize different steps of the analysis.
The development of an analysis tool was successful. There are indications for an oscillation with an angular frequency of $\omega \sim 0.95$. However this result is unstable with respect to the selected time range. There is not sufficient evidence which justifies the rejection of M_0 in favor of M_1.

6.3.2 Possible improvements

A possibility to further improve the analysis technique would be to apply it to different data sets like the 2006 or the 2010 experiment and compare the results. Some adjustments to the current procedure would be required since the different ion species have different frequency ranges and charge states. The latter influences the Schottky signal directly which is proportional to the charge (i.e. the number of mirror charges induced in every revolution). So the higher charges of $^{140}Pr^{48+}$ and $^{142}Pm^{60+}$ yield a better S/N ratio. Hence in general it is to be expected that an adjusted automated analysis would yield better results for $^{142}Pm^{60+}$ and $^{140}Pr^{48+}$ than for $^{122}I^{52+}$.
One could investigate the possibility of detecting multiple decays within daughter traces. With a larger data set the amplitude distributions of daughter traces could be characterized sufficiently well to unambiguously identify multiple decay cases and identify multiple decay times therein. This would increase the usable data set significantly.
As discussed in [8] a new 245 MHz resonator detector was developed which has a much higher sensitivity than the older capacitive pickup device. This leads to a more precise determination of decay times since even the cooling traces can be observed. Furthermore it allows the identification of multiple decays within one file which makes the selection of single-decay cases dispensable.

The last experiment to this date about an oscillatory behavior of a decay curve was performed in October 2014 with $^{142}Pm^{60+}$. The data analysis is

ongoing. While it will not be able to verify the results obtained in this thesis it may clarify unambiguously whether the oscillation exists in the Promethium case, once the data analysis is completed.
If the oscillations are confirmed it may lead to future experiments investigating this staggering effect.

References

[1] Yu. A. Litvinov et al. "Measurement of the β^+ and Orbital Electron-Capture Decay Rates in Fully Ionized, Hydrogenlike, and Heliumlike ^{140}Pr Ions". In: *Physical Review Letters* 99 (26 Dec. 2007), p. 262501 (cit. on pp. 1, 12).

[2] Yu.A. Litvinov et al. "Observation of non-exponential orbital electron capture decays of hydrogen-like 140Pr and 142Pm ions". In: *Physics Letters B* 664.3 (2008), pp. 162–168 (cit. on pp. 1, 7, 9, 12–16, 20, 24, 62, 72).

[3] B. Schlitt et al. "Schottky mass spectrometry at the heavy ion storage ring ESR". In: *Hyperfine Interactions* 99.1 (1996), pp. 117–125 (cit. on p. 5).

[4] B. Schlitt. "Schottky Mass Spectrometry at the Heavy Ion Storage Ring ESR". PhD thesis. Heidelberg: Ruprecht-Karls-Universität Heidelberg, Naturwissenschaftlich-Mathematische Gesamtfakultät, June 1997 (cit. on p. 5).

[5] M. Steck et al. "Cooled heavy ion beams at the ESR". In: *Nuclear Physics A* 626.1–2 (1997), pp. 495–498 (cit. on p. 5).

[6] D.R. Atanasov et al. "Half-life measurements of stored fully ionized and hydrogen-like 122 I ions". In: *The European Physical Journal A* 48.2 (2012) (cit. on p. 6).

[7] Richard B. Firestone and Virginia S. Shirley. *Table of Isotopes*. 8th ed. John Wiley & Sons,Inc, 1996 (cit. on p. 7).

[8] P. Kienle et al. "High-resolution measurement of the time-modulated orbital electron capture and of the decay of hydrogen-like 142Pm60+ ions". In: *Physics Letters B* 726.4–5 (2013), pp. 638–645 (cit. on pp. 7, 10, 12, 16, 18, 19, 74).

[9] Yuri A Litvinov and Fritz Bosch. "Beta decay of highly charged ions". In: *Reports on Progress in Physics* 74.1 (2011), p. 016301 (cit. on p. 8).

[10] URL: https://www.gsi.de/fileadmin/_migrated/pics/GSI_3D_03_05_300dpi_neu.jpg (visited on 10/03/2014) (cit. on p. 8).

[11] H. Geissel et al. “The GSI projectile fragment separator (FRS): a versatile magnetic system for relativistic heavy ions”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 70.1–4 (1992), pp. 286–297 (cit. on p. 9).

[12] Bernhard Franzke. “The heavy ion storage and cooler ring project ESR at GSI”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 24–25, Part 1 (1987), pp. 18–25 (cit. on p. 9).

[13] H. Geissel et al. “Ions penetrating through ion-optical systems and matter — non-liouvillian phase-space modelling”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 282.1 (1989), pp. 247–260 (cit. on p. 9).

[14] Zygmunt Patyk et al. “Orbital electron capture decay of hydrogen- and helium-like ions”. In: *Physical Review C* 77 (1 Jan. 2008), p. 014306 (cit. on p. 12).

[15] G T Emery. “Perturbation of Nuclear Decay Rates”. In: *Annual Review of Nuclear Science* 22.1 (1972), pp. 165–202 (cit. on p. 20).

[16] A. N. Ivanov and P. Kienle. “Time Modulation of the K-Shell Electron Capture Decay Rates of H-like Heavy Ions at GSI Experiments”. In: *Phys. Rev. Lett.* 103 (6 Aug. 2009), p. 062502 (cit. on p. 20).

[17] A. Merle. “The GSI oscillation mystery”. In: *Progress in Particle and Nuclear Physics* 64.2 (2010), pp. 445–447 (cit. on p. 20).

[18] H. Burkhardt et al. “Oscillations in the GSI electron capture experiment”. In: (2008). arXiv: 0804.1099 [hep-ph] (cit. on p. 20).

[19] Carlo Giunti. “The GSI Time Anomaly: Facts and Fiction”. In: *Nuclear Physics B - Proceedings Supplements* 188 (2009), pp. 43–45 (cit. on p. 20).

[20] F. Giacosa and G. Pagliara. “(Oscillating) non-exponential decays of unstable states”. In: *PoS* BORMIO (2012), pp. 28–37 (cit. on p. 20).

[21] A. W. F. Edwards. *Likelihood: an account of the statistical concept of likelihood and its application to scientific inference*. Cambridge Univ. Press, 1972 (cit. on p. 44).

[22] URL: http://seal.web.cern.ch/seal/snapshot/work-packages/mathlibs/minuit/ (cit. on p. 52).

[23] Kenneth P. Burnham and David R. Anderson. *Model selection and multimodel inference: a practical information-theoretic approach*. 1st ed. New York: Springer, 1998 (cit. on p. 71).

[24] S. Kullback and R. A. Leibler. “On Information and Sufficiency”. In: *The Annals of Mathematical Statistics* 22.1 (1951), pp. 79–86 (cit. on p. 71).

Curriculum Vitae

Christoph Klaushofer

2012 -	**University of Vienna**
	MSc Physics
2008 - 2012	**University of Vienna**
	BSc Physics
2007 - 2008	**Sportmedizinisches Institut, Salzburg**
	Civilian Service
1999 - 2007	**Musisches Gymnasium, Salzburg**
	Matura