Understanding VMI & MBB Group Spectrometer

J. D. Pickering

July 21, 2020

Chapter 1

Understanding VMI

1.1 Ion Imaging

Ion imaging, in the broadest sense, is a technique whereby ions are focussed onto a 2D detector screen such that a 2D image of the ions is created. Judicious choice of how this focussing is achieved can lead to two different outcomes, spatial-map imaging (SMI) or velocity-map imaging (VMI). Simply put, SMI focusses ions produced from the same *spatial position* within the ionising source to the same position on the 2D image (irrespective of their momenta in the detection plane); whereas VMI focusses ions produced with the same *momentum* in the ionising source to the same position on the 2D image (irrespective of the position of ionisation within the source). Both of these techniques are widely used within the physical sciences. SMI is most often used in imaging mass spectrometry (for example, to build up a picture of a surface sample based on the positions of the recorded ions), whereas VMI finds an important niche in studies of chemical dynamics, as it allows the velocity of ions produced via photochemical processes to be easily and accurately measured. Indeed, this was *the* reason for the creation of the VMI technique by Eppink and Parker - and their spectrometer design is still widely used today in laboratories worldwide.

In the MBB group, we are primarily focussed on VMI - as the ultimate goal is to characterise ultrafast processes in molecules using femtosecond laser-induced Coulomb Explosion. The goals here are two-fold, one possibility is to characterise the structures of neutral molecules that are difficult to interrogate spectroscopically (as the Coulomb Explosion is a non-resonant strong field process, so is not subject to many of the selection rules that govern 'conventional' optical spectroscopy). This is a technique which has been proven to work for 'simple' molecules in prior work by the Brouard Group, Michael, and the author. Extending this to more complex systems is challenging, but potentially fruitful. A second possibility is to use the ultrafast imaging as a tool, coupled with inducing a photochemical process (e.g. photodissociation) on an ultrafast timescale, such that

the process of interest can be monitored in real-time. This is the 'femtochemistry' that won Ahmed Zewail the Nobel Prize in 1999, and the creation of a true 'molecular movie' of a chemical process (e.g. covalent bond formation), has so far proved elusive (despite what people doing ultrafast diffractive imaging try to peddle at conferences).

The purpose of this document is firstly to give a background of the working principles of VMI, together with some (hopefully) interesting insights into the nature of velocity-focussing and drawing numerous analogies to classical 'light' optics. Secondly it is to document the development of the spectrometer from conception to creation. Having known almost nothing about electrostatic lenses or use of SIMION and Lua before I started this (and with very minimal assistance), I think a document explaining the process and identifying many of the pitfalls I fell into will be useful for future students.

1.2 Velocity Map Imaging

As stated, VMI is a process whereby the momenta of ions from a source are mapped onto a 2D screen. Ions with the same *momenta* in the detector plane are mapped onto the same *spatial* position on the detector, regardless of their *initial* spatial position in the source region. This is illustrated in Figure 1.1.

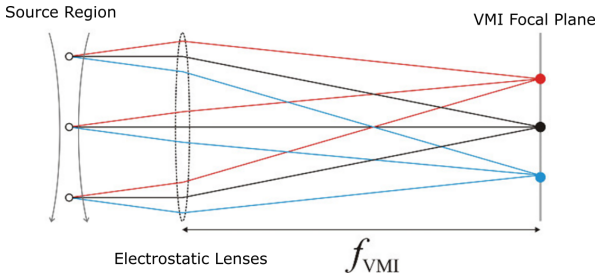

Figure 1.1: Illustration of how VMI is achieved using an electrostatic lenses. Ions with the same momenta have the same colour. Figure adapted from Harb *et al* (doi: 10.1063/1.3505799).

Before we proceed and look at some actual spectrometer designs, it is useful to define some coordinate conventions. These are shown in Figure 1.2.

The laboratory frame (Figure 1.2) is defined to coincide with the internal coordinates used in the SIMION software package, such that the X-axis is the time-of-flight (ToF) axis, and the detector plane is the YZ-plane. The laser is assumed to propagate along the Z-axis, and therefore the polarisation plane is the XY-plane. The molecular beam pulses approach along the X-axis. Note that

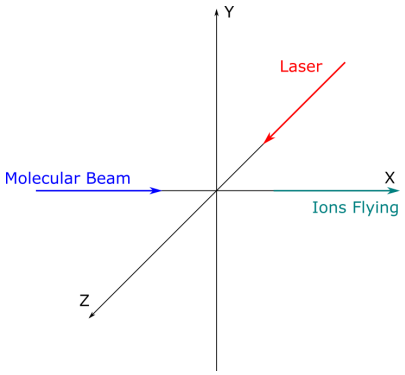

Figure 1.2: Coordinate conventions, consistent with SIMION's internal coordinates. The detector is in the YZ-plane.

we maintain the convention whereby the laboratory frame is referred to using upper-case letters, as is common in literature.

