IDENTIFICATION OF SINGLE BARIUM ATOMS WITH RESONANCE IONIZATION MASS SPECTROSCOPY FOR THE nEXO NEUTRINOLESS DOUBLE BETA DECAY EXPERIMENT

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Scott Kravitz
July 2017

This dissertation is online at: http://purl.stanford.edu/ys086xx4482

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Giorgio Gratta, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Leo Hollberg

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Hari Manoharan

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

The mechanism for neutrino mass generation is an open question in particle physics, with implications for new physics at higher mass scales. If neutrinos are their own antiparticle, meaning they have Majorana masses, this property can be measured through a rare process called neutrinoless double beta decay ($0\nu\beta\beta$). EXO-200 is an experiment searching for $0\nu\beta\beta$ of ^{136}Xe using a liquid xenon time projection chamber. The first part of this work describes a search for the two-neutrino double beta decay of ^{136}Xe to an excited state of ^{136}Ba with EXO-200 data using machine learning techniques. A measurement of the half-life of this decay can provide valuable input to nuclear models and reduce systematic uncertainties of the $0\nu\beta\beta$ process. The second part of this work describes efforts to greatly improve the sensitivity of a future experiment similar to EXO-200 through the identification of the ^{136}Ba daughter produced in the double beta decay of ^{136}Xe. In the technique explored here, barium ions are adsorbed onto a substrate, where they neutralize, and are then transported to a separate identification chamber. They are subsequently desorbed using laser-induced thermal desorption (LITD) and selectively ionized using resonance ionization spectroscopy (RIS), followed by identification using a time of flight mass spectrometer. Improvements in understanding of the LITD and RIS processes, backgrounds to the technique, and substrate cleaning methods are presented, bringing the technique closer to single Ba atom sensitivity.

Acknowledgments

I have many people to thank for the support, guidance, and friendship which have enabled me to complete this thesis work and made my time at Stanford truly exceptional. The best part of my graduate experience has been the community of motivated, welcoming, and genuine people here and I will miss seeing them regularly.

Thanks to Giorgio for giving me the room to pursue this work as I saw fit while also insisting that my experiments be as thorough as possible. His excitement for the prospects of this work has been contagious. Thanks also to Bill Fairbank for being a constant source of knowledge on barium tagging and for always being an enthusiastic collaborator. Thanks to the nEXO collaboration, which has shown me a microcosm of the particle physics community, and has provided many creative and useful ideas over the years. I am grateful to all of the staff in Varian, from the machinists to the front office to Marcia, for doing all I have asked of them and much more.

The members of the Gratta group have helped me develop as a scientific researcher and have been perpetually willing to help me troubleshoot issues or come up with new ideas. Karl Twelker in particular was an excellent mentor and taught me how to do real, honest experimental physics despite my inexperience with Conflat flanges. Dave Moore and Igor Ostrovskiy taught me the right approach to data analysis, which always includes at least one more cross-check. Dan Fudenberg has always been there to scheme up the next experiment with me or to commiserate on the frustrations of barium tagging.

Thanks to all of my friends, at Stanford and elsewhere, who constantly remind me that there are things more important than research. I look forward to continuing to be a part of your lives from across the Bay or across the country. To everyone at Goggles Optional: it has been so much fun podcasting with you, and I hope to continue spreading knowledge through bad jokes with my spare time.

Special thanks to my parents, who have always believed in me and pushed me to be my best - but also to avoid getting caught up in the little things that go wrong. And also to Kenna: spending time with you is always the brightest part of my day, and I cannot wait to share that brightness with you every day.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Neutrinos are electrically-neutral subatomic particles which pervade our Universe. They are present as relics of the early Universe [1]; as messengers from supernovae, carrying even more energy than the light from these explosions [2]; as escapees from deep beneath the crust of the Earth [3]; and as radiation from the Sun. The flux of neutrinos from the Sun at the surface of the Earth is on the order of $10^{11}/\text{cm}^2/\text{s}$ [4], yet neutrinos remain one of the least-understood particles of the Standard Model of particle physics due to their exceptionally weak interactions with normal matter. A better understanding of neutrinos can improve models of astrophysical and cosmological processes, as well as provide clues to unresolved questions that go beyond the Standard Model, such as dark matter and the hierarchy problem.

To date, three types or "flavors" of neutrinos have been experimentally-detected: ν_e, ν_μ, and ν_τ, corresponding to the three generations of charged leptons, e, μ, and τ [5]. There are also three different neutrino masses, m_1, m_2, and m_3, though the neutrino flavor eigenstates do not correspond to the mass eigenstates ν_1, ν_2, and ν_3. Neutrino oscillation experiments have measured the splittings between the neutrino masses [6], while beta decay measurements [7, 8] and astrophysical experiments have set upper bounds on the sum of the neutrino masses [9], but the absolute neutrino mass scale remains unknown. These limits establish neutrinos as peculiarly-light particles, with masses at least five orders of magnitude below the next lightest particle, the electron. In addition, the mass-generation mechanism for neutrinos has not been

determined. Hence, it is possible for neutrinos to be either Dirac particles, as are all other matter in the Standard Model, or Majorana particles, which are their own antiparticles [10]. This latter case is particularly well-motivated, as it could arise from a natural mechanism for producing the low neutrino mass scale while also pointing to new, heavy particles.

Both of these questions, the absolute neutrino mass scale and the Majorana or Dirac nature of neutrinos, can be answered through the observation of a rare process known as neutrinoless double beta decay ($0\nu\beta\beta$). Chapter 2 of this dissertation provides more detail on the underlying theory of Majorana neutrinos and $0\nu\beta\beta$. EXO-200 and its future successor, nEXO, are experiments searching for $0\nu\beta\beta$ in ^{136}Xe. The work presented in this dissertation has two primary components: a search for the double beta decay of ^{136}Xe to an excited state of ^{136}Ba with EXO-200 data, which can provide valuable input to models of the $0\nu\beta\beta$ process (Chapter 3); and development of a system to greatly improve the sensitivity of nEXO or a similar future experiment through the identification of the ^{136}Ba daughter produced in the $0\nu\beta\beta$ of ^{136}Xe (Chapters 4-9).

Chapter 2

Neutrino Mass and Double Beta Decay

2.1 Neutrino Mass

The Standard Model (SM) of particle physics was originally developed under the assumption that neutrinos are massless particles. Recently, neutrino oscillation measurements have established that neutrinos do have mass, earning the experimenters the 2015 Nobel Prize in Physics [6]. However, the masses of the neutrinos are at least five orders of magnitude smaller than that of the electron, the lightest of all other massive particles in the SM, which may point to an entirely new mass generation mechanism. In fact, there is no way to generate neutrino masses in the SM, including the standard Higgs mechanism that gives quarks and charged leptons their masses, without coupling neutrinos to additional fields. These two points - the surprisingly-small masses of the neutrinos and the requirement of new fields to generate them -

establish the study of neutrino masses as an excellent means of probing physics at higher energies.

Complicating the study of neutrino masses is the fact that the eigenstates of their mass matrix do not correspond to the eigenstates of the weak interactions they participate in. A neutrino produced by a weak interaction such as β decay is created as a flavor eigenstate ν_e, ν_μ, or ν_τ, which is a linear superposition of the mass eigenstates ν_1, ν_2, and ν_3. The mass eigenstates are what govern the propagation of the neutrino, so the different mass components of the initial flavor eigenstate each propagate differently, picking up phases that give rise to neutrino oscillation [11]. The Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix U is the matrix which transforms the mass basis into the flavor basis, and hence describes the particular linear combination of mass eigenstates corresponding to any given flavor eigenstate:

$$|\nu_\alpha\rangle = \sum_i U^*_{\alpha i} |\nu_i\rangle \tag{2.1}$$

where $\alpha = e, \mu, \tau$ are the flavor indices and $i = 1, 2, 3$ are the mass indices.

The PMNS matrix can be conveniently parametrized in terms of three mixing angles $(\theta_{12}, \theta_{23}, \theta_{13})$ and three phases $(\delta, \alpha_1, \alpha_2)$ [12] as given below:

$$
\begin{aligned}
U &= \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix} \\
&= \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{bmatrix} \\
&= \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \begin{bmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{bmatrix}
\end{aligned}
\tag{2.2}
$$

where $c_{ij} = \cos(\theta_{ij})$, $s_{ij} = \sin(\theta_{ij})$, and the matrix has been factorized to highlight its structure as a mixing of each of the three pairs of neutrinos plus the CP-violating Dirac phase δ and two Majorana phases which reduce to the identity matrix ($\alpha_1 = \alpha_2 = 0$) unless the neutrino is a Majorana particle as described in Sec. 2.2.

For the sake of illustrating the phenomenon of neutrino oscillation, consider a simplified case with only two neutrino flavor or mass eigenstates, in which case there is a single mixing angle θ and $U = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}$. Each mass eigenstate ν_i picks up a phase $e^{-iE_it/\hbar}$ during its propagation over time t, so the probability of a neutrino starting out in flavor α and being detected as the other flavor β is given by

$$
\begin{aligned}
P\left(\nu_\alpha \to \nu_\beta\right) &= \left|\langle \nu(t)|\nu_\beta\rangle\right|^2 \\
&= \left|\left(\cos(\theta)e^{iE_1 t/\hbar}\langle\nu_1| + \sin(\theta)e^{iE_2 t/\hbar}\langle\nu_2|\right)\left(-\sin(\theta)|\nu_1\rangle + \cos(\theta)|\nu_2\rangle\right)\right|^2 \\
&= \cos(\theta)^2\sin(\theta)^2\left|e^{iE_2 t/\hbar} - e^{iE_1 t/\hbar}\right|^2 \\
&= \sin^2(2\theta)\sin^2\left(\frac{(E_1 - E_2)t}{2\hbar}\right) \\
&= \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right)
\end{aligned} \tag{2.3}
$$

where $\Delta m^2_{ij} = m^2_i - m^2_j$ and we've assumed the neutrino to be in the ultrarelativistic regime where $E_i \approx E + \frac{m_i^2}{2E}$ for an average neutrino energy E and the oscillation distance $L \approx t/c$. This form only depends on the mass squared splitting Δm^2_{ij}, so oscillation experiments are in general not sensitive to the sign of this splitting nor to the absolute scale of the neutrino masses. With this in mind, the current best estimates of the neutrino oscillation parameters, averaged from multiple oscillation experiments [5], are given in Table 2.1. Experiments have yet to determine the CP-violating phase δ, the mass ordering (sign of Δm^2_{32}), and the absolute mass scale (equivalently, the minimum neutrino mass), as well as the Majorana phases which do not affect neutrino oscillation.

Parameter	Best Fit	3σ
Δm^2_{21} $[10^{-5}\text{eV}^2]$	7.37	6.93 – 7.97
$\lvert\Delta m^2\rvert$ $[10^{-3}\text{eV}^2]$	2.50 (2.46)	2.37 – 2.63 (2.33 – 2.60)
$\sin^2\theta_{12}$	0.297	0.250 – 0.354
$\sin^2\theta_{23}$	0.437 (0.569)	0.379 – 0.616 (0.383 – 0.637)
$\sin^2\theta_{13}$	0.0214 (0.0218)	0.0185 – 0.0246 (0.0186 – 0.0248)

Table 2.1: Neutrino oscillation parameters. By convention, the masses are labeled so that ν_1 and ν_2 have the smallest mass squared splitting and $m_1 < m_2$. Values (values in parentheses) correspond to the normal ordering $m_1 < m_2 < m_3$ (inverted ordering $m_3 < m_1 < m_2$). Because $\Delta m^2_{31} \approx \Delta m^2_{32}$, an average splitting is defined: $\Delta m^2 \equiv m^2_3 - (m^2_2 + m^2_1)/2$.

In addition to oscillation experiments, measurements of the endpoint of tritium β decay [7, 8] have set limits on an average neutrino mass

$$m_\beta \equiv \sqrt{\sum_i |U_{ei}|^2 m_i^2} < 2.05 \text{ eV}, \ 95\% \text{ CL} \tag{2.4}$$

while cosmological measurements of baryon acoustic oscillations from the Planck Collaboration [9] set limits on the sum of neutrino masses

$$\sum_i m_i \ < \ 0.23 \text{ eV}, \ 95\% \text{ CL} \tag{2.5}$$

As previously noted, these limits establish neutrinos as abnormally light as compared to the other massive particles in the SM. It remains a mystery as to why this is (if there is indeed a natural explanation for their low masses), as well as how these masses are generated. One possible mechanism for generating neutrino masses within the SM framework is via the standard Higgs mechanism, which produces a Dirac mass term in the Lagrangian after electroweak symmetry breaking, as with the masses of all other quarks and leptons in the SM:

$$\mathcal{L}_D = -m_D \overline{\Psi}\Psi = -m_D \left(\overline{\Psi_L}\Psi_R + \overline{\Psi_R}\Psi_L \right) \tag{2.6}$$

where the usual four-component spinor Ψ has been separated into its left-handed and right-handed chirality components $\Psi_{L,R} \equiv [(1 \pm \gamma_5)/2]\Psi$.[1] However, the SM only contains the left-handed component of the neutrino field, as experiments have only definitively-established two degrees of freedom for each neutrino flavor - typically simply referred to as neutrinos (produced in *e.g.* β^+ decay) and their weak interaction opposites referred to as antineutrinos (from *e.g.* β^- decay). A minimal extension involves adding three right-handed analogues to the left-handed neutrinos we observe [13]. To preserve the symmetries of the SM, these new right-handed neutrinos must be sterile - they have no gauge quantum numbers, meaning that they do not participate in

[1] Note that $\overline{\Psi_L}$ has right-handed chirality, but for a general SM field is considered to be "left-handed" in the sense of being charged under the left-handed SU(2) gauge. For consistency, I will use the term "left-handed chirality" when I want to refer to Ψ_L or $\overline{\Psi_R}$, and simply "left-handed" to refer to the components charged under the SU(2) gauge Ψ_L and $\overline{\Psi_L}$.

the electromagnetic, weak, or strong interactions. However, this requires the Yukawa coupling strength to the Higgs field (a free parameter in this minimal extension) to be several orders of magnitude smaller than the equivalent coupling parameter for electrons. While there is nothing in principle to stop such an unnatural pattern of coupling strengths, alternative mechanisms exist which provide a more elegant and natural explanation for the low neutrino masses - mechanisms which require neutrinos to be Majorana particles.

2.2 Majorana Neutrinos

An alternative form for the mass term in the Lagrangian to Eq. 2.6 is called a Majorana mass:

$$\mathcal{L}_M = -\frac{1}{2} m_L \left[\overline{(\Psi_L)^c} \Psi_L + h.c. \right] \tag{2.7}$$

where Ψ^c is the charge conjugate of Ψ. A similar term can also be constructed from only the right-handed parts of the field, Ψ_R, though in principle only two degrees of freedom are needed to construct a Majorana mass.

For the sake of generality, let us consider an extension to the SM which includes both left-handed neutrinos (three of which have been observed in nature) and right-handed neutrinos [10]:

$$\nu_L \equiv \begin{bmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \\ \vdots \end{bmatrix}, \ \nu_R \equiv \begin{bmatrix} \nu_{eR} \\ \nu_{\mu R} \\ \nu_{\tau R} \\ \vdots \end{bmatrix}, \ n_L \equiv \begin{bmatrix} \nu_L \\ (\nu_R)^c \end{bmatrix} \tag{2.8}$$

Using these fields, the most general mass term we can construct is given by

$$\mathcal{L}_{M+D} = -\frac{1}{2} \overline{(n_L)^c} \mathcal{M} n_L + h.c., \ \mathcal{M} = \begin{pmatrix} \mathcal{M}_L & \mathcal{M}_D^T \\ \mathcal{M}_D & \mathcal{M}_R \end{pmatrix} \tag{2.9}$$

If the eigenvalues of the mass matrix $\mathcal{M}$ are not degenerate, we can always find a

basis where $\mathcal{M}$ is diagonal with positive entries m_k and the mass term takes the form

$$\mathcal{L}_{M+D} = -\frac{1}{2}\sum_k m_k \overline{\phi_k}\phi_k \tag{2.10}$$

and the new mass eigenstates $\phi_k = \phi_k^c$. This condition, where the mass eigenstates are equal to their own charge conjugates, is often referred to as the Majorana condition. This condition is also frequently paraphrased as saying that Majorana neutrinos are their own antiparticles. Note that it is the mass eigenstates ϕ_k that satisfy this condition (and have no well-defined chirality), while the left-handed interaction eigenstates are formed from the left-handed components of these mass eigenstates.

To preserve the gauge symmetries of the SM Lagrangian, we require $\mathcal{M}_L = 0$ or else the introduction of modifications to *e.g.* the Higgs field [14]. However, $\mathcal{M}_R$ will in general be non-zero. Only in the special case where $\mathcal{M}_L = \mathcal{M}_R = 0$ and the (positive) mass eigenvalues become degenerate do we get neutrinos with a Dirac mass - so we see that Majorana neutrinos are in some sense the more general expectation. In this special Dirac case, there is an additional U(1) symmetry corresponding to lepton number conservation (and the lack of Majorana phases in the mixing matrix). Because there is no need for the SM to respect this global symmetry, however, this is often referred to as an accidental symmetry of the model. If it turns out that neutrinos are Majorana particles, this symmetry is broken and replaced with the approximate symmetry of helicity, which is violated at the level of $(m_\nu/E_\nu)^2$. In other words, for most neutrinos where $E_\nu \gg m_\nu$, this violation of lepton number is very small, and measuring such an effect is quite difficult. It is only because the neutrino masses are so small that the question of their Dirac or Majorana nature is still unanswered.

If we consider a different case where the eigenvalues m_R of $\mathcal{M}_R$ satisfy $m_R \gg m_D$, then we are in the regime of the see-saw mechanism [10]. In this case, we get heavy sterile neutrinos plus light Majorana neutrinos with mass of order $m \approx m_D^2/m_R$. We expect m_D to be comparable to the Higgs vacuum expectation value when generated from the usual Yukawa coupling term, in which case the masses of the sterile neutrinos must be on roughly the order of the GUT scale (where the SM gauge couplings are approximately equal) for the light neutrinos to have the masses they do. Hence, this

scenario explains the light masses of the neutrinos by positing sterile neutrinos at a much higher mass scale.

One final reason this model of neutrino masses is well-motivated arises if we take a more model-independent approach, where the SM is treated as an effective field theory of a higher-energy complete theory. In this view, non-renormalizable terms are allowed in the Lagrangian and all non-SM fields are integrated out. The only term of mass dimension 5 (the lowest-order correction) consistent with the SM gauge symmetries is then one which gives rise to Majorana masses for the neutrinos [15].

If neutrinos are indeed Majorana particles, then lepton number is not exactly conserved in nature, which means lepton-number-violating processes such as neutrinoless double beta decay can be experimentally searched for.

2.3 Double Beta Decay

Double beta decay is a rare, second-order process where an atomic nucleus undergoes two simultaneous beta decays - that is to say that two of its neutrons are converted into protons through the emission of two electrons and two electron antineutrinos, as depicted in Fig. 2.1. While double beta decay (hereafter $2\nu\beta\beta$) can occur as long as the final daughter nucleus has lower energy than the initial nucleus, in practice it is only observable if single beta decay is also energetically forbidden, as the rate of $2\nu\beta\beta$ is otherwise dwarfed by the beta decay signal. Despite its half-life of 10^{18} yr or more, $2\nu\beta\beta$ has been observed in several isotopes including ^{136}Xe, which was first measured by the EXO-200 experiment [16] (described further in Chapter 3).

An alternative, even more rare version of double beta decay is neutrinoless double beta decay ($0\nu\beta\beta$), in which no neutrinos are emitted. This hypothetical process violates lepton number conservation by two units due to the lack of electron antineutrinos. There are multiple possible channels through which $0\nu\beta\beta$ can occur [17, 18, 19], but light neutrino exchange (Fig. 2.1) is the most commonly considered, and is motivated by extensions to the SM which require neutrinos to be Majorana particles. Even if light neutrino exchange is not the primary contributor to the $0\nu\beta\beta$ decay rate,

however, observation of $0\nu\beta\beta$ would necessitate that neutrinos have a non-zero Majorana mass [20]. In fact, $0\nu\beta\beta$ is widely considered to be the most practical means of determining the Majorana or Dirac nature of neutrino masses.

If we assume that light neutrino exchange dominates the $0\nu\beta\beta$ process, then the half-life for this process is given by [10]

$$[T^{0\nu}_{1/2}]^{-1} = G_{0\nu}(Q, Z) \left| M_{0\nu}(Z, A) \right|^2 \langle m_{\beta\beta} \rangle^2 \tag{2.11}$$

where $G_{0\nu}(Q, Z)$ is a phase space factor that depends on the nucleus itself and the difference in initial and final nucleus energies $Q = E_f - E_i$; $M_{0\nu}(Z, A)$ is a nuclear matrix element which depends on the wavefunction overlap of the initial and final nuclear states, which is challenging to calculate due to the intricacies of nuclear structure models; and $\langle m_{\beta\beta} \rangle^2$ is the effective Majorana mass defined as

$$\langle m_{\beta\beta} \rangle \equiv \left| \sum_i U_{ei}^2 m_i \right| \tag{2.12}$$

The proportionality of the half-life to this effective Majorana mass can be seen from the Feynman diagram in Fig. 2.1, where each mass eigenstate contributes an amplitude proportional to U_{ei}^2 (one factor of the mixing matrix for each neutrino vertex, to go from electron flavor eigenstate to mass eigenstate ν_i) and a factor of m_i from the Majorana mass vertex. Note that even in the case of Majorana neutrinos, because the U_{ei}^2 factors carry phases and hence can cancel each other out, it is possible for the rate of $0\nu\beta\beta$ to be arbitrarily small.

Oscillation parameters provide some information on the allowed U_{ei} values, which in turn set constraints on the possible effective Majorana mass as a function of the minimum neutrino mass (Fig. 2.2). The most valuable contribution of $0\nu\beta\beta$ experiments is to search for new physics through measuring, or setting limits on, the half-life of this process. Multiple experiments have searched for $0\nu\beta\beta$ in several different isotopes, with some recent limits summarized in Table 2.2. In addition, as Fig. 2.2 demonstrates, if $0\nu\beta\beta$ experiments succeed in measuring the process's half-life, they can also potentially say something about the absolute neutrino mass scale and the mass ordering.

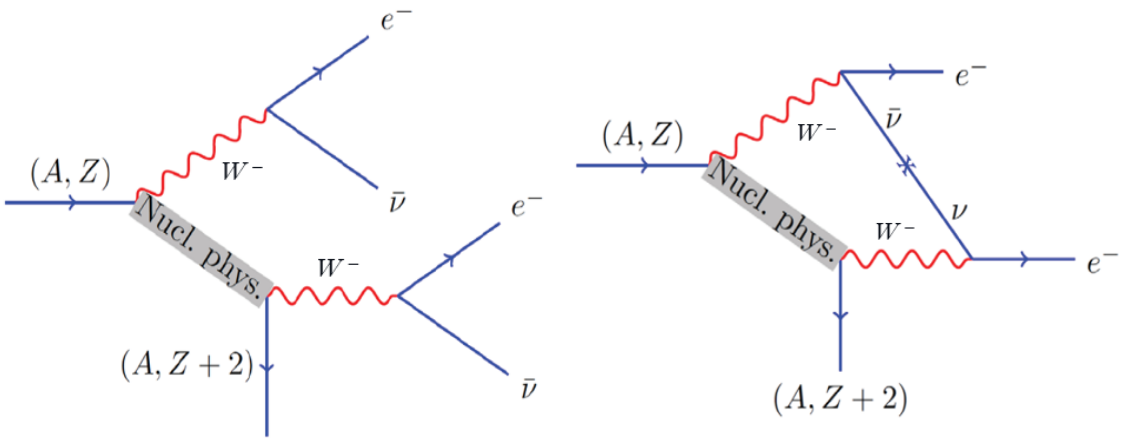

Figure 2.1: Feynman diagrams for two-neutrino double beta decay (left) and neutrinoless double beta decay via light neutrino exchange (right).

Isotope	$T^{0\nu}_{1/2}$ Limit (10^{25} yr)	Experiment	Reference
^{136}Xe	10.7	KamLAND-Zen	[21]
^{136}Xe	1.1	EXO-200	[22]
^{76}Ge	2.1	GERDA	[23]
^{130}Te	0.4	CUORE	[24]
^{100}Mo	0.11	NEMO-3	[25]

Table 2.2: Limits on the $0\nu\beta\beta$ half-life at 90% confidence for several experiments and isotopes.

Because the lack of neutrinos in $0\nu\beta\beta$ alters the kinematics of the decay, it can be distinguished from $2\nu\beta\beta$ by very precisely measuring the energy of the electrons that are emitted. In fact, because these electrons are orders of magnitude lighter than the nuclei they come from, they carry away essentially all the energy of the decay, Q, so that the signature of $0\nu\beta\beta$ in the electron energy spectrum is a monoenergetic peak at the endpoint of the $2\nu\beta\beta$ spectrum, with a width only limited by the energy resolution of the detector (Fig. 2.3). Measuring the energy of the emitted electrons is hence the primary means of searching for $0\nu\beta\beta$ in all current detectors. However, because this process is so rare, any signal in a real detector will be completely masked

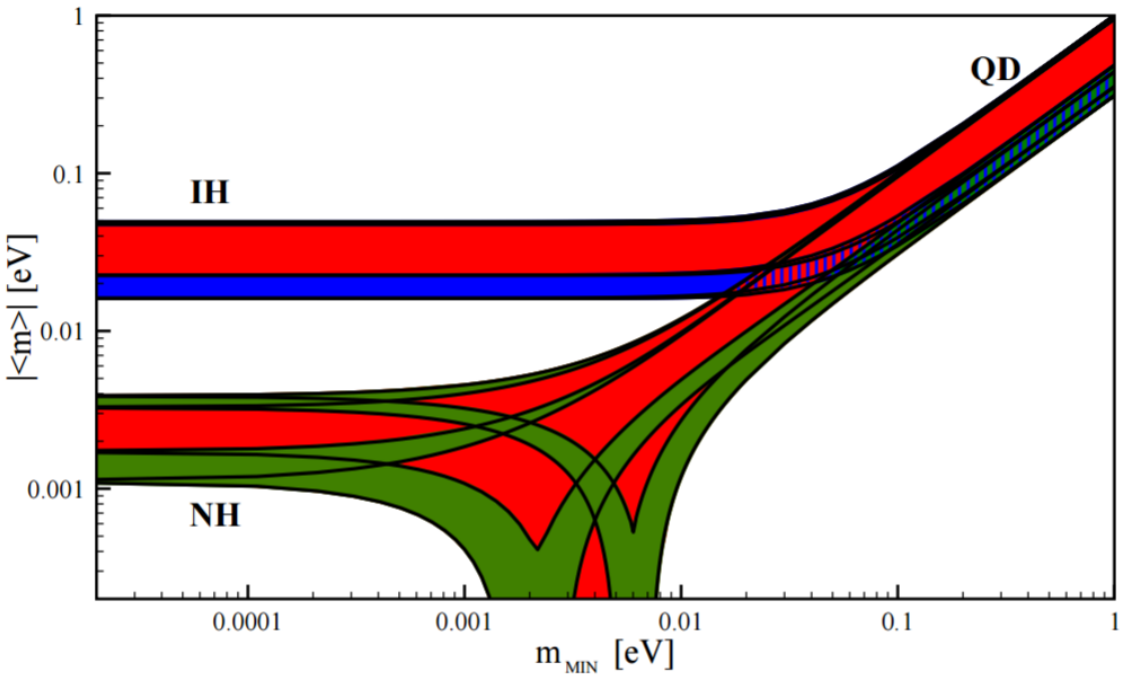

Figure 2.2: Allowed ranges for the effective Majorana mass $|\langle m \rangle|$, as a function of the minimum neutrino mass. Contours are drawn using best fit values and 2σ confidence intervals for neutrino oscillation parameters for the inverted mass ordering (IH), normal ordering (NH), and quasi-degenerate ordering (QD). Red regions correspond to Majorana phases that are CP-violating, while the blue and green regions have CP-conserving Majorana phases. Taken from [5].

by background events if care is not taken toward minimizing them. Successful observation is likely to require a combination of large isotope mass and run time, good energy resolution, exceptionally low backgrounds (both from trace radioactivity in detector components and cosmogenics), and more clever means of distinguishing between the signal and background using additional information such as event position. Chapter 3 gives a brief introduction to EXO-200, an experiment searching for $0\nu\beta\beta$ in ^{136}Xe. It then describes one way of taking advantage of the additional information beyond event energy available to EXO-200 to search for a rare version of $2\nu\beta\beta$ using machine learning methods. Chapters 4–9 describe a more ambitious method to eliminate backgrounds altogether through the identification of the daughter ^{136}Ba atom.

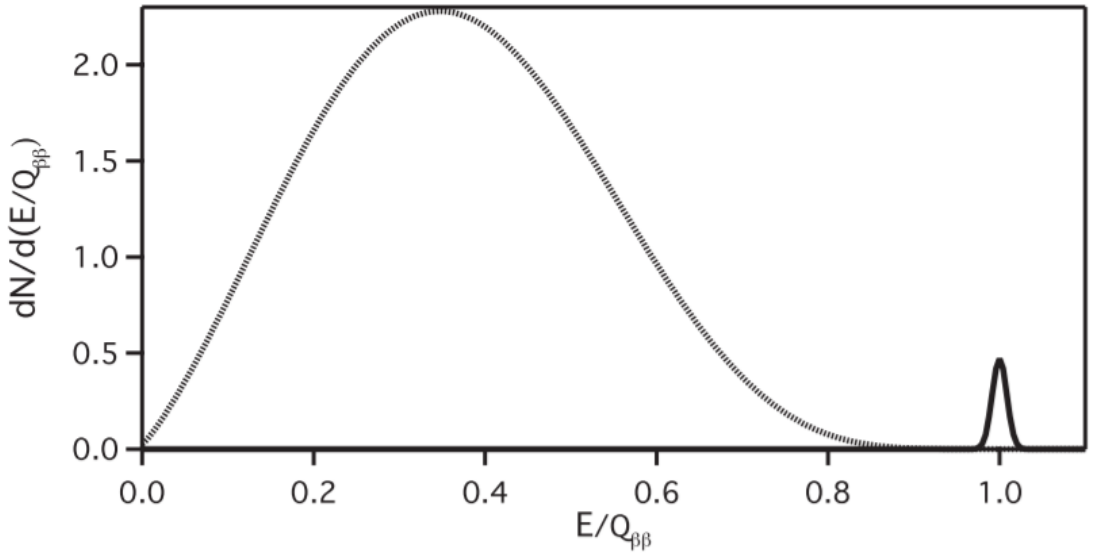

Figure 2.3: Distribution of summed electron energies in $2\nu\beta\beta$ (dashed curve) and $0\nu\beta\beta$ (solid curve) assuming the rate for the latter is 1% of the former (a substantial exaggeration, for visibility) and a 1σ energy resolution of 2%. Taken from [10].

Chapter 3

Search for $2\nu\beta\beta$ decay of ^{136}Xe to the 0_1^+ excited state of ^{136}Ba

3.1 Liquid Xenon Detector Advantages

Among the different $0\nu\beta\beta$ candidate isotopes, ^{136}Xe has several advantages. ^{136}Xe has a fairly high natural abundance of 9% which can be enriched to 80% or more with little difficulty compared to solid state isotopes such as ^{76}Ge. Xe can also be continuously recirculated through purifiers to remove dissolved radioisotopes and electronegative impurities. Its Q value is relatively high, which means more abundant lower-energy backgrounds are not a concern. When condensed into liquid xenon (LXe), its high Z prevents γ-rays from penetrating more than 10 cm or so, which makes the inner volume of large, monolithic LXe detectors capable of very low backgrounds.

This background reduction power is greatly augmented by use in a time projection chamber (TPC), which allows for reconstruction of the full three-dimensional position of an energy deposit in the detector. With this capability, the outer regions of a LXe TPC can be used to measure backgrounds by fitting the spectrum of events with respect to both energy and position within the detector. The energy resolution of this technique is significantly better than what is achievable in liquid scintillator detectors, making backgrounds from $2\nu\beta\beta$ negligible in the $0\nu\beta\beta$ energy region of interest (ROI) even for large detectors. In the event of a $0\nu\beta\beta$ discovery, the same detector can be filled with Xe depleted of ^{136}Xe to ensure the signal events are not in fact unexpected backgrounds. Lastly, the fluid nature of Xe makes probing it for its daughter isotope, ^{136}Ba, plausible as a powerful means of background rejection (Chapter 4).

3.2 EXO-200 Detector Description

The EXO-200 detector is a cylindrical, single-phase LXe TPC, 40 cm in diameter and 44 cm in length, filled with xenon enriched to 80.6% in ^{136}Xe. Two drift regions are separated by a cathode at the center. A detailed description of the detector can be found in [26]. Radioimpurities in the detector components were minimized by a careful screening process [27] and a detector design optimized to use a minimal amount of materials. External radioactivity is reduced by $\geq$25 cm thick lead walls on all sides and additional passive shielding of $\geq$50 cm of high purity cryogenic fluid,

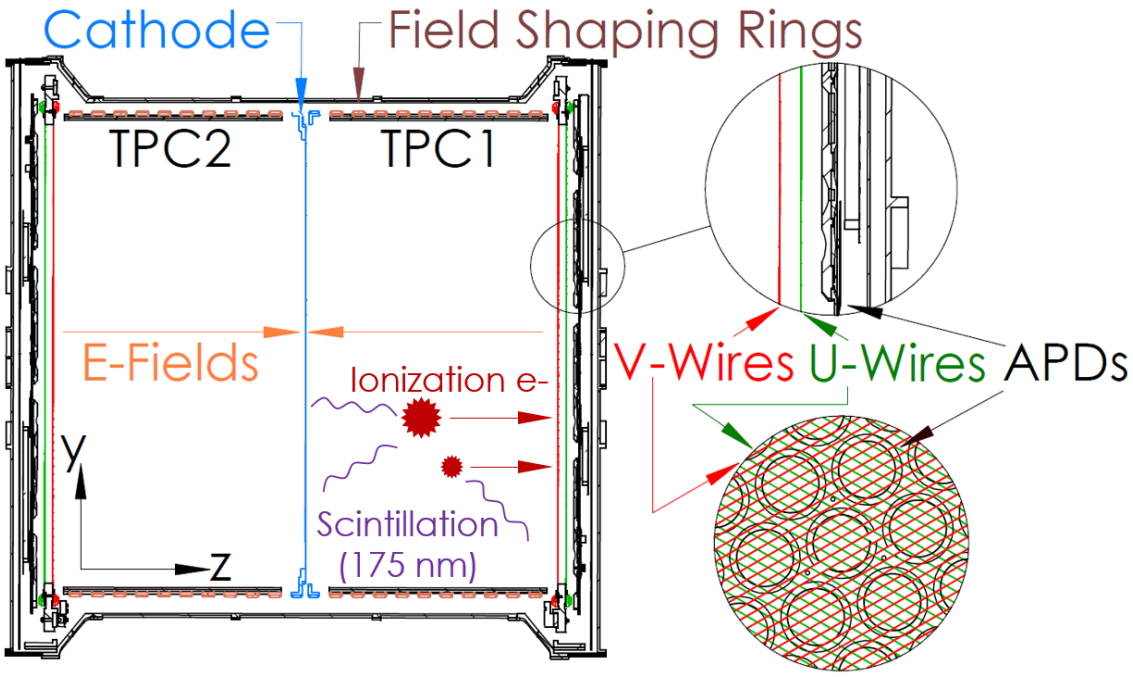

Figure 3.1: Cut out schematic of the EXO-200 LXe TPC. The top inset shows the wire grids and APDs in more detail while the bottom inset shows the same section in the X/Y plane.

HFE-7000 [28]. The detector is located inside a clean room at the Waste Isolation Pilot Plant (WIPP) in Carlsbad, NM, USA, under an overburden of 1585^{+11}_{-6} meters water equivalent [29]. The remaining cosmic ray flux is detected by an active muon veto system consisting of plastic scintillation panels that cover the clean room on four sides.

Energy depositions by ionizing radiation create free electrons and scintillation light in the LXe that are registered by anode wire grids and arrays of avalanche photodiodes (APDs), respectively (Fig. 3.1). Two sets of wire grids form each anode. From the cathode to the anodes, charges will pass by V-wires (induction) first before being collected on the U-wires (charge collection). The U- and V- wire grids are offset by 60° to allow for two-dimensional (2D) reconstruction in the plane perpendicular to the axis of the TPC. The time difference between the prompt scintillation light and drifted charge collection allows for the position of the event in the drift direction

(Z) to be determined. Charge deposits (clusters) in a given event that are spatially separated by ∼1 cm or more can be individually resolved. An event can then be classified as single-site (SS), or multisite (MS), depending on the number of observed charge clusters. The total energy of an event is determined by combining the charge and scintillation signals, allowing improved energy resolution from the anticorrelation between these channels [30]. Radioactive γ sources are periodically deployed at several positions near the TPC to characterize the detector response.

3.3 Theory of $2\nu\beta\beta$ Decays to Excited States

The Standard Model allows $2\nu\beta\beta$ decays to the first 0^+ excited state, denoted hereafter as 0_1^+, of the daughter nucleus if this state is energetically accessible. These decays are suppressed relative to their ground-state counterparts and are generally accompanied by the emission of de-excitation γs, creating a signature that is distinct from typical γ backgrounds and $2\nu\beta\beta$ decays to the ground state.

Measurements of decays to excited states may provide additional constraints on the nuclear matrix elements (NMEs) relevant to $\beta\beta$ decay. Using the NME for the excited state decay in a ratio between it and the ground-state decay would allow any shared uncertainties in the NMEs for these transitions to be canceled. Better knowledge of these NMEs could lead to a more precise determination of the effective Majorana neutrino mass from $0\nu\beta\beta$ half-life measurements [31]. Searches for decays to excited states may also test exotic theories of alternate $\beta\beta$ decay mechanisms.

The first investigation of $2\nu\beta\beta$ decay to excited states was performed by Fiorini *et al.* in 1977 with ^{76}Ge [32]. The first dedicated search was the Milano experiment in 1982, also with ^{76}Ge [33], while the first positive signal was obtained in 1995 for the $2\nu\beta\beta$ decay mode of ^{100}Mo to the 0_1^+ excited state of ^{100}Ru [34]. This decay was then precisely measured by the NEMO-3 experiment where all β and γ tracks were observed [35]. $2\nu\beta\beta$ decays to excited states have also been observed in ^{150}Nd in 2004 [36] and again in 2014 where the γ coincidence was explicitly measured [37]. More recently, current $0\nu\beta\beta$ experiments such as Gerda [38] and KamLAND-Zen [39] have searched for decays to excited states using ^{76}Ge and ^{136}Xe, respectively.