Ions are produced by an ionisation source, in this case a focussed laser pulse, and are assumed to be created approximately isotropically[1]. The source region is defined by the overlap of the focussed laser beam, and (generally) a gas pulse from a molecular beam source. This region is located between two electrodes, referred to as the 'repeller' and 'extractor' electrodes, as shown in Figure 1.3. The voltage applied to these electrodes is such that there is an electric field gradient[2] between them, and the nascent cations are accelerated down a potential 'hill', away from the repeller electrode. After the extractor electrode, there is a grounded electrode to ensure that the region after the extractor electrode is essentially field free all the way until the detector.

[1] This is, of course, also a simplification as the polarisation of the laser field will induce some anisotropy (dependent on the molecular target).

[2] Many times weaker than the electric field of the laser pulse, so it does not perturb the electronic structure of the molecule significantly/at all.

Most spectrometers (and the spectrometer designed here) are cylindrically symmetric around the X-axis, and the electrodes are simply annular discs positioned orthogonal to the ToF axis (in the YZ-plane). The repeller plate has a small hole in it to allow the molecular beam to pass through, and the extractor plate has a larger hole, and the third lens has an even larger hole. As mentioned, the voltages applied to the repeller and extractor plates create an electric field gradient that accelerates ions towards the detection region. **The form of this electric field determines the focussing behaviour.**

Figure 1.3 shows a model of a typical spectrometer (in the XZ-plane), with some realistic voltages applied to the repeller and extractor. Plotted in red are lines of *equipotential*, that is, lines where the potential due to the electric field is a constant. Between the repeller and extractor plates, the equipotential lines are approximately parallel. The region where the field extends through the extractor electrode into the field-free region is clearly not composed of parallel lines - this

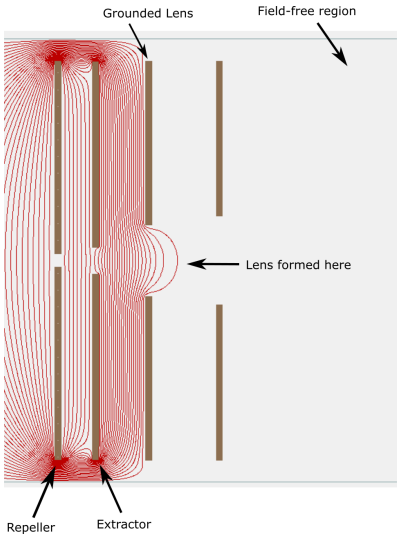

Figure 1.3: Screenshot from SIMION showing a simple VMI interaction region. The electrostatic lens is formed in the region where the equipotentials (red lines) bulge - inhomogeneous electric field.

'bulge', or inhomogeneity in the electric field, creates an **electrostatic lens** - this lens is what gives rise to the VMI focussing[3].

[3]The shape of the equipotentials look rather like a spherical glass light lens - this is not a coincidence. In a light lens the curvature means that different light rays are deviated by different amounts depending on where they hit the lens - leading to focussing behaviour.

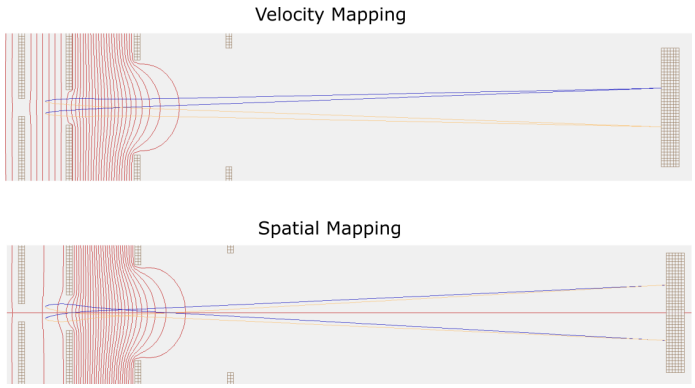

Figure 1.4: Diagram showing the difference between velocity mapping (VMI) and spatial mapping (SMI). The two diagrams show the same geometry and ions, but with different voltages applied to the repeller and extractor electrodes (see text). Ions of a common colour have the same velocity in the detector plane.

As the ions pass through this area of inhomogeneous field ('through the lens'), they undergo focussing. Figure 1.4 shows that the ions come through the lens, fly through a field-free region, and then hit a 2D detector screen. The two spectrometers shown in Figure 1.4 have the same geometry, but with different voltages

applied to the electrodes (see below). The blue and orange trajectories correspond to ions of the same energy, but different momenta in the detector plane. The blue ions have an initial momentum directed upwards, whereas the orange ions have momentum directed downwards. There are several points to note about this:

- Ions of a common colour are created at different points within the interaction region, but with the same momentum in the detector plane. It can be seen that in the diagram marked 'Velocity Mapping', the ions with a common colour are focussed to the same point on the detector - this is the point at which the ions are termed to be 'velocity focussed'. Ions of a common velocity are focussed to the same point, regardless of their initial positions within the source region.
- In the diagram marked 'Spatial Mapping', it can be seen that ions formed from the same position are mapped to the same point on the detector - this is the point at which the ions are 'spatially focussed'.
- The voltages applied to the electrodes in both cases are similar - the repeller is held at $+2\,\mathrm{kV}$, and the extractor is held at $+1.58\,\mathrm{kV}$ (VMI) and $+1.85\,\mathrm{kV}$ (SMI).
- In general, it is impossible to make an electrostatic lens that is diverging (which is possible with a light lens). It may seem as though the ions are diverging having passed through our lens, but they are actually *converging* relative to their initial velocities (as can be seen most clearly in the SMI example).

The positions of the two focal planes can be controlled by adjusting the voltages applied to the two plates, or by adding a third lens after the extractor that has a voltage applied to it to 'fine tune' the velocity focussing.

Generally speaking, in VMI, the voltage ratios between the different plates are what determine the focussing behaviour - rather than an 'absolute' voltage. As such it is common to see voltages given as a repeller/extractor/third lens ratio, rather than absolute numbers. The overall size of the image on the detector is controlled by the repeller voltage, and the VMI condition will be maintained as the repeller voltage is varied, provided that the voltages on the other electrodes vary with it (maintaining the same overall ratio). Higher voltages on the repeller lead to stronger focussing, so higher energy ions hit smaller radii on the detector. As the detector is a finite size, this can be useful to ensure that the ion channel of interest fills as much of the detector as possible. Controlling these voltages in synchrony (using a computer) gives a rudimentary form of magnification of the VMI image. The highest energy ion it is possible to focus is generally dictated by the voltage that can be applied to the repeller plate without arcing[4] - around $9\,\mathrm{kV}$ is generally possible with standard power supplies.

[4] Where a spark bridges the gap between plates - you will know about it if this happens.

1.2.1 Why VMI?

An obvious question is why this technique is called *velocity* map imaging, when in reality it is not the ion velocities that are focussed, but the ion momenta. That is, two ions with the same momentum will hit the same point on the detector - whereas two ions of very different masses will hit different points even if their velocities are the same. Therefore, we see that really ion momenta are mapped onto the detector. Sometimes people often erroneously say that kinetic energy is focussed rather than velocity - but this is even more incorrect than saying velocity is focussed, as kinetic energy is a scalar quantity and therefore an image produced where only kinetic energy is focussed would have no angular resolution.

The terminology, I believe, arises as follows. Clearly, for a given mass, ions of different velocities will hit different points on the detector. These experiments are normally performed by gating the detector such that ions of a specific m/z ratio are the only detected ions. Within that specific mass channel, ions hitting different points on the detector necessarily have different velocities - hence the name *velocity* map imaging.

1.3 Electrostatic and Light Optics

Having understood an overview of how the VMI process works, it is now useful to spend more time thinking about the nature of electrostatic optics, and particularly which analogies can be made to light optics. The analogies are many, but there are also areas where electrostatic optics are quite different to light optics.

1.3.1 Electrostatic Lenses

The simplest form of electrostatic lens is a single-element charged cylinder. In the example previously, the cylinder was very short and with thick walls, so that it was better described as an annular disc, but it was still a cylinder. While a detailed discussion of the behaviour of many different types of electrostatic lenses is beyond the scope of this document, below is a summary of the more useful underlying physics to aid comprehension of the behaviour of a spectrometer. For a more complete treatment see standard texts on electrostatic optics (such as Helmut Liebl's *'Applied Charged Particle Optics'*, or D.W.O Heddle's *'Electrostatic Lens Systems'*).