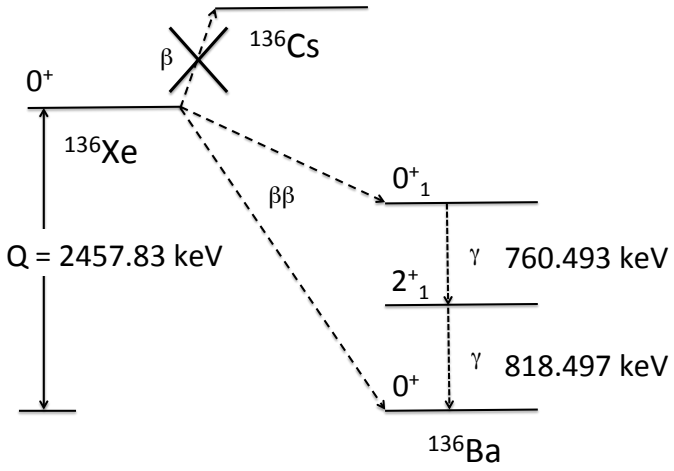

Figure 3.2: Level scheme of the $\beta\beta$ decay of ^{136}Xe. Decay to the excited state of 0_1^+ will result in the emission of two γs during the de-excitation to the ground state. The energy levels of the excited states are taken from [40]. Only the 0^+ and 2^+ levels relevant for this search are shown.

Fig. 3.2 shows the energy level scheme of the $\beta\beta$ decay of ^{136}Xe. A typical $\beta\beta$ decay with a Q value of 878.8 keV transitions from the ground state of ^{136}Xe to the 0_1^+ state of ^{136}Ba. The de-excitation from the 0_1^+ state results in the emission of two γs with energies of 760.5 keV and 818.5 keV. The two de-excitation γs are emitted with a preferential angular correlation to be aligned or antialigned [31]. Although the 2_1^+ excited state of ^{136}Ba is at lower energy than the 0_1^+ state, the direct decay to it from the parent 0^+ ground state of ^{136}Xe is highly suppressed from angular momentum. The intermediate 2_1^+ excited state has a half-life of 1.930 ps [40], which is not resolved temporally in most practical detectors. EXO-200 has the ability to identify both the $\beta\beta$ and de-excitation γs in an excited state decay.

For decays to excited states, $2\nu\beta\beta$ decay to the 0_1^+ state is expected to be the dominant decay mode, although $2\nu\beta\beta$ decay to the 2_1^+ state or $0\nu\beta\beta$ decays to either 0_1^+ or 2_1^+ are also possible. While EXO-200 can conduct searches for these other decays, this analysis focuses on the decay to the 0_1^+ excited state.

The $2\nu\beta\beta$ decay rate for $0^+ \rightarrow 0^+$ transitions, including to the 0_1^+ excited state, can be written in the form,

$$[T_{1/2}^{2\nu}]^{-1} = G^{2\nu}(E_0, Z) \left| M_{GT}^{2\nu} - \frac{g_V^2}{g_A^2} M_F^{2\nu} \right|^2 , \tag{3.1}$$

where $T_{1/2}^{2\nu}$ is the half-life, $M_{GT}^{2\nu}$ and $M_F^{2\nu}$ are the Gamow-Teller and Fermi nuclear transition matrix elements, respectively, $G^{2\nu}$ is the phase space integral that depends on the Q value (E_0) and the atomic number of the daughter nucleus (Z), and g_V and g_A are the vector and axial-vector weak interaction coupling constants, respectively. The NMEs, $M^{2\nu}$, and phase space factor (PSF), $G^{2\nu}$, are expected to differ for decays to the ground and excited states. The PSF term results in a suppression of the decay rate because of the smaller Q value for the transition to the excited state.

The expected rate for the $2\nu\beta\beta$ decay of ^{136}Xe to the first excited state is calculated from the measured $2\nu\beta\beta$ decay rate to the ground state and the theoretical values for the PSF and NMEs of both the ground and excited states. For the ^{136}Xe $2\nu\beta\beta$ decays, the ratio of the PSF between the ground state and the excited state using the Schenter-Vogel Fermi function approximation and the Wilkinson correction for nuclear size is 3915 [41]. Using the PSF from [42], a suppression ratio of 3956 is found. The two calculated values are within $\sim$1% of each other, indicating that the theoretical uncertainty on the PSF suppression is small.

The NMEs for $\beta\beta$ decays have been calculated using the microscopic interacting boson model (IBM-2) with a method for isospin restoration and Argonne short-range correlations in the closure approximation in [43]. The ratio of the NMEs between the ground state and the excited state for ^{136}Xe is 1.7. From Eq. (3.1), this leads to an additional suppression factor of 2.9 for the decay rate to the excited state. The calculated NME values are model dependent and have larger theoretical uncertainty than the PSF. Including both suppression factors and applying them to the ground state half-life measured by EXO-200 of 2.16×10^{21} yr [44], the expected $2\nu\beta\beta$ decay half-life of ^{136}Xe to the 0_1^+ excited state of ^{136}Ba is estimated to be 2.5×10^{25} yr. While this prediction puts the decay beyond the sensitivity of the present EXO-200 dataset, uncertainties in the NMEs may produce an overestimate of the half-life, as

in the cases of ^{100}Mo and ^{150}Nd.

3.4 Data Set and Methodology

This search uses the same data set ("low-background data") and event selection criteria as recent EXO-200 searches for $0\nu\beta\beta$ decay [22, 45]. A total of 477.60±0.01 live days of data were collected between September 22, 2011 and September 1, 2013. Events consistent with noise, coincident with the muon veto, with more than one scintillation signal, or within 1 s of other events in the TPC are removed.

The fiducial volume is hexagonal with an apothem of 162 mm. Only regions within this hexagonal volume that are >10 mm from the cathode and anode wire planes are included. This geometry corresponds to a ^{136}Xe mass of 76.5 kg, or 3.39×10^{26} atoms of ^{136}Xe. The total exposure is 100 kg·yr (736 mol·yr). As in previous analyses, an energy range (summed over all charge clusters in an event) of 980−9800 keV is used. Finally, all events are required to have fully reconstructed U-, V-, and Z-positions.

A Geant4-based [46] Monte Carlo (MC) simulation of the detector and shielding (described in detail in [44]) is used to model the detector response. The simulation of the $2\nu\beta\beta$ decay of ^{136}Xe to the 0_1^+ excited state of ^{136}Ba (hereafter referred to simply as "excited state events") accounts for the smaller Q value and angular correlation between the de-excitation γs. This MC is used to estimate the detection efficiency, determined by the percentage of these MC excited state events that survive all event selection cuts. The resulting efficiency for the excited state signal is $(23.2 \pm 2.0)\%$, with the dominant losses in efficiency arising from the fiducial volume and full reconstruction cuts. Errors in the estimate of this efficiency are accounted for in § 3.6.

Based on periodic calibrations using γ sources (^{228}Th, ^{60}Co, ^{226}Ra, and ^{137}Cs), the energy scale and resolution are determined by fitting the full shape of the energy spectra observed in calibration data to MC simulations. In particular, the ^{60}Co source produces events with multiple γs of energies similar to those produced in the excited state decay. These calibration events show good agreement with MC simulations across all energies. The energy scale calibration and resolution are determined separately for SS and MS events.

To search for the $2\nu\beta\beta$ decay to the excited state, a binned maximum-likelihood (ML) fit is performed simultaneously over the SS and MS events using probability density functions (PDFs) in two dimensions: energy and an excited state "discriminator" variable. The PDFs are generated using the MC simulation and smeared with the appropriate energy resolution function. The same background components used in previous analyses [22, 45] are also used here. The discriminator is specifically optimized to search for the unique event structure of the decay to the 0_1^+ excited state using machine learning techniques to be detailed in § 3.5. Using a 2D fit in energy and the machine learning discriminator improves the sensitivity to excited state events by more than a factor of three over the more generic technique in [22, 45], as detailed in § 3.5.5.

3.5 Machine Learning

This analysis uses machine learning techniques to create a variable ("discriminator") which, for each event, indicates how "signal-like" (+1) or "background-like" (-1) it is. The machine learning software TMVA [47] (part of the ROOT data analysis framework [48]) is used to "train" an algorithm, using simulated signal and background data sets, to construct an optimized discriminator from a set of input variables and characterize the discriminator's performance.

3.5.1 Input variables

Prior to applying this algorithm to the low-background data, the input variables used to build the discriminator were finalized by optimizing the expected sensitivity to excited state events (as determined by the method described in § 3.5.5), while minimizing systematic errors resulting from disagreement between data and MC (discussed in § 3.6). By waiting to perform the fit until the method to determine the discriminator variable is finalized, potential biases from tuning the algorithm to the data set are minimized. The number of input variables is limited to reduce the sensitivity of the discriminator to systematic differences between the data and MC simulation.

The majority of the discriminating power is provided by several event variables used in previous EXO-200 analyses [22, 45]. These include the number of charge clusters in the event (multiplicity), the total event energy determined from the ionization and light response (energy), and the minimum separation of the charge deposits in the event from the anode wire plane or cylindrical walls of the TPC (standoff distance). The agreement between data and MC for these variables was studied in previous analyses [44].

Although energy and multiplicity are included as inputs to the discriminator, the fit procedure also uses energy as the second fit dimension and separates the fit into SS and MS events. Including these additional dimensions allows the fit to better constrain individual background model components, while the discriminator is primarily useful for distinguishing between excited state events and all other backgrounds.

In addition, variables designed to take advantage of the energy deposits from the ^{136}Ba de-excitation γs are used. These variables, γ_1, γ_2, and γ_{sum}, are defined as:

$$\gamma_{\mathrm{i}} \equiv \min_{j \in S} \{|E_j - \epsilon_i|\}\,, \tag{3.2}$$

where S is the set of charge clusters in an event, E_j is the energy of cluster j, and ϵ_i are the de-excitation energies, $\epsilon_1 = 760.5$ keV, $\epsilon_2 = 818.5$ keV, and $\epsilon_{sum} = 1579$ keV. Hence, in the case where an event has a cluster of energy close to ϵ_i, the corresponding variable γ_i will be close to zero. Because only the total scintillation energy of an event is measured by the APDs, the γ_i variables are determined using only the energy reconstructed from the charge signal of each cluster.

Additional variables were also considered which attempted to estimate which charge clusters came from multiple Compton scatters of the same γ ("superclustering"). One method of superclustering, called k-means, groups charge clusters together by proximity in such a way that there are three superclusters at the end (unless there were fewer than three charge clusters in the first place). Another superclustering method uses the known angular form of Compton scattering, assuming one cluster as the origin (from the $\beta\beta$) from which the two initial γs with the de-excitation energies ϵ_i are emitted, and groups together charge clusters to minimize the difference between

the Compton angle and the angle from data. These superclustering techniques are then used to define a new γ_i of the same form as Eq. 3.2 but where S is the set of superclusters and E_j is the summed energy of the clusters within a supercluster. None of these variables incorporating superclustering methods showed an improvement in the ultimate sensitivity, so they were not included in the final training. This lack of improvement from superclustering suggests that Compton scattering may erase most of the relevant event information which is not already contained in the γ_i variables without superclustering.

3.5.2 Training data set

The training data consist of a signal class (excited state events) and a background class (all others). The dominant background to this search is $2\nu\beta\beta$ decays to the ground state, because they are the primary component of the low-background data. The background class was drawn from MC simulations, using a list of background components in quantities determined by the best fit model from a prior analysis of this data [45], to ensure that the training circumstances accurately reflect the low-background data set to which the discriminator will be applied. An equal number of excited state events were simulated to produce the signal class. All event selection cuts were applied prior to building this data set. In addition, both MC data classes were split in half, with one half used for training and the other for testing, to ensure that any patterns found in the training data set were not the result of statistical fluctuations. The distributions of each of the input variables used for this training data set are shown in Fig. 3.3.

3.5.3 Machine learning algorithm

Several different machine learning algorithms were investigated, including boosted decision trees (BDTs), multilayer perceptrons (MLPs, a class of neural network), k-nearest neighbors (KNN), and simple rectangular cuts on the inputs. In preliminary tests of these algorithms using a cut on the discriminator variable to separate signal from background in the testing data set, the BDT provided the greatest discriminating

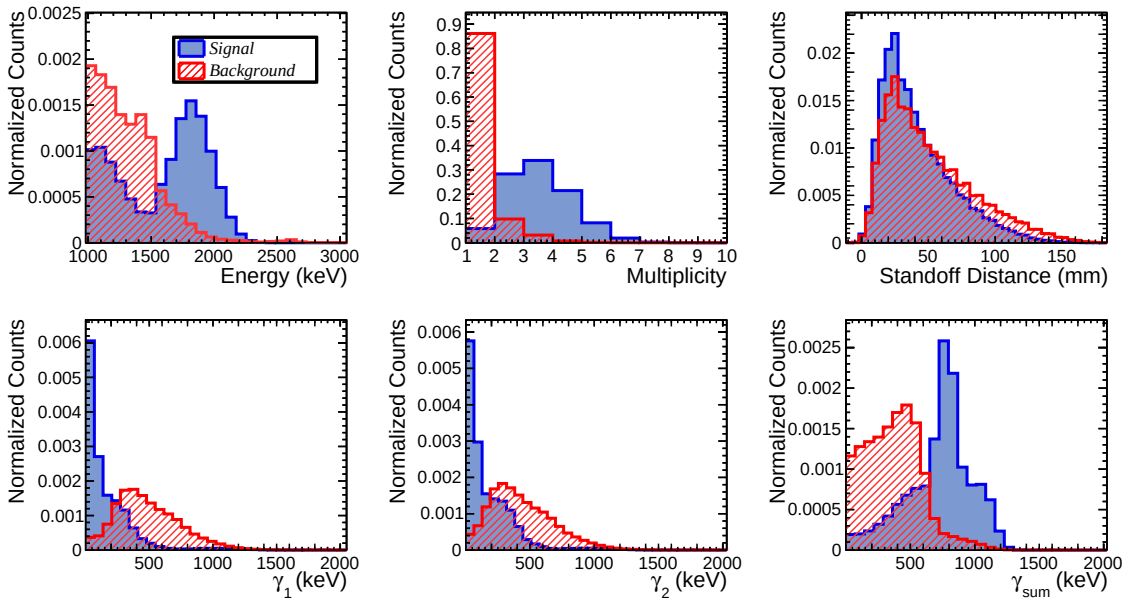

Figure 3.3: Distributions of the input variables to the discriminator based on MC simulations. Signal (excited state events) is shown in solid blue and background (all other events) in hatched red.

power. Adjustments to the training parameters of the BDT did not indicate significant improvement by this metric. More detailed studies of the sensitivity and systematic errors of the three most promising of these algorithms (MLP and two BDT variants) were also performed, as described in § 3.5.5 and 3.6. These studies provided further evidence that the BDT algorithm is optimal among those considered and that changes to its training parameters have a small effect on its performance: The sensitivity for the MLP was roughly 50% that of the BDT, and changing the BDT training parameters led to a 2% increase in sensitivity at the cost of a 5% increase in the systematic error.

3.5.4 Trained decision tree results

After training the BDT, the effectiveness of individual variables in deciding whether an event is signal-like or background-like can be determined. The BDT consists of many individual decision trees, each of which performs a series of binary cuts on the

input variables, with the final nodes of the cuts each assigned to either signal (+1) or background (-1). The cuts at each intermediate stage are chosen by finding the cut from all of the input variables which maximizes the separation of signal from background for the decision tree's set of training data that have made it past all the previous cuts. Once a decision tree is made, the events it misclassifies are weighted more strongly in the training of the subsequent decision tree.

The discriminator variable is then given by a sum of the individual trees' assignments, weighted by their classification performance. Further description of the BDT algorithm can be found in [47]. The ranking for any given variable is a measure of the fraction of decision tree cuts which use that variable, and is given in Table 3.1. As expected, multiplicity is an effective discriminator between the decays to the excited state, which are largely MS, and $2\nu\beta\beta$ decays to the ground state, which are primarily SS, with most of the additional information contained in the energy variables.

Rank	Variable	Importance
1	Multiplicity	0.28
2	Energy	0.27
3	γ_{sum}	0.14
4	Standoff distance	0.12
5	γ_1	0.10
6	γ_2	0.09

Table 3.1: Importance ranking of the input variables in the final boosted decision tree. The "importance" of each variable denotes the fraction of decision tree cuts which use that variable.

The trained BDT is then applied to low-background data, calibration data, and MC to determine the value of the discriminator variable for each event. In particular, this allows for comparison between data and MC for calibration sources, which can be used to quantify systematic effects. These comparisons are done for ^{60}Co, ^{226}Ra, and ^{228}Th, both for SS and MS events. The primary γ event backgrounds consist of MS events, for which a representative comparison using ^{228}Th is shown in Fig. 3.4. While the data and MC distributions typically agree within 10%, the remaining deviations are accounted for as systematic errors, to be discussed in § 3.6.

Because $2\nu\beta\beta$ decays to the ground state constitute one of the largest backgrounds

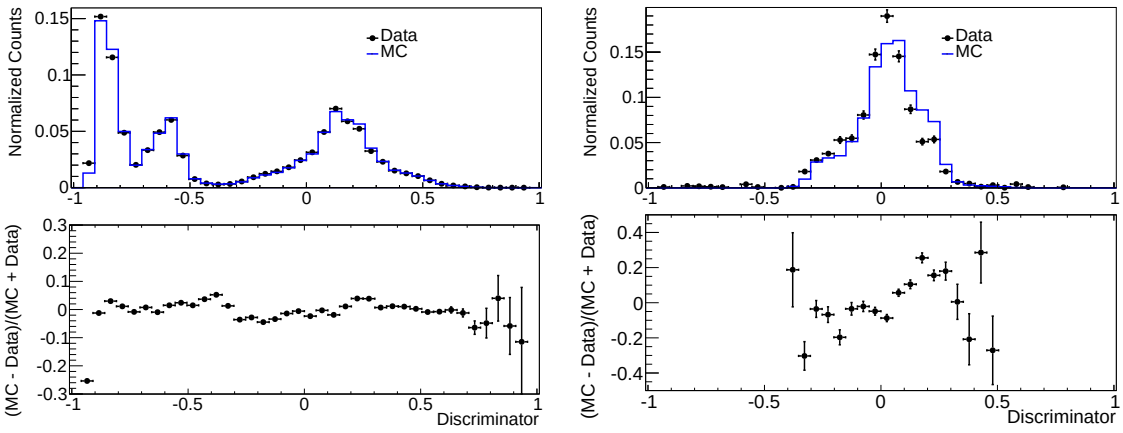

Figure 3.4: Discriminator agreement between data and MC for MS ^{228}Th (left) and MS $2\nu\beta\beta$ decays to the ground state (right). The $2\nu\beta\beta$ data are generated by subtraction of the non-$2\nu\beta\beta$ components, in amounts given by a prior analysis, from the low-background data. Residuals are shown in the bottom panels, ignoring those bins with insufficient counts.

to this search, it is important to be able to quantify the difference in the shape of the discriminator variable distribution for data and MC for this component. To determine the $2\nu\beta\beta$ spectrum, each non-$2\nu\beta\beta$ component (from MC) is subtracted from the low-background data, in amounts given by a fit from a prior analysis [45]. The final subtracted $2\nu\beta\beta$ distribution is then compared to the MC $2\nu\beta\beta$ component in the discriminator variable (Fig. 3.4). The results indicate good SS agreement, with some notable differences in the MS spectrum, which are accounted for in § 3.6.

3.5.5 Sensitivity and significance estimates

Before the final fit to the data, the discriminator sensitivity was estimated by performing a two-dimensional profile likelihood fit on toy MC data sets. These MC data sets are created using the best-fit model from a previous analysis of the data [45], with no excited state events included. A fit in energy and discriminator is applied to find the 90% confidence level (CL) upper limit on the number of excited state events, as determined by a profile of the negative log likelihood (NLL). An expected sensitivity of 43 events, corresponding to $T_{1/2}^{2\nu} = 1.7 \times 10^{24}$ yr, is given by the median upper limit

obtained from toy MC data sets. This demonstrates an improvement by more than a factor of three over a simple fit to energy and standoff distance (as in [22, 45]), which has a sensitivity of 170 events, corresponding to $T_{1/2}^{2\nu} = 4.2 \times 10^{23}$ yr.

The toy MC data sets with zero injected excited state events can also be used to calculate the significance of a nonzero final fit to data. The distribution of ΔNLL values for fits with 0 excited state events is calculated; then, the fraction of toy MC data sets with ΔNLL values greater than that found by the fit to data determines the level of compatibility with the null hypothesis. The full sensitivity and significance distributions are shown in § 3.7.

3.6 Systematic Uncertainties

Systematic uncertainties in this analysis can be divided into those deriving from the event selection and background model, previously evaluated in [22], and those unique to this analysis. In all cases, these uncertainties are accounted for by applying Gaussian constraints on the fit parameters.

Unique to this analysis is an excited state event normalization term that accounts for discrepancies between MC and data in the discriminator and energy distributions of background components. To calculate this, toy MC data sets with a nonzero number of excited state events are generated from PDFs that have been skewed by the relative differences between data and MC, as measured in calibration source data (Fig. 3.4). These toy data sets are then fit to the standard, un-skewed PDFs, and the resulting difference between the fitted and simulated number of excited state events is determined. This fractional difference is accounted for as a systematic error which is applied in the normalization term specific to the excited state component.

While the deviations between calibration data and MC measured for the ^{228}Th, ^{226}Ra, and ^{60}Co sources are used to skew the corresponding background components, the $2\nu\beta\beta$ ground state decay is skewed by deviations between MC and the background-subtracted data set. Differences between calibration sources are small, and the majority of the skewing error comes from the $2\nu\beta\beta$ component, so the exact choice of which calibration source is most similar to a given background component is inconsequential.

Because of statistical errors resulting from the background subtraction, the skewing cannot be precisely measured in the tail of the $2\nu\beta\beta$ distribution (*i.e.* values of the discriminator >0.5 or <-0.4 in the MS data). To prevent artificially distorting the distribution by errors resulting from the background subtraction, the skewing is forced to go to zero for bins with too few (< 20) events. Several methods for suppressing the skewing in the low statistics tails were tested, with negligible effects on the resulting error. If an excited state decay signal were present in the data, its incomplete subtraction could lead to a slight overestimation of this systematic; however, while the estimation of this error may be conservative, this background-subtracted data set is only used for this study, so the procedure does not bias the final result. The final error found by this study is 15%.

The additional Gaussian constraints are quantified in Table 3.2. The SS fraction constraint is applied to the ratio of the number of SS events to the total number of events [SS/(SS+MS)] for each component, and is determined by the deviations in this ratio between data and MC for the calibration sources. A normalization constraint common to all components, accounting for the uncertainty in detection efficiency, is calculated from studies of the event reconstruction efficiency and the rate of events from calibration sources of known activity. A further normalization constraint is applied to account for uncertainties in the location of degenerate backgrounds. The relative fractions of neutron-capture related backgrounds, coming from cosmic ray muons and radioactive decays in the salt surrounding the experiment, are constrained according to MC studies and data coincident with muon-veto-panel events. The activity of radon in the liquid xenon is measured by observing characteristic ^{214}Bi to ^{214}Po coincidences in the detector, and the uncertainty on this measurement is translated into a constraint on the normalization of the radon components.

The final systematic accounts for uncertainties in the "β-scale", which describes the possible difference in energy scales of β-like and γ-like events. Because the majority of the energy in the excited state events is deposited by the de-excitation γs, these events are calibrated as γ-like components in the MC simulations. The β-scale variable is defined as an energy-independent ratio of γ to β energy scales and is allowed to float as a free parameter in the profile likelihood fit. This work finds the β-scale to be

0.9943 ± 0.0006.

Among the systematics, the SS fraction error had the largest effect on the final upper limit of the search. This uncertainty allows the largest background, $2\nu\beta\beta$ decays to the ground state, to shift from SS to MS, making them more signal-like.

Constraint	Error (%)
Excited state normalization	15
SS fraction	4
Common normalization	8.6
Background normalization	20
Neutron capture fractions	20
Radon in liquid xenon	10

Table 3.2: Gaussian constraints applied to fit parameters to account for systematic uncertainties. These errors are explicitly included as input to the final fit to the low-background data, and are not the systematic uncertainties on the final result.

3.7 Results

Plots of the discriminator variable and energy spectra of MS events for the fit to the data are shown in Figs. 3.5 and 3.6. These plots illustrate the relative shapes and best fit quantities for all components, but do not contain all information used in the full fit, such as constraints on the SS fraction. The best fit values for background components are compatible with prior results, with the calculated half-life of the $2\nu\beta\beta$ decay to the ground state agreeing within systematic error [44]. This fit was also checked for robustness against hypothetical backgrounds ^{88}Y and ^{110m}Ag, considered because of their mixed β and γ composition that could mimic an excited state signal. Both backgrounds were found to have a $< 1\sigma$ effect on the fit result, and prior analyses found no evidence of either component, so they are not included in the final result. The corresponding profile likelihood scan in the number of excited state events is shown in Fig. 3.7. The profile finds a best-fit value of 43 events, with a 90% CL upper limit of 104 events, assuming Wilks' theorem [49, 50] holds. Taking this result as a limit on the half-life of the decay to the 0_1^+ excited state gives a constraint of

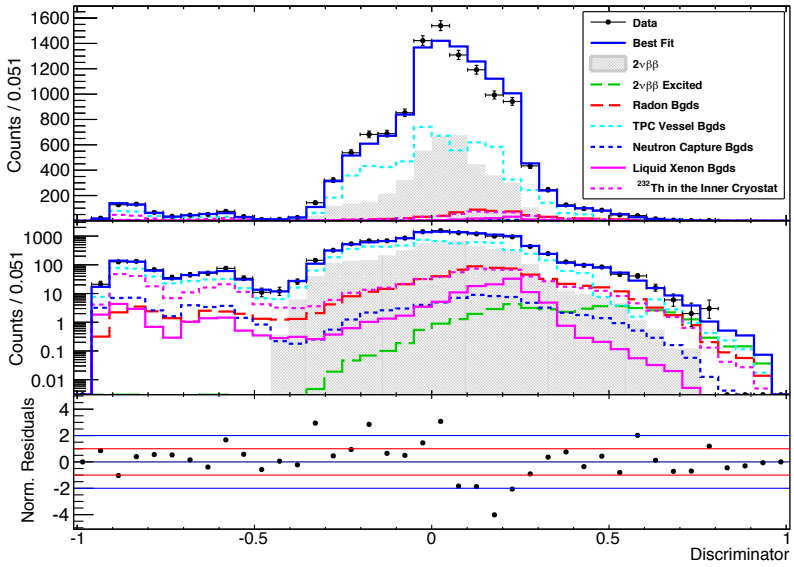

Figure 3.5: Fit results for the discriminator variable spectrum of MS events. The excited state event distribution is given by the dashed green line, concentrated toward positive discriminator values. Data points are shown in black and residuals between data and the best fit, normalized to the Poisson error, are presented, ignoring bins with 0 events.

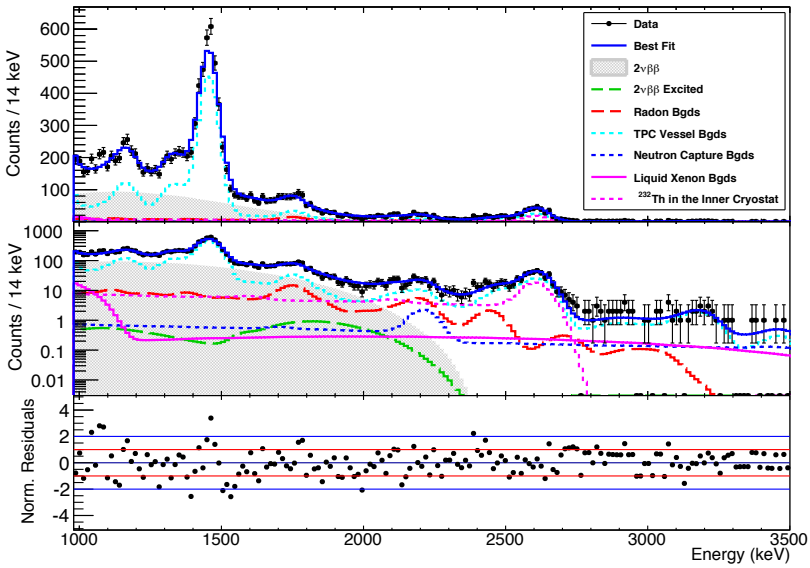

Figure 3.6: Fit results for the energy spectrum of MS events. Data points are shown in black and residuals between data and the best fit, normalized to the Poisson error, are presented, ignoring bins with 0 events.

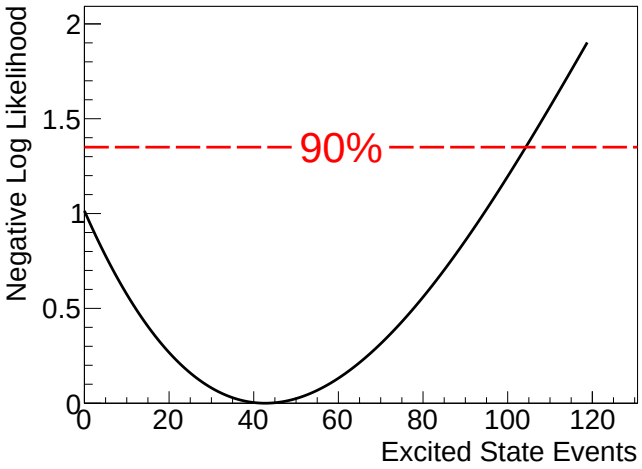

Figure 3.7: NLL profile in the number of excited state events. The best fit value is 43 events, while the fit with 0 events has a ΔNLL value of 1.0. The 90% upper limit (dashed line) is 104 events.

$T^{2\nu}_{1/2} \geq 6.9 \times 10^{23}$ yr at 90% CL.

The upper limit obtained from the data is competitive with a recent result from the KamLAND-Zen experiment [39] despite being roughly a factor of two weaker than the median expected sensitivity of $T^{2\nu}_{1/2} = 1.7 \times 10^{24}$ yr calculated in § 3.5.5. The calculation of significance, as described in § 3.5.5, indicates that the result from data is consistent with the null hypothesis at 1.6σ (Fig. 3.8). Thus, this analysis does not find statistically significant evidence for a nonzero component of $2\nu\beta\beta$ decays to the excited state.

This limit is consistent with theoretical calculations which predict, with substantial uncertainty, a half-life of order $\sim 10^{25}$ yr; hence, sensitivity to this decay may be within reach of future analyses. Since the publication of these results [51], improvements through more data and relaxation of the requirement that all charge clusters have a reconstructed three-dimensional position have shown indications that the sensitivity can be nearly doubled. Additional data taken after a detector upgrade with better energy resolution promise to provide further improvement.

This work also demonstrates the merit of applying machine learning algorithms to EXO-200 data, and shows that such an approach is compatible with careful accounting

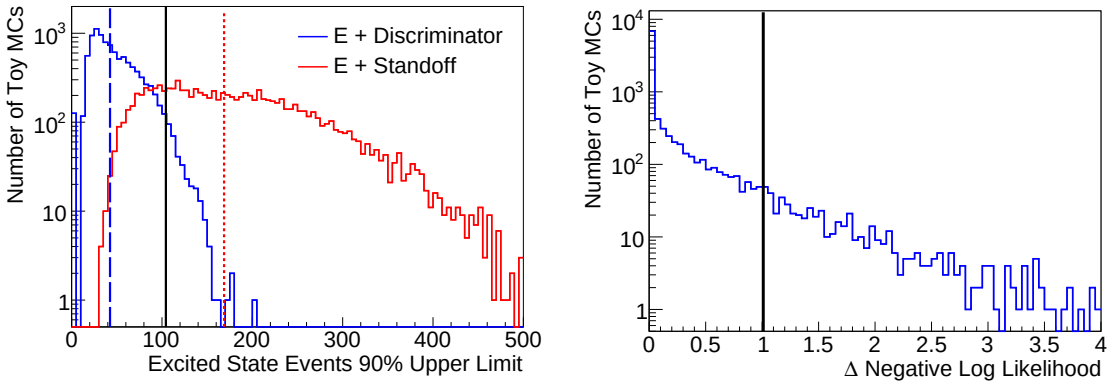

Figure 3.8: Left: Distribution of 90% CL upper limits on the number of excited state events from toy MC data sets with no excited state events, using fits to energy and either standoff distance (red) or discriminator (blue). Median values are drawn with dashed lines, with the limit from data at 104 events drawn as a solid line. Right: Distribution of ΔNLL values from toy MC data sets with no excited state events. The ΔNLL for 0 events from data is shown as a solid line at 1.0.

of systematic errors. The framework developed here is currently being extended to allow it to be tailored to any signal of interest with the appropriate choice of input variables. This includes the possibility of performing the primary $0\nu\beta\beta$ search using machine learning by developing a discriminator based off of event topology beyond what is captured in the current SS/MS categorization.

Chapter 4

nEXO and Barium Tagging

4.1 nEXO

Pushing the limits of sensitivity to $0\nu\beta\beta$ by orders of magnitude will require a larger, more sensitive experiment than EXO-200. Its successor, nEXO (currently in the design phase), aims to be capable of discovering $0\nu\beta\beta$ with a half-life in excess of 10^{27} yr. In the case of $0\nu\beta\beta$ caused by light neutrino exchange, this corresponds to a sensitivity to Majorana masses covering the full experimentally-allowed range for the inverted mass ordering. In the case of a normal mass ordering, nEXO is also expected to have a good chance at discovery of $0\nu\beta\beta$ through light neutrino exchange [52, 53], while alternative mechanisms are in fact more favorable for a normal mass ordering. The design of nEXO will build on the successful LXe TPC design of EXO-200, while increasing in sample mass to 5 tonnes of enriched LXe and implementing several improvements to its detection capabilities. A side-by-side comparison of the two TPCs is shown in Fig. 4.1, along with a depiction of the experiment in its potential site at the SNOLAB cryopit in Sudbury, Ontario.

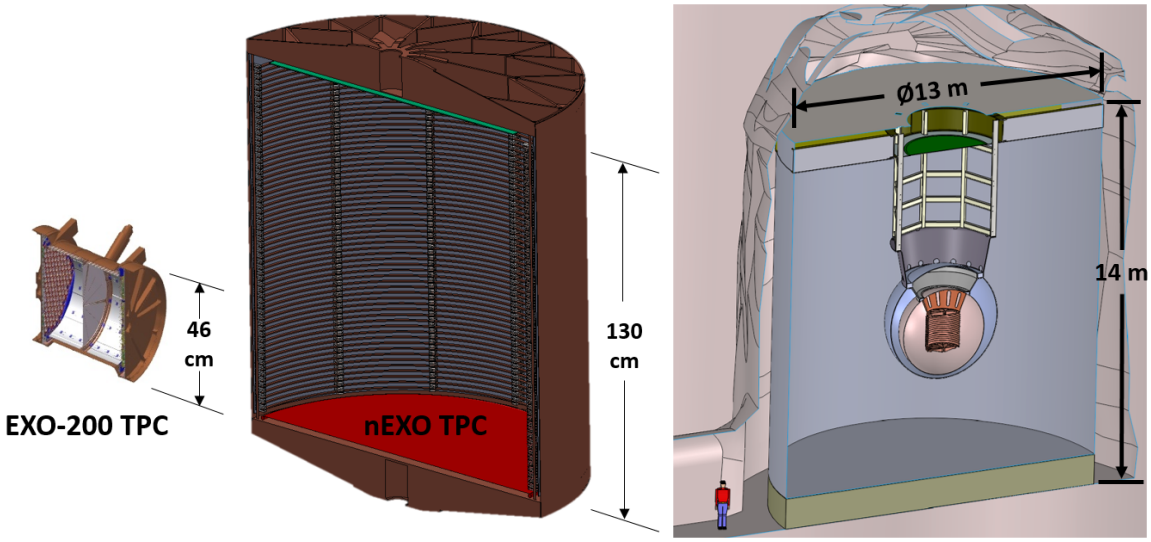

Figure 4.1: Schematics of the EXO-200 and planned nEXO TPCs, to scale (left), and an artist's rendition of nEXO on-site at the SNOLAB cryopit (right).

Among the changes that will improve nEXO's performance relative to EXO-200 are the use of silicon photomultipliers placed around the barrel of the TPC, crossed 2D charge readout strips, and a single drift volume with no central cathode. The self-shielding and background tagging properties of Xe make such a monolithic detector design highly effective at reducing γ backgrounds from the central volume of the detector.

Even with the combined benefits of radiopure materials, self-shielding, and dedicated analysis techniques, however, it is beyond current capabilities to eliminate backgrounds entirely throughout the volume of such a large detector. For this, a more ambitious technique is under development: identification of the ^{136}Ba daughter ion produced by double beta decay events, hereafter referred to as Ba tagging. In this scheme, when a potential $0\nu\beta\beta$ event is observed, the origin of the event is searched for any ^{136}Ba – an isotope which is not produced by γ backgrounds. In the ideal case with perfect ^{136}Ba detection efficiency and no false positives, Ba tagging results in a detector which is background-free in the energy region of interest (ROI) near the Q

value of the $0\nu\beta\beta$ decay. The only exception is events from the tail of the $2\nu\beta\beta$ distribution, which also produce ^{136}Ba – but with a realistic energy resolution of $\sim 1\%$, fewer than one event is expected in the ROI during ten years of operation. If a Ba tagging technique can be established to reliably work with high efficiency, it will be included as an upgrade to nEXO or in a future experiment.

When in a statistics-limited regime, as expected when searching for only a few $0\nu\beta\beta$ signal events in multiple years of data, removal of backgrounds can dramatically improve sensitivity. As a simplified illustration, in a simple counting experiment with backgrounds that must be substracted statistically, the sensitivity to Majorana mass scales as the fourth root of exposure:

$$\langle m_{\beta\beta} \rangle \propto \frac{1}{\sqrt{T^{0\nu}_{1/2}}} \propto \frac{1}{\sqrt[4]{Nt}} \tag{4.1}$$

where N is the number of ^{136}Xe atoms in the analysis volume and t is the livetime for the experiment. This is purely due to the statistical improvement on the knowledge of background rates. For comparison, in the background-free case, the sensitivity scales as the square root of exposure:

$$\langle m_{\beta\beta} \rangle \propto \frac{1}{\sqrt{T^{0\nu}_{1/2}}} \propto \frac{1}{\sqrt{Nt}} \tag{4.2}$$

From this expression, it is clear that the most substantial benefits of background removal come for large detector volumes running for long times, as is the case for nEXO and potential future-generation experiments.

More quantitative simulations of the nEXO experiment have been produced to investigate the potential improvement in sensitivity achieved by Ba tagging (Fig. 4.2). For 100% Ba tagging efficiency, the sensitivity to the half-life of the decay can by improved by almost a factor of five over five years of livetime, which translates to an improvement in the sensitivity to the effective Majorana mass by more than a factor of two.

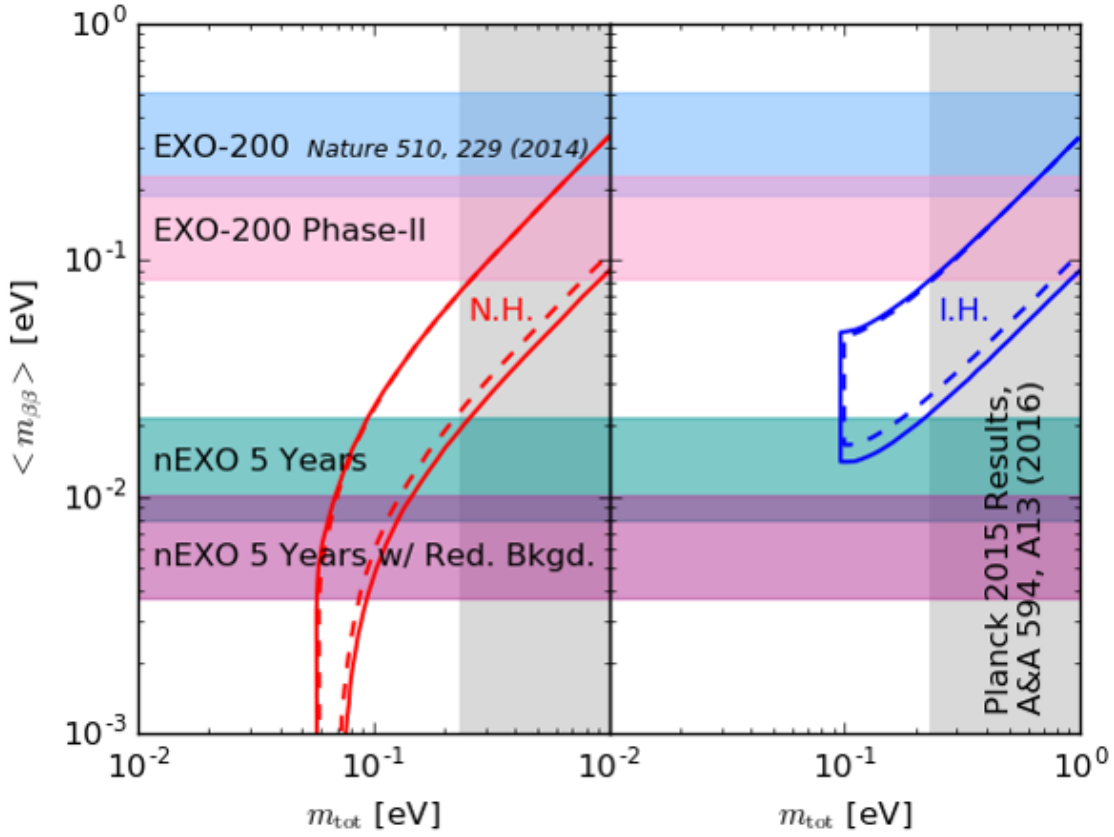

Figure 4.2: Projected sensitivity of nEXO to Majorana mass as a function of total neutrino mass after 5 years, with and without Ba tagging ("reduced background"). This assumes the light neutrino exchange mechanism. Limits on the allowed parameter space from oscillation experiments are also drawn, for the normal neutrino mass ordering (N.H., left) and inverted ordering (I.H., right). Limits on the sum of the neutrino masses from the Planck experiment are also shown (shaded region). The width of the sensitivity bands is due to uncertainties in the nuclear matrix elements (see Sec. 2.3). This assumes 100% efficiency for Ba tagging.

4.2 Ba Tagging from Liquid Xenon

Due to the inherent challenges of detecting a single atom out of a macroscopic volume, several different techniques for Ba tagging are being developed in parallel [54, 55]. The remainder of this dissertation, however, presumes a technique where the daughter Ba ion is produced in a LXe detector environment, electrostatically attracted and

adsorbed onto the surface of a probe, and then extracted to a separate chamber for identification.

In this scenario, the first step of the Ba tagging process is the production of a candidate event in the LXe TPC. This event must be recognized as a candidate through real-time reconstruction, to know when and where to search for Ba ions. Here, a candidate event is defined as one with a single energy deposit of energy within the $\pm 2\sigma$ energy resolution of the Q value, after appropriate quality cuts such as restricting the event position to some fiducial volume and vetoing cosmic ray events. At this real-time reconstruction stage, simple estimators of the energy and position from raw waveforms should suffice.

Next, a probe with the appropriate substrate for Ba adsorption must be introduced to the detector near the reconstructed position of the event. Development of such a probe which is capable of accessing the full detector volume is ongoing, but is beyond the scope of this work. For electrostatic attraction to work, any Ba produced must be in the form of an ion. The initial double beta decay will result in a Ba^{++} ion, but several processes can change this. First, fast processes such as interaction of the Ba^{++} ion with the cloud of secondary electrons produced by the decay can change its charge state. These are the least well-understood, although empirical measurements from EXO-200 suggest that single β decay of ^{214}Pb in LXe results in a $^{214}Bi^{+}$ ion roughly 75% of the time, so one might reasonably expect a similarly-high ion fraction for double beta decays in LXe [56]. In addition, there can be charge exchange with the LXe and with its impurities. Because the first ionization energy (IE) of Ba is 5.2 eV while the second IE is 10.0 eV, and the band gap of LXe is 9.2 eV, Ba^{++} is expected to charge exchange with the LXe to Ba^{+}, but not to neutral Ba [57, 58]. Furthermore, due to the exceptionally low levels of electronegative impurities required for nEXO, neutralization from these impurities is not a concern. Hence, the Ba is expected to be in the Ba^{+} charge state most of the time when the probe is deployed, though a non-negligible fraction of the time it may instead be a neutral Ba atom.

There are also some relevant restrictions on how long this extraction can take. If the electric field of the TPC is not shut off during extraction, the drift of the Ba^{+} ion will eventually remove it from the detector. The mobility of Ba^{+} is $\sim$

$2 \times 10^{-4}\text{cm}^2\text{V}^{-1}\text{s}^{-1}$ [59], so for an electric field on the order of 400 V/cm (the field used in EXO-200), the Ba^+ will drift about 1 mm/s. Hence, there are a few minutes to adsorb the Ba^+ onto the probe substrate before it drifts a substantial fraction of the 130 cm drift length. The drift distance will have to be corrected for when setting the location of the probe, but the uncertainty on this (from diffusion and currents in the LXe caused by its steady recirculation) should be sufficiently small. Moreover, the rate of candidate events in the detector is expected to be on the order of one per day, so the dead time introduced by the probe insertion process (during which normal data taking would not occur) is negligible.

More significantly, there is the risk of a false positive due to background Ba^+ ions already in the LXe from its steady rate of $2\nu\beta\beta$ (unrelated to the candidate event being tested). However, the electric field will cause these to drift through the TPC and eventually plate onto the cathode or some other surface. This can be modeled as a balance between the steady production of Ba due to $2\nu\beta\beta$ at every point in the detector and the removal rate of Ba due to its drift. The production rate is given by

$$r_+ = \frac{f_{136}\rho_{LXe}}{m_{Xe}\tau^{2\nu}} \tag{4.3}$$

where f_{136} is the fraction of LXe which is ^{136}Xe (∼85% for EXO-200), ρ_{LXe} is the mass density of LXe (~ 3 g/cm^3), m_{Xe} is the average mass of a Xe atom in the LXe ($\sim 2 \times 10^{-22}$ g/atom), and $\tau^{2\nu}$ is the lifetime of the $2\nu\beta\beta$ decay ($\sim 2.9 \times 10^{21}$ yr [44]). This leads to a production rate of $\sim 10^{-7}$ Ba cm^{-3} s^{-1}. In equilibrium, the flux of Ba removed for a column extending the full drift length of the detector is

$$\Phi = N_{col} v_{drift} / L_{drift} \tag{4.4}$$

where N_{col} is the average equilibrium density of Ba atoms per area in the column, $v_{drift} \approx 1$ mm/s is the drift velocity, and $L_{drift} = 130$ cm is the full detector length.

In equilibrium, the rate of Ba production in this column equals the rate leaving it ($r_+ L_{drift} = \Phi$), so $N_{col} \approx 0.017$ Ba/cm^2 and the average Ba density throughout this column is $\rho_{Ba} \approx 10^{-4}$ Ba/cm^3. In fact, the density increases linearly with distance from the anode (since all Ba produced closer to the anode will pass through that

volume), so the maximum Ba density in equilibrium is twice this, or 2×10^{-4} Ba/cm^3. To get a false positive rate of $<1\%$ then requires capturing only those Ba^+ ions within a sphere of roughly 2 cm radius. While such a capture region should be achievable by setting the appropriate bias voltage for the substrate, this has not been tested.

It is assumed that neutral Ba is not a concern for this process as it will not be attracted to the biased substrate – hence, only motion of Ba neutrals due to diffusion and the continuous LXe recirculation might cause an interaction with the substrate.

The production of Ba from $2\nu\beta\beta$ also carries with it a considerable benefit: it provides a reliable calibration source which allows periodic measurements of the efficiency of the entire Ba tagging process. The false positive rate can similarly be measured by performing the tagging process without a candidate event signal. These measurements will ensure that systematic errors from tagging imperfections are quantitatively accounted for in a final analysis.

These concerns regarding how to integrate such a probe in a real LXe TPC such as nEXO are not the subject of this dissertation, however. The following chapters are devoted to the more fundamental process of adsorbing Ba onto a substrate (in either a vacuum or Xe environment), removing it via infrared thermal desorption, and identifying it via resonance ionization mass spectroscopy.

Chapter 5

Laser-Induced Thermal Desorption and Resonance Ionization Spectroscopy

Any successful Ba tagging scheme must satisfy two crucial properties: it must be highly efficient, since each candidate event produces only one Ba atom, and it must be exceptionally selective, to separate that atom from the macroscopic volume or surface surrounding it. Arguably the best techniques for achieving such high selectivity are

from the realm of atomic physics. Specifically, spectroscopy takes advantage of each atom's unique spectral lines for identification. However, spectroscopy of an atom on or in a substrate is generally impractical due to the broadening of spectral lines caused by interactions of the atom with its neighbors[1]. Therefore, once a Ba atom has been adsorbed onto a substrate for transport, it must first be removed to a vacuum environment before attempting spectroscopic identification.

A schematic overview of this Ba tagging technique after Ba adsorption onto a surface is shown in Fig. 5.1. First, adsorbed atoms (adatoms) are removed via laser-induced thermal desorption (LITD). In this step, an infrared laser is incident on the substrate, which causes rapid heating sufficient to thermally dissociate bonds between adatoms and the substrate. This results in a gas phase plume of atoms directly in front of the surface. Next, two lasers tuned to Ba spectral transitions are passed through the plume. These lasers excite the Ba through a two-step resonant process to an autoionizing state which relaxes through the emission of an electron. This process, called resonance ionization spectroscopy (RIS), causes selective ionization of only the Ba atoms in the plume. Finally, the Ba^+ ion is accelerated out of the laser interaction region by a series of electrodes to its final detection. In this work, the Ba^+ is drifted through a time-of-flight mass spectrometer (TOF-MS) to separate by mass all ions created, and then detected by a channel electron multiplier (CEM) with high efficiency for detecting single ions. This is convenient for evaluating what ions are ultimately produced during this process. However, this technique could also straightforwardly couple ions to an ion trap, where laser-induced fluorescence (LIF) of Ba could be performed for an additional, non-destructive, stage of identification [60, 61].

[1] A notable exception is spectroscopy inside a transparent, weakly-interacting matrix such as Xe ice, which is itself a Ba tagging technique currently being pursued[55].

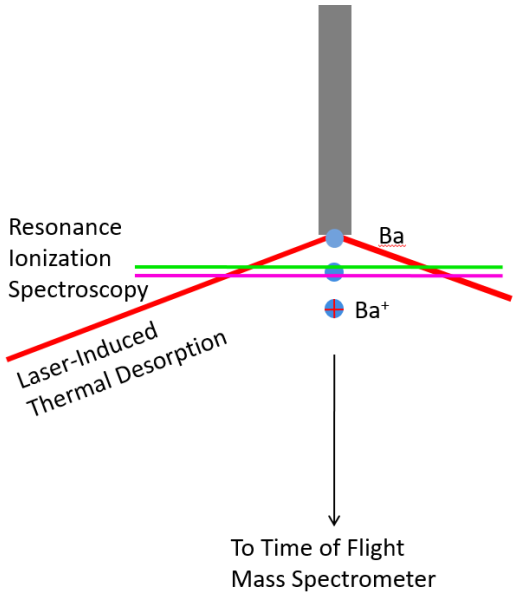

Figure 5.1: Ba tagging technique employed in this work.

5.1 Laser-Induced Thermal Desorption

5.1.1 Adsorption and Desorption

The desorption of an adatom depends strongly on the chemical interaction between it and the underlying substrate. A typical adatom bound through a chemical bond (chemisorption) experiences a potential energy curve like that shown in Fig. 5.2. This potential is primarily characterized by a single parameter, the desorption energy E_{des}, which is equal to the depth of the potential well. Typically, measurements of E_{des} are available for situations where a full atomic layer of adatoms (monolayer) or more is present on top of a substrate. This is chemically very different from the case of a single Ba atom adsorbed onto a bare substrate, so for many substrates the best estimates are given by *ab initio* calculations using density functional theory [62, 63, 64, 65].

However, these calculations are designed for the case of an adatom binding with a single, atomically-smooth crystal of just the substrate atoms, and are not guaranteed to match empirical measurements of real surfaces. The binding energy in general depends on what other adatoms might be nearby, the local crystal orientation, the

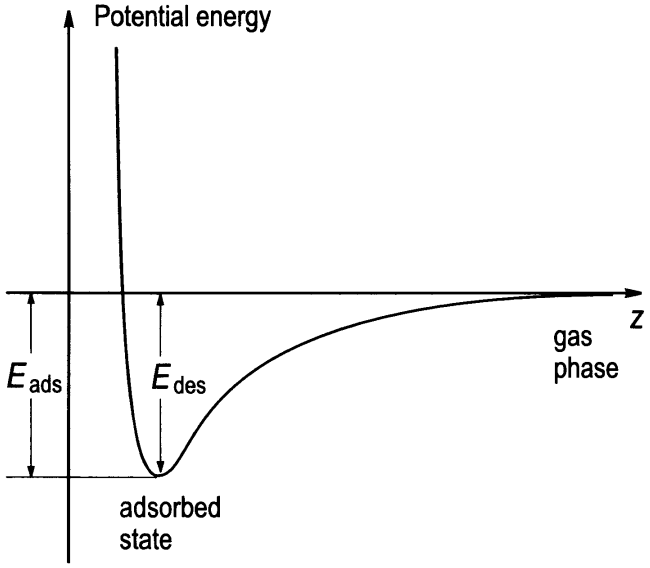

Figure 5.2: Potential energy for a typical adatom. From [66].

surface topography, and whether the surface is truly the bare substrate material or, more commonly, some oxide or other compound, so that it is quite difficult to estimate E_{des} accurately. Accordingly, a Ba atom deposited in one location on the substrate may have a different E_{des} than in a different location, unless the substrate is uniform at the scale of a few atoms.

Thermal desorption offers little in the way of atomic selectivity, as distinguished from quantum desorption mechanisms such as desorption induced by electronic transitions (DIET) or resonant laser ablation (RLA) [67]. However, it is much simpler to achieve for an arbitrary adsorbate. For a single atom of a given E_{des}, the rate of thermal desorption r_{des} is given by the Arrhenius equation [66]:

$$r_{\text{des}} = f \exp(-E_{\text{des}}/k_B T) \tag{5.1}$$

where the rate factor f is typically on the order of 10^{12} s^{-1}, roughly the frequency of atomic vibrations in the lattice [68]. For a laser that heats the surface to a time-dependent temperature $T(t)$, the total probability of desorption from purely thermal

processes for a single pulse is then

$$P_{\text{des}} = \int_{t_0}^{t_f} f \exp(-E_{\text{des}}/k_B T(t)) dt. \tag{5.2}$$

5.1.2 Laser Heating and Desorption Rate

In this work, the laser pulses used for desorption are at 1064 nm and are ∼7 ns in duration. The energy in each pulse is absorbed by free charges in the substrate but then rapidly transferred to heat in the lattice with a timescale of roughly $10^{-13}\,\text{s}^{-1}$. Therefore, it is expected that the atomic motion of the lattice (which governs the desorption process) is approximately in thermal equilibrium with the electronic motion throughout the process.

The absorption of each laser pulse's energy in the substrate depends on the material's optical attenuation length (or absorption coefficient) α. The absorbed power density as a function of depth, z, for a laser pulse which is Gaussian in time is then given by [69]

$$\dot{q}(z,t) = (1-R)\,\alpha \frac{F}{\sqrt{2\pi}\sigma_t} e^{-\frac{(t-t_0)^2}{2\sigma_t^2}} e^{-\alpha z} \tag{5.3}$$

where R is the surface reflectivity, F is the laser fluence (energy per area), and σ_t is the time width of the laser pulse. Because the beam size is typically large compared to the thermal diffusion length, the temperature as a function of position parallel to the sample surface is not significantly affected by lateral thermal diffusion, so it is appropriate to treat the heat flow as one-dimensional (into the substrate). In addition, the thermal diffusion length is typically large compared to the optical penetration depth, so the energy deposited can be treated as a heat source at the surface. The effects of heat loss due to surface re-radiation are orders of magnitude below the incident pulse intensity, so can be safely ignored. With these approximations, the temperature of the substrate can be modeled using the one-dimensional heat equation

[70]:

$$\rho c_p \frac{\partial T(z,t)}{\partial t} = \frac{\partial}{\partial z}\left[K \frac{\partial T(z,t)}{\partial z}\right] + \dot{q}(z,t) \tag{5.4}$$

where T is the temperature, ρ is the mass density, c_p is the heat capacity, and K is the thermal conductivity. This differential equation can be solved numerically, given the boundary conditions that the whole sample starts at room temperature ($T(t = 0) = 300$ K) and that the temperature deep in the solid is left unchanged throughout the process ($T(z \to \infty) = 300$ K).

A solution to this equation for an incident pulse with a maximum intensity of 20 MW/cm^2 is presented in Fig. 5.3, using properties for heavily-doped Si at 1100 K [71]. This simple model has some limitations, most notably that it does not include the temperature dependence of the substrate's thermal and optical properties. This can in fact have a considerable effect on the temperature profile, depending on the material. It is particularly relevant when the substrate is heated above the melting temperature, in which case the heat of fusion must be properly accounted for. The numerical solution shown does not reach the melting temperature of Si (1687 K), but some of the data in this work is expected to reach or exceed melting temperatures for the substrate in question. For these reasons, the calculations presented here should be considered a good guide for the main qualitative features of LITD in general, but not a complete, quantitative model of the process. A calculation for Si at 694 nm which takes these effects into account is presented in [72]. Such a detailed treatment is in practice only useful under conditions where the surface is highly uniform, the desorption energies are well-known, and all substrate thermal and optical properties are established.

With this in mind, the solution to Eq. 5.3 and 5.4 can be combined with Eq. 5.1 and 5.2 to get the instantaneous desorption rate versus time and the integrated desorption probability for a single laser shot (Fig. 5.4). There are a few important features of this desorption curve. First, there is a strong, non-linear dependence of the desorption rate on the surface temperature, so that most of the desorption happens in about 10 ns near the time of the peak surface temperature. This also

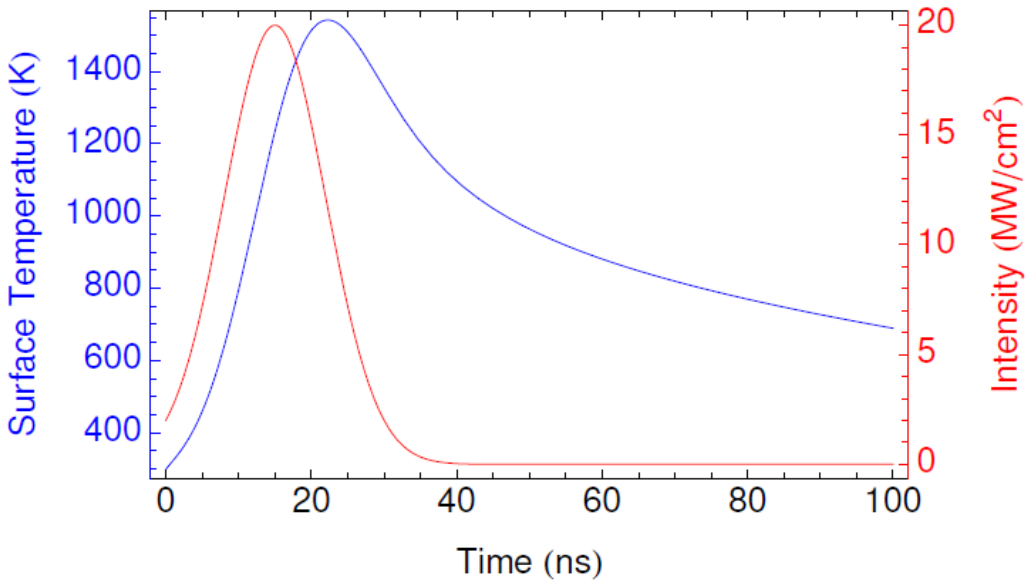

Figure 5.3: Solution to the one-dimensional heat equation for the surface temperature of heavily-doped Si due to a 1064 nm Gaussian laser pulse of 7 ns duration and 20 MW/cm^2 peak intensity. The peak intensity occurs at $t_0 = 15$ ns. Does not account for temperature dependence of Si optical and thermal properties.

means that the desorption happens primarily from the center of the laser beam, with tails that fall off at a shorter distance scale than naïvely expected from the beam's spot size. Additionally, there is a very strong dependence on E_{des}. This means that running at a given laser power (peak temperature) mostly probes adsorbates with a fairly narrow range of E_{des}, as adsorbates with lower E_{des} will quickly burn off unless there is a continual source of them, *e.g.* the bulk of the substrate. Hence, such a system exhibits hysteresis in that each increase in desorption laser power leads to different adsorbates being removed, but returning to lower power does not lead to the reappearance of adsorbates of lower E_{des}.

The strong dependence on E_{des} means that it is crucial to find a substrate material and preparation which allows for desorption in a practical number of laser shots. This leads to an approximate range of acceptable desorption energies: a desorption energy below about 1 eV is weak enough that it may break spontaneously at room temperature while moving the probe (about 1% probability over 15 min at 1 eV), while an adsorbate with a desorption energy above about 2.5 eV would take thousands of

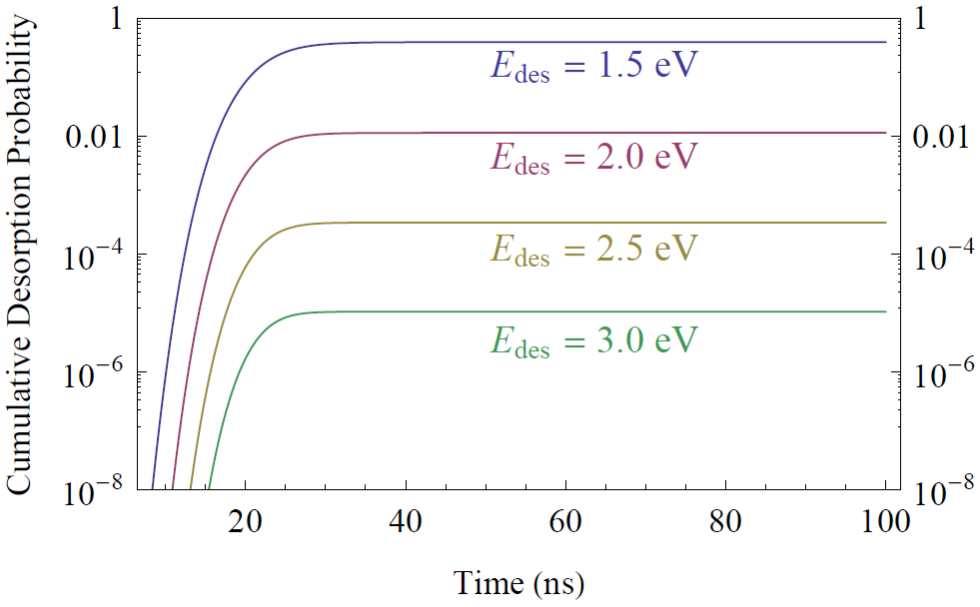

Figure 5.4: Cumulative desorption probability during a single laser pulse versus time, for an adsorbate of different desorption energies. The surface temperature profile is the same as in Fig. 5.3. A rate factor of $f = 10^{12}$ s^{-1} is used.

laser shots or more to be removed with high probability if the substrate reaches about 1500 K. This upper limit can be pushed up if higher temperatures are reached, either through higher laser power or a more easily-heated substrate, but this runs the risk of damaging the surface through ablation or creating excessive backgrounds due to desorption of ions.

5.1.3 Ablation and Desorbed Ions

Often a distinction is made between laser desorption – a thermal process where small fractions of a monolayer are evaporated per laser shot, as described above – and ablation, which is more akin to a small surface explosion, heating the surface well beyond its vaporization temperature, removing many layers of substrate, and leaving behind a crater. The processes for laser ablation are generally more challenging to model than desorption, so descriptions in the literature tend to focus on empirical results over careful analytical treatment. More complete models which attempt to

account for the heat of vaporization and surface recession, phase explosion (explosive vaporization), and shielding of the substrate from the laser via plasma formation are available for aluminum [69, 73] and may be valid for other materials as well [74].

Though there is no clear separating line between these two regimes, laser fluences below $\sim 10^6$ W/cm^2 are considered in the desorption regime, while fluences above $\sim 10^9$ W/cm^2 clearly cause ablation [75]. While ablation is a more violent process than desired for the case of removing a single adsorbed atom, removal of adsorbates with a high E_{des} may require laser fluences which are high enough to begin causing some of the effects associated with ablation. In this work, laser intensities are in an intermediate regime of $10^6 - 10^7$ W/cm^2, with fluences below 0.5 mJ/cm^2. This should be well below the threshold for phase explosion of about 5.0 mJ/cm^2 [74], but may in some cases cause surface damage, particularly if there is non-uniform heating due to surface structures or deviations from a Gaussian laser profile.

It is also possible at high laser intensities to remove adsorbates as ions directly, without any photoionization processes such as RIS. In a mass spectrometer such as that employed in this work, such ions may constitute a background to the resonantly-ionized Ba signal if they are close in mass to Ba, especially if produced from a high-power laser pulse which can lead to long-time tails to each mass peak's time of flight distribution. A first-order approximation to the desorbed ion fraction, which is valid when the desorbed material is in thermal equilibrium, is given by the Saha-Langmuir equation, here assuming ionization states above singly-ionized are negligible and that the degeneracies of the ground state and ionized state are equal:

$$\frac{n_+}{n_0} = \frac{2}{n_e \lambda_e^3} \exp\left(-\frac{E_i}{k_B T}\right) \tag{5.5}$$

where n_+ is the density of ions, n_0 is the density of neutral atoms, n_e is the density of electrons, $\lambda_e \equiv h/\sqrt{2\pi m_e k_B T}$ is the electron's de Broglie wavelength, T is the temperature of the desorbed material, and E_i is the first ionization energy.

In practice, the desorbed material does not thermalize after its emission, so it may or may not have a well-defined temperature T. If the desorption happens thermally to begin with, however, the thermodynamics of ion removal will depend more on the

difference in energy between the adsorbed atom and the final state of the desorbed ion plus the substrate with an extra electron. In other words, the way electron(s) are shared in the chemical bond with the substrate is important in determining the ion fraction. Qualitatively, the form of Eq. 5.5 demonstrates that the ion fraction will depend exponentially on some effective energy for desorption as an ion. Hence, the backgrounds from the desorption process will heavily favor those atoms (or molecules) which are most readily ionized.

5.2 Resonance Ionization Spectroscopy

After the LITD step, the neutral Ba can be separated from the other neutral components by taking advantage of its unique spectral transitions via RIS. With sufficient laser power and the right choice of level scheme, each transition can be saturated, allowing for ionization efficiencies which approach 100% and hence detection of single atoms [76]. In addition, with lasers of wavelengths above 320 nm (corresponding to the first ionization energy of cesium [58]), a single photon has insufficient energy to ionize any atom. This means that atoms other than the atom of interest can only be ionized by highly-suppressed processes such as two-photon non-resonant absorption to the ionization continuum. Consequently, the only backgrounds that can be produced in the RIS step are from photoionization of molecules and from atoms with nearby spectral lines, if the laser linewidths are too broad.

5.2.1 RIS Scheme

The level scheme for Ba used in this work is shown in Fig. 5.5. Any atomic Ba desorbed is expected to be in its $6s^2\ {}^1S_0$ ground state, since its lowest-lying excited state is 1.1 eV from the ground state [77], and is hence expected to be strongly suppressed for thermal desorption due to the Boltzmann factor of Eq. 5.1. The initial transition is from the ground state to the 6s6p $^1P_1$ excited state via single photon absorption at 553.548 nm (553.702 nm) in air (vacuum). The Ba in this state then transitions to the 5d8d $^1P_1$ autoionizing state via absorption of a photon at

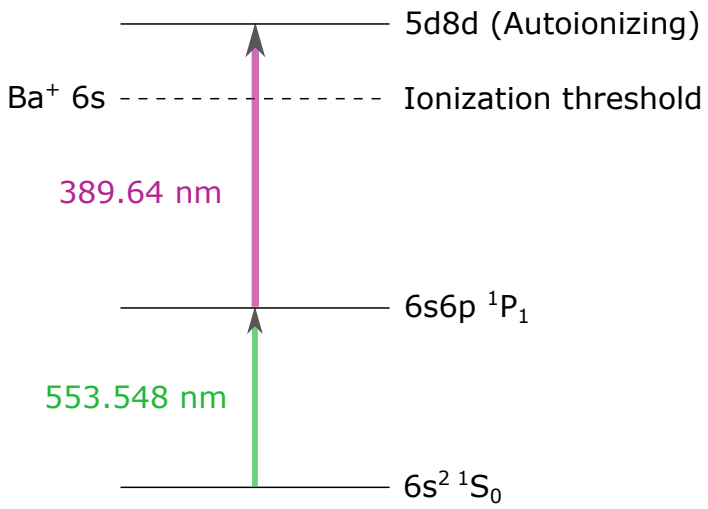

Figure 5.5: RIS scheme for this work. Wavelengths listed are in air.

389.64 nm (389.75 nm) in air (vacuum). Autoionizing states are present in atoms with more than one electron, where discrete atomic states involving two or more excited electrons interfere with continuum ionization states [78]. The probability of an autoionizing state decaying via electron emission (ionizing) is orders of magnitude higher than via photon emission. Such states are especially convenient for RIS because they carry relatively high cross sections for photoabsorption compared to ionization directly to the continuum.

This scheme was chosen due to its high cross sections, which allow for the possibility of 100% ionization efficiency with intensities available with commercial dye lasers: the 553.5 nm (green) transition has the highest cross section of all Ba ground state transitions [77], while the 389.6 nm (UV) transition has the highest cross section of autoionizing states from the 6s6p $^1P_1$ excited state [79]. In addition, the presence of other atomic lines should not significantly affect this scheme: the lifetime for the decay of the 6s6p $^1P_1$ excited state to other states not in this scheme (6s5d states) is 40 μs or more, which is much longer than the 10 ns timescale relevant for the RIS process. Furthermore, there are no transitions from either the ground state or the 6s6p $^1P_1$ excited state which are nearby the RIS lasers in use, as the nearest such transition is from the ground state to the 5d6p $^3P_1$ state at 388.93 nm (air) [77] which

is well outside the linewidth of the UV dye laser when tuned to 389.64 nm.

5.2.2 Green Transition

Let us now consider in more detail each of these two transitions. In the case of the transition from the ground state, the transition linewidth is narrow relative to the 553.5 nm (green) dye laser linewidth of $\Delta\nu_{beam} = 0.0012\,\text{nm}$ (1.2 GHz in frequency). Hence, the fraction of the laser power which is relevant for this transition scales with the inverse linewidth of the laser for a Gaussian or Lorentzian lineshape, and the absorption cross section for a given transition is given by [76]

$$\sigma = \frac{g_2}{g_1}\frac{\lambda^2}{4\pi^2}\frac{\Gamma_{21}}{\Delta\nu_{beam}} \tag{5.6}$$

where g_i is the degeneracy factor for the ground state ($i = 1$, $g_1 = 1$) or the excited state ($i = 2$, $g_2 = 3$ for unpolarized light), λ is the transition wavelength, and $\Gamma_{21} = 1.19 \times 10^8\,\text{s}^{-1}$ [77] is the spontaneous transition rate from the excited state to the ground state. For this transition and the laser linewidth used, the absorption cross section as calculated by Eq. 5.6 is $\sigma = 2.36 \times 10^{-11}\,\text{cm}^2$.

Such a transition is said to be saturated when (in the absence of other lasers) the excited state population is in equilibrium with the ground state population when accounting for degeneracy factors: $n_2 \approx \frac{g_2}{g_1+g_2} n_1 = f_2 n_1$ with n_i the density of atoms in state i and $f_2 \equiv \frac{g_2}{g_1+g_2}$. For this to occur, two conditions must be satisfied:

1) The flux of photons must be sufficient to excite ground state atoms faster than their rate β of removal from the excited state: $\sigma F > \beta$, where F is the flux of photons.

2) The fluence (time-integrated energy density) of photons must be sufficient for each ground state atom to have many chances to absorb a photon: $f_2\phi\sigma \gg 1$, where ϕ is the fluence of photons.

The flux condition is easily satisfied with pulsed dye lasers which have very high instantaneous fluxes. Because the transition rate out of the 6s6p $^1\text{P}_1$ excited state is dominated by the decay to the ground state with rate Γ_{21}, the lifetime for decay of

the excited state is about $1/\Gamma_{21} \approx 8\,\text{ns}$, so a given atom may decay and be re-excited during the comparable $\sim 10\,\text{ns}$ duration of the laser pulses.

Hence, saturation is achieved for an interaction region where the local energy fluence $\Phi(\mathbf{r}) = (h\nu)\,\phi(\mathbf{r})$ satisfies the condition

$$\Phi(\mathbf{r}) \gg \Phi_{sat} \equiv \frac{h\nu}{f_2\sigma} \tag{5.7}$$

which defines the saturation fluence, Φ_{sat}. More commonly, the fluence is averaged across the beam and an effective interaction area A is defined, which leads to a condition on the total beam energy $E \gg E_{sat} \equiv (h\nu)\,A/f_2\sigma$. This is useful when the density of atoms to be excited is expected to be fairly uniform, as in the case of a thermal vapor cell where experimental determinations of E_{sat} are often made, but for the case of thermal desorption, the atom density of the cloud must in general be taken into account. Fortunately, the very high cross section for this transition means that for this work which employs a beam of roughly 2 mm diameter, $E_{sat} < 10\,\text{nJ}$, whereas beam energies of order $100\,\mu\text{J}$ are typically used. Therefore, the entire interaction region is easily saturated even during the early- and late-time tails of the pulse, and any efficiency losses due to incomplete ionization can be attributed to the transition to the autoionizing state.

5.2.3 UV Transition

In the case of the autoionizing transition, mixing with continuum ionization states causes the linewidth to be considerably wider than the dye laser beam, $\Delta\nu_I \approx 0.05\,\text{nm} \gg \Delta\nu_{beam} \approx 0.001\,\text{nm}$ (2.0 GHz). Under these conditions, the absorption cross section is approximately independent of linewidth, and has been measured to be $\sigma \approx 5.50 \times 10^{-16}\,\text{cm}^2$ [79, 80]. These cross sections are calculated in Ba vapor cells where the density is very high and collisions between atoms are expected to equalize the populations of all three degenerate m_J sublevels of the 6s6p $^1\text{P}_1$ excited state. In these conditions, no effect of the relative polarization of the two resonant lasers is observed on the cross section.

This need not be the case for thermal desorption, where atomic densities are much

lower and collisions may be negligible, depending on the incident laser power. In this case, the linear polarization of the green laser excites atoms to states with $m_J = 0$ along the polarization axis. The transition from this $J = 1$, $m_J = 0$ state to the autoionizing state with $J = 1$ is then forbidden if the UV is polarized parallel to the green laser, since this polarization will leave m_J unchanged. There is significant mixing between states, which means the transition is not purely one with $\Delta J = \Delta m_J = 0$, and moreover hyperfine interactions for odd isotopes can cause depolarization effects that make this resonance still present even for parallel polarizations. The lasers are also not perfectly linearly polarized. However, there is still a substantial difference in the cross section for parallel versus perpendicular RIS laser polarization in the absence of interactions to equalize the m_J populations [81, 82].

With its lower absorption cross section, the autoionizing transition is then the one most likely to contribute to ionization efficiency losses. Here, it is useful to have a more complete model of the ionization efficiency. A full description involves solving the system of differential rate equations for each possible transition, as done in *e.g.* [76]. However, if we assume that the green transition is saturated for the duration of the UV pulse, *i.e.* the ground state and the 6s6p ^{1}P$_1$ excited state are at all times kept in the ratio $f_2 = g_2/(g_1 + g_2)$, then the fraction of atoms which are ionized in a region of local fluence $\Phi(\mathbf{r})$ is simply [83]

$$f_i(\mathbf{r}) = 1 - e^{-f_2 \sigma \Phi(\mathbf{r})/h\nu} = 1 - e^{-\Phi(\mathbf{r})/\Phi_{sat}} \tag{5.8}$$

and we recall that $f_2 = 1/2$ for linearly polarized light with no interactions that equalize the m_J sublevels, while $f_2 = 3/4$ when the m_J sublevels are equally populated. This assumption of saturating the green transition at all times is easily verified by checking that the ionization rate is unchanged (within experimental error) by increases in the green laser intensity. Such tests in this work have validated this assumption for green laser energies above $\sim 10\,\mu$J.

From this calculation, we can get the total ionization fraction f_I of all Ba atoms removed by LITD into the ground state by integrating along with the density of Ba

atoms $\rho(\mathbf{r})$:

$$f_I = \int \rho(\mathbf{r}) \left(1 - e^{-\Phi(\mathbf{r})/\Phi_{sat}}\right) d^3\mathbf{r} \tag{5.9}$$

where $\rho(\mathbf{r})$ is normalized to 1 and is assumed to be effectively constant for the duration of the RIS pulses. This is expected to be true, as typical thermal velocities at LITD temperatures are of order $1\,\mathrm{mm}/\mu\mathrm{s}$ so the transit time across the roughly mm-sized RIS beams is of order $1\,\mu\mathrm{s}$, while the RIS duration is of order 10 ns. This also assumes that Φ_{sat} is constant in time, which may not be true in the presence of collisions which can alter f_2.

5.2.4 Saturation and Detuning

Prior work has utilized a Ba oven to produce a sufficiently large density of Ba atoms to get a reliable measure of the linewidth (Fig. 5.6) and average saturation energy (Fig. 5.7) of each transition using the same lasers as in this work [84]. In the linewidth measurements, the intensity of the laser being detuned was set low to minimize power broadening effects. Doppler broadening, caused by the motion of the atoms, broadens the natural linewidth by a factor of $\Delta\lambda = \lambda_0\sqrt{\frac{8k_BT\ln 2}{mc^2}}$. This is small compared to the measured linewidth of both transitions for the prior measurement at room temperature. At LITD temperatures of a few thousand Kelvin, the broadening of the green transition is significant, but unimportant due to its high absorption cross section. The broadening of the UV transition is small compared to its inherent linewidth even at LITD temperatures - $\Delta\lambda \approx 0.001\,\mathrm{nm}$ at 3000 K. The saturation measurements confirm that the green transition is readily saturated, while the UV transition approaches saturation at $1\,\mathrm{mJ}$/pulse for the conditions used. Due to differences in the Ba density profile (uniform and of effectively infinite extent for this prior work) and potentially the magnetic sublevel population (if atomic collisions played a role in the prior work), the saturation curve for this work is expected to be similar but not identical.