It can be shown that for a particle travelling through a field with potentials V_1 and V_2 on either side, the ratio of the angle of incidence α_1 to the angle of refraction α_2 is given by Equation 1.1

$$\frac{\sin(\alpha_2)}{\sin(\alpha_1)} = \sqrt{\frac{V_1}{V_2}}\left(= \frac{n_1}{n_2}\right) \tag{1.1}$$

The term on the right hand side is the ratio of the refractive indices of medium 1 and 2, and by comparison with Snell's law (in brackets) it can be seen that $\sqrt{V}$ plays the role of refractive index in electrostatic optics. As should be known from any introductory optics course, the refractive index determines the degree to which a light ray entering a medium is refracted, or 'bent' upon entering the medium. Clearly then, within electrostatic optics, a particle entering a region of higher field will experience a greater force and be deviated ('refracted') by the field more. This should intuitively make sense, at least on a qualitative level.

We can continue deriving equations of motion for the particle in a uniform field, as is done rigorously in Liebl or Heddle, and arrive at various expressions for the focussing properties of various electrostatic lenses. We will not discuss all of these here, as it would detain us unnecessarily, but there are some useful points to note that we will directly make reference to in several places.

Considering the generic cylindrical lens, it is clear that focussing effect of a lens is larger (for a given lens geometry) if the voltage applied to the lens is higher. Looking back to our picture with the 'bulging' field (in Figure 1.3), a larger 'bulge' will result in stronger lensing. Think of the equipotentials as forming the shape of a conventional light lens[5] - the more strongly curved the lens, the stronger the focussing power (and the shorter the focal length). This arises because for a very curved lens, the angles α are correspondingly larger than for a very flat lens. Decreasing the aperture diameter whilst holding the aperture voltage constant results in a larger 'bulge', and therefore a stronger lens. This is illustrated in Figure 1.5.

[5] Analogous as we have seen that potential V is proportional to n^2.

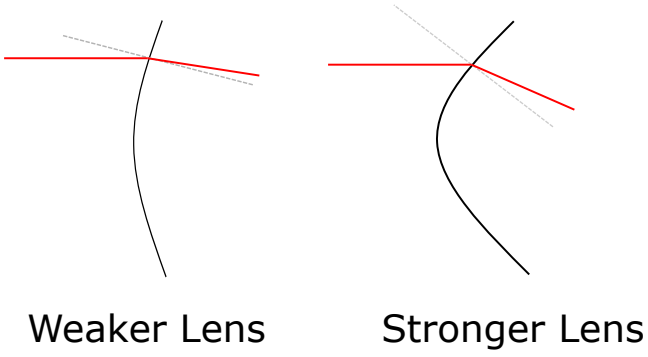

Figure 1.5: Illustration of lens action - more curved lenses are stronger as the angle of the incoming beam (red) to the surface normal (grey dashed) is larger than for a less curved lens.

Strictly speaking, the focussing power of a lens (focussing power $\propto \frac{1}{f}$ where f is the focal length), depends on the *difference in potential on either side of the lens*. In most cases, the lens is placed in a field-free region, so the potential around the lens is produced by the voltage applied to the lens. Clearly from the contour plots in Figure 1.3, the stronger the field, the stronger the lensing effect - because the difference between the potential caused by the lens and the field-free region (zero potential) is greater. Electrostatic lenses are always converging, as (in most cases) it is not possible to produce an equipotential pattern that is convex - as would be required for a diverging lens.

Generally, electrostatic lens systems are made up of a number of different lenses - in the same way that a camera lens contains many different individual lenses. There is one type of lens that it is commonly used in ion imaging, the so-called *einzel lens* or *unipotential lens*. This is normally a three-element lens, where the potentials of the outer two electrodes are zero (grounded), and the middle electrode has some voltage applied to it. This lens has the useful property that it focusses the beam of charged particles without changing their energy. This kind of lens is always converging, for the examples we will use. There are many examples and designs of einzel lenses that can be found online or in books.