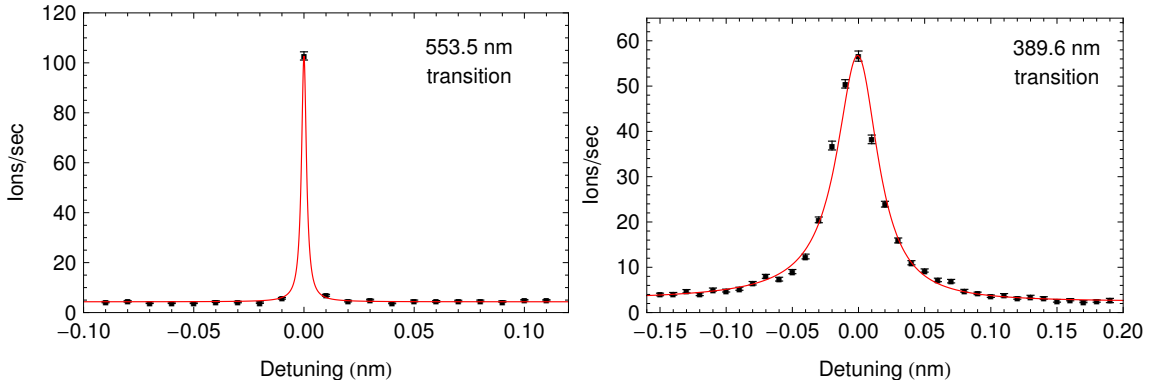

Figure 5.6: Ba^+ detection rate as a function of wavelength detuning from each resonance. The lineshape is a convolution of the inherent linewidth and the linewidth of the dye laser. Errors shown are statistical only. Left: The green transition, fit to a Lorentzian centered at 553.548 nm. Right: The UV transition, fit to a Fano profile (accounting for mixing of the autoionizing state and the continuum) centered at 389.64 nm.

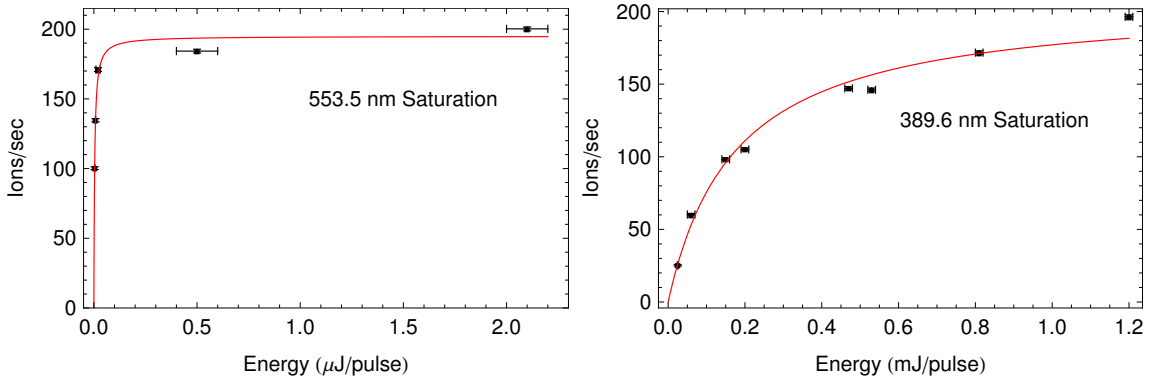

Figure 5.7: Ba^+ detection rate as a function of energy for each laser. Fits were performed assuming a constant average laser intensity, so they do not account for increasing ionization from the tails of the beam as the intensity is increased.

Chapter 6

Experimental Design

The goal of the Ba tagging system developed in this work is to demonstrate deposition and recovery of Ba ions from a liquid xenon (LXe) environment. Ideally, the technique pursued here will allow for detection of a single Ba ion with high efficiency

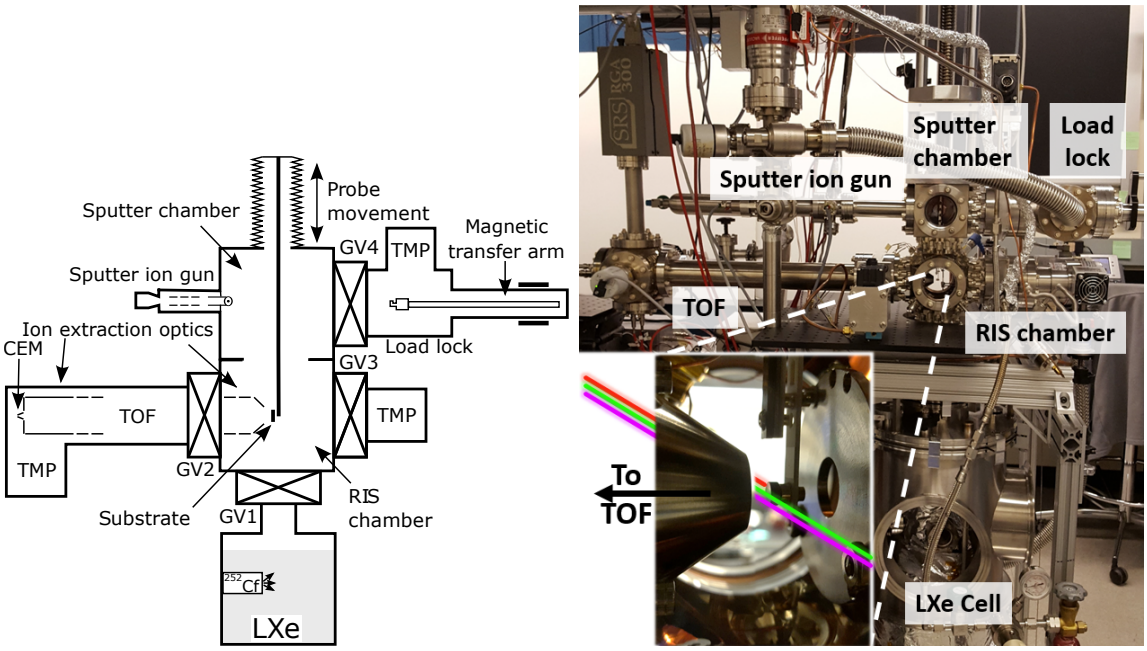

Figure 6.1: Simplified diagram (left) and photograph (right) of the Ba tagging system.

when fully characterized. This includes the steps of drifting ions to a substrate, adsorption onto the substrate, removal of the substrate to vacuum, desorption from the substrate, ionization of desorbed neutrals, and finally identification of ions by mass. Integration with a full TPC including event reconstruction and probe delivery to any region in the TPC is not part of the scope of this project.

A diagram of the full system is shown in Fig. 6.1, with detailed descriptions of its components and their operation in the following sections. Additional description of the system is available in [85] and [86], with more recent additions to the system described in §6.1-6.2. The full Ba deposition and recovery process is summarized as follows:

1) The deposition substrate, which is mounted on the end of a moveable probe, is prepared *in situ* as desired; this may include *e.g.* cleaning by IR laser exposure or ion sputtering.

2) A gate valve (GV1) between the RIS chamber and LXe cell is opened and the probe is lowered into the LXe cell where a Ba production source is present. This

cell may be operated in vacuum, Xe gas, or LXe as desired and appropriate for the Ba production source of choice.

3) Ba from the source is adsorbed onto the substrate for a duration corresponding to the desired level of Ba deposition. The substrate may be electrically-biased to promote collection of Ba produced in ionic form.

4) The probe is returned to the RIS chamber. If operating in a Xe environment, GV1 can be closed and any Xe gas in the chamber is pumped away to return the RIS chamber to vacuum.

5) Ba on the substrate is desorbed by LITD and ionized by RIS lasers as described in Chapter 5. To better ensure complete removal of deposited Ba, the desorption laser is scanned across the surface of the sample, hitting each location for many laser shots.

6) Any ions created are then guided by electrodes through the time of flight (TOF) mass spectrometer to a final channel electron multiplier (CEM) capable of detecting single ions.

6.1 Load Lock Chamber and Substrate Mounting

The most recent design for substrate transfer and mounting is shown in Fig. 6.2. Pin-shaped substrates are held in place by a spring-loaded receptacle which is part of a ceramic hybrid mount. This ceramic hybrid fits onto a transfer key at the end of a magnetically-coupled translation arm which allows the substrate to move from the load lock chamber to the main chamber while under vacuum, as shown in Fig. 6.1. The substrate mount then mates to a second hybrid circuit attached to the tip of the movable probe by pushing two receptacles on the substrate mount onto two pins on the probe tip. This provides a mechanical connection for the probe tip to the substrate mount while ensuring that the orientation and position of the substrate relative to the probe tip is fixed and repeatable from one substrate to the next. It also makes two separate electrical connections to the substrate and a conductive

pad directly above the substrate via conductive traces that run along the back of the substrate mount. This allows both to be biased during desorption to shape the electric field for ion collection and sets the drift voltage for the TOF. The pins on the probe tip ceramic hybrid are connected to feedthroughs at the top of the probe via Kapton-coated wires which pass inside the probe tubular shaft. After mating the two hybrid circuits, the transfer arm is disengaged by pulling the probe tip up and drawing the transfer key back through the slot in the substrate mount. Substrate removal is achieved by reversing the entire process.

The probe shaft is a thin-walled stainless steel tube 1.27 cm in diameter and 101 cm long, sufficient to allow travel from the sample transfer position to the Ba-deposition position in the LXe cell. The thin-walled design minimizes the thermal mass of the probe to avoid excess boiloff during immersion in LXe. Its vertical movement is actuated through a long bellows by an external stepper motor with a nominal position accuracy of 3 μm. The transverse motion of the probe is constrained by pairs of spring-loaded rollers, one at the bottom of the bellows and one near the top of the LXe cell, which progressively engage the shaft as it is lowered. These features ensure that the position during both Ba deposition and recovery are repeatable to better than 100 μm.

This design has a number of useful qualities for the full Ba deposition and recovery process. The materials used for the probe shaft, probe tip, and substrate mount are all ultra-high vacuum (UHV) compatible and operable from room temperature to LXe temperatures. This includes the fluxless solder used to connect the gold-plated pins and receptacles to the conductive traces on the ceramic hybrids. The rollers are made of Vespel. The compact design for the substrate mounting system is required for the probe tip to pass through the rollers. Grabbing the substrate from the back, with its face offset from the ceramic, reduces the likelihood of ion production from the desorption laser hitting components other than the substrate. It also eliminates the possibility of contaminants from the mount migrating to the substrate face, a concern for previous designs utilizing clips to hold the substrate in place.

The addition of the load lock chamber allows the RIS chamber to always remain under vacuum (or filled with very pure gas Xe), even while replacing substrates. This

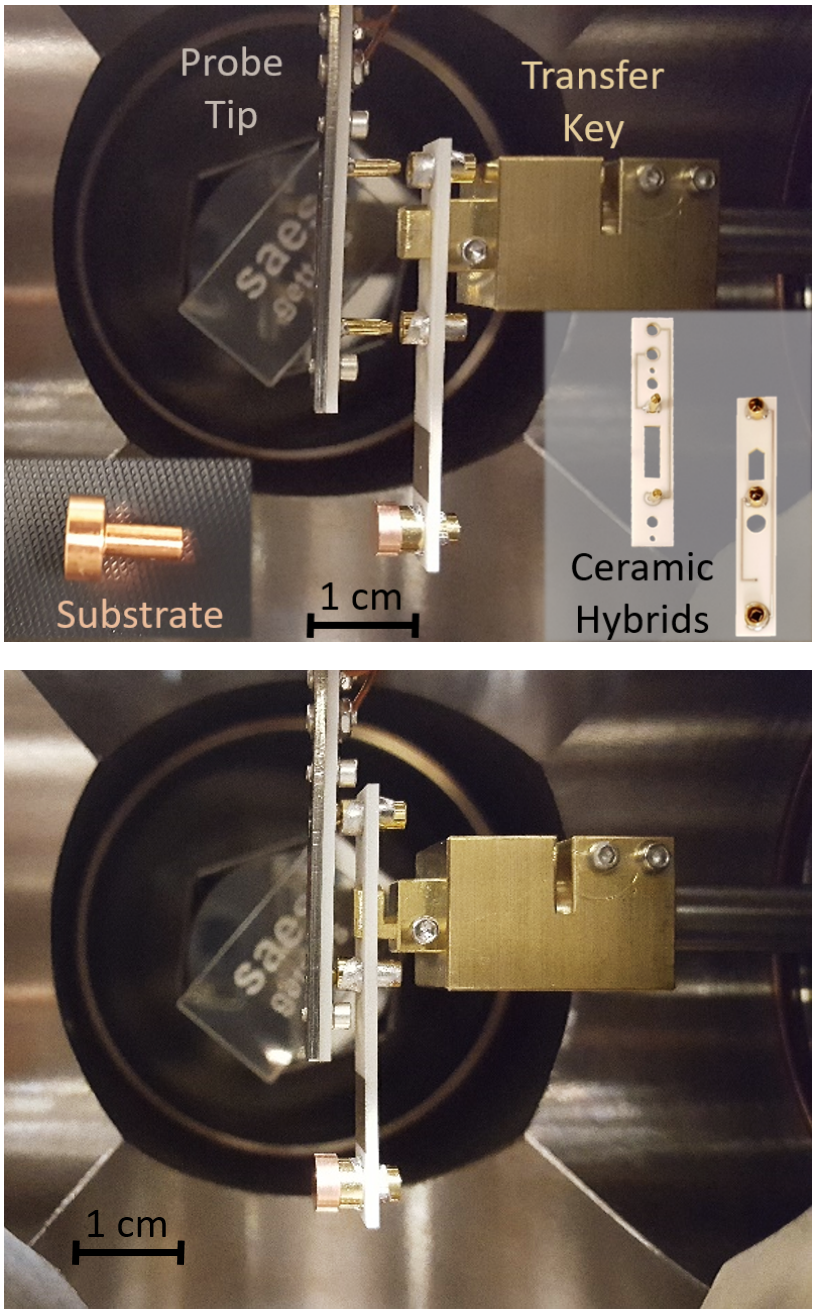

Figure 6.2: Top: Substrate transfer system just before mating the probe tip to the substrate mount. Insets show a copper substrate and the two ceramic hybrid circuits with conductive traces visible. Bottom: Substrate transfer system during transfer key disengagement.

allows the system to be baked once to 200 C to remove water vapor adsorbed onto the vacuum chamber walls, permanently reducing the vacuum chamber pressure. This is valuable as residual vacuum pressure will cause the development of adsorbate layers on the substrate. Such adsorbates might promote formation of Ba molecules such as BaO, which the RIS process would not be sensitive to unless the Ba-O bond is broken during the desorption step. A non-evaporable getter (NEG) pump[1] is also included in the main chamber, which is heated during bakeout to further reduce the pressure of components such as H_2O and H_2. Such reductions exceed what is achievable from the turbopumps without the risk from ion pumps of producing background ions in the TOF or substrate contamination. After both bakeout and NEG activation, the total RIS chamber pressure is reduced below 1×10^{-8} mbar, with the H_2O partial pressure declining by a factor of >200.

6.2 Argon Ion Source

The chamber substrates are transferred into also houses an argon ion source (Fig. 6.1). This allows *in situ* preparation of substrates by ion sputtering to remove adsorbate layers as well as shape the sample topography. It is positioned at normal incidence for equal sputtering in both dimensions parallel to the substrate surface; for ease of alignment relative to the substrate; and for ready sputtering of selected sections of the substrate to better test the effects of sputtering. The ion beam is roughly Gaussian in intensity, with a width during normal operation of $\sigma_x \approx \sigma_y = 0.2$ mm at a typical working distance of 45 mm. The beam current is tunable up to 1 μA and the beam can be automatically rastered across a rectangular area of the substrate by two pairs of electrostatic deflection plates at the beam aperture. Ion energies are adjustable in the 1-5 keV interval.

The beam current and alignment relative to the substrate is verified by measuring the current through an ammeter connected in series to substrate and ground as the beam is deflected. The profile of the current versus beam position as it crosses the

[1] SAES Getters CapaciTorr D400-2, 400 L/s H_2 pump speed, 300 L/s H_2O pump speed

substrate edge can be fit to an error function to get the beam size, and the half-current position can be used to set a point on the edge of the substrate. Scanning both vertically and horizontally at several positions produces a rough outline of the substrate, which can then be used to find the substrate center and the proper scaling from deflection voltage to transverse distance at the substrate. Using this method, the center of the substrate has been determined to be stable relative to the ion beam to better than 0.1 mm, including shifts between substrates.

The sputter yield for 5 keV Ar ions in Si is about 1-1.5 atoms/ion and in Cu is about 3-5 atoms/ion [87, 88]. When running at the full 1 μA current and rastered across a 1 mm^2 area of the substrate, this translates to an average sputter rate of 0.12-0.18 nm/s in Si and 0.2-0.4 nm/s in Cu. Hence, removal of several atomic layers across a large fraction of the substrate is achievable in minutes, while removal of a few μm of substrate material is possible by sputtering overnight.

6.3 Xenon System

A copper cell with a volume of about 1 L is used for the condensation of Xe gas when a LXe environment is desired (Fig. 6.3). Because of the narrow temperature window for the liquid phase of Xe, the temperature of the cell is monitored and maintained within 1 K or less of 165 K. Cooling is provided by a copper heat strap connecting the bottom of the cell to a cold reservoir of liquid nitrogen (LN_2). The high thermal conductivity of this strap allows for cooling of the cell to 165 K in several hours, but also necessitates heating of tens of Watts through resistive heaters to maintain the operating temperature. The cell and dewar are surrounded by a vacuum jacket to reduce thermal conduction with surrounding air and covered in reflective superinsulation to reduce radiative heat transfer. Temperatures at several locations on the cell are measured using resistive thermal devices. These readings are used as feedback to a computer-controlled PID loop which adjusts the heating power to maintain the cell temperature. The Ba source is mounted onto one of the CF flanges brazed onto the cell.

Xe is delivered to the system from a Xe gas cylinder through a regulator which

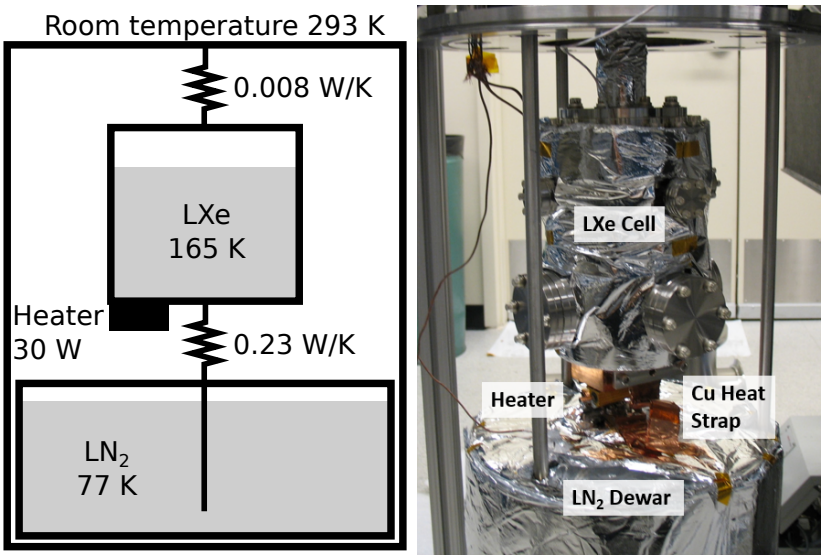

Figure 6.3: Diagram of the thermal flow for the LXe cell (left) and photograph (right). During operation, the system is kept under vacuum insulation and the heat strap is covered in reflective superinsulation.

feeds into a purifier to remove impurities such as oxygen and water vapor. The pressure inside the system is monitored by Baratron capacitance manometers when Xe is present and by cathode gauges when under vacuum. Xe recovery is achieved by immersing a second gas cylinder in LN_2 such that any Xe gas present will condense onto its walls (cryopumping). Opening a valve from the RIS chamber to the cryopump bottle allows for recovery of Xe from over 900 Torr down to under 1 Torr in ~ 1 min. The remaining Xe can be safely pumped out to the atmosphere through a turbomolecular pump (TMP).

6.4 Barium Ion Sources

The ultimate goal of identifying depositions of single Ba ions with high efficiency requires that Ba sources under consideration have very low production rates which are also reliably known. For these reasons, and to avoid the risk of contaminating the vacuum chamber and its components with large amounts of Ba, Ba ovens and similar sources as used in prior work [84] were not considered. The use of radionuclides as a

means of Ba production is a convenient alternative due to their compact nature, the ease of determining their activity, and the ability to produce only a few atoms at a time. Two such sources are used: a ^{148}Gd source for use in vacuum environments for initial testing and a ^{252}Cf source for use in high-pressure Xe gas and, in the future, in LXe.

6.4.1 ^{148}Gd-driven Source

This Ba source is described in detail in [89]. It consists of a layer of ^{148}Gd electrodeposited onto a silicon substrate with a thin layer of BaF_2 evaporated on top. When the ^{148}Gd α decays, the recoiling nucleus carries enough energy to knock loose a few atoms from the BaF_2 layer. The source used in this work has an activity of 173 Bq with an estimated Ba production of roughly 5 atoms per forward recoil calculated using the SRIM simulation package [90], for a total production rate of roughly 430 Hz. The relative production of Ba^+ ions versus Ba atoms is estimated in [89] but has not been measured. This source has been partially characterized by placing it at the entrance to the TOF, allowing some of its lower-energy ions to pass through and be identified by mass. These measurements show roughly equal amounts of Ba^+ and BaF^+, though little is known about the majority of the emitted Ba which is expected to be neutral. Most Ba is emitted with relatively low energy (< 1 keV) and hence is not expected to embed itself into the substrate material, though the tail of the energy distribution extends out to 89 keV. For the same reason, this Ba source is unsuitable for use in a Xe environment, as neutral Ba atoms emitted will be stopped by the Xe rather than being collected onto the substrate.

6.4.2 ^{252}Cf Fission Fragment Source

To accomplish Ba deposition in Xe, a higher-energy source of Ba is required. ^{252}Cf decays primarily through the emission of αs, but 3% of decays proceed by fission fragmentation into two neutron-rich isotopes and their subsequent decay chains. These two initial fragments are most likely to be around masses 110 u and 140 u [91], so that there is significant production of Ba isotopes, which are close in mass to this second

peak. The high energies of these decays ($\sim$ 70 MeV on average) allow the fission fragments to travel $\sim$ 30 μm in LXe according to SRIM [90], far enough to readily escape their image charge from the source when the substrate is biased negatively relative to the source. In addition, the stopping range in gas Xe at 1 bar is on the order of 1 cm, so a similar biasing of the substrate should allow operation in gas Xe provided the separation from the source is sufficient.

The half-life of ^{252}Cf is 2.65 yr [40], which is sufficient for use over an extended period, but also produces a sufficient decay rate for observation of Ba deposition during the initial phase where this technique is not yet sensitive to single atoms. A complete calculation of the population of each decay isotope versus time can be done by solving a series of coupled differential equations including for each isotope its daughter(s) and parent(s). In the case of Ba, its production and subsequent decay is entirely through β decay processes, so its differential equations are of the form:

$$\dot{N}_{^{A}\mathrm{Ba}}(t) = N_0 \lambda_0 Y^{ind}_{^{A}\mathrm{Ba}} + N_{^{A}\mathrm{Cs}}(t)\lambda_{^{A}\mathrm{Cs}} - N_{^{A}\mathrm{Ba}}(t)\lambda_{^{A}\mathrm{Ba}} \tag{6.1}$$

where $N_{^{A}\mathrm{Ba}}$ is the population of Ba isotopes of mass number A, $\lambda_i = 1/\tau_i = \ln(2)/T_i^{1/2}$ is the decay rate for the corresponding isotope, and $Y^{ind}_{^{A}\mathrm{Ba}}$ independent fission yield for the isotope in question. Similar equations can be written down for all other parents which yield Ba isotopes. A numerical solution to these equations for all Ba isotopes produced in significant quantities by the ^{252}Cf source is presented in Fig. 6.4 for a source activity of 17 kBq (the activity, as of April 2017, of the source used in this work). A summary of these results and the half-lives of the Ba products is presented in Table 6.1.

This calculation assumes that there are no fission products present in the Xe to begin with, though the validity of this assumption depends on what happens to the fission products after they are released. If the source is biased to re-adsorb fission fragment ions (when the substrate is not in the deposition location) and there is not significant neutralization of fission fragments in the Xe environment, this assumption should hold.

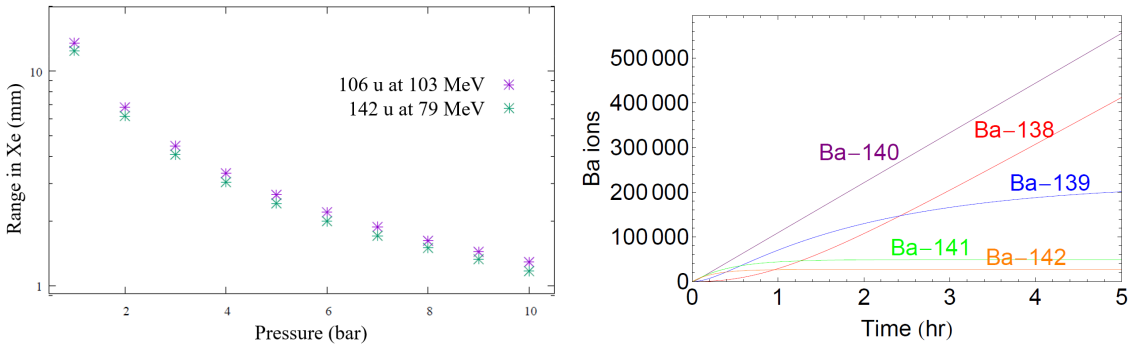

Figure 6.4: (Left) Stopping range of ^{252}Cf fission fragments in high pressure Xe gas at the two fission peaks around mass 106 u and 142 u at their most common energies. Plot courtesy T. Brunner. (Right) Accumulation of Ba isotopes produced in significant quantities versus time for a 17 kBq ^{252}Cf fission source, assuming no fission products to start.

Isotope	Half-life	$N(t = 30\,\text{min})$	$N(t = 5\,\text{hr})$
^{138}Ba	Stable	6400	410000
^{139}Ba	83.1 min	30000	200000
^{140}Ba	12.8 day	53000	560000
^{141}Ba	18.3 min	33000	49000
^{142}Ba	10.6 min	23000	27000

Table 6.1: Summary of Ba isotope half-lives and production from a 17 kBq ^{252}Cf source.

6.5 Laser Optics

The typical optical setup used in this work is presented in Fig. 6.5. The output of a single Nd:YAG pulsed laser is passed through a frequency doubling (tripling) crystal as a pump beam to provide the energy for the green (UV) dye laser. A second Nd:YAG laser provides the IR desorption pulse. Both lasers are Q-switched by an external computer-controlled fast timing signal which repeats at approximately 10 Hz. This constant repetition rate provides maximum power stability, while the fine timing of the pulses can be adjusted to set the relative delay $\delta t = t_{RIS} - t_{IR}$ between the desorption and RIS pulses and affect which lasers are present during data taking. A typical run alternates repeatedly between a shot with the pulses combined ($\delta t \approx 1\,\mu$s), and with

just the IR ($\delta t \approx -100\,\mu$s). Because the data-taking window for the TOF is no more than $70\,\mu$s, this ensures that any ion counts coming during this IR-only configuration are from the desorption laser, while counts from the combined configuration are from both the desorption and RIS lasers. This allows for determining which TOF peaks are produced by direct ionization from the desorption laser versus ionization from the RIS lasers. The exact delay for the combined configuration can be adjusted to maximize the signal for resonantly-ionized Ba by improving the overlap of the desorbed Ba and the RIS beams. Two mechanical shutters ensure that the substrate is only exposed to the laser beams during data taking.

A series of beam pickoffs and photodiodes (PDs) provide shot-by-shot timing information as well as rough power measurements for each beam, all of which is recorded by the computerized data-taking program. The most reliable power measurements, however, come from a separate power meter which is used to determine the laser powers before each data-taking run. The desorption laser power is set by an internal power knob. The RIS pump beam is always set to maximum power for improved stability, but the individual green and UV dye laser beams' powers are set by attenuating the beam through a polarizing beamsplitter (BS). Because the dye laser beams are linearly polarized, a half-wave plate can be rotated to set the incoming polarization relative to the BS, which determines the relative power that passes through. The beams' polarization remains fixed by the orientation of the respective BS. A dichroic mirror allows for the green and UV beams to be collinear as well as optimally spatially overlapped, to ensure efficient resonant ionization.

6.5.1 Desorption Optics

The desorption laser has a Gaussian spatial profile before passing through any optical elements. Because the desorption rate is highly non-linear in temperature (and hence laser intensity), a more uniform intensity profile is preferable, to improve desorption uniformity across the beam spot. The uniformity of the profile is improved by passing the beam through a pinhole so that the central, flattest region of the Gaussian is preserved, as in some ablation experiments [92, 93]. The minimum size of the pinhole

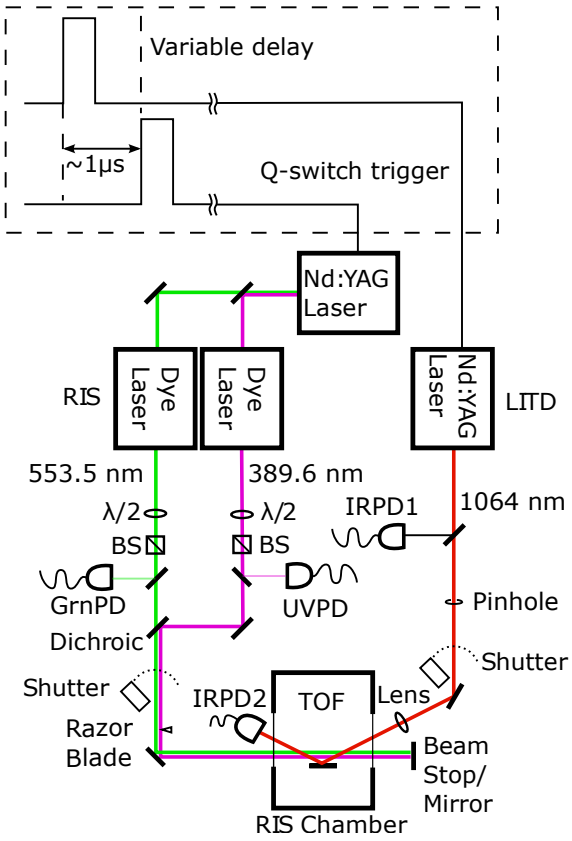

Figure 6.5: Typical laser optics configuration and timing.

is limited, however, by requirements on the beam power, which makes the difference from a Gaussian beam modest.

The beam is then passed through a lens to set the spot size of the beam on the sample. The final mirror and lens are mounted on a translation stage with horizontal and vertical degrees of freedom to allow the beam to be rastered across the surface of the substrate. Smaller beam sizes allow for better spatial resolution, which can be helpful in identifying mass peaks in the TOF (*e.g.* by spatial correlations), for distinguishing uniform features from localized ones (*e.g.* small particles on the substrate surface), etc. This comes at the expense of requiring more positions to be probed to cover the full substrate area, and hence more time for fixed laser intensity. A reasonable balance is achieved with a lens that focuses the beam 25 cm away, whereas the sample is 14 cm from the lens.

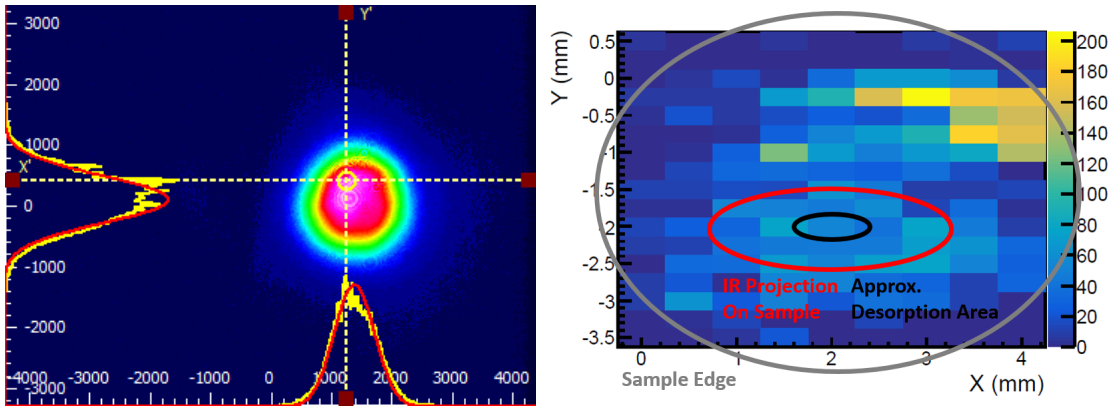

Figure 6.6: (Left) Intensity profile of the desorption laser after the iris and lens, at roughly the same distance from the focus as the sample. Slices along the horizontal and vertical axis at the point of highest intensity are shown on the coordinate axes along with Gaussian fits. Units are in μm. (Right) Integrated counts for a specific TOF mass peak as a function of desorption laser position. The sample geometry is overlaid in grey and the projected laser spot size for one particular desorption position is overlaid in red. Note that the scale of variation (overlaid in black) is much smaller than the beam size due to the non-linear dependence of desorption on surface temperature.

The spatial profile of the beam under these conditions at a similar distance to the focus as the substrate is shown in Fig. 6.6. Due to the beam's angle of incidence of $\sim 70°$, its projection on the substrate is roughly three times wider along the horizontal axis. The distance from the lens to the substrate also changes slightly as the translation stage moves, so that the spot size is roughly 5% smaller in each dimension at the far edge of the substrate compared to the near edge 5 mm away. This corresponds to a slight preference for desorption from the far edge, though this effect is typically smaller than changes due to local surface features.

A summary of the typical running parameters for the desorption laser is given in Table 6.2. The range of powers used here corresponds to a peak intensity $I_0 = P/(2\pi\sigma_x\sigma_y) = (4-18)$ MW/cm^2 (where P is averaged over the 7 ns pulse duration), which is in an intermediate regime between desorption and ablation, as described in §5.1.3. Often a data-taking run involves rastering the desorption beam across the full substrate surface – this both allows for recovery of Ba from the entire deposition area

Laser	Duration (ns)	Power (mJ/pulse)	Horiz. size (mm)	Vert. size (mm)
IR	7	1-5	2.5	1.0
Green	8.5	0.2	1.3	1.3
UV	6.0	2.0	0.40	1.8

Table 6.2: Summary of typical laser operating parameters. Spot size is $w = 2\sigma$ from a Gaussian fit to the profile. Duration is $\pm 2\sigma$ from a Gaussian fit to the temporal profile as measured by a fast PD. For the RIS lasers, the power and spot sizes are measured without the razor blade in place.

as well as provides a TOF mass spectrum which averages over local surface features, making comparison of one run to the next more reliable. Such a scan typically takes on the order of 30 min to an hour, allowing relatively rapid data analysis. A sample plot of the signal from one TOF mass peak versus desorption laser position is given in Fig. 6.6. Its features are significantly smaller than the laser spot size due to the non-linear dependence of desorption on surface temperature. Though some datasets may include locations where the desorption beam center is off of the sample edge, the analysis region can easily be restricted to a smaller area if desired, after the fact.

The alignment of the desorption beam relative to the sample is set by measuring the amount of light reflected by the sample with a separate PD (IRPD2) as the translation stage is moved. The image produced by this signal is a convolution of the beam profile with the sample surface shape, and its center (judged by eye) corresponds to the translation stage position where the beam center coincides with the sample center. Note that because of deviations from a Gaussian beam, the peak power position (which corresponds to the region of maximum desorption) does not exactly coincide with the beam center; this difference is ~ 0.5 mm on the 5 mm substrate, which makes the choice of alignment center unimportant as long as it is consistent for a given substrate.

6.5.2 RIS Optics

Typical beam parameters for both RIS lasers are presented in Table 6.2, along with profile measurements in Fig. 6.7. The asymmetry in the UV profile ensures that it

covers the full substrate desorption region (vertical extent) while concentrating the beam power close to the surface where the desorbed plume is expected (horizontal extent). Without further shaping, the beam centers must be offset from the substrate surface to avoid ablating it, which reduces the spatial overlap of the RIS beams with the desorbed plume. A simple solution to this is achieved with the introduction of a razor blade to produce a sharper rising edge for the RIS beams, as shown in Fig. 6.8. The position of the razor blade relative to the center of the beam can be adjusted to maximize the RIS signal. Though the exact position will depend on the UV power, RIS delay, and the size of the desorbed plume, for typical operation a good balance between retaining most of the UV power and achieving a sharp rising edge is found when 60–75% of the beam's energy is transmitted past the razor. For 75% transmission, the rise distance (10% to 90% of the max intensity) is approximately $250\,\mu$m.

Once the razor position is set, the alignment of the RIS beams relative to the substrate surface can be set using the final RIS mirror's tilt as well as the precision translation stage it is mounted on. This is performed each time a new substrate is used. The vertical alignment is set by tilting the mirror and aligning the beams to the center of the substrate's shadow by eye. This rough alignment suffices due to the tall vertical profile of the UV beam. The horizontal angular alignment is set by tilting the mirror into the substrate surface so that a reflection appears and then tilting it back away (adjusting the translation stage accordingly) until the reflection disappears and the beams run parallel to the substrate surface. The horizontal linear alignment is set using the translation stage. The RIS efficiency has been observed to be somewhat sensitive to this alignment, so it is done periodically for any given substrate. Because the spatial overlap with the desorbed plume is generally enhanced by keeping the RIS beams as close to the substrate surface as possible, alignment is performed by shifting the beams toward the substrate until ion counts are observed in the TOF spectrum with the desorption laser blocked. Though the threshold fluence for production of background ions from this process depends on the substrate material and its surface features, measurements with this system set conservative limits of $F_{ion,\,\mathrm{UV}} < 50\,\mu\mathrm{J/mm^2}$ and $F_{ion,\,\mathrm{grn}} < 5\,\mu\mathrm{J/mm^2}$. Because the position where the

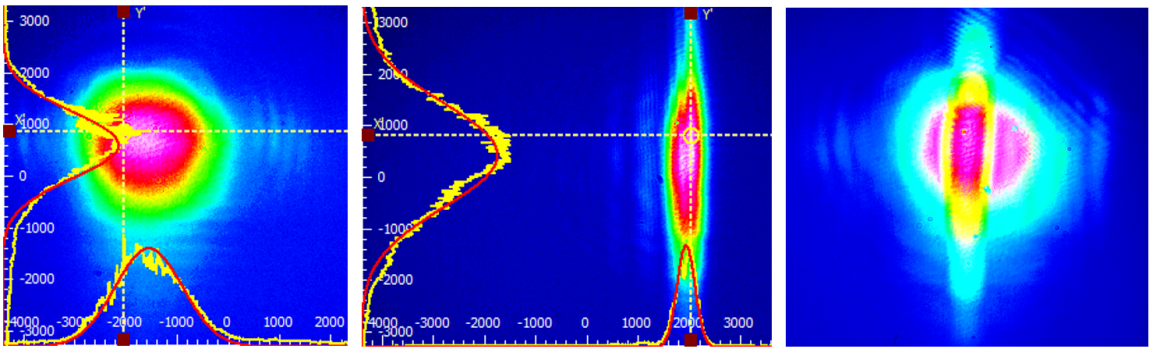

Figure 6.7: Intensity profile of the green (left) and UV (center) lasers, and their spatial overlap (right) without the razor blade in place. Slices along the horizontal and vertical axis at the point of highest intensity are shown on the coordinate axes along with Gaussian fits. Scale is in μm.

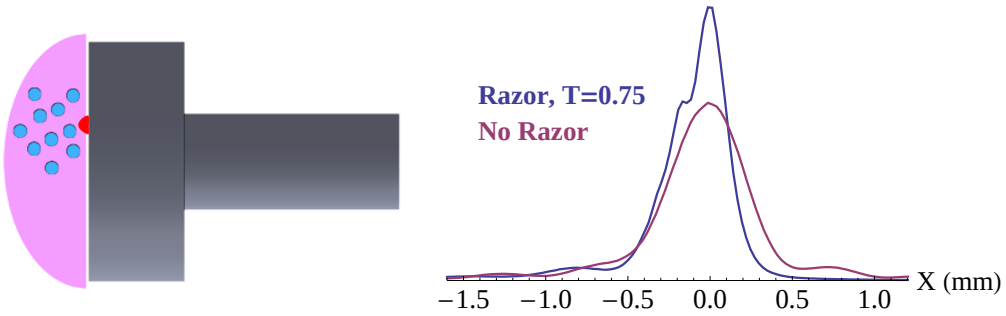

Figure 6.8: (Left) Desired shape of RIS beams (purple) for maximum spatial overlap with atoms desorbed from across the substrate. Desorbed atoms (blue) are shown as they might come from one particular desorption location (red). (Right) Measured UV intensity profile (arbitrary height) projected onto the axis normal to the substrate surface, with the razor blade to the right transmitting 75% of the UV energy (blue), and without the razor blade (purple). The two profiles are normalized to have equal total energies.

profile exceeds either of these thresholds changes as a function of RIS beam power, this alignment is typically performed anytime the power is adjusted.

The wavelengths of both beams are measured daily with a wavemeter, to ensure that their central wavelengths are within 0.001 nm of the transition wavelengths, as described in §5.2. A summary of the RIS wavelengths, peak fluences, and estimated saturation fluences is provided in Table 6.3. As noted previously, the green transition

is easily saturated over a very wide region, so the ionization efficiency is primarily governed by the UV beam's power and spatial overlap with the desorbed plume. This spatial overlap depends both on the beam shape and position as well as the desorbed plume shape at the time of the RIS beams' transit. Because the plume shape evolves in time, generally increasing in size, overlap can be improved by adjusting the IR to RIS time delay δt, with an optimal value typically in the range of $0.5 - 2.0 \mu s$. The default delay is $1\,\mu$s. The relative timing of the green and UV beams is also relevant, since the UV transition can only occur for atoms previously excited by the green laser. Measurements with a fast PD indicate that the center of the UV beam appears $\sim 2\,$ns before the green, but due to the $\sim 8.5\,$ns duration of the green beam and its high absorption cross section, the early tail of the green beam saturates the full interaction region for the duration of the UV pulse. This is verified by increasing the green beam's power and noting the lack of increase in RIS rate above $\sim 20\,\mu$J.

Wavelength (nm)	Peak Fluence (μJ/mm^2)	Saturation Fluence (μJ/mm^2)
553.548	79	1.5×10^{-4}
389.64	1700	9.3

Table 6.3: Summary of RIS laser wavelengths and fluences under typical operating conditions. Wavelengths are in air. Saturation fluences are calculated from [84], and may not be accurate under these conditions, so should be considered rough guidelines only.

6.6 Data Acquisition

Detection of ions by the CEM is destructive, necessitating that all relevant data be recorded for each desorption laser pulse, at the 10 Hz repetition rate of the lasers. During data taking, signals from the PDs and the CEM must be digitized and read out quickly. For the PDs, signals are digitized at 60 MS/s (17 ns resolution), with the peak amplitude and time written for each laser shot. The CEM signal is passed through a high-bandwith preamp and then digitized at 250 MS/s (4 ns resolution) for a $70\,\mu$s time window beyond the trigger time (tied to the laser trigger). This includes

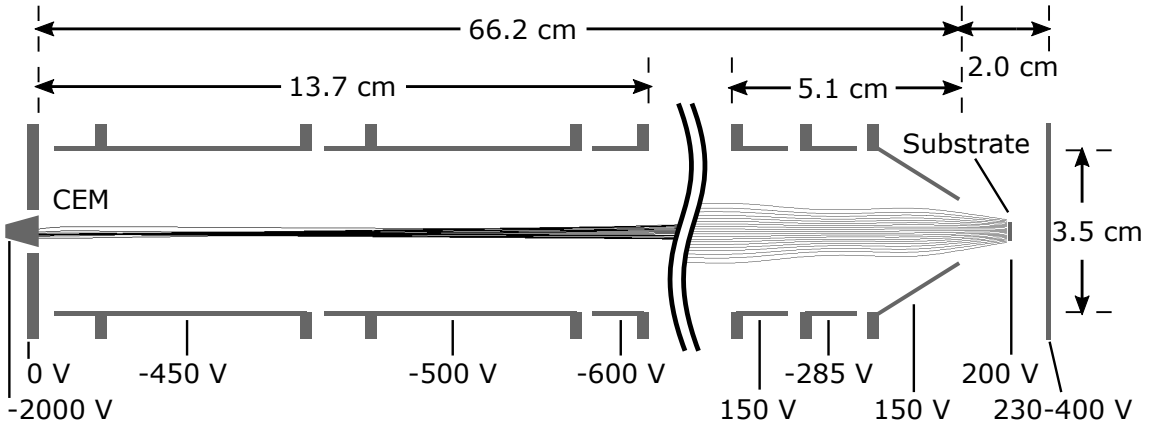

Figure 6.9: Design of the time of flight spectrometer. Simulated trajectories for ions generated uniformly across the substrate with no initial kinetic energy are shown to demonstrate the effects of the electrodes. The exact position of the substrate varies slightly with different substrate mounts, but is ~ 1 cm from the cone-shaped electrode in front of it. To account for this, the voltage of the backing plate behind the substrate is adjusted. An additional conductive pad is present above the substrate to help shape the electric field for initial ion collection in some designs.

masses of up to 300 amu/e. The 4 ns resolution is required to separate pulses from multiple ions of similar mass within each laser shot. The digitized signal is then searched in real time for pulses which exceed a voltage threshold, set so that few true pulses will be rejected and the count rate from noise during the 70 μs time window is negligible. The amplitude and time of each of these pulses is then recorded. Finally, the position of the desorption laser translation stage is recorded for each shot, so that data from different desorption locations on the substrate can be compared.

6.7 Time of Flight Mass Spectrometer

The design of the TOF is shown in Fig. 6.9. The principle behind the TOF is that any ions created at a fixed electric potential V relative to ground will share the same kinetic energy when they reach the ground potential (assuming they are singly-charged). When they drift at ground across a fixed distance, their arrival time

is later for more massive species:

$$t_i = v_i/L = \sqrt{2V\frac{q_i}{m_i}}/L \tag{6.2}$$

where t_i is the time of flight, m_i/q_i is the mass-to-charge ratio, and L is the drift length. Differences in path length due to *e.g.* non-uniform electric fields create a range of arrival times for any given mass and hence limit the mass resolution of the spectrometer, though this effect is reduced for longer drift lengths.

In the design here, ions created by the LITD and RIS processes are initially collected and accelerated by an electric field of about 50 V/cm set up by electrodes near the substrate. This is primarily determined by the voltage of the substrate and the electrode in front of it, which is cone-shaped to allow better optical access for the desorption laser. Because the initial collection is crucial to the transport efficiency and the mass resolution of the TOF, the field is further shaped by an electrode behind the substrate. Some substrate mount designs also include a conductive pad just above the substrate for further control (§6.1). Ions then pass through a triplet of electrodes forming an Einzel lens to start focusing them to a smaller spot, and then a long drift length at ground. Finally, a series of far electrodes re-focus the ions to the CEM for detection.

6.7.1 Time of Flight Simulation

To optimize the voltages, a full electrostatic simulation of the ion transport process is performed using SIMION [94]. To model the transport of resonantly-ionized ^{138}Ba (the most common isotope of natural Ba), ions are:

1) Seeded uniformly across the substrate surface, to model desorption at all locations;

2) Given initial velocities drawn from a (forward-only) Maxwell-Boltzmann distribution at some input temperature relevant for LITD;

3) Allowed to drift freely for 1 μs according to their initial velocities, simulating

the delay between the LITD and RIS pulses;

4) And finally fed to the electrostatic simulation as ions with the positions and velocities determined above.

Because the RIS process is intended to provide most of the separation power, transport efficiency is valued above mass resolution. Using the model above, the electrode voltages are iteratively adjusted in simulation to maximize an objective function which includes both an efficiency and a mass resolution term, given reasonable initial guesses for the voltages. The approximate optimized voltages are shown in Fig. 6.9. The exact results depend on the substrate mount used and the desorption temperature, but for the scheme with a conductive pad above the substrate described in §6.1 and a somewhat high desorption temperature of 3000 K, the simulated transport efficiency is >99% and the mass resolution is $m/\Delta m > 100$, with Δm defined as the FWHM of the TOF distribution in mass. This is not sufficient to distinguish between different isotopes of Ba, but does allow for separation from peaks a few atomic mass units away, such as ^{133}Cs. The estimated TOF relative to t_0 from the desorption laser pulse for resonantly-ionized ^{138}Ba is 39 μs.

These simulations suggest that the distribution of initial ion positions contributes substantially to the spread in TOFs. The initial velocities are also relevant, due both to the extent that they adjust the initial position during the 1 μs drift time and due to the difference in trajectory after ionization. The effect of the desorption temperature on the TOF distribution is demonstrated in Fig. 6.10. There is significant broadening for higher desorption temperatures, as well as a delay of about 1 μs between peaks produced by the desorption laser ("directly-ionized") and the RIS lasers. This TOF delay can be somewhat larger than the delay between the desorption and RIS lasers due to the fact that the desorbed Ba first drifts away from the substrate to a lower potential before being resonantly-ionized; the magnitude of this effect increases both with desorption temperature and RIS delay. In addition, the desorption itself occurs over a finite period, with long-time tails due to the temperature profile as described in §5.1.2. This can broaden the directly-ionized TOF distribution, though will have a significantly smaller effect on the width of the resonantly-ionized distribution.

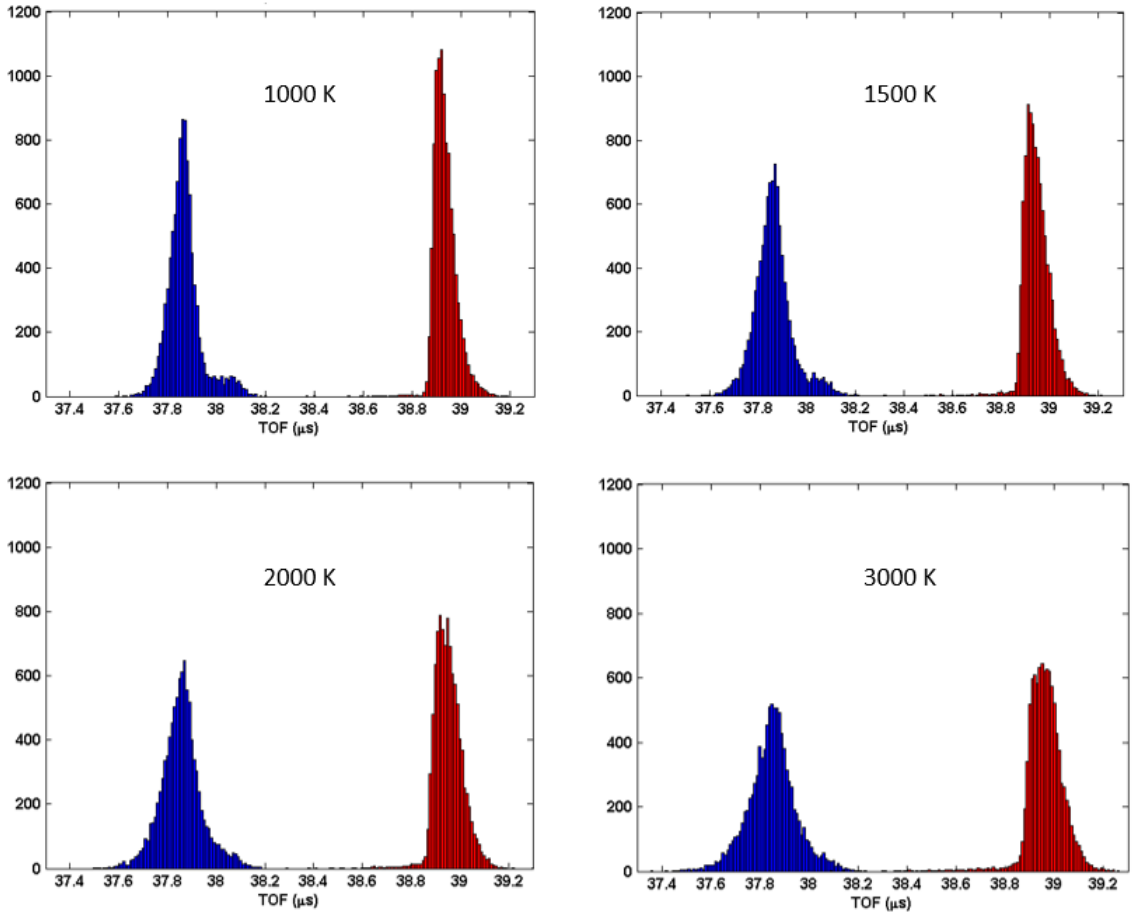

Figure 6.10: TOF distributions for ^{138}Ba ions produced during the LITD process (blue) and from resonant ionization after a 1 μs delay (red), according to SIMION. Equal numbers of the two distributions are generated. The different panels correspond to different desorption temperatures, for Maxwell-Boltzmann distributions.

6.7.2 Time of Flight Features, Calibration, and Corrections

A sample TOF spectrum taken using the voltages described in the previous section is shown in Fig. 6.11. Experimental results confirm the approximate RIS Ba TOF from simulations, ranging from 37-41 μs depending on the substrate mount used, RIS delay, desorption location, etc., though SIMION systematically predicts slightly longer TOFs. The mass resolution is somewhat worse than in simulations, again varying significantly depending on experimental conditions but typically reaching an

$m/\Delta m$ between 50-100 as judged by a Gaussian fit to ^{133}Cs which is at a similar mass to Ba and contains only a single isotope.

The absolute transport efficiency is not measured directly. However, the relative transport efficiency can be inferred by comparing the rate of CEM hits ("counts") in the TOF spectrum for a single desorption location as voltages are varied. There is substantial intrinsic variation in the count rate, in part due to continual changes to the substrate surface as a result of the desorption process. Because of this, experimental optimization of the voltages is impractical. However, the optimum voltage configuration found by simulations has been verified to correspond to a local maximum in transport efficiency measured in this manner, demonstrating that deviations on the order of 10 V for some of the more crucial electrodes can cause noticeable reductions in the transport efficiency.

As is typical of TOF spectra taken with this system, the data in Fig. 6.11 show several peaks other than directly-ionized Ba and RIS Ba, including a strong component of alkali ions. When interpreting TOF spectra, it is important to keep in mind that only ions are observed, so that those elements and compounds which are most readily-ionized are strongly favored, as described in §5.1.3. Hence, the peak ratios are not in general representative of the relative amounts of different elements or compounds on the surface, and in particular the alkalis (with the lowest ionization energies) tend to dominate regardless of the substrate material. Although these peaks may be considered backgrounds, they are easily separated from Ba by mass, which makes their ubiquity quite convenient for calibrating the TOF mass scale.

To set the mass calibration for a particular dataset, the peak TOF t_c for a calibration peak of known mass-to-charge ratio m_c/q_c is found and used to set the value

$$\alpha \equiv \frac{t_c}{\sqrt{m_c/q_c}} \tag{6.3}$$

where it is typically assumed that the calibration peak is singly-charged, so $q_c = e$. In the case of the simple formula from Eq. 6.3, $\alpha = L/\sqrt{2V}$. There are additional complexities related to variation in the flight path and creation voltage which make a constant value of α for all masses no longer strictly correct. However, the validity

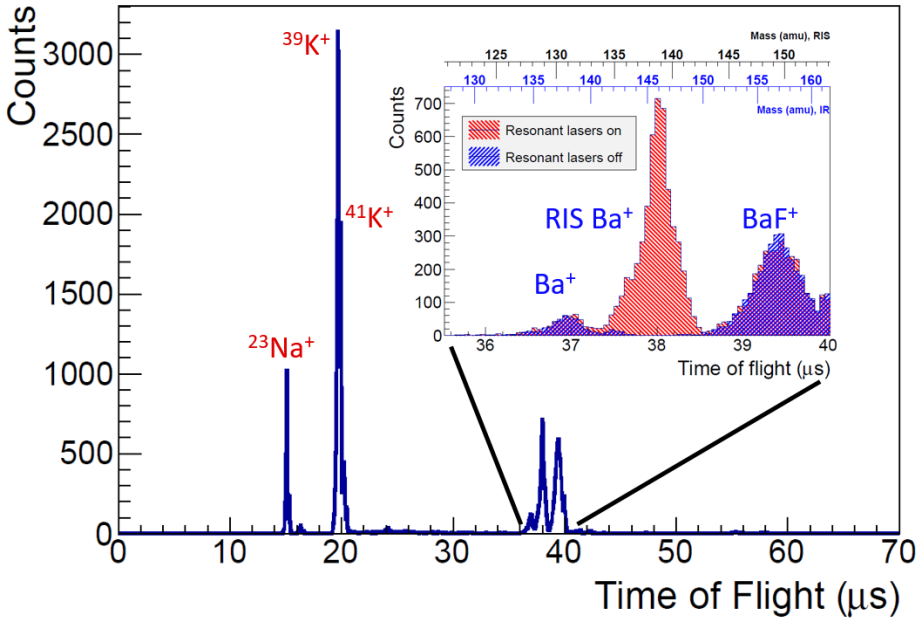

Figure 6.11: Sample TOF spectrum from a single desorption position. Inset shows the Ba region of interest in both TOF and mass. Because of the time delay between desorption and RIS lasers, there are two different mass scales. Alkalis appear in almost all TOF spectra, so can be used as a convenient mass calibration peak.

of this approach from well-characterized peaks ^{7}Li up to ^{138}Ba has been repeatably established to within the mass resolution of the TOF. Given this single calibration value α, the TOF scale can then be converted to a mass scale, as shown for example in Fig. 6.11.

Here it is crucial to use the correct t_0 to determine the true TOF. For directly-ionized peaks, this is the desorption time (within about 10 ns of the time the desorption laser hits the surface on average), while for ions produced by the RIS lasers it is the RIS time. Because the laser times are measured by PDs, relative to the same trigger for the TOF data acquisition, counts from a particular laser shot can be shifted in TOF to correct for jitter in the laser timing, typically on the order of 50 ns for the desorption laser and 100 ns for the RIS lasers. It is impossible to know for any particular count whether the ion was directly-ionized or produced by the RIS lasers when all lasers are present, so the full TOF spectrum must be corrected the same way. Data presented in this work use the jitter-corrected IR time as t_0, with two different

mass scales corresponding to direct ionization and resonant ionization, as exemplified in Fig. 6.11. This improves the mass resolution of directly-ionized peaks, while the alternate procedure using the RIS time as t_0 will improve the resolution for RIS peaks. Masses are then assigned to peaks as a whole (rather than individual counts) by judging whether they are present when the RIS lasers are off and then using the corresponding mass scale. As noted in the previous section, the RIS peaks also shift slightly due to being ionized at a different potential; this effect generally depends on the particular peak, and is difficult to model due to the unknown desorption velocity distribution, so it is not typically corrected for in a systematic way. It rarely shifts the mass estimates by more than 1-2%.

The position of the desorption laser is also recorded for each laser shot, which means that the TOF can be calibrated with an independent α_j for each position j, provided an appropriate calibration peak is present. A typical scan covering a large region of the substrate contains over 100 positions which can vary substantially in the mass peaks present, so only the most reliable peaks, Na and K, have sufficient statistics to be used for this position-dependent calibration. This procedure is rarely performed because in practice it has been shown to have a small effect on the mass resolution for most scans; however, it is quite useful for data with particularly large variations in TOF scale, as in the case of a poor electrical connection for the substrate bias voltage.

Finally, when precise quantitative estimates of count rates are required, the quantum efficiency of the CEM must be taken into account. This is typically on the order of 50% when the rate of ions is small (<1 ion per μs), but decreases for higher rates due to charge depletion - the bias voltage source cannot supply electrons quickly enough to replace those lost during the avalanche process, so the gain of the CEM is decreased and some counts may have insufficient signal amplitudes to be detected. This means that for high count rates (as seen for high desorption power runs), the quantum efficiency varies with TOF, and hence ratios of TOF peak heights may not represent the ratios of ions reaching the CEM.

An example of CEM charge depletion for a dataset with high count rate is shown in the left panel of Fig. 6.12. For TOFs below Na, around 14 μs, the count rate

is low, and the mean CEM amplitude is far enough away from the count threshold that losses are negligible. The very high-count-rate K peaks around 20 μs cause a major charge depletion effect, which continues out to large TOFs. To correct for this, the distribution of CEM amplitudes for a given peak of interest can be fitted to a Gaussian (Fig. 6.12, right panel), and the fraction of the fitted distribution below the detection threshold voltage is the efficiency loss. This method is accurate as long as the peak of the distribution is below the threshold - if not, the fit will be unreliable. It is rare for peaks near the Ba TOF to exhibit efficiency losses from CEM charge depletion above 50%, due to having sufficient recovery time from the K peaks, though scans with the highest desorption power have shown this.

It is worth noting that the CEM amplitude distribution is only Gaussian for a TOF window with a fixed CEM gain (*i.e.* sufficiently narrow). Most mass peaks are sufficiently narrow for a single Gaussian to fit the data reasonably well, but this can be more difficult for TOF regions of rapidly-changing CEM gain, such as near the K peaks. This changing CEM gain also means that peaks with high count rates are biased toward lower TOFs, making them appear less massive. In practice, this is usually only relevant for the K peaks, and just means that they are less reliable for mass calibration than Na.

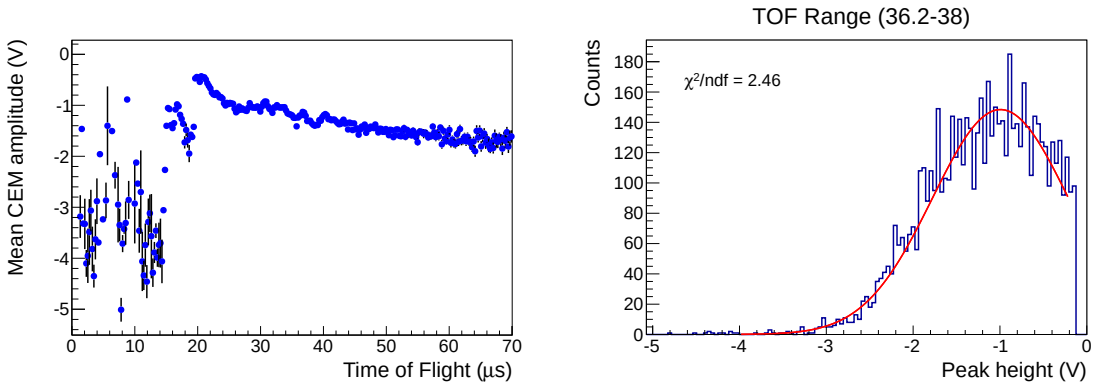

Figure 6.12: (Left) Mean amplitude of CEM counts versus TOF bin, with statistical error bars. Lower voltages correspond to stronger signals. CEM counts in TOF ranges with high rates, such as near K at 20 μs, tend to have weaker signals due to CEM charge depletion. (Right) Distribution of CEM amplitudes for the Ba TOF region of the same dataset, displaying partial loss in count rate due to charge depletion. Gaussian fit shown in red.

Chapter 7

Barium Deposition and Recovery in Vacuum

7.1 Deposition Tests with Radionuclide-Driven Source

7.1.1 Experimental Conditions

Due to the additional complexity and time required to attempt Ba recovery from a LXe environment, many initial tests have been done in a vacuum environment, using the ^{148}Gd-driven Ba source described in §6.4.1. The tests described in this section were performed prior to the upgrades to the vacuum system and installation of the argon ion source described in §6.1-6.2. Several different substrate materials and preparation techniques were attempted, with the goal of finding an optimal surface for Ba desorption. Such a surface would possess low E_{des} (see §5.1), minimal Ba contamination in its bulk, and a chemical state which promotes desorption of Ba in atomic (rather than molecular) form. Initial tests from single-crystal Si of various preparations and crystal orientations, SiC, Ni, and refractory metals such as Ta and W have been performed and summarized in [86]. Such tests did not provide conclusive data for the best material or surface preparation for Ba recovery.

Prior to use, all substrates are cleaned in an ultrasonic bath of acetone followed by ethanol for at least 15 min each. Further attempts at cleaning and chemical preparation of the substrates, both before use and *in situ* are detailed in §8.2. Typical data taking for a new substrate begins with low desorption power (0.5-1.0 mJ/pulse) to avoid damaging the substrate surface. When probing the substrate with the desorption laser at higher power than it has previously been exposed to, the rate of ion counts in the TOF is typically much higher than at the previous power, then decays with continued laser exposure. This is attributed to removal of adatoms with values of E_{des} which are readily desorbed at the higher power but were not removed by prior laser exposure as well as changes to the surface which increase E_{des} or decrease the local heating rate (surface conditioning) - *e.g.* smoothing of sharp surface features which could enhance desorption. Regardless of the mechanism, this means that the effects of laser exposure are generally not reversible.

Because each new data taking run occurs under different surface conditions, repeatability can be challenging. To help with this, the desorption laser is typically rastered across the full substrate surface and the TOF is summed across all locations, as described in 6.5.1. Once the position-averaged TOF is reasonably stable from one run to the next, Ba deposition and recovery is attempted. This proceeds with trials of increasing desorption power until reaching a threshold where detection of Ba (or other TOF peaks of interest) occurs. Such a threshold effect is typical of the desorption process [67], with Ba detection thresholds for this system typically ranging from 1-5 mJ/pulse.

7.1.2 Initial Tests

Despite the range of experimental conditions, TOF data taken at sufficient desorption power (above the Ba desorption threshold) show several common features as exemplified in Fig. 7.1: many peaks produced by the desorption laser ("directly-ionized"), including alkali ions such as Na and K (used for mass calibration, see §6.7.2), metals such as Al and Ni, and other unidentified intermediate-mass peaks; Ba-related peaks including resonantly-ionized Ba and Ba compounds; and a broad background of counts that starts around the ^{39}K TOF and steadily decreases out to large TOFs.

However, despite the presence of both directly-ionized Ba and RIS Ba in many TOF spectra, these peaks have not been reliably shown to come as a result of exposure to the Ba source. Consequently, it must be determined where the "background" Ba (not from the Ba source) in the TOF spectra originates (investigated in detail in Chapter 8) and what happens to the Ba produced by the source.

There are some peaks which clearly appear (or increase) as a result of exposure to the Ba source, most notably one at roughly 157 amu/e which is believed to be BaF^+ (Fig. 7.2). This peak has been reliably seen as a result of source exposure on Si, Ni, and W substrates. It also decays quickly relative to other TOF peaks, suggesting a relatively modest amount of it is present in a readily-desorbed form, as expected for a surface deposition. As described in §6.4.1, the Ba source is known to produce BaF^+ as well as Ba^+, so the presence of BaF^+ in the TOF spectra is not

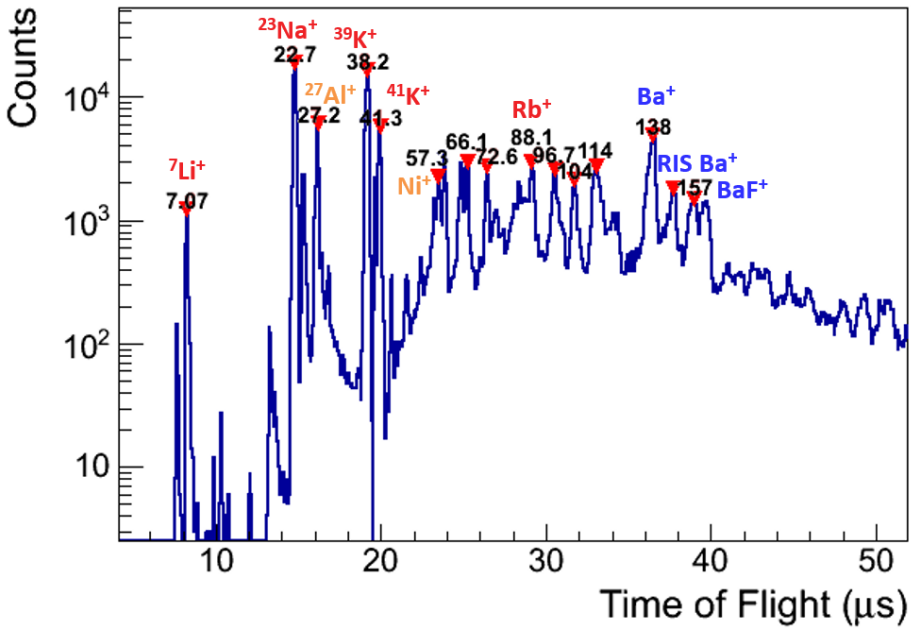

Figure 7.1: Sample TOF spectrum in log scale from a Si substrate at moderate desorption power (3.0 mJ/pulse), summed across multiple desorption positions. Peaks are automatically labeled by mass-to-charge ratio (amu/e) according to the procedure detailed in §6.7.2 and some common elements and compounds are identified. Almost all peaks are ionized by the desorption laser rather than the RIS lasers.

surprising. Relative production rates of BaF and Ba from this source are estimated to be comparable to each other at the time the source was constructed based off of TOF measurements of ions [89]. The estimated efficiency of Ba recovery when the 157 amu/e peak increases as a result of source exposure typically varies from 0.1–0.5%. This assumes all increases in the 157 amu/e peak are from BaF deposition and a deposition rate given by the Ba yield estimated at the time of source production (430 Hz into the forward hemisphere including all charge states and molecular forms) corrected for the solid angle that the substrate subtends (about 1/16 of the forward hemisphere). Attempts to reduce backgrounds through higher-power IR exposure necessitate increasingly-high-power future scans for desorption of deposited BaF. This suggests that this cleaning technique reduces desorption efficiency for all peaks, and may be due to surface conditioning as described in §7.1.1.

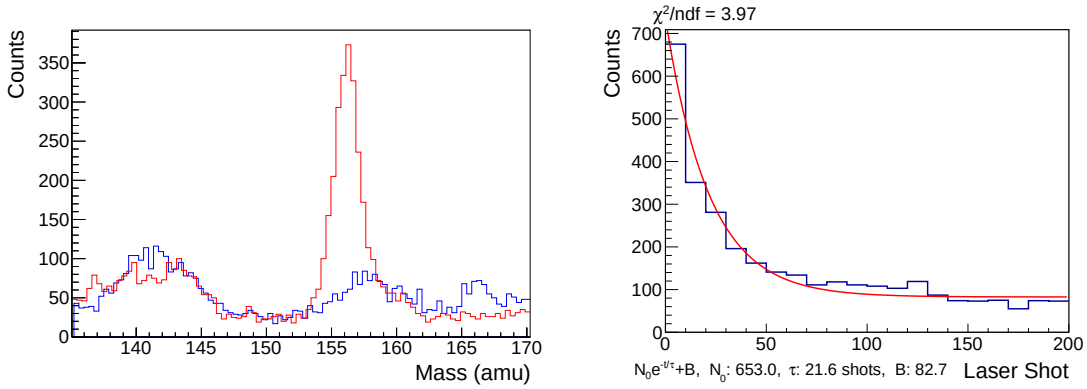

Figure 7.2: (Left) TOF from a Si substrate in the Ba mass region before and after 15 hours of exposure to the Gd-driven Ba source. A peak at 157 amu/e, believed to be BaF^+, appears after this exposure. (Right) Count rate for the 157 amu/e peak versus laser shot from the same dataset, averaged over all locations in the scan. The count rate is fit to a decaying exponential plus a constant background term for counts coming from nearby peaks. The calculated χ^2 uses Poisson errors, which underestimate the true fluctuations in count rate.

The absence of an increase in the directly-ionized Ba and RIS Ba peaks is hypothesized to be due to deposited Ba binding to surface adsorbates and subsequently being desorbed in molecular form. In particular, the Si is expected to have an SiO_2 surface layer while other substrates are likely to have oxygen present from residual O_2 and H_2O in the vacuum chamber, so that formation of BaO is likely. The 3 amu mass difference between BaF^+ and BaO^+ is close to the resolution of the TOF spectrometer, making the identification of the peak as BaF^+ favored but not with high confidence. If BaO^+ is indeed absent from the TOF spectra, a plausible explanation is the relative ionization energies of the two molecules: 6.5 eV are required to ionize BaO compared to 4.6 eV to ionize BaF. If the Ba source produces similar quantities of BaF and Ba, and the Ba binds to the surface as BaO while the BaF binds weakly to the surface compared to the Ba–F bond strength, the LITD process would heavily favor production of BaF^+ over BaO^+ in the TOF. This hypothesis motivated in part the improvements to the vacuum system described in §6.1-6.2.

Despite the clear deposition of this 157 amu/e peak from the Ba source across several substrates spanning years of data taking, later data found such deposition

increasingly difficult to observe. This is likely due to steady chemical degradation of the source over time, as indicated by a comparison of the TOF spectrum at different points during its use (Fig. 7.3). These data were taken by placing the source directly in front of the TOF spectrometer and using the signal from an α detector the source was plated onto as the trigger (t_0). This measurement is not sensitive to production of neutrals, nor to ions with high energy which are not efficiently transported by the TOF spectrometer. However, it clearly indicates both a decline in the ratio of Ba^+ and BaF^+ to other non-Ba peaks (most notably Na^+) as well as a decline in the absolute rate of production of both Ba^+ and BaF^+ of roughly an order of magnitude. This could be a result of chemical changes to the surface which cause fewer of the outgoing Ba to be ionized (while leaving the total production rate the same), or something more severe such as the accumulation of a thicker layer of material which could lower the yield of Ba from the source. Regardless of the explanation, this source is considered unreliable for more recent data where deposition of the 157 amu/e peak is not seen.

The difficulty of Ba deposition and recovery using the Gd-driven source did not, however, prevent evaluation of the LITD and RIS processes for background Ba. Above a certain threshold intensity, such a background was seen for all substrates which were thoroughly investigated. A summary of the threshold desorption intensity required for reliable observation of resonantly-ionized background Ba from different substrate materials is presented in Table 7.1[1].

The peak listed as RIS Ba in the TOF measurements mentioned above is identified both by its mass assignment - consistent with 138 amu after accounting for the delay between the desorption and RIS lasers and ionization at a slightly lower potential - and its dependence on both of the resonant lasers. When either of the lasers is blocked, the signal from this peak vanishes, as demonstrated in Fig. 7.4. However, this signal decays with laser exposure on many substrates, so that statistical fluctuations and the lack of a steady count rate makes precise measurements of the RIS process, such

[1]Ni is from Goodfellow and has >99.99% purity. Si substrates are machined single crystals with intrinsic dopants, of unknown crystal orientation. Cu substrates are electroformed and very high purity as described in §8.2.3. Ta foil is 99.9% purity from Goodfellow. Graphene substrates are produced as described in §7.1.3.

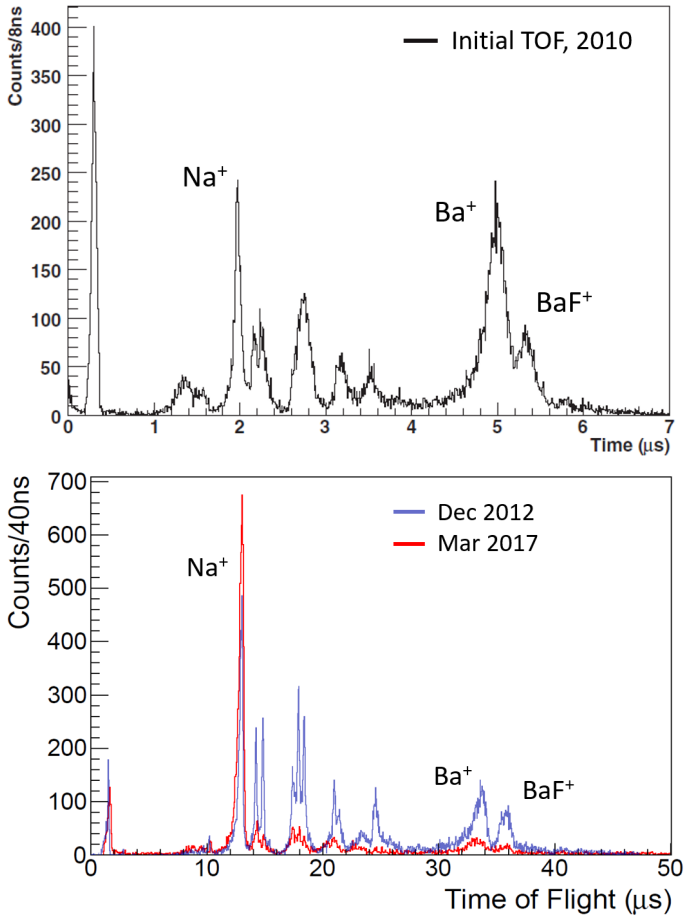

Figure 7.3: TOF spectrum of the Gd-driven source at different points during its use. A different TOF spectrometer with different voltages (and hence TOFs) was used for the initial data (top). Due to its different collection efficiency, a comparison of absolute count rates to more recent data is not meaningful. Data from 2012 and 2017 (bottom) use the TOF spectrometer described in this work and are normalized to the same exposure time, for direct comparison of the emission rate. Both the absolute rate of Ba^+ and BaF^+ ions as well as the rate relative to the Na^+ peak have decreased.

as detuning, difficult. Resistively heating some of the substrates to several hundred degrees C was shown to boost the background RIS Ba signal temporarily, but not enough to overcome these issues. The heating method and a discussion of its effect on the background Ba in the TOF spectra are described in more detail in §8.1.1.

Substrate	Des. Threshold Power (mJ/pulse)	Des. Threshold Max Intensity (MW/cm^2)
Ni	4	15
Si	3	11
Cu	3	11
Ta foil	0.3	6
Graphene on Ni	1.5	6
Graphene on Cu	1	4

Table 7.1: Summary of desorption threshold intensities for "background" RIS Ba for different substrate materials. Thresholds depend strongly on local surface features and the presence of nearby peaks which might obscure a small RIS Ba signal, so be considered past threshold, a noticeable RIS Ba signal must be present in the TOF summed across the entire $2\,\text{mm} \times 2\,\text{mm}$ desorption region. They are considered rough indications of the relative ease of LITD and RIS of Ba from the bulk of the material. Peak intensity is given by $I_0 = P/(2\pi\sigma_x\sigma_y)$ (where the pulse power P is averaged over the 7 ns pulse duration). Data from Ta foil were taken with a reduced spot size.

7.1.3 Results From Graphene

Single-layer graphene is known to be less chemically reactive than many materials, making it an attractive candidate for easy deposition and desorption of Ba. Its simple structure also makes calculations of bond energies tractable and fairly reliable, with density functional theory calculations estimating the bond energy for Ba on single-layer graphene to be 0.67 eV [64], much lower than for materials tested previously. It is also possible that its chemical properties may prevent the buildup of adsorbates such as O atoms, reducing the formation of Ba molecules and hence promoting desorption in atomic form.

For these reasons, Ba deposition and recovery tests were performed using commercial graphene grown via chemical vapor deposition (CVD) which was subsequently transferred onto either a copper or nickel substrate. This sheet is described by the manufacturer as over 95% single-layer graphene, with the remainder being two layers thick. There are some possible concerns to keep in mind for this approach. First, a binding energy of 0.67 eV is likely too low, allowing desorption at room temperature in seconds. However, the presence of the underlying substrate certainly changes the

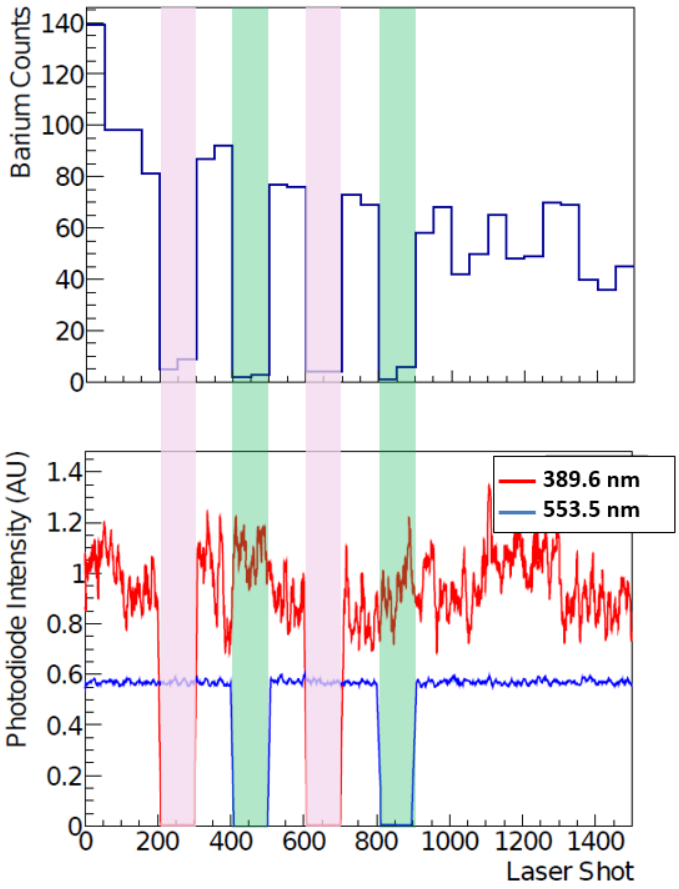

Figure 7.4: Count rate for the RIS Ba peak while alternately blocking one of the RIS lasers, as shown by the associated photodiode reading dropping to zero. Blocking either laser reduces the count rate to a negligible level (the background from nearby peaks), demonstrating the two-step nature of the ionization process.

chemical properties of the graphene, possibly making it more reactive. Additionally, the energy barrier to migration between adsorption sites is estimated to be exceedingly low for Ba, of order 0.05 eV [64], allowing for the possibility of deposited Ba to migrate to the substrate edges or beyond. This may be prevented by topographical features of the underlying substrate (which is not atomically-smooth) which could pin the Ba at local surface defects. There is also a risk of laser damage to the graphene. One experimental estimate of the damage threshold for graphene produced in a similar manner puts it at below 15 mJ/cm^2 for an 800 nm laser of 50 fs duration [95].

Keeping the peak fluence below this corresponds to pulse energies below 0.6 mJ/pulse for the desorption spot sized used in this work, though the damage threshold may be higher at the 1064 nm wavelength and 7 ns duration of the desorption laser.

Initial tests with a graphene on Cu substrate showed a clear TOF spectrum at a desorption power of 0.3 mJ/pulse, considerably lower than required on other substrates for similar TOF count rates. Attempts at Ba deposition and recovery showed no signs of significant TOF change, including a lack of the 157 amu/e peak that often came as a result of source exposure on earlier substrates. There was a small RIS Ba signal from isolated locations on the substrate seen when increasing the desorption power to 0.6 mJ/pulse. A further test of Ba deposition and desorption at this power did not show an increase in this signal, however, suggesting that it most likely came from background Ba. At 0.7 mJ/pulse, peaks with masses that could fit C^+ and Cu^+ began to appear in the TOF spectra coming from the edge of the substrate, suggesting possible surface damage from the laser removing C atoms from the graphene layer and perhaps even Cu from the underlying substrate. At 1.0 mJ/pulse, a significant RIS Ba peak was observed at locations across the full substrate surface (Fig. 7.5), but this peak did not increase after exposure to the Ba source.

Results from graphene on Ni were qualitatively very similar, though slightly higher desorption power was required for comparable count rates. A comparison of TOF spectra for 1 mJ/pulse desorption power for Cu and Ni with and without graphene is shown in Fig. 7.6. Despite the possibility of laser damage to the graphene at this power, it is clear that the graphene substrates show significantly higher count rates. This is interpreted as evidence for lower desorption energies from the graphene substrates in general, though the surface may also be more effectively heated for a given laser power. Another alternative explanation is the graphene substrates could possess more surface contaminants to be desorbed. The slightly lower desorption thresholds from graphene on Cu compared to graphene on Ni suggest that the underlying substrate material plays a role, though the differences are much less than between bare Cu and Ni.

The presence of background Ba on these graphene substrates is not well-understood. Initial contamination on the surface of the graphene may be part of the explanation,

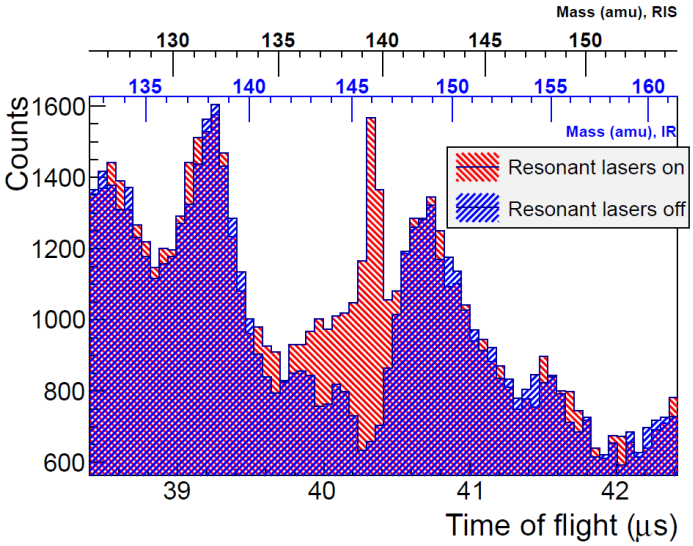

Figure 7.5: TOF spectrum for graphene on Cu at 1 mJ/pulse desorption power. The mass calibration for the RIS scale is slightly high due to ionization at lower potential than directly-ionized peaks, as detailed in §6.7.1.

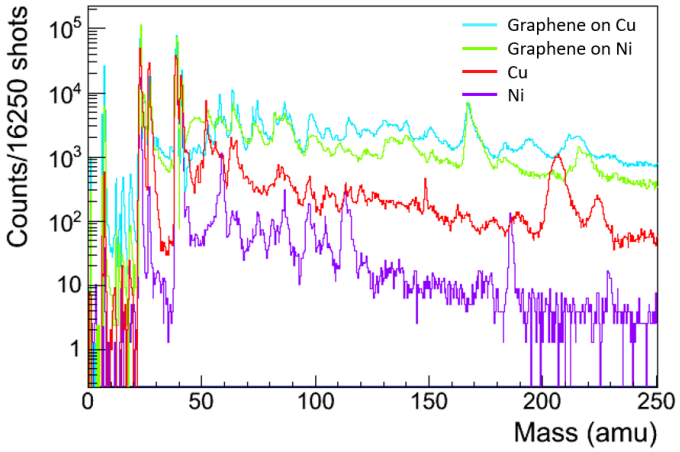

Figure 7.6: TOF spectra in calibrated mass for materials with and without graphene, during their first exposure to IR at 1 mJ/pulse. Data are summed across many locations on the substrate and counts are normalized to 16250 laser shots.

and the signal does decay with time as expected for a source of Ba that may be depleted over time. There may also be Ba in the first few atomic layers of the underlying substrate which diffuses up through the graphene or through areas of the graphene

which were damaged by the desorption laser. If graphene is used in the future, further investigation of this may be warranted.

As an additional check, one of the graphene on Ni substrates was exposed to very high laser power (30 mJ/pulse), expected to completely remove the graphene layer. Subsequent high-power desorption scans (9 mJ/pulse) showed a TOF spectrum with almost exclusively alkalis and Ni clusters. After exposure to the Ba source, the TOF showed an increase in the 157 amu/e peak and a peak around 74 amu/e (both seen on prior substrates), as shown in Fig. 7.7. There was also an increase around 143 amu/e which could be ^{144}Sm which is the α decay daughter of ^{148}Gd. All peaks which increased after deposition decayed away after a few hundred laser shots, while the others (presumably from the bulk of the Ni substrate) did not. The recovery efficiency from the 157 amu/e peak, assuming it is BaF^+ from the source, was 0.18%, similar to previous tests. No increase was observed in directly-ionized or resonantly-ionized Ba.

These results suggest the chemical properties of graphene prevent recovery of the 157 amu/e peak - perhaps it is not BaF^+ but some other molecule of Ba which does not form on graphene. It seems less likely that BaF^+ was in fact deposited but that its desorption threshold was not reached before removing the graphene, as both directly-ionized and resonantly-ionized Ba were seen reliably on the substrate at 1 mJ/pulse and other substrates have shown similar desorption thresholds for BaF^+ and directly-ionized Ba. The lack of directly-ionized Ba or RIS Ba after such high power cleaning could be due to removal of surface Ba contamination or damage to the substrate which inhibits desorption for bulk contaminants.

Another graphene on Ni substrate was tested with desorption power ramping up to 6 mJ/pulse, and showed a significant background RIS Ba signal with very little decay over time. It seems likely that there was significant damage to the graphene at such high power, possibly removing it entirely. However, the background RIS Ba signal from a pure Ni substrate with similar desorption power was much weaker and decayed fairly quickly, so either there is some graphene left which improves the RIS Ba signal or else this particular Ni substrate happens to have significantly more background Ba for unknown reasons. Regardless of the source of the background Ba,

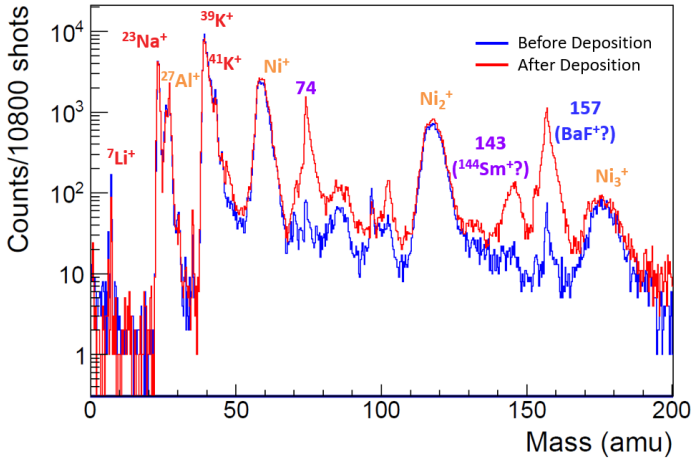

Figure 7.7: TOF spectra in calibrated mass before and after exposure to the Ba source. The substrate is graphene on Ni and the IR power is 9 mJ/pulse. Data are summed across many locations on the substrate and counts are normalized to 10800 laser shots.

having a steady RIS Ba signal allowed for a series of tests to verify and optimize the LITD and RIS process, at least for background Ba under these specific conditions.

7.2 RIS Characterization

7.2.1 Wavelength Detuning and Saturation

Using TOF data from a single desorption location of the graphene on Ni substrate described in the previous section, the dependence of the RIS Ba count rate on each of the RIS wavelengths was checked for consistency with previous measurements. For low powers, where power broadening is not significant (Fig. 7.8, left), the central wavelengths and approximate linewidths listed in §5.2 match well with the measurements from this system. Closer to the typical operating powers (Fig. 7.8, right), there is significant power broadening of the linewidths, especially for the UV transition, though the central wavelengths remain unchanged.

Saturation curve measurements were also performed to better understand possible RIS efficiency losses due to insufficient power. As described in §5.2.3, the green

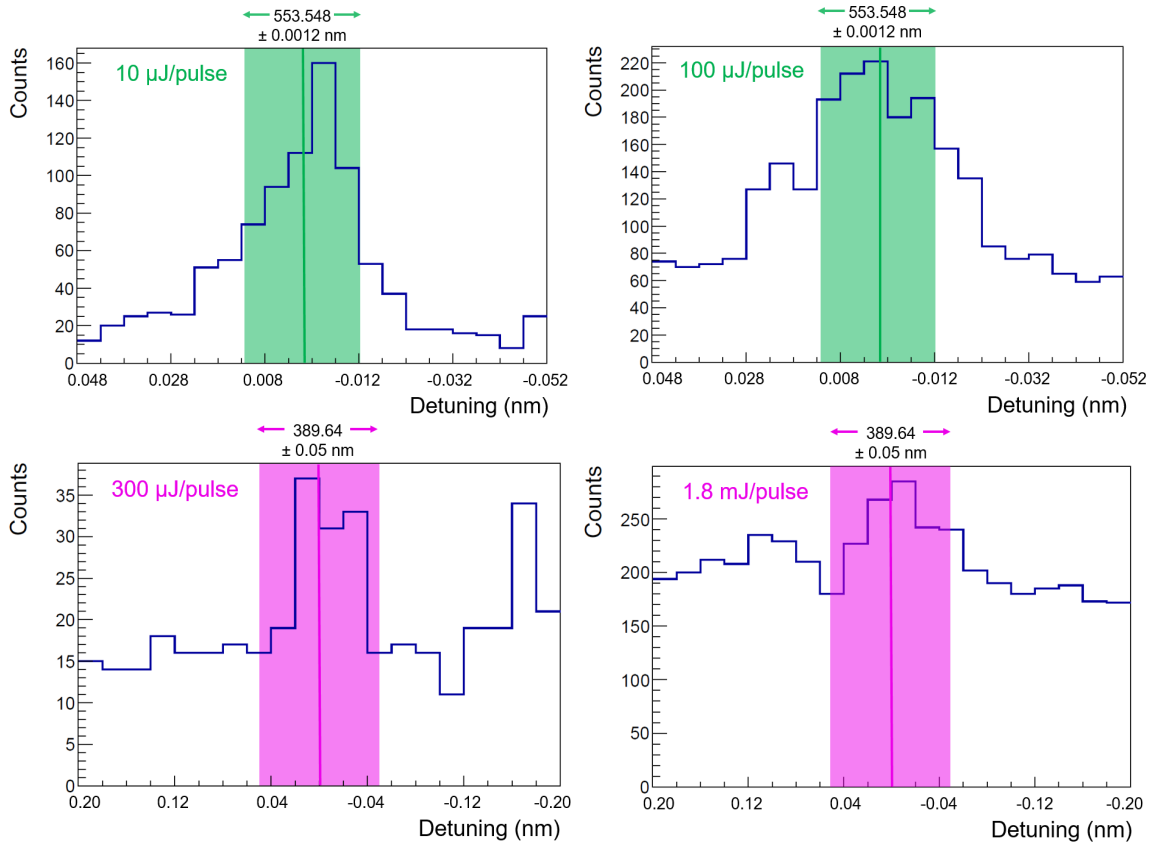

Figure 7.8: RIS Ba counts versus laser wavelength for data from a graphene on nickel substrate at high desorption power (6 mJ/pulse) and a 1.2 μs delay between the desorption and RIS lasers. Colored bands correspond to the central wavelengths and widths from prior measurements far from saturation of the wavelength being detuned (see Fig. 5.6). Fluctuations in the desorption rate for neutral Ba, as judged by the directly-ionized Ba count rate, are responsible for some of the features other than the peak at the central wavelength.

transition's high absorption cross section makes it readily saturated across the entire interaction region even without careful alignment. The RIS Ba count rate was verified to be effectively independent of green power above 1 μJ/pulse for low UV power (under 100 μJ/pulse) and above 15 μJ/pulse for high UV power (2 mJ/pulse, not far from the maximum dye laser output power). The dependence on UV power is expected, as higher UV powers saturate a wider region of the desorbed Ba plume, requiring higher green powers to saturate the same area for the initial transition. This makes typical

operation at a green power of 200 μJ/pulse more than sufficient.

Saturation measurements for the UV transition are shown in Fig. 7.9. Recall that the RIS beams are shaped by a razor to avoid background counts in the TOF from the beams impinging on the substrate directly and producing ions (§6.5.2). Although the RIS count rate will depend on the overlap of the UV beam with the desorbed Ba plume as described in §5.2.3, a simplified model using an effective beam size and assuming a uniform spatial density of Ba atoms throughout the interaction region can be used to roughly estimate the RIS efficiency. This is done by fitting the RIS Ba count rate to the function $r_0\left(1-e^{-E/E_{\text{sat}}}\right)$, where r_0 is an overall scaling factor, E is the total UV energy, and E_{sat} is the average saturation energy as defined in §5.2.2. The fit finds $E_{\text{sat}} = 0.7\,\text{mJ/pulse}$, which corresponds to an ionization efficiency at the typical 2.0 mJ/pulse UV power of 95%. It is expected that the count rate will continue to rise with UV power faster than the simple exponential form used here due to increased ionization from the tails of the beam, but this establishes that the RIS step is unlikely to be the primary loss channel for deposited Ba, and that the typical operating parameters are likely compatible with the ultimate goal of near 100% Ba detection efficiency.

The saturation curve here can also be compared to that from prior work as described in §5.2.4. The functional form of the fit extracting E_{sat} there is slightly different (see [84]), and extracts a saturation energy of $E_{\text{sat}} = 0.4\,\text{mJ/pulse}$, smaller than seen here but still within a factor of two. It is unsurprising that the effective saturation energy is different for this system, where the Ba density profile is different and the magnetic sublevel populations may be different as well.

7.2.2 Timing Measurements

Another parameter which can significantly affect the RIS efficiency is the delay between the desorption and RIS lasers. Prior measurements from Si substrates have shown evidence for an optimal delay of roughly 1 μs at IR powers ranging from 3–5 mJ/pulse. Using the stronger RIS Ba signal from the graphene on nickel substrate, this was tested more thoroughly. The RIS Ba rate from a single desorption location

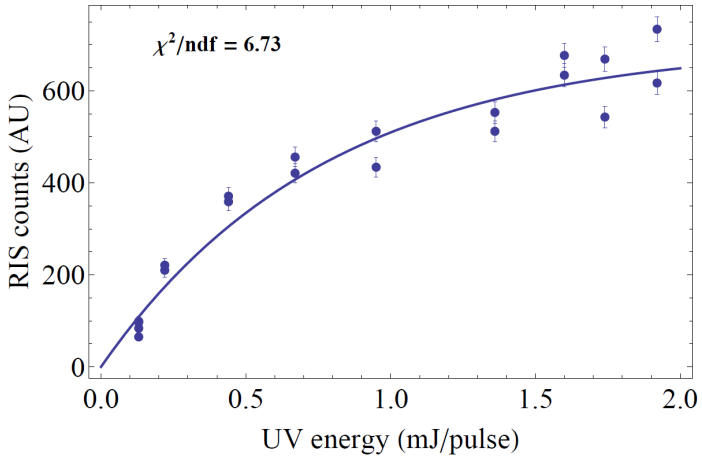

Figure 7.9: RIS Ba counts versus UV power for data from a graphene on nickel substrate at high desorption power (6 mJ/pulse) and a 1.2 μs delay between the desorption and RIS lasers. The solid curve is a fit of the form $r_0\left(1-e^{-E/E_{\text{sat}}}\right)$, with the best fit value of $E_{\text{sat}} = 0.7\,\text{mJ/pulse}$. The energy listed is the total beam energy, before being shaped by a razor blade which is positioned to transmit 60% of the beam energy, as described in §6.5.2. Data points are taken first while increasing the UV power then while decreasing it, to reduce systematic effects from a possible overall decrease in desorbed Ba versus laser exposure. The errors and χ^2 shown are Poissonian, which underestimates the true fluctuations in count rate.

was averaged over many laser shots at each of several different delays (Fig. 7.10). To factor out the effects of changing Ba desorption rate, the RIS Ba counts were divided by the counts in the directly-ionized Ba peak. This assumes the ratio of desorbed Ba^+ to neutral Ba to be constant. Because the desorption rate is high, the data show evidence of CEM charge depletion effects (see §6.7.2) which change depending on the RIS delay. To account for that, data were also corrected for the estimated relative CEM efficiency for each dataset and delay.

Both the CEM-corrected and uncorrected data agree with previous measurements showing the highest RIS efficiency for a delay of 1.0–1.2 μs. The effects of changing the delay between 0.5 μs and 2.0 μs are relatively modest, however. This is good evidence that the RIS lasers are saturating a fairly wide region in front of the substrate, so that the RIS efficiency is not very sensitive to changes in the time during which the plume expands from its initial desorption location.

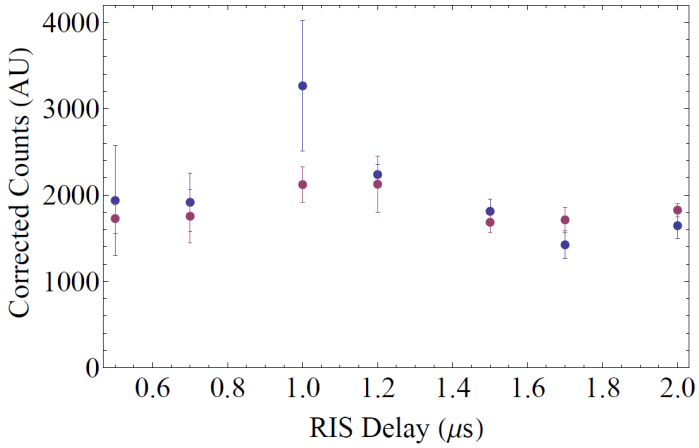

Figure 7.10: RIS Ba counts versus delay between the desorption and RIS lasers for data from a graphene on nickel substrate at high desorption power (5 mJ/pulse) and 700 μJ/pulse UV power, with the razor blade used to shape the RIS beams transmitting 60% of the UV energy. Data in blue are divided by the estimated rate of Ba desorption at each delay as judged by the counts in the directly-ionized Ba peak. Data in purple are corrected for the estimated relative CEM detection efficiency for RIS Ba at each delay, then scaled by the (CEM-corrected) counts in the directly-ionized Ba peak. Error bars are estimated by the standard error from multiple measurements at each delay.

7.2.3 Polarization Dependence

Because of the substantial difference in absorption cross section between UV light polarized perpendicular to the angular momentum of the excited Ba atom versus parallel to it, the RIS rate may in principle depend on the relative polarization between the green and UV lasers (§5.2.3). This was tested by varying the angle of the UV beamsplitter, while adjusting the UV half-wave plate to keep the UV power constant. The effects of this adjustment are shown in Fig. 7.11. These data suggest that the polarization has no effect on the RIS efficiency at the level of 10% or more. This could be due to collisions of the excited Ba atoms with the desorbed plume which randomize the orientation of the Ba, though the density of desorbed atoms is expected to be low enough that collisions during the roughly 10 ns RIS laser duration should be rare. A more plausible explanation is that the UV intensity is sufficient to saturate the transition regardless of the initial orientation of the excited Ba atom, in

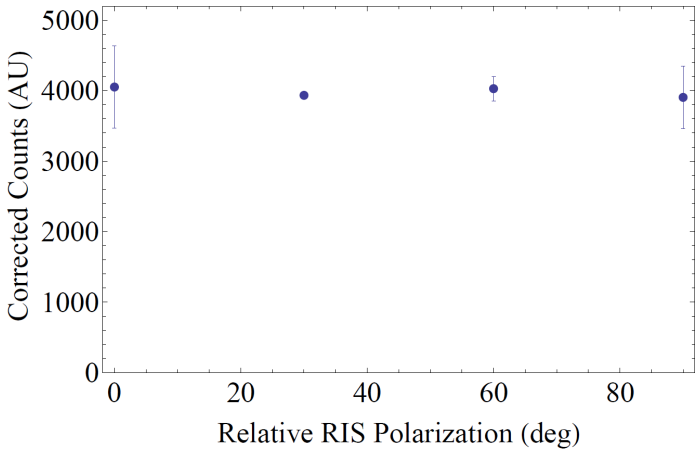

Figure 7.11: RIS Ba counts versus polarization angle between the green and UV lasers for data from a graphene on nickel substrate at high desorption power (6 mJ/pulse) and 600 μJ/pulse UV power, with the razor blade transmitting 60% of the UV beam energy. Data are divided by the estimated rate of Ba desorption at each delay as judged by the counts in the directly-ionized Ba peak. Error bars are estimated by the standard error from multiple measurements at each angle.

which case the RIS rate is not expected to show any polarization dependence.

7.2.4 Desorbed Plume Measurements

Although measurements such as those in §7.2.1 suggest that the RIS efficiency is typically quite high, this will depend on the overlap of the desorbed plume and the RIS beam profiles. The RIS Ba rate is typically not sensitive to adjustments to the UV vertical alignment of less than a few millimeters due to the UV beam's vertically-skewed aspect ratio, but has shown a clear dependence on the UV position normal to the substrate surface which relates to the extent of the desorbed plume at the RIS time. Typically the RIS rate is maximized when the UV is as close to the surface as possible (while avoiding background counts from directly hitting the substrate), motivating the alignment procedures described in §6.5.2. A better understanding of the optimal alignment and beam shape, as well as a complete model of the RIS efficiency, requires knowledge of the extent of the desorbed plume.

A fair amount of literature exists on the velocity profiles produced by laser desorption and ablation. Most models treat desorption as instantaneous [67], occurring at the time of the peak surface temperature. This is typically valid due to the strong non-linear dependence of desorption rate on temperature and the assumption that the time of flight when measuring the velocities is $\gtrsim 100\,\mathrm{ns}$. At low fluences (desorption rather than ablation), the velocity distribution function is commonly modeled as a forward Maxwell-Boltzmann distribution [96]:

$$f(v_x, v_y, v_z)\mathrm{d}^3v \propto e^{-m(v_x^2+v_y^2+v_z^2)/2k_BT}\ \mathrm{d}v_x\mathrm{d}v_y\mathrm{d}v_z,\ v_x \geq 0 \tag{7.1}$$

where m is the mass of the desorbate, v_x is normal to the substrate surface, v_y and v_z are in the plane of the substrate, and T is the desorption temperature. The desorption temperature may be equal to the maximum surface temperature for low laser fluences [70], though only adsorbates from the high-energy tail of the thermal distribution at the surface are desorbed, so this need not be the case [67]. For ablation, collisions modify this distribution. In this case, the desorbed plume is separated into three regions: the Knudsen layer directly in front of the substrate where collisions cause both a forward streaming velocity $\mathbf{u_{KL}}$ for the exponential factor ($\boldsymbol{v}^2 \rightarrow (\boldsymbol{v} - \boldsymbol{u_{KL}})^2$) and backward velocities ($v_x < 0$) are allowed; a region of adiabatic expansion where the forward streaming velocity increases and the temperature decreases; and a region of free expansion [97]. There is typically also a forward-focused angular dependence, commonly modeled to have the form $\cos^n(\theta)$ for an integer n which depends on the kinetics of the process.

For the laser fluences used in this work, formation of a Knudsen layer is not expected. However, experiments have demonstrated that an average of only a few collisions per desorbate is required for this formation [97], so it is not out of the question. A few additional caveats are warranted here. First, some data have shown evidence of two temperature components (fast and slow desorbates), so the simple single-temperature models above may be incomplete [67]. In addition, velocity distribution data is typically inferred by time of flight well after the free streaming region,

whereas the position distribution near the substrate surface is important for resonance ionization. Lastly, measurements of time of flight or local desorbate density tend to be fairly smooth and featureless – this means that even though they may be reasonably fit to a Maxwell-Boltzmann distribution, the extracted fit parameters may not be good measures of the underlying physical properties such as temperature. This is frequently seen in a degeneracy between desorption temperature and streaming velocity for fits to ablation data [67].

Fortunately, the most relevant measure of the desorbed plume for the RIS process is its extent normal to the surface. To gauge this, the RIS Ba rate was measured as a function of RIS distance from the substrate surface when horizontally-focused using a cylindrical lens to a $\pm 2\sigma$ width of $<150\,\mu$m for the UV beam. The green beam was kept wide to ensure saturation for a larger region than the UV might ionize from and the razor was removed to avoid diffraction effects. Data from this procedure, for four different RIS delay times, are presented in Fig. 7.12. There is evidence of a slight decay in count rate with laser exposure, so at each delay, the RIS position was moved away from the substrate and then back toward it to reduce the correlation between RIS position and laser exposure. The distance of closest approach, at the threshold of ion count production from the UV laser, is estimated to be 0.1 mm or less using the UV beam parameters and the estimated count production fluence of $50\,\mu$J/pulse (§6.5.2).

Because the UV horizontal spot size is small compared to the plume's spatial extent, the count rate is representative of the local desorbed Ba atom density. The data are fit to Maxwell-Boltzmann distributions integrated over the transverse axes (along the major axis of the UV profile, which is a few mm tall, and the laser propagation direction). The non-uniform UV vertical profile will lead to a slight reduction in overall normalization for longer delays, but this effect is small and degenerate with the drop in count rate over time. This model also assumes free streaming ($x = v_x t$) and instantaneous desorption at the desorption laser time. These fits give a range of temperatures from about 2800–4600 K and mild evidence of increasing temperature with time. They also show decreasing plume integrals at lower delays taken later in time, consistent with the decay in count rate from increased laser exposure.

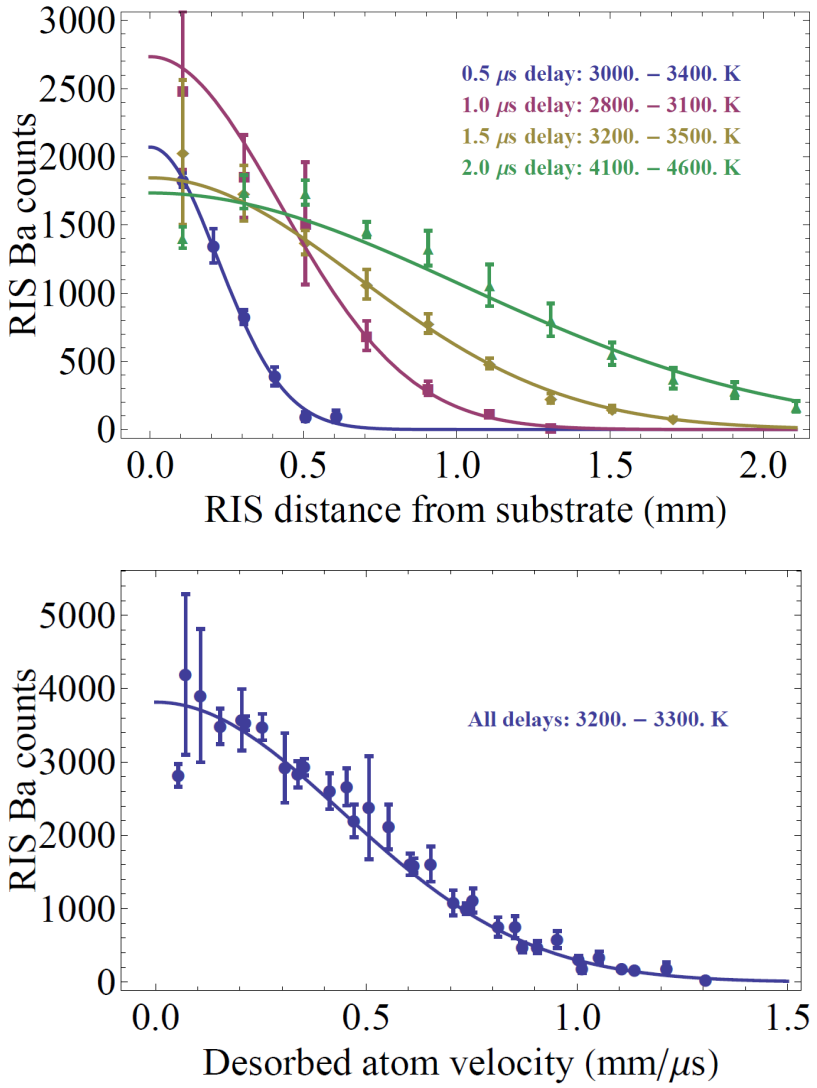

Figure 7.12: (Top) RIS Ba counts versus RIS beam distance from the substrate for data from a graphene on nickel substrate at high desorption power (6 mJ/pulse) and 1 mJ/pulse UV power, for different RIS delay times. The UV beam is horizontally-focused to a $\pm 2\sigma$ width of $<150\,\mu$m. Error bars are estimated by the standard error from multiple measurements at each position and delay. Data at each delay are fit to a Maxwell-Boltzmann distribution, and estimated temperature ranges are extracted from these fits. Data were taken in decreasing order of delay and the count rate decays slightly over time, causing a reduction in integrated counts for shorter delays. (Bottom) Same data, with counts and error bars for each delay scaled to have equal integrals based off the fits above and with velocity $v_x = x_{RIS}/t_{delay}$.

The apparent increase in temperature with delay may be due to degeneracy between desorption temperature and streaming velocity, if a Knudsen layer is present. More complex fits, such as adding a streaming velocity term, gave inconsistent streaming velocities across delays and an unphysical negative streaming velocity for the 1.5 μs delay data, so it seems likely models with additional parameters are overfitting the data. A single combined fit in desorption velocity ($v_x = x_{RIS}/t_{delay}$) is also given in Fig. 7.12, setting the integrals for each set of delay data equal for direct comparison. It shows good consistency across delays, suggesting that a single Maxwell-Boltzmann distribution near 3300 K with no collisions is a reasonable model for the desorbed plume. The extracted temperature estimates are not especially reliable (they on the order of 50% between delays and are sensitive to the minimum separation from the substrate surface, especially for small delays), but the shape of the plume at each delay, and the corresponding velocity distribution which tails off at around 1 mm/μs, are valuable inputs to estimates of the total RIS efficiency.

Additional data at slightly lower desorption power (5 mJ/pulse) were also taken using the same method. They showed larger fluctuations and are generally of lower quality, but also fit reasonably well to Maxwell-Boltzmann distributions and show some evidence of increasing temperature with delay. The combined fit at this lower power gives a slightly lower best fit temperature of 2800 K.

7.2.5 RIS Efficiency Estimates

Using the known extent of the desorbed plume, as well as the measured RIS beam profiles, the efficiency of resonant ionization can be fully modeled. The efficiency is calculated as the local ionization fraction times the local plume density, integrated over the full beam profile, as given in Eq. 5.9. This assumes that the green transition is fully saturated for the duration and spatial extent of the UV pulse, which is readily achievable. The desorbed plume profile is generated assuming a point-like desorption source and a Maxwell-Boltzmann distribution with a temperature given by the fits from the previous section. This corresponds to a uniform angular profile, which is conservative since the alternative is a forward-focused profile which, for a fixed

(measured) plume extent normal to the surface, means less transverse spreading and hence better overlap with the beam profile.

The results of numerical integration using the measured beam profile and this desorbed plume model are presented in (Fig. 7.13). For the parameters given there, a beam energy of 2 mJ/pulse and a saturation fluence of 10× the theoretical value (a better fit to prior empirical measurements) gives an estimated RIS efficiency of 83–84% for desorption temperatures from 2000–5000 K and for desorption from the center of the substrate. Desorption from near the edge at 2 mm vertically offset from the center of the substrate gives 49–54% ionization efficiency for 2000–5000 K. These numbers represent the most accurate estimates for typical operation.

The model above assumes the UV beam is aligned as close to the substrate as possible without creating background counts: the fluence of the measured UV profile at the surface of the substrate (at its maximum along the vertical axis) is equal to the conservative count production threshold from §6.5.2, $F_{ion,UV} = 50\,\mu\text{J/mm}^2$. The desorbed plume and estimated ionization fraction profiles are shown overlaid in Fig. 7.14. As illustrated there, the efficiency loss comes primarily from the desorbed atoms closest to the substrate surface, where the rising edge of the UV profile has not yet reached the saturation threshold. This wide rising edge and periodic bumps further out are caused by diffraction from the razor blade. Increasing the RIS delay allows the plume to expand further, so that fewer desorbed atoms are in the rising edge region, but also increases the transverse spread of the atoms, leading to an overall decrease in efficiency. Scaling the UV profile down along just the horizontal axis can improve the ionization efficiency: reducing the profile by a factor of 0.7 ($\sigma_x = 0.14$ mm) and shifting the profile inward to account for the new count threshold location results in estimated efficiencies of 96% for desorption from the center and 64% for desorption from 2 mm vertically offset.

To validate this estimate, the RIS beams were passed through a cylindrical lens at different locations to adjust the horizontal extent of the beams. The optimal signal was indeed seen at roughly 0.7× scale, with a signal increase of about 10% over the standard beam size. Further improvement in efficiency (particularly for desorption from the edges of the substrate) will require a steeper rising edge than the razor

blade can provide. Initial tests attempted to use a one-to-one telescope to image the razor blade at the substrate surface (and undo diffraction effects). The focused beams caused damage to the RIS mirrors when the telescope was placed in its optimal location (imaging the razor blade at the surface), precluding RIS measurements. Placement of the telescope elsewhere validated its ability to reduce the extent of the rising edge at the image location substantially, but did not lead to a measured improvement in the RIS rate due to its poor positioning. Future work using high-power mirrors to allow optimal telescope placement may be able to boost the estimated efficiency to near 100% for the full desorption area.

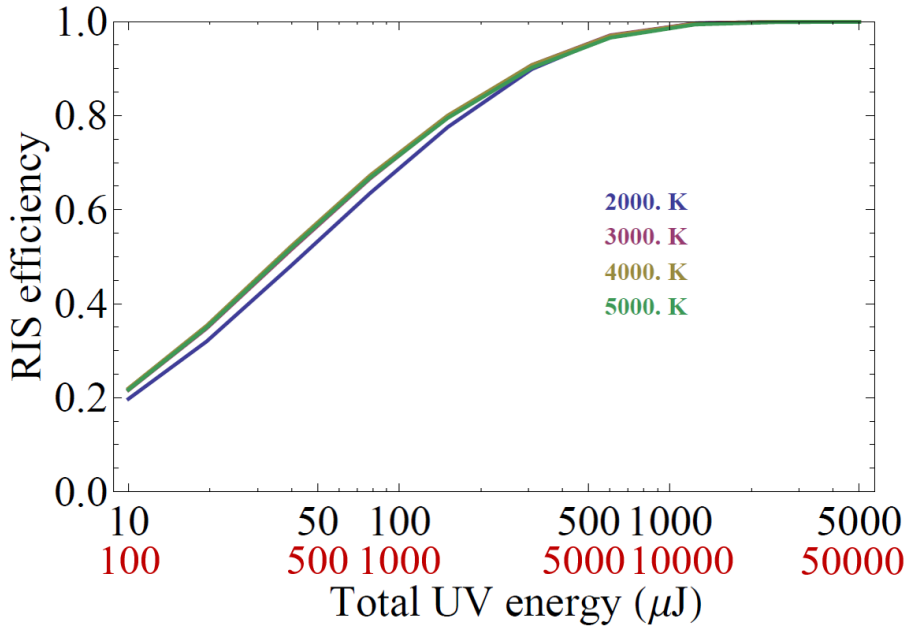

Figure 7.13: Estimated total ionization efficiency versus UV beam energy for Ba desorbed from a single point at the center of the substrate at different desorption temperatures. The measured UV beam profile with the razor transmitting 75% of the beam energy is used. A 1 μs RIS delay and an estimated offset of 550 μm between the peak beam power and the substrate surface (at the count production threshold) are used. The theoretical saturation fluence of 9.3 μJ/mm^2 is assumed for the energy scale in black, while the red energy scale uses 10$\times$ the theoretical saturation fluence, which is more in line with prior empirical measurements. This model depends only on the ratio of beam energy to saturation fluence. Similarly, the temperature affects the desorption length scale as the square, so quadrupling the temperature is equivalent to doubling the desorbed plume size (or the RIS delay).

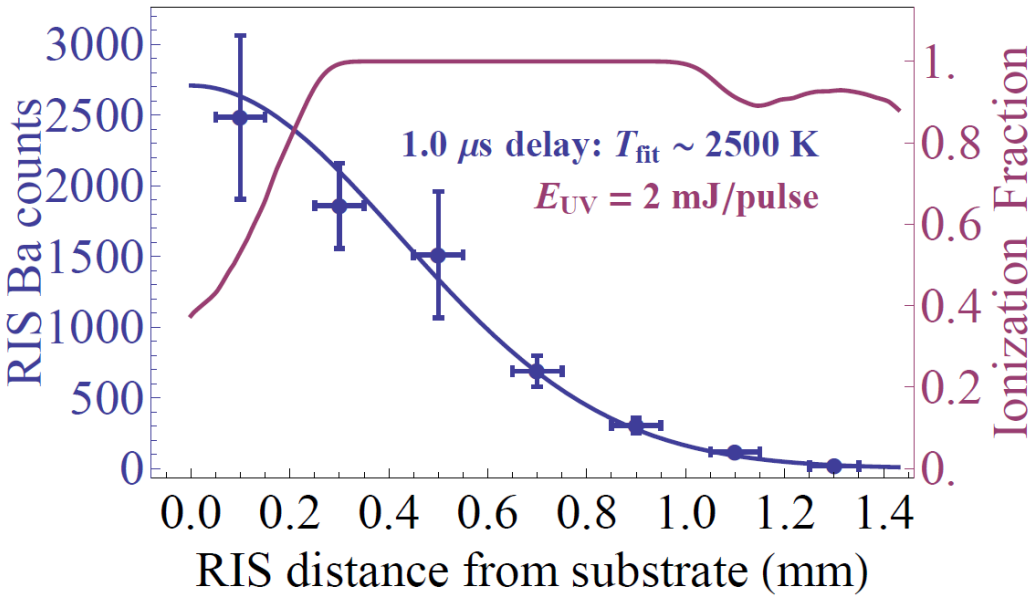

Figure 7.14: Desorbed plume data at 1 μs delay, as from Fig. 7.12, with the local ionization fraction (averaged along the vertical direction) overlaid. The same beam parameters and 550 μm offset as used in Fig. 7.13 are used here. The ionization fraction assumes a saturation fluence of 10$\times$ the theoretical value, which is more in line with prior empirical measurements. The measured beam profile is used, with its gradual rising edge and the dip near 1.2 mm caused by diffraction from the razor.

Chapter 8

Background Characterization and Reduction

8.1 Barium Contamination Measurements

Complete characterization of this technique requires an understanding of the background counts in the TOF mass spectra, most notably "background Ba" – Ba that

is not deliberately deposited from the Ba source. Unlike other ions of mass near that of Ba, background Ba is an issue even in a system that identifies ions without the use of mass, such as fluorescence in an ion trap. In general, it is beneficial to reduce background counts as much as possible. However, as the TOF spectra are not representative of the surface adsorbates present, a reduction in background counts is not necessarily an indication of cleaning (removal of contaminants). Such reductions may instead be due to changing the chemical or topographical features of the surface, which could affect how readily desorbates are removed as ions or their absolute desorption rate. These changing surface features may be unfavorable for a Ba signal, making it difficult to judge the optimal surface preparation without a reliable Ba deposition signal. Nevertheless, attempts at measuring and removing backgrounds have given a better indication of their source and what methods may be effective at cleaning, regardless of impact on a potential Ba deposition signal.

8.1.1 Substrate Heating

Resistive heating of surfaces, *i.e.* passing electrical current through them, has been shown to effectively remove surface contaminants on Si and SiC substrates through thermal desorption [98]. This was tested as a cleaning technique for Si and SiC substrates, as described in more detail in [86]. Rather than reducing background counts, these tests showed increases in many TOF peaks, including Ba and, even more so, RIS Ba (Fig. 8.1). Such increases occurred with or without prior direct exposure to the Ba source. A separate method of heating metal substrates, such as W and Ta, by electron impact at keV of energy was also pursued, and also resulted in Ba and RIS Ba increases.

This effect may be due to diffusion of impurities from the bulk to the surface (as happens to dopants during heating [99, 100]); from changing chemical properties of the surface (*e.g.* removal of oxides and breakup of BaO); or from transferring Ba from nearby components which also heat up. These increases persisted even after exposing the substrates to air at atmospheric pressure, which would likely cause any removed oxides to reform.

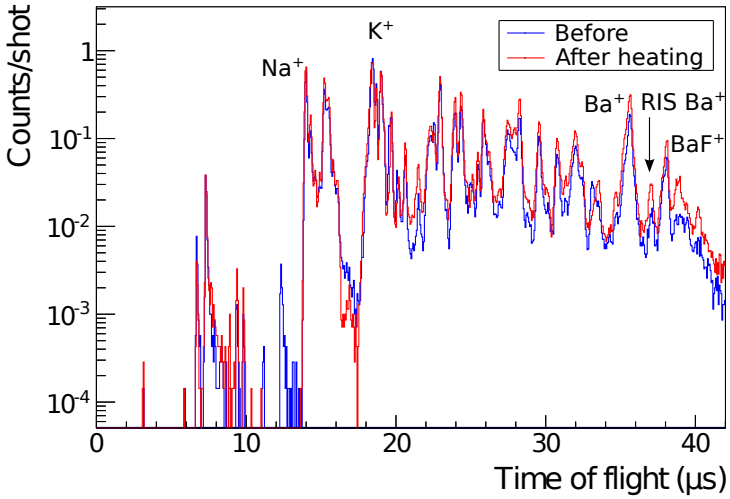

Figure 8.1: TOF spectrum before and after resistive heating of a Si substrate. The cleaning procedure involved heating for 1 hour at 700 C and then increasing to 1050 C for 30 seconds to remove oxides. After heating there is slightly more Ba^+, BaF^+, and RIS Ba^+. From [86].

8.1.2 Environment and Tools

If background Ba such as that seen from substrate heating were coming from the Ba source, it could be due to Ba accumulation in the vacuum chamber environment. However, replacing the substrate mount and nearby electrodes showed no effect on the background Ba signal, before or after substrate heating. The windows were etched in HCl, which is expected to be the most effective acid at removing BaF_2 [101], and again no effect on the background Ba was seen. In case the tools used to cleave and mount the Si substrates were somehow contaminated by the Ba source, a substrate was prepared using entirely new tools in a separate lab, to no effect.

No correlation was seen between time since substrate insertion in the vacuum chamber and the background Ba rate, further suggesting that the Ba is not from the chamber itself. Substrates resistively heated in a separate stand-alone vacuum chamber (never exposed to the Ba source) and then transferred to the system showed similar levels of Ba and RIS Ba. These trials all suggest that the background Ba is present independently of the environment or handling tools, likely from some combination of

initial surface contamination and intrinsic contamination of the bulk material.

8.1.3 Secondary Ion Mass Spectrometry

Quantitative measurements of the Ba contamination of different substrates require an independent, calibrated device. Secondary ion mass spectrometry (SIMS) removes and ionizes surface atoms via sputtering with high-energy incident ions followed by separation by mass. Incident ions possess much higher energy than that of a typical bond, making this process much less sensitive to the chemical state of atoms at the surface than thermal desorption.

This technique was applied[1] to Si substrates with various preparations and used to measure concentrations of contaminants as a function of depth (Fig. 8.2) as well as surface position (Fig. 8.3). These measurements indicate that background Ba is present at the surface of the substrate even without use in the Ba tagging system, with somewhat more Ba available at the surface for resistively-heated substrates. This background Ba seems to decrease with sputter depth, suggesting a higher concentration at the very surface. Moreover, the SIMS images demonstrate the presence of Ba across the entire surface, with some micron-scale spots of higher concentration that may be attributed to "dust" particles. This supports the hypothesis that a substantial amount of the background Ba is concentrated at the surface and can be removed with appropriate cleaning techniques such as ion sputtering. In addition, there is weak evidence that heating increases the surface concentration of Ba, as might occur as a result of thermal diffusion of contaminants from the bulk to the surface.

Though a measurement of the absolute amount of Ba present on the surface requires knowing the ion yield for Ba on Si, a lower limit can be set by supposing that all surface Ba atoms are detected by the SIMS measurement. Using the SIMS images for Ba and the known surface density of Si atoms, a surface concentration of >100 atomic ppb of Ba is measured. This assumes all detected Ba comes from the surface monolayer; this lower limit is reduced if Ba is removed from deeper within the substrate, though the ion beam is scanned sufficiently rapidly so as not to remove more

[1]Part of this work was performed at the Stanford Nano Shared Facilities (SNSF), supported by the National Science Foundation under award ECCS-1542152.

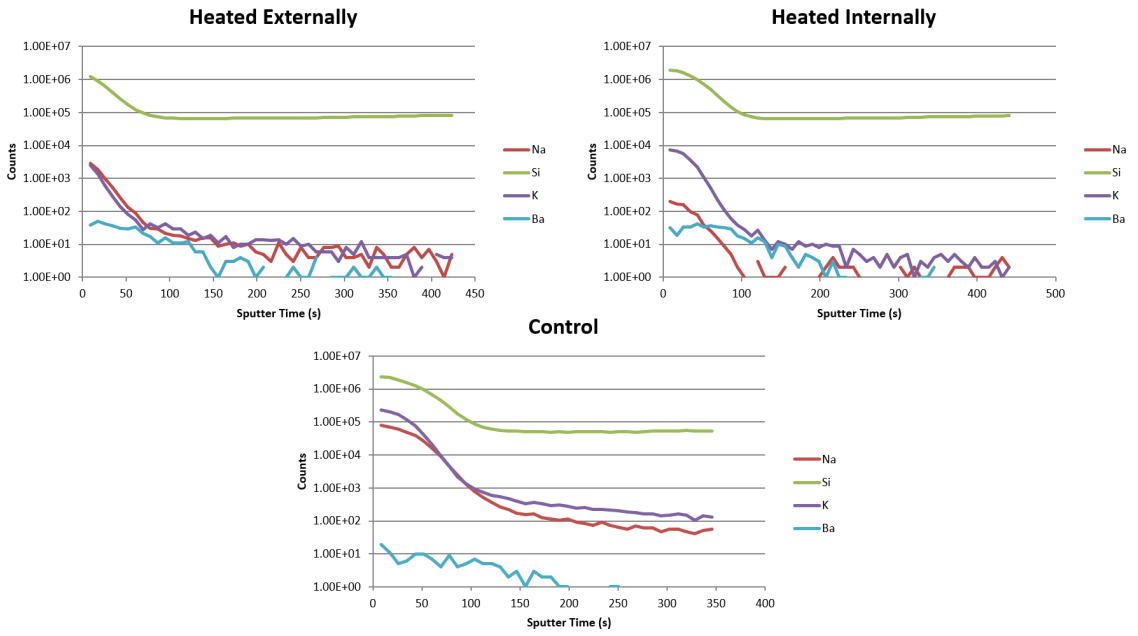

Figure 8.2: Depth profiles of several atomic species using secondary ion mass spectrometry on three Si substrates. One was resistively heated inside of the system ("heated internally"), another in a separate vacuum chamber ("heated externally"), and a third was measured as received from the manufacturer ("control"). The approximate sputter rate is 8 nm every 100 s. Initial decreases in signal are related to buildup of the sputter ions in the surface which reduces the ion yield – only when the underlying Si signal stabilizes are changes in the count rates indicative of concentration.

than a couple of monolayers of material in any given location. Measurements of a Ta substrate were also performed, with roughly 100× more surface Ba observed and a similar reduction in the Ba signal with sputter time; however, these results are less reliable due to difficulties obtaining an appropriate Ba calibration.

8.1.4 Inductively-Coupled Plasma Mass Spectrometry

The Ba signal from SIMS measurements above was just above the detection threshold of the device, despite being relatively easily measured by desorption and RIS. A more sensitive but still calibrated measurement technique, inductively-coupled plasma mass spectrometry (ICPMS), allows for measurement of trace contaminants in the bulk of

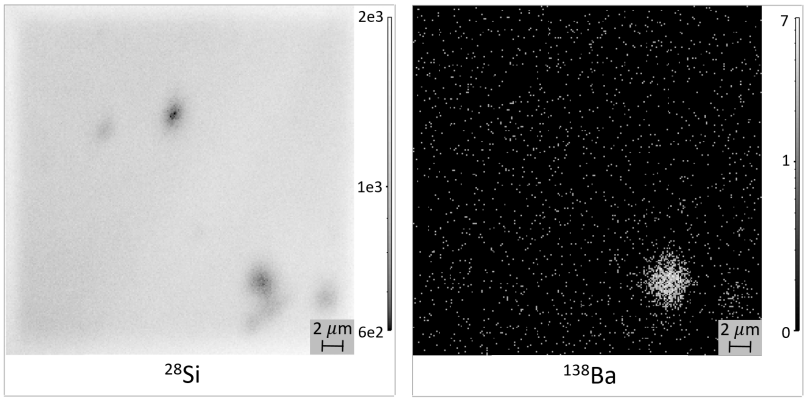

Figure 8.3: Secondary ion mass spectrometry images of Si and Ba from a Si substrate. Low concentrations of Ba atoms are present across the surface, with occasional micron-scale spots of higher concentration. These are interpreted as "dust" particles which have higher Ba concentration than the underlying Si substrate.

a material. A substrate is first dissolved in acid, then a thin stream is sprayed into a chamber where it forms a plasma as a result of inductive heating. This plasma is then separated by mass and ion count rates for a given mass window are measured. By comparing to acid standards with a known concentration of a given atom, a calibrated concentration in the analyte is achieved.

This method was performed on the Si substrates measured with SIMS above in Fig. 8.2 as well as a few others. Measurements of Ba concentration, averaged over the full substrate mass, are given in Table 8.1. They indicate that all Si substrates investigated posses Ba contamination at the level of a few ppb, independent of the particular Si wafer or manufacturer. However, those substrates which were heated (inside or outside of the Ba tagging system) may have shown systematically lower Ba concentrations by roughly a factor of two. This provides conclusive evidence that resistive heating does not result in the accumulation of Ba on the substrate. Instead, it suggests that heating may cause diffusion of Ba to the surface, both increasing the amount of Ba available for IR desorption and reducing somewhat the total Ba concentration by removing some Ba. It is also noteworthy that the bulk concentrations estimated by ICPMS are a factor of 10 or more lower than the lower limit on the surface concentrations estimated by SIMS. This provides further evidence that

Substrate	Ba conc. (ng/g)
SIMS control	22.0 ± 1
Si control 1	17.5 ± 1
Si control 2	9.0 ± 0.5
SIMS heated ext.	11 ± 0.8
SIMS heated int.	8.6 ± 0.6
Si used 1	5.5 ± 0.5
Si used 2	9.8 ± 0.7

Table 8.1: Bulk Ba concentration measurements of Si substrates using ICPMS, in ng of Ba per g of Si. This includes all natural Ba isotopes. Concentration by atom is given by multiplying by the ratio of atomic weights, roughly $28/138 \approx 0.2$. Control substrates were never exposed to the Ba tagging system or heated, while "used" substrates were heated and exposed to the desorption laser inside the system.

sputtering away surface contaminations could substantially reduce background Ba.

8.2 Substrate Cleaning Methods

The primary goal of the cleaning techniques investigated here is to reduce background Ba while avoiding modifications to the surface that may hinder desorption of Ba source deposits. To improve repeatability between substrates, it is also desirable for cleaning techniques to produce a well-defined chemical state of the surface. Considerable effort in the field of surface science has been devoted to development of techniques to prepare atomically-clean surfaces, here defined as those with surface contamination below the level of a few percent of a monolayer. It is believed that an atomically-clean surface is favorable for Ba tagging, as this should reduce the likelihood of Ba molecule formation and such a simple surface chemistry is easier to understand. However, in principle contaminants may reduce the desorption energy for Ba at the surface, so other surface preparations may be considered in the future.

Techniques for atomic cleaning vary between substrate materials and include pulsed laser cleaning [102], thermal annealing, and ion sputtering [98]. Due to the issues of increased background Ba from resistive heating described in §8.1.1, thermal annealing is not considered further here despite its utility in developing clean surface

reconstructions. Common to these techniques is the requirement of low pressure inside the vacuum chamber where the cleaning is performed, typically below 1×10^{-9} mbar. At a pressure of 1×10^{-9} mbar, a monolayer of contamination takes roughly 20 min to form, assuming that all incident atoms stick to the surface. Unfortunately, such low pressures are not currently achievable in the system considered here due to its large surface area and components which cannot be baked above 200 °C. A truly atomically-clean surface is therefore unlikely to last for the duration of a data-taking run (30 min), but atomic cleaning methods may still improve the Ba tagging process. Reductions in the vacuum chamber pressure are described further in §8.2.4.

8.2.1 Laser Cleaning

Pulsed laser irradiation has been shown to produce an atomically-clean Si surface preparation under the right conditions through thermal desorption of surface contaminants and diffusion into the bulk depending on the contaminant [102]. Unlike resistive heating, laser irradiation is not expected to cause diffusion of bulk contaminants to the surface due to the shallow depth of heating and the thermal gradient which promotes diffusion deeper into the substrate. With sufficient power, laser cleaning also causes annealing of the surface which may produce ordered surface reconstructions.

At laser fluences of roughly 20 mJ/mm^2 for a 694 nm laser [103] and 8 mJ/mm^2 at 248 nm [104], atomically clean surfaces have been produced with ~ 10 laser shots. Visible surface damage was observed at slightly higher fluence, but the gap between the fluence needed for significant cleaning and the damage threshold may be widened by reducing the thickness of the surface oxide layer through chemical etching prior to use in the vacuum system [104].

In this work, all attempted laser cleaning is done with the 1064 nm desorption laser, which is likely to produce different results due to its higher absorption length in Si. Nevertheless, laser cleaning in this manner has been shown to be both the easiest and most-effective method for reducing background counts in the TOF spectra. An example of the effects of laser cleaning at higher power than used for data-taking scans is shown in Fig. 8.4. Reductions in background Ba of two orders of magnitude

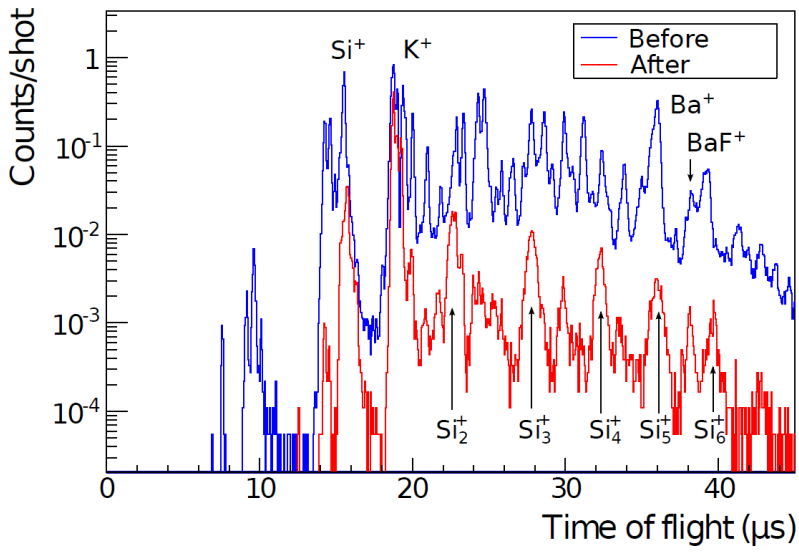

Figure 8.4: Laser cleaning of a Si substrate. Data were taken at 15 mJ/cm^2 and 200 laser shots per location, while the cleaning scan in between was at 23 mJ/cm^2 and 1200 laser shots per location. All peaks decrease in magnitude, though the presence of Si_n^+ clusters in the TOF suggests possible surface damage. From [86].

or more appear to be possible with this technique, though there is also a risk of irreversible damage to the substrate, as indicated by the appearance of Si clusters in the TOF spectra. Similar reductions in TOF count rate have been observed on every substrate material investigated.

As described in 7.1.1, each exposure of the substrate to higher laser power results in increased TOF counts at first followed by a rapid decay and an eventual plateau. This is likely due to some combination of removal of surface atoms and conditioning of the surface to reduce desorption rates. It is believed that some level of laser cleaning is appropriate to reduce backgrounds without harming potential Ba deposition signals excessively, though quantitatively finding the optimal laser cleaning power is challenging without a more reliable deposition signal.

8.2.2 Chemical Etching

Acid treatment of Si samples can be used to remove the surface oxide layer and, for some preparations, lead to a hydrogen-terminated surface which is less chemically-reactive. As noted in §8.2.1, this may also improve prospects for laser cleaning without surface damage. Several Si samples were treated using the chemical etching procedure recommended by the Stanford Nanofabrication Facility:

1) Initial dip in H_2O_2 and H_2SO_4 solution

2) 10 min in 5:1:1 (by volume) H_2O:H_2O_2:NH_4OH mixture at 50 °C

3) 30 s in 50:1 H_2O:HF mixture at room temperature

4) 10 min in 5:1:1 H_2O:H_2O_2:HCl mixture at 50 °C

This procedure is expected to remove organics and the surface oxide layer as well as metallic surface impurities such as Ba. Soon after chemical treatment, the substrate was transferred to the vacuum chamber to avoid chemical changes due to long-term exposure to atmospheric pressure.

While there is some mixed evidence that this procedure reduces background Ba in data taken at low desorption power, it does not prevent the appearance of background Ba after resistive heating, nor does it seem to reduce background Ba at more typical operating powers. This is consistent with the background Ba from resistive heating and medium- to high-power desorption being primarily from bulk contamination of the substrates, while background Ba from initial low-power desorption scans may be from Ba at the surface layers. If this is true, laser cleaning may already be effective at removing surface Ba, without the need for ion sputtering.

Despite this operating hypothesis, use of Si substrates from different manufacturers and production methods showed similar levels of background Ba at medium to high desorption power. It may be that Si wafer production and cleaning methods (including the nominaly-cleaner float-zone production method) do little to the intrinsic bulk Ba contamination, primarily removing more common and lighter contaminants such as C and O. This hypothesis also matches ICPMS measurements which show similar Ba contamination (at the few ppb level) across several Si substrates (§8.1.4).

8.2.3 Ion Sputtering

Initial tests of the Ar ion source were performed before bakeout of the vacuum system using copper substrates, as sputtering is known to be effective at cleaning metal surfaces. Ultrapure electroformed copper was provided by Pacific Northwest National Lab for use as substrate material. Measurements of the bulk Ba concentration of similar substrates using glow discharge mass spectrometry gave limits of < 0.1 ng/g to < 1 ng/g depending on the substrate. These limits are lower than the ICPMS measurements of Ba in earlier Si substrates, suggesting the possibility of reduced background Ba after sputtering away surface contaminants. Although the primary goal of these sputtering tests is removal of backgrounds, sputtering away oxide layers and other surface contaminants (which may or may not be present in the TOF spectra) may also improve repeatability across substrates and reduce formation of Ba molecules.

Before extensive laser cleaning, sputtering was observed to cause reductions in the count rate of the TOF spectra. However, the "baseline" TOF spectrum at a relatively low IR power of 1 mJ/pulse (where repeated data taking runs show almost identical spectra) is very similar whether or not the substrate has been cleaned via sputtering (Fig. 8.5). This suggests that in fact laser cleaning alone may be sufficient to remove most surface contaminants including Ba (as suggested in §8.2.2), with ion sputtering providing a relatively small further reduction in background counts.

Once the baseline TOF has been reached with laser cleaning or a mixture of sputtering and laser cleaning, further sputtering does not seem to permanently affect the TOF spectrum. Instead, relatively short sputter scans result in slight increases to the TOF count rate at all masses, followed by a decay to the previous count rate in around 100 laser shots at 1 mJ/pulse. Longer sputter scans may cause larger temporary increases, but are reversed by laser cleaning at an intermediate desorption power of 3 mJ/pulse (Fig. 8.6).

The effects of sputtering on the TOF spectra have been studied in further detail, including possible explanations for the temporary increase in count rates after sufficient prior laser cleaning. Two primary classes of changes to the spectra are observed. First, the TOF count rate increases for desorption locations across the

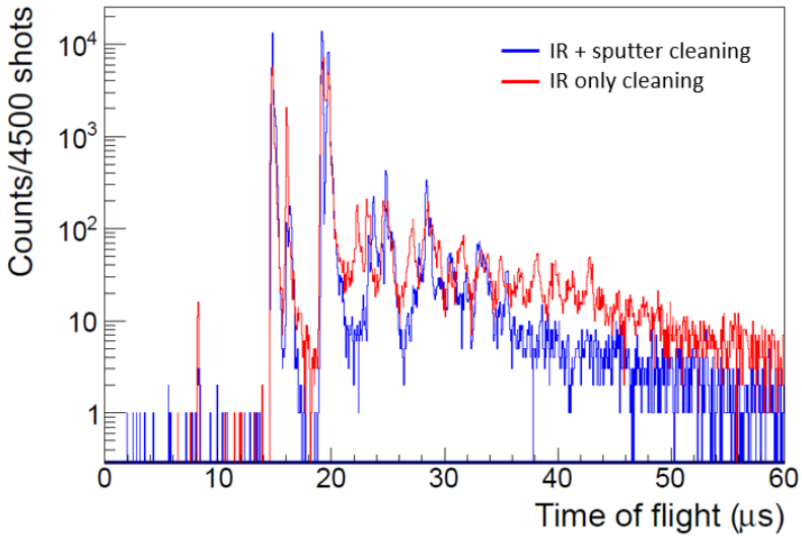

Figure 8.5: TOF spectra from two different Cu substrates, one cleaned by alternating sputtering and data taking/IR cleaning at 1 mJ/pulse and the other only laser cleaned at the same power. Sputtering does not significantly reduce the background count rate beyond what is achievable by IR cleaning alone.

substrate, including those which are not sputtered. For brief sputter scans, this effect is rapidly reversed by even low power laser cleaning. This is believed to be the result of redeposition of a small fraction of the sputtered atoms back onto the substrate surface, possibly up to a few mm away from the sputter location. These redeposited atoms may be adsorbed with a lower desorption energy than if they remained in the bulk, explaining their rapid decay in the TOF spectra. Extended sputtering of the conductive pad a few mm above a Cu substrate during initial positioning tests resulted in visible depositions on the top edge of substrate, lending plausibility to the redeposition hypothesis.

Second, longer sputter scans seem to more strongly affect the TOF count rate from the sputtered region, and require more extensive laser cleaning to reverse. This may be due to changing topographical features of the sputtered surface area. SEM images of a Cu substrate which was sputtered extensively in one region are shown in Fig. 8.7, indicating the presence of flat regions with steep walls which may promote desorption.

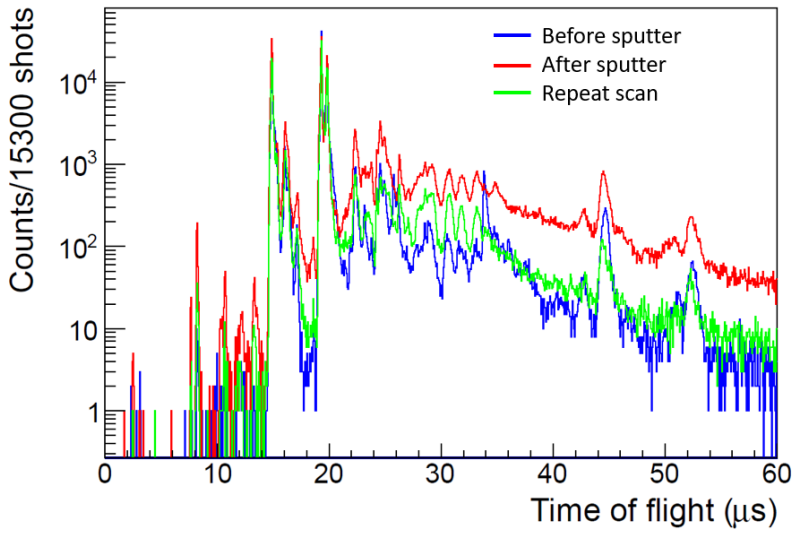

Figure 8.6: Transient effects of extended sputtering on the TOF spectrum. Data are taken at 3 mJ/pulse IR power, with a 1 hr sputter scan at 1 μA beam current and 5 keV beam energy in between. The expected sputter depth is on the order of 0.5 μm for the sputtered region. Though the TOF count rate does increase significantly after sputtering, the spectrum returns to quite close to the pre-sputtering spectrum after a second data-taking run.

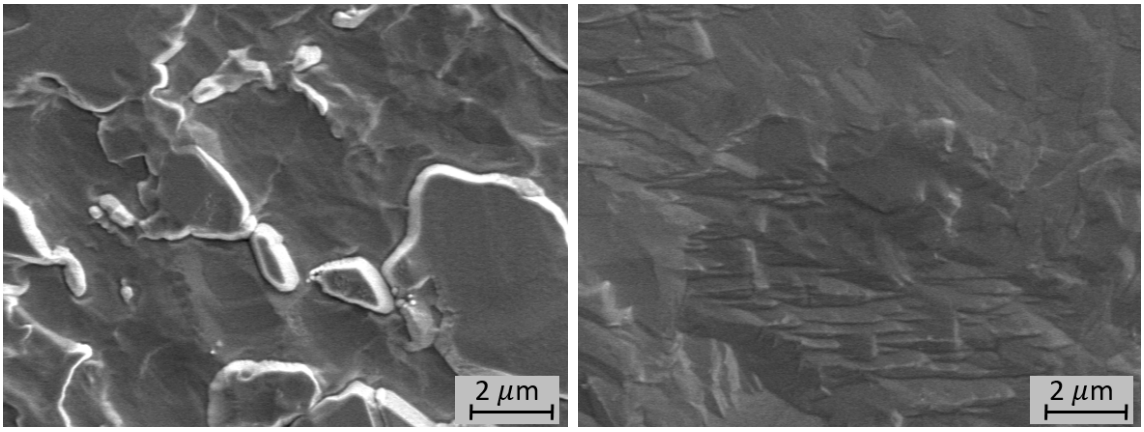

Figure 8.7: SEM images of a sputtered (left) and bare (right) Cu substrate. A 3 mm x 1.5 mm area of the substrate was sputtered for 3.5 hr at 5 keV beam energy and 1 μA beam current, corresponding to an estimated 2.5 μm average sputter depth. Sputtering seems to promote the formation of relatively flat plateaus with steep walls, which may increase desorption rates.

An alternative explanation for TOF count rate increases is implantation of contaminants from the Ar ion beam itself. However, sputtering typically increases the count rate of all mass peaks, and it seems unlikely that the contaminants in the Ar would match those on the substrate prior to sputtering.

As extended sputtering through as much as 1 μm of the substrate does not cause any changes to the TOF spectrum which cannot be reversed by IR cleaning, these cleaning techniques demonstrate removal of Ba from initial surface contamination. The remaining mass peaks in the TOF spectra must have as their source the bulk of the substrate and/or adsorbed molecules from the vacuum chamber pressure. Hence, further reductions in background rates are likely to be achieved primarily through cleaner bulk substrate materials or reduced vacuum chamber pressure - more extensive laser cleaning will reduce background rates, but at the likely cost of reduced desorption overall, including of depositions from the Ba source.

8.2.4 System Bakeout

To improve the vacuum chamber pressure, the full system was wrapped in heater tape and aluminum foil and baked to 150-200 °C for 11 days. During this period, all filaments (such as that in the hot cathode pressure gauge) were heated to degas them and the non-evaporable getter pump was conditioned at high temperature to chemically activate its getter material. Afterward, the pressure in the RIS chamber as measured by the hot cathode gauge was below 1×10^{-8} mbar. Measurements from a residual gas analyzer at the end of the TOF region were used to estimate the relative fractions of different residual gasses. After applying cathode gauge correction factors, the partial pressure of H_2O is estimated at 5×10^{-9} mbar, a reduction by more than a factor of 200, with the majority of the remaining pressure from H_2 and a small amount of N_2. This corresponds to a roughly 10 min monolayer formation time for H_2O assuming a sticking fraction of 1.

TOF spectra taken at 1 mJ/pulse IR power from Cu samples before and after bakeout of the system are shown in Fig. 8.8. The two spectra look very similar, suggesting that the substantial reduction in vacuum pressure had very little effect on

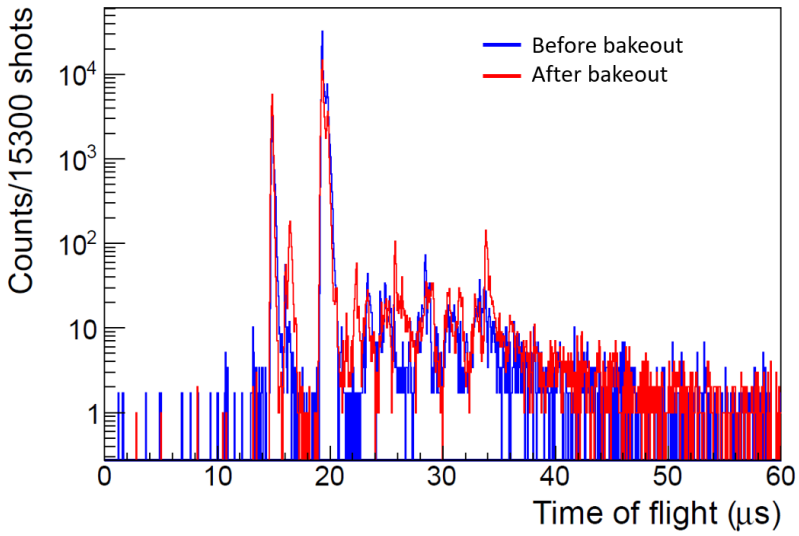

Figure 8.8: Effect of bakeout on TOF backgrounds for two Cu substrates. Both were laser cleaned at 2 mJ/pulse and then data were taken at 1 mJ/pulse IR power. The two spectra do not differ substantially, suggesting that TOF backgrounds are unaffected by the lower pressure.

background counts. It is possible that the relatively short monolayer formation time of 10 min means that such a pressure reduction has little effect on the equilibrium surface state. However, it seems most likely that the source of the TOF backgrounds is the bulk of the substrate itself.

After bakeout, Si substrates were also used both for background studies and attempts at Ba deposition and recovery. The effects of sputtering the Si substrates are qualitatively similar, in that improvements over laser cleaning alone are minimal and subsequent sputtering causes TOF count rates to temporarily increase. At low IR power (1 mJ/pulse), the Cu substrates show more background counts, whereas at higher power (3 mJ/pulse), the Si substrates have more counts (Fig. 8.9). There is not enough information to properly model which material is cleaner, as the contamination levels and desorption energies likely change with laser exposure/cleaning. In addition, even for a fixed contamination level, higher count rates could mean either a lower average desorption energy or a higher surface temperature (for a fixed laser

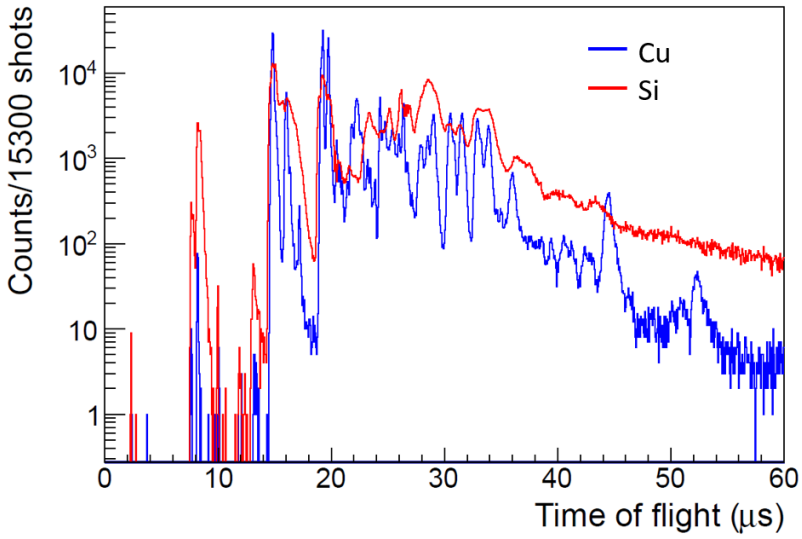

Figure 8.9: TOF backgrounds for Si and Cu substrates at 3 mJ/pulse IR power. These are "baseline" spectra after repeated datasets at the same power.

power). However, the simple Arrhenius model presented in Eq. 5.1 (with a single average desorption energy) suggests that the lower count rate at low power and higher count rate at higher power of Si corresponds to both a higher average desorption energy on Si and a higher total level of contamination. This is consistent with the GDMS and ICPMS measurements of Ba contamination in the two materials which indicated a lower bulk Ba concentration in the Cu, making it potentially better as a substrate material.

Attempts at Ba deposition and recovery after bakeout were performed for both substrate materials. Trials were performed at desorption powers ranging from 1-5 mJ/pulse with and without sputtering the sample immediately before exposure to the Ba source, for durations from 15 min to several days. However, these trials did not show convincing evidence of recovery of Ba from the source. It may be that the reduced vacuum pressure and different surface chemistry/topography from sputtering do not reduce Ba molecule formation or otherwise improve the Ba recovery process. However, these tests are inconclusive due to the degradation of the Gd-driven Ba source, as previously described in §7.1.2. Further attempts at Ba recovery using the ^{252}Cf fission fragment source in gas xenon are described in the following chapter.

Chapter 9

Tests in Gas Xenon and Future Work

9.1 Ba Deposition in Gas Xenon

Use of the ^{252}Cf fission source (6.4.2) for Ba deposition avoids the chemical degradation issues of the previous ^{148}Gd-driven source and also allows for tests in gas or liquid Xe, more closely matching the environment of the final application. For faster iteration, the tests here have been performed in gas Xe rather than liquid Xe. As shown in Fig. 6.4, ^{252}Cf fission fragments in 1 bar of gas Xe have an average range of just over 1 cm. The original source mount places it at 1 cm from the substrate during deposition, so it is expected that some fission fragments will implant in the substrate, while others are stopped in the gas and drifted to the negatively-biased substrate surface. Because the source changing procedure requires breaking vacuum, and modifications to the source mount to optimize the distance to the sample were expected, subsequent bakeout of the system was postponed. Trials presented here are

therefore done in an environment of higher chamber pressure.

Si substrates are used to test deposition. First, repeated data-taking scans at a fixed laser power are performed until a stable baseline TOF is achieved. Then the vacuum chamber is filled with 1 bar of Xe gas and the substrate is moved to the deposition location. During deposition, the substrate is biased to −200 V to attract positively-charged fission fragments. After deposition, all voltages are returned to ground and the Xe is cryopumped out of the system. Data is then taken at the same desorption power as the pre-deposition scans within an hour of deposition to minimize effects of fission fragment decay or slow changes to the substrate surface.

9.1.1 Ba Recovery at Low Desorption Power

Initial tests at 1 and 1.5 mJ/pulse desorption power showed signs of increases in several TOF peaks including Ba and RIS Ba (Fig. 9.1). These increases were stronger from some desorption locations than others, possibly due to differences in the local surface chemistry and topography or from non-uniform deposition on the substrate. Subsequent scans at the same desorption power showed greatly-reduced counts in many peaks, including Ba, which is consistent with the increases coming from loosely-bound surface deposits.

If the difference in counts in the Ba and RIS Ba peaks before and after deposition are from the fission source, a recovery efficiency can be estimated. The data at 1.5 mJ/pulse are used for this calculation, as the increase is larger for this higher power. Using the calculations of Fig. 6.4 extended to 23 hours, the estimated total production of ^{140}Ba is 2.5 million atoms and of ^{138}Ba is 2.3 million atoms. The efficiency of collection onto the substrate surface is unknown, but cannot be higher than 50%, corresponding to the fraction of fission fragments traveling into the forward hemisphere. That corresponds to no more than 2.4 million total Ba atoms collected. The Ba peak increased by roughly 5000 counts (1000 counts for the RIS Ba peak), corresponding to a minimum recovery efficiency of 0.2% (0.04%).

Though this is well below the desired efficiency, it may underestimate the true efficiency of recovery from the surface considerably if the initial collection is much

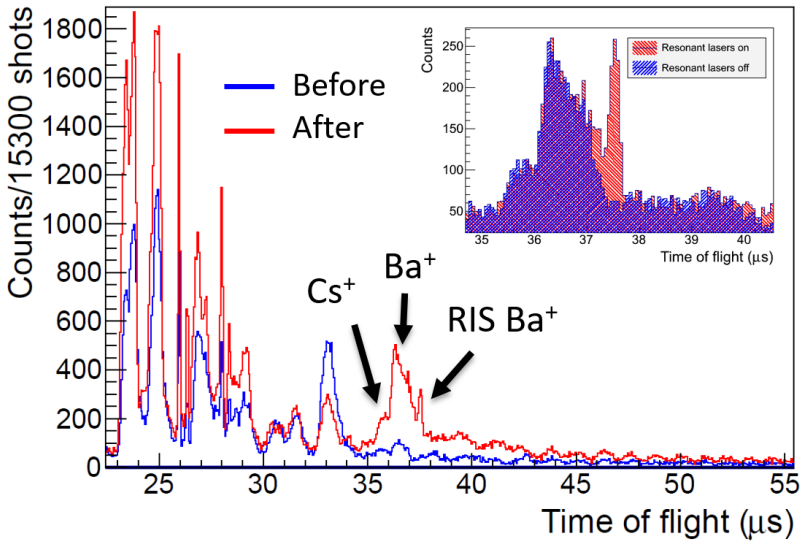

Figure 9.1: TOF before and after 23 hours of deposition from the ^{252}Cf source. Taken at 1.5 mJ/pulse desorption power. The inset shows the Ba TOF region, including the RIS Ba peak in red. Increases in Cs, Ba, and RIS Ba are believed to be a direct result of deposition. The decrease at 33 μs, corresponding to $\sim$ 114 amu/e, is not currently understood.

less than 50% efficient. Moreover, recovery at the optimal desorption power may be significantly more efficient, due to the non-linear dependence of the desorption process on surface temperature.

To ensure the increases were not due to some other part of the deposition process, such as exposure to the Xe environment or movement of the substrate, a repeat trial was done with the substrate 4 cm above the fission source - far enough removed to considerably reduce or eliminate deposition onto the substrate from the source. The TOF spectrum observed after this test was nearly identical to the baseline TOF immediately preceding it, implying that deposition from the source may be the cause of the prior TOF count increases.

However, attempts at duplicating the deposition signal on the same substrate afterward were not successful, including trials at higher desorption laser power. This means that the check for TOF count increases with the substrate 4 cm above the fission source was not conclusive. The lack of a deposition signal could be due to surface damage from repeated exposure to the desorption laser at 1.5 mJ/pulse or

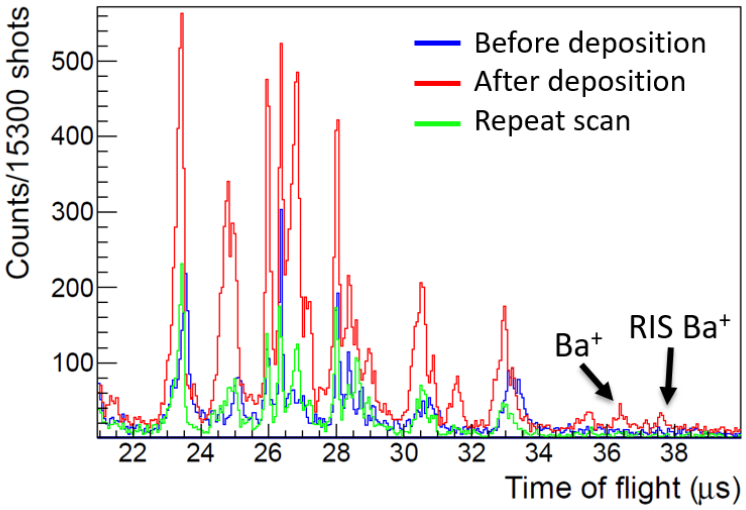

Figure 9.2: Duplication of deposition results from the ^{252}Cf source on a second Si substrate. Taken at 1 mJ/pulse desorption power. The TOF spectrum mostly reverts to its pre-deposition state after a second scan.

higher.

To test this, similar trials were performed using a new Si substrate. Data taken at 1 mJ/pulse showed TOF count increases which may be signs of deposition and recovery (Fig. 9.2), though the increases were somewhat smaller and were not observed at higher desorption power. These results are interpreted to indicate successful Ba deposition and recovery from a gas Xe environment, though the conditions for doing so reliably are not yet known. These conditions may depend not only on the substrate preparation, such as the history of laser exposure (the results of which may be irreversible), but also on differences between substrates of the same material.

Further trials were aimed at finding conditions for better Ba deposition and recovery. Sputtering directly before deposition did not result in an improved deposition signal. Because of the possibility of fission fragment implantation in the substrate surface, the fission source mount was adjusted to put it at 2 cm separation from the substrate at the deposition location, as well as to allow the source itself to be biased to better shape the electric field for collection. Neither the increased distance nor the source bias showed evidence of enhanced deposition, though the unreliable nature of

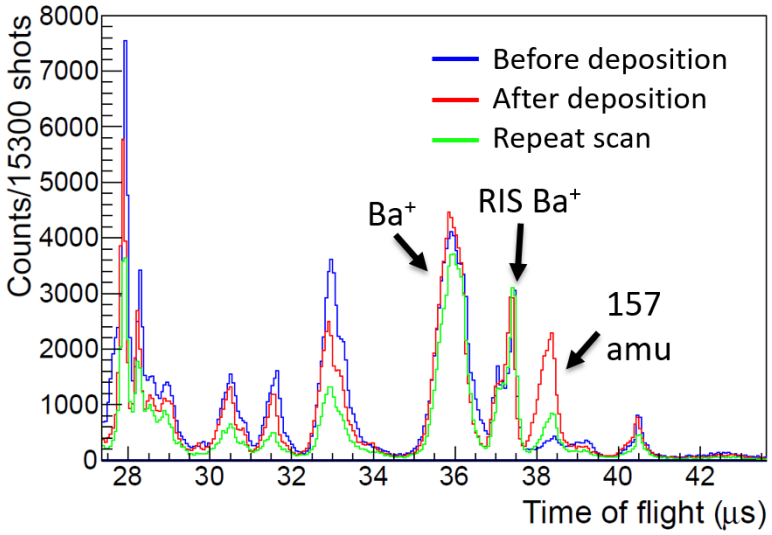

Figure 9.3: Deposition of a peak at 157 amu/e from 14 days of exposure to the ^{252}Cf source. Taken at 4 mJ/pulse desorption power. This peak decays after a second scan.

the deposition signal makes such results somewhat inconclusive.

9.1.2 Ba Recovery at High Desorption Power

Trials were also performed at higher desorption laser power, with the intent of desorbing Ba in different chemical forms or at different surface sites than what was seen at 1 mJ/pulse. At 4 mJ/pulse, a large RIS Ba peak is present, believed to be coming from Ba in the bulk of the Si. Despite this considerable background, an extended deposition over 14 days was attempted. Afterward, the TOF spectrum showed a considerable increase in a peak near 157 amu/e, which rapidly decayed afterward (Fig. 9.3). The Ba peak also showed signs of a slight increase, though this could be due to deposition of ^{137}Cs which is also produced by the ^{252}Cf source. The increases in the 157 amu/e peak came from the same desorption locations where the Ba and RIS Ba signals were strongest, and its counts were also correlated with the Ba and RIS Ba signals from shot to shot, *i.e.* laser shots which produced counts in the 157 amu/e TOF window also tended to produce counts in the Ba TOF window.

A repeat of this test for a shorter deposition time of 4 days produced similar results,

with a smaller increase. The increases in the 157 amu/e TOF window for the two trials were 15400 (14 days) and 3600 (4 days), compared to expected Ba depositions of fewer than 31 million atoms (14 days) and 10 million atoms (4 days), calculated using 50% collection efficiency. These signals correspond to minimum efficiencies of recovery from the surface of 0.049% (14 days) and 0.036% (4 days). Given that the signal in the 157 amu/e peak had not yet decreased to its baseline value before the 4-day deposition, simply subtracting the pre-deposition counts underestimates the increase in the peak from deposition. Hence, the data are compatible with a deposition signal that scales as expected with exposure to the ^{252}Cf source, possibly with a small reduction in recovery efficiency after further laser exposure. A further trial of deposition over 4 days with the source and substrate bias at opposite polarity, to repel positive ions from the substrate, showed no signs of such an increase.

The 157 amu/e peak's rapid decay after source exposure and its dependence on both the exposure time and the substrate bias suggest that this peak derives from atoms deposited by the source which are produced as positive ions in the Xe gas. This peak's spatial correlation with the Ba and RIS Ba signals is weak evidence that it may be a Ba molecule, with the common desorption locations possessing a surface chemistry favorable for Ba desorption. If this peak is indeed a singly-charged molecular ion including Ba from the ^{252}Cf source, the rest of the molecule should be in the 17-19 amu mass range, assuming the Ba isotopes match those of the ^{252}Cf source, primarily ^{138}Ba and ^{140}Ba. The molecule could be BaF, as was presumed to be the result of loading from the Gd-driven BaF source (§7.1.2), though the ^{252}Cf source does not produce any F, so there would have to be significant F at the Si surface for this to be the case. BaH_3O would have a similar mass assignment, and has been observed in a similar mass spectrometer system in the presence of significant H_2O vapor pressure [105]. BaO is another possibility, though for 14 days of exposure, there is expected to be slightly more ^{138}Ba than ^{140}Ba, so the mass comes out slightly below the calibrated estimate of 157 amu/e. This difference is comparable to the 2-3 amu/e FWHM mass resolution of the spectrometer, so the data are only in slight tension with this molecular assignment.

It may instead be a molecule of ^{137}Cs, such as CsOH, though this is even more in

tension with the calibrated mass estimate. Unfortunately, the limited mass resolution does not allow for a meaningful distinction between a single peak and a sum of two peaks separated by 2 amu, so distinguishing between molecules involving Cs and those involving Ba from the ^{252}Cf source is difficult with the current system. Further work is needed to establish with certainty the identity of this peak and its source.

9.2 Conclusions and Future Work

Significant progress has been made toward an apparatus capable of identifying Ba atoms with high sensitivity and toward an understanding of the underlying processes, but further work is needed to reach the single atom level. "Background Ba", which is not intentionally deposited on the substrate, has been useful for understanding LITD and RIS in this system, but is an obstacle to the ultimate goal of single atom sensitivity. Analysis of the background Ba signal has led to a complete characterization of the RIS process, including measurements of the desorbed atom profile. The estimated RIS ionization efficiency using this model exceeds 80%, making the RIS step close to its ultimate goal of near 100% efficiency. Achieving this goal will likely require more sophisticated shaping of the RIS beams to bypass diffraction effects from the present razor blade beam shaper. This may be accomplished in the future using a one-to-one telescope.

Despite its use as a diagnostic tool, background Ba must ultimately be eliminated from the TOF mass spectra to reach the desired sensitivity. Ba contamination of substrates has been measured using SIMS and ICPMS, finding both Ba concentrated at the surface as well as within the bulk of all substrates tested. While surface contamination is readily removed using laser cleaning and/or ion sputtering, greatly reducing the background Ba signal, many peaks in the TOF remain, including Ba, which are believed to come from the bulk of the material.

Further reduction in the backgrounds will likely require cleaner bulk materials or more exotic substrates. Work is currently underway at Brookhaven National Lab to produce a suspended graphene sheet for use in this system. This would allow removal of surface contaminants by either running current through the graphene or by cleaning

it with the desorption laser at low power. There is a risk of damage to the graphene from laser exposure, but the relatively low desorption energy of Ba on graphene may allow for its removal at laser powers below the damage threshold. Similarly, thin membranes as used in microelectromechanical systems have the potential to be effectively cleaned throughout their volume via laser heating and desorption.

Relatively minor backgrounds to the RIS Ba signal are also present from nearby peaks in the TOF spectra produced by the desorption laser alone ("directly-ionized"). These, and all directly-ionized peaks, can be eliminated by pulsing the substrate voltage during the delay between the desorption and RIS lasers to return the ions to the substrate surface. This would have the added benefit of greatly reducing charge depletion of the CEM from the high ion count rate, potentially boosting the RIS Ba signal significantly when running at very high desorption laser power. However, as the directly-ionized peaks are often the most reliable (in some cases the only) measure of the surface state, removing only the dominant low-mass peaks is preferable. This could be achieved through the addition of a Bradbury-Nielsen gate further down the TOF path: a series of parallel wires with a time-dependent bias deflect ions arriving earlier (with lower mass), while allowing later ions to pass through [106].

The lack of a reliable Ba deposition signal remains the most pressing challenge, and may require a better understanding of the underlying surface chemistry for Ba desorption on various substrates. Early trials using the Gd-driven Ba source showed signs of deposition of a peak believed to be BaF, but no increases in the directly-ionized Ba or RIS Ba peaks in the TOF were seen as a result of exposure to this source. It is believed that this is due to binding of deposited Ba to surface contaminants, causing formation of molecules such as BaO which the RIS process is no longer sensitive to. Hence, desorption of neutral Ba molecules may be a significant channel for Ba loss without detection. If the chemical state of the adsorbed Ba can be changed to favor desorption as an atom, the high RIS efficiency should allow for detection of almost all desorbed Ba.

Toward that aim, bakeout of the vacuum system resulted in significant reductions in vacuum pressure, including H_2O, but did not lead to the appearance of a deposition signal of atomic Ba from the Gd-driven source. Similarly, no deposition signal

was observed after ion sputtering of surfaces to remove oxides. However, chemical degradation of the Gd-driven source is believed to have greatly reduced its Ba output. Future efforts using the ^{252}Cf fission fragment source after bakeout, including sputtering trials, have the potential for better results. In particular, substrates such as graphene and Ta, which showed no successful deposition with the Gd-driven source at higher vacuum pressure, may prove promising under these new conditions. These materials are especially interesting due to their empirically-low Ba desorption thresholds and the ability of Ta to reduce BaO, which may promote its desorption in atomic form.

Tests using the ^{252}Cf fission source for deposition on Si before bakeout have already shown some promise, with increases in the directly-ionized Ba and RIS Ba signals appearing at relatively low desorption power. Though deposition and recovery under these conditions has been repeated a few times, the surface properties required for doing so reliably are not yet known. Using the calculated production of Ba from the ^{252}Cf source and a maximum collection efficiency of 50%, the minimum efficiency of Ba recovery from the surface under favorable conditions is estimated at 0.25%. This is well below the ultimate efficiency goal, but the strong dependence of the deposition signal on desorption power, surface conditioning from the laser, and the particular Si substrate suggest that the current efficiency is not an inherent limit of the technique.

There is also evidence of deposition and recovery of another peak in the TOF spectra, which may be a Ba molecule, as a result of exposure to the ^{252}Cf source, when higher desorption power is used. Its appearance is more reliable than the atomic Ba signal at lower desorption power, though if resulting from Ba deposited from the ^{252}Cf source, its efficiency of recovery is significantly lower (above 0.05%). This may be evidence that Ba from the source can be desorbed in at least two forms, though further work is required to definitively prove that the peak comes from the source and not some other part of the deposition procedure. If it does, tests of its relative strength under different conditions may provide valuable insights on how to improve the atomic Ba signal as well.

Lastly, this technique must be tested in a LXe environment. Because the apparatus was designed around this goal, this is expected to be relatively straightforward.

LXe has been stably maintained in the xenon cell previously, so the Xe cryogenic system's successful operation has been verified. Deposition from the ^{252}Cf source in LXe should work equally well as in gas xenon. The substrate mount was designed to account for differential thermal contraction when submerged in LXe, though care will have to be taken during submersion to minimize thermal shock. It is possible that Ba from the ^{252}Cf source might adsorb differently onto the substrate surface in LXe than in gas Xe, in which case further studies may be necessary. If the full deposition and recovery process is tested in LXe and its efficiency measured, then those results can be used to evaluate whether it is worthwhile to pursue integrating the technique with a working LXe TPC.

Bibliography

[1] J. Lesgourgues and S. Pastor, *Massive neutrinos and cosmology*, *Physics Reports* **429** (2006) 307 – 379.

[2] K. Hirata, T. Kajita, M. Koshiba, M. Nakahata, Y. Oyama, N. Sato et al., *Observation of a neutrino burst from the supernova SN1987A*, *Phys. Rev. Lett.* **58** (Apr, 1987) 1490–1493.

[3] T. Araki, S. Enomoto, K. Furuno, Y. Gando, K. Ichimura, H. Ikeda et al., *Experimental investigation of geologically produced antineutrinos with KamLAND*, *Nature* **436** (Jul, 2005) 499–503.

[4] J. N. Bahcall, *Neutrinos from the sun*, *Scientific American* **221** (July, 1969) 28–37.

[5] Particle Data Group collaboration, C. Patrignani et al., *Review of Particle Physics*, *Chinese Physics C* **40** (2016) 100001.

[6] F. Capozzi, E. Lisi, A. Marrone, D. Montanino and A. Palazzo, *Neutrino masses and mixings: Status of known and unknown 3ν parameters*, *Nuclear Physics B* **908** (2016) 218 – 234. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.

[7] C. Kraus, B. Bornschein, L. Bornschein, J. Bonn, B. Flatt, A. Kovalik et al., *Final results from phase II of the Mainz neutrino mass search in tritium β decay*, *The European Physical Journal C - Particles and Fields* **40** (2005) 447–468.

[8] V. N. Aseev, A. I. Belesev, A. I. Berlev, E. V. Geraskin, A. A. Golubev, N. A. Likhovid et al., *Upper limit on the electron antineutrino mass from the Troitsk experiment*, *Phys. Rev. D* **84** (Dec, 2011) 112003.

[9] Planck collaboration, Ade, P. A. R., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C. et al., *Planck 2015 results - XIII. cosmological parameters*, *A&A* **594** (2016) A13.

[10] F. T. Avignone, S. R. Elliott and J. Engel, *Double beta decay, Majorana neutrinos, and neutrino mass*, *Rev. Mod. Phys.* **80** (Apr, 2008) 481–516.

[11] B. Kayser, *Neutrino Mass*, ch. 1: Neutrino Mass, Mixing, and Flavor Change, pp. 1–24. Springer Berlin Heidelberg, September, 2003. 10.1007/978-3-540-44901-0_1.

[12] J. Bernabeu, *On the history of the PMNS Matrix... with today's perspective*, *Nuovo Cim.* **C037** (2014) 145–154.

[13] M. Gonzalez-Garcia and M. Maltoni, *Phenomenology with massive neutrinos*, *Physics Reports* **460** (2008) 1 – 129.

[14] R. Zukanovich Funchal, B. Schmauch and G. Giesen, *The Physics of Neutrinos*, 1308.1029.

[15] S. Weinberg, *Baryon- and lepton-nonconserving processes*, *Phys. Rev. Lett.* **43** (Nov, 1979) 1566–1570.

[16] EXO collaboration, N. Ackerman, B. Aharmim, M. Auger, D. J. Auty, P. S. Barbeau, K. Barry et al., *Observation of two-neutrino double-beta decay in* ^{136}Xe *with the EXO-200 detector*, *Phys. Rev. Lett.* **107** (Nov, 2011) 212501.

[17] M. Hirsch, H. Klapdor-Kleingrothaus and O. Panella, *Double beta decay in left-right symmetric models*, *Physics Letters B* **374** (1996) 7 – 12.

[18] M. Hirsch, H. Klapdor-Kleingrothaus and S. Kovalenko, *On the SUSY accompanied neutrino exchange mechanism of neutrinoless double beta decay*, *Physics Letters B* **372** (1996) 181 – 186.

[19] H. Klapdor-Kleingrothaus and U. Sarkar, *Neutrinoless double beta decay with scalar bilinears*, *Physics Letters B* **554** (2003) 45 – 50.

[20] J. Schechter and J. W. F. Valle, *Neutrinoless double-β decay in SU(2)$\times$U(1) theories*, *Phys. Rev. D* **25** (Jun, 1982) 2951–2954.

[21] KamLAND-Zen collaboration, A. Gando, Y. Gando, T. Hachiya, A. Hayashi, S. Hayashida, H. Ikeda et al., *Search for Majorana neutrinos near the inverted mass hierarchy region with KamLAND-Zen*, *Phys. Rev. Lett.* **117** (Aug, 2016) 082503.

[22] EXO-200 collaboration, J. B. Albert et al., *Search for Majorana neutrinos with the first two years of EXO-200 data*, *Nature* **510** (2014) 229–234, [1402.6956].

[23] M. Agostini, M. Allardt, A. Bakalyarov, M. Balata, I. Barabanov, L. Baudis et al., *Search of neutrinoless double beta decay with the GERDA experiment*, *Nuclear and Particle Physics Proceedings* **273** (2016) 1876 – 1882.

[24] CUORE collaboration, K. Alfonso, D. R. Artusa, F. T. Avignone, O. Azzolini, M. Balata, T. I. Banks et al., *Search for neutrinoless double-beta decay of ^{130}Te with CUORE-0*, *Phys. Rev. Lett.* **115** (Sep, 2015) 102502.

[25] NEMO-3 collaboration, R. Arnold, C. Augier, J. D. Baker, A. S. Barabash, A. Basharina-Freshville, S. Blondel et al., *Results of the search for neutrinoless double-β decay in ^{100}Mo with the NEMO-3 experiment*, *Phys. Rev. D* **92** (Oct, 2015) 072011.

[26] M. Auger et al., *The EXO-200 detector, part I: Detector design and construction*, *JINST* **7** (2012) P05010, [1202.2192].

[27] D. S. Leonard et al., *Systematic study of trace radioactive impurities in candidate construction materials for EXO-200*, *Nucl. Instrum. Meth. A* **591** (2008) 490–509, [0709.4524].

[28] "3M HFE-7000." http://www.3m.com/.

[29] E.-I. Esch, T. J. Bowles, A. Hime, A. Pichlmaier, R. Reifarth and H. Wollnik, *The cosmic ray muon flux at WIPP*, *Nuclear Instruments and Methods in Physics Research A* **538** (Feb., 2005) 516–525, [astro-ph/0408486].

[30] E. Conti et al., *Correlated fluctuations between luminescence and ionization in liquid xenon*, *Phys. Rev. B* **68** (2003) 054201, [hep-ex/0303008].

[31] M. Duerr, M. Lindner and K. Zuber, *Consistency test of neutrinoless double beta decay with one isotope*, *Phys. Rev. D* **84** (2011) 093004, [1103.4735].

[32] E. Fiorini, *Proceedings, Int. Conf. NEUTRINO '77* **2** (1978) 315–320.

[33] E. Bellotti, E. Fiorini, C. Liguori, A. Pullia, A. Sarracino and L. Zanotti, *An experimental investigation on lepton number conservation in double beta processes*, *Lett. Nuovo Cim.* **33** (1982) 273–283.

[34] A. S. Barabash et al., *Two neutrino double-beta decay of ^{100}Mo to the first excited 0^+ state in ^{100}Ru*, *Phys. Lett. B* **345** (Feb., 1995) 408–413.

[35] M. F. Kidd, J. H. Esterline, W. Tornow, A. S. Barabash and V. I. Umatov, *New Results for Double-Beta Decay of Mo-100 to Excited Final States of Ru-100 Using the TUNL-ITEP Apparatus*, *Nucl. Phys. A* **821** (2009) 251–261, [0902.4418].

[36] A. S. Barabash, F. Hubert, P. Hubert and V. I. Umatov, *Double-beta decay of Nd-150 to the first 0+ excited state of Sm-150*, *Phys. Atom. Nucl.* **67** (2004) 1216–1219.

[37] M. F. Kidd, J. H. Esterline, S. W. Finch and W. Tornow, *Two-neutrino double-β decay of ^{150}Nd to excited final states in ^{150}Sm*, *Phys. Rev. C* **90** (2014) 055501, [1411.3755].

[38] GERDA collaboration, M. Agostini et al., *$2\nu\beta\beta$ decay of ^{76}Ge into excited states with GERDA Phase I*, *J. Phys. G* **42** (2015) 115201, [1506.03120].

[39] KamLAND-Zen collaboration, K. Asakura, A. Gando, Y. Gando, T. Hachiya, S. Hayashida, H. Ikeda et al., *Search for double-beta decay of 136Xe to excited states of 136Ba with the KamLAND-Zen experiment*, *Nuclear Physics A* **946** (2016) 171 – 181, [1509.03724].

[40] A. A. Sonzogni, *Nuclear Data Sheets* **95** (2002) 837. Data extracted from the ENSDF database, Oct. 9, 2015, http://www.nndc.bnl.gov/.

[41] P. Vogel. Private communication, 2012.

[42] J. Kotila and F. Iachello, *Phase space factors for double-β decay*, *Phys. Rev. C* **85** (2012) 034316, [1209.5722].

[43] J. Barea, J. Kotila and F. Iachello, *$0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements in the interacting boson model with isospin restoration*, *Phys. Rev. C* **91** (2015) 034304, [1506.08530].

[44] EXO-200 collaboration, J. B. Albert et al., *Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector*, *Phys. Rev. C* **89** (2014) 015502, [1306.6106].

[45] EXO-200 collaboration, J. B. Albert et al., *Search for Majoron-emitting modes of double-beta decay of ^{136}Xe with EXO-200*, *Phys. Rev. D* **90** (2014) 092004, [1409.6829].

[46] J. Allison et al., *Geant4 developments and applications*, *IEEE Trans. Nucl. Sci.* **53** (2006) 270.

[47] A. Hocker et al., *TMVA - Toolkit for Multivariate Data Analysis*, *PoS (ACAT 2007)* **ACAT** (2007) 040, [physics/0703039].

[48] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81–86.

[49] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, *Annals Math. Statist.* **9** (1938) 60–62.

[50] G. Cowan, *Statistical Data Analysis*. Oxford Science Publications, Clarendon, 1998.

[51] EXO-200 collaboration, J. B. Albert, D. J. Auty, P. S. Barbeau, D. Beck, V. Belov, M. Breidenbach et al., *Search for $2\nu\beta\beta$ decay of ^{136}Xe to the 0_1^+ excited state of ^{136}Ba with the EXO-200 liquid xenon detector*, *Phys. Rev. C* **93** (Mar, 2016) 035501.

[52] A. Caldwell, A. Merle, O. Schulz and M. Totzauer, *A Global Bayesian Analysis of Neutrino Mass Data*, 1705.01945.

[53] M. Agostini, G. Benato and J. Detwiler, *Discovery probability of next-generation neutrinoless double-β decay experiments*, 1705.02996.

[54] T. Brunner, D. Fudenberg, V. Varentsov, A. Sabourov, G. Gratta, J. Dilling et al., *An RF-only ion-funnel for extraction from high-pressure gases*, *International Journal of Mass Spectrometry* **379** (2015) 110 – 120.

[55] nEXO collaboration, B. Mong, S. Cook, T. Walton, C. Chambers, A. Craycraft, C. Benitez-Medina et al., *Spectroscopy of Ba and Ba^+ deposits in solid xenon for barium tagging in nEXO*, *Phys. Rev. A* **91** (Feb, 2015) 022505.

[56] EXO-200 collaboration, J. B. Albert, D. J. Auty, P. S. Barbeau, D. Beck, V. Belov, M. Breidenbach et al., *Measurements of the ion fraction and mobility of $\alpha-$ and β-decay products in liquid xenon using the EXO-200 detector*, *Phys. Rev. C* **92** (Oct, 2015) 045504.

[57] N. Schwenter, E.-E. Koch and J. Jortner, *Electronic excitations in condensed rare gases*, ch. 4: Electronic excitations in rare-gas liquids, pp. 63–82. Springer-Verlag Berlin Heidelberg, 1985. 10.1007/BFb0111641.

[58] W. Haynes, *CRC Handbook of Chemistry and Physics, 97th Edition*. CRC Press, 2016.

[59] S.-C. Jeng, W. M. F. Jr and M. Miyajima, *Measurements of the mobility of alkaline earth ions in liquid xenon*, *Journal of Physics D: Applied Physics* **42** (2009) 035302.

[60] J. Wodin, *Single ion trapping in a buffer gas and liquid Xe energy resolution studies for the EXO experiment*. PhD thesis, Stanford University, February, 2007.

[61] M. P. Green, *On ion probes and traps, barium tagging for the EXO double beta decay detector*. PhD thesis, Stanford University, June, 2010.

[62] J. Wang, J. A. Hallmark, D. S. Marshall, W. J. Ooms, P. Ordejón, J. Junquera et al., *Bonding and diffusion of Ba on a Si(001) reconstructed surface*, *Phys. Rev. B* **60** (Aug, 1999) 4968–4971.

[63] Z. G. Wang and X. T. Zu, *First principles calculations of adsorption and diffusion of Ba on a reconstructed Si(001) surface*, *Surface Review and Letters* **13** (2006) 365–368.

[64] K. Nakada and A. Ishii, *Migration of adatom adsorption on graphene using DFT calculation*, *Solid State Communications* **151** (2011) 13 – 16.

[65] L. A. Hemstreet and S. R. Chubb, *Electronic structure of c(2×2) Ba adsorbed on W(001)*, *Phys. Rev. B* **47** (Apr, 1993) 10748–10753.

[66] K. Oura, M. Katayama, A. V. Zotov, V. G. Lifshits and A. A. Saranin, *Surface Science: An Introduction*, ch. 12: Elementary Processes at Surfaces I. Adsorption and Desorption, pp. 295–323. Springer Berlin Heidelberg, 2003. 10.1007/978-3-662-05179-5.

[67] P. Hess, *Photoacoustic, Photothermal and Photochemical Processes at Surfaces and in Thin Films*, ch. 3: Time-of-Flight Analysis of IR and UV Laser-Induced Multilayer Desorption and Ablation, pp. 55–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 1989. 10.1007/978-3-642-83945-0_3.

[68] P. Redhead, *Thermal desorption of gases*, *Vacuum* **12** (1962) 203 – 211.

[69] D. Marla, U. V. Bhandarkar and S. S. Joshi, *A model of laser ablation with temperature-dependent material properties, vaporization, phase explosion and plasma shielding*, *Applied Physics A* **116** (2014) 273–285.

[70] D. Burgess Jr., P. C. Stair and E. Weitz, *Calculations of the surface temperature rise and desorption temperature in laser-induced thermal desorption*, *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films* **4** (1986) 1362–1366.

[71] G. E. Jellison Jr. and D. H. Lowndes, *Optical absorption coefficient of silicon at 1.152 μ at elevated temperatures*, *Applied Physics Letters* **41** (1982) 594–596.

[72] B. Koehler and S. George, *Laser-induced desorption of H_2 from Si(111)*7×7, *Surface Science* **248** (1991) 158 – 172.

[73] A. H. A. Lutey, *An improved model for nanosecond pulsed laser ablation of metals*, *Journal of Applied Physics* **114** (2013) 083108.

[74] B. Wu and Y. C. Shin, *Modeling of nanosecond laser ablation with vapor plasma formation*, *Journal of Applied Physics* **99** (2006) 084310.

[75] R. E. Russo, *Laser ablation*, *Applied Spectroscopy* **49** (1995) 14A–28A.

[76] G. S. Hurst, M. G. Payne, S. D. Kramer and J. P. Young, *Resonance ionization spectroscopy and one-atom detection*, *Rev. Mod. Phys.* **51** (Oct, 1979) 767–819.

[77] J. J. Curry, *Compilation of wavelengths, energy levels, and transition probabilities for Ba I and Ba II*, *Journal of Physical and Chemical Reference Data* **33** (2004) 725–746.

[78] U. Fano, *Effects of configuration interaction on intensities and phase shifts*, *Phys. Rev.* **124** (Dec, 1961) 1866–1878.

[79] M. A. Kalyar, M. Rafiq, Sami-ul-Haq and M. A. Baig, *Absolute photoionization cross section from the 6s6p $^{1,3}P_1$ excited states of barium*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **40** (2007) 2307.

[80] B. Willke and M. Kock, *Measurement of photoionization cross sections from the laser-excited Ba I (6s6p) $^1P_1^0$ state*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **26** (1993) 1129.

[81] R. P. Wood, C. H. Greene and D. Armstrong, *Photoionization of the barium 6s6p $^1P^\circ_1$ state: Comparison of theory and experiment including hyperfine-depolarization effects*, *Phys. Rev. A* **47** (Jan, 1993) 229–235.

[82] L.-W. He, C. E. Burkhardt, M. Ciocca, J. J. Leventhal, H.-L. Zhou and S. T. Manson, *Correlation effects in the photoionization of Ba(6s6p $^1P_1$): Determination of cross sections for production of specific final J states*, *Phys. Rev. A* **51** (Mar, 1995) 2085–2093.

[83] W. M. Fairbank Jr., *A seat-of-the pants approach to RIS theory*, *AIP Conference Proceedings* **388** (1997) 457–464.

[84] M. Montero Díez, *Development of a resonance ionization spectroscopy ion-transport probe for the Enriched Xenon Observatory*. PhD thesis, Stanford University, May, 2012.

[85] K. Twelker, S. Kravitz, M. Montero Díez, G. Gratta, W. Fairbank Jr., J. B. Albert et al., *An apparatus to manipulate and identify individual Ba ions from bulk liquid Xe*, *Review of Scientific Instruments* **85** (2014) 095114.

[86] K. Twelker, *Surface adsorption and resonance ionization spectroscopy for barium identification in neutrinoless double beta decay experiments*. PhD thesis, Stanford University, August, 2014.

[87] P. C. Zalm, *Energy dependence of the sputtering yield of silicon bombarded with neon, argon, krypton, and xenon ions*, *Journal of Applied Physics* **54** (1983) 2660–2666.

[88] G. D. Magnuson and C. E. Carlston, *Sputtering yields of single crystals bombarded by 1- to 10-keV Ar+ ions*, *Journal of Applied Physics* **34** (1963) 3267–3273.

[89] M. M. Díez, K. Twelker, W. F. Jr., G. Gratta, P. S. Barbeau, K. Barry et al., *A simple radionuclide-driven single-ion source*, *Review of Scientific Instruments* **81** (2010) 113301.

[90] J. F. Ziegler, M. Ziegler and J. Biersack, *SRIM – the stopping and range of ions in matter (2010)*, *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **268** (2010) 1818 – 1823. 19th International Conference on Ion Beam Analysis.

[91] T. England and B. Rider, *Evaluation and compilation of fission product yields*, Tech. Rep. LA-UR-94-3106, Los Alamos National Laboratory (LANL), October, 1994.

[92] R. Noll, *Laser-induced breakdown spectroscopy*, ch. 7: Material ablation, pp. 97–118. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. 10.1007/978-3-642-20668-9.

[93] K. K. Anoop, X. Ni, M. Bianco, D. Paparo, X. Wang, R. Bruzzese et al., *Two-dimensional imaging of atomic and nanoparticle components in copper plasma plume produced by ultrafast laser ablation*, *Applied Physics A* **117** (2014) 313–318.

[94] D. A. Dahl, *SIMION for the personal computer in reflection*, *Int. J. Mass Spectrom.* **200** (Dec., 2000) 3–25.

[95] M. Currie, J. D. Caldwell, F. J. Bezares, J. Robinson, T. Anderson, H. Chun et al., *Quantifying pulsed laser induced damage to graphene*, *Applied Physics Letters* **99** (2011) 211909.

[96] R. Kelly and R. Dreyfus, *On the effect of Knudsen-layer formation on studies of vaporization, sputtering, and desorption*, *Surface Science* **198** (1988) 263 – 276.

[97] R. Kelly, A. Miotello, A. Mele and A. G. Guidoni, *Laser Ablation and Desorption*, vol. 30 of *Experimental Methods in the Physical Sciences*, ch. 5: Plume Formation and Characterization in Laser-Surface Interactions, pp. 225 – 289. Academic Press, 1997. 10.1016/S0076-695X(08)60397-X.

[98] R. Musket, W. McLean, C. Colmenares, D. Makowiecki and W. Siekhaus, *Preparation of atomically clean surfaces of selected elements: A review*, *Applications of Surface Science* **10** (1982) 143 – 207.

[99] J. Kubby and J. Boland, *Scanning tunneling microscopy of semiconductor surfaces*, *Surface Science Reports* **26** (1996) 61 – 204.

[100] F. G. Allen, J. Eisinger, H. D. Hagstrum and J. T. Law, *Cleaning of silicon surfaces by heating in high vacuum*, *Journal of Applied Physics* **30** (1959) 1563–1571.

[101] A. E. Smirnov and A. A. Urusovskaya, *Selective etching and dissolution of BaF_2 crystals*, *Journal of Materials Science* **15** (1980) 1183–1188.

[102] S. M. Bedair and H. P. S. Jr., *Atomically clean surfaces by pulsed laser bombardment*, *Journal of Applied Physics* **40** (1969) 4776–4781.

[103] D. Zehner, *Surface studies of pulsed laser irradiated semiconductors*, *Semiconductors and Semimetals* **23** (1984) 405 – 470.

[104] R. Larciprete and E. Borsella, *Excimer laser cleaning of Si(100) surfaces at 193 and 248 nm studied by LEED, AES and XPS spectroscopies*, *Journal of Electron Spectroscopy and Related Phenomena* **76** (1995) 607 – 612. Proceedings of the Sixth International Conference on Electron Spectroscopy.

[105] Y. Ito. Private communication, 2017.

[106] N. E. Bradbury and R. A. Nielsen, *Absolute values of the electron mobility in hydrogen*, *Phys. Rev.* **49** (Mar, 1936) 388–393.