A useful approach to the mathematics of electrostatic lenses is made by using *transfer matrices*. These matrices are analogous to the $ABCD$ matrices that may be familiar from light optics. Essentially the idea is that each element of an optical system can be described by a matrix, which effects a transformation on a vector which describes the initial position r (perpendicular to the optical axis) and propagation angle α (relative to the optical axis) of a particle, for example as:

$$\begin{pmatrix} r_{\mathrm{Out}} \\ \alpha_{\mathrm{Out}} \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} r_{\mathrm{In}} \\ \alpha_{\mathrm{In}} \end{pmatrix} \tag{1.2}$$

This kind of formalism should be entirely familiar from courses on transformations of vectors by matrices. In this case, the input and output vectors are obivously labelled, and the matrix $ABCD$ is effecting some transformation - most texts on electrostatic optics will give lists of transfer matrices for various kinds of optical elements.

The real power in the transfer matrix method is that it allows the effects of multiple lenses to be decoupled from each other. For example, I would have an input vector $\boldsymbol{v}$ that passes through several optical elements denoted by transfer matrices $\mathrm{M}_1, \mathrm{M}_2, \mathrm{M}_3$. Then my output vector $\boldsymbol{v'}$ would be given by Equation 1.3.

$$\boldsymbol{v'} = \mathrm{M}_3\mathrm{M}_2\mathrm{M}_1\boldsymbol{v} \tag{1.3}$$

Then, the action of individual elements, or simply removing one element, can be easily calculated by changing which matrices are applied. The key point is that it **decouples** focussing elements from each other. In our case, we will therefore treat

the extractor-repeller region as distinct from any lenses further down the flight tube. This simplifies analysis a great deal, but is also clearly a simplification, but will generally hold provided that the fields from adjacent elements don't interfere with one another.

The derivations above have all been made under the **paraxial approximation** - this may be familiar from light optics, but it essentially is an assumption that any ray (or particle trajectory) makes a small angle to the **optical axis** - defined here as the ToF (X) axis. Generally this angle is small enough such that $\sin(\alpha) \approx \alpha$. When the angle is larger, the paraxial approximation is not valid and generally aberrations (imaging errors) can become severe. However, we will see that many VMI spectrometers currently in use clearly break this rule of thumb - in general you will find that many texts and papers on electrostatic optics are vastly more concerned with aberration control than we seem to be in doing VMI. This is because in VMI, we are generally only looking at a single ion, or a very small range, and so the effect of various aberrations (see later) is often not severe over the image we look at[6], or are actively desirable REF EPPINK AND PARKER. Most work on electrostatic optics is in fields like electron microscopy or ion beam manipulation, where aberrations have a much more deletrious effect on the experiments than in VMI.

[6] For an analogy to light optics, if you're only looking at the blue parts of an image, chromatic aberration will never seem to be an issue.

This said, it is obviously best to avoid undesired aberrations if you can. The next section details some common aberrations and how they appear in electrostatic optics. For now, here is a set of 'rules-of-thumb' when designing ion optics (with thanks to Steve Thompson!):

- Ensure that the beam energy (i.e. particle charge $\times$ repeller voltage) is greater than any lens voltages - or you will create an electrostatic mirror, rather than a lens!

- Keep the beam angle to below 10 mrad where possible, to ensure the paraxial equation is valid and to minimise aberrations. This is often not possible in VMI, however.

- Generally try to keep the lens thick in relation to the spacing between the lenses - i.e. $L < T$ where L is the distance between adjacent lenses and T is the lens thickness. This ensures that focussing happens inside the lenses, and not outside (JDP IS THIS RIGHT?)

- Try to keep the ions to within around 10% of D, where D is the aperture diameter, to control aberrations. This is entirely analogous to light optics, as any photographer will attest to - the edges of the lens are where it performs worst.

We will see that VMI ends up breaking a lot of these rules, as sometimes the aberrations are actually desirable - but it's important to have them in your mind if trying to design improvements (even if you end up breaking them all eventually..).

1.3.2 Aberrations

Having discussed aberration control previously, it makes sense to define what aberrations actually are. Simply, an aberration is just an 'imaging error' caused by imperfections in the optics. There are two aberrations that it is useful to know about in this context, and only one that we will really focus on. These are **spherical aberration** and **chromatic aberration**. You may already be familiar with these from light optics, and their behaviour in electrostatic optics is pretty much entirely analogous to this.

Spherical aberration can be simply described as *'rays travelling through different parts of the lens focus to different points'* - and is illustrated in Figure 1.6. In electrostatic optics, this is entirely analogous, except rather than a light ray we use a particle trajectory. Clearly, then, ensuring that most of the trajectories travel through the middle of the lens would be 'best practice' and minimise spherical aberrations.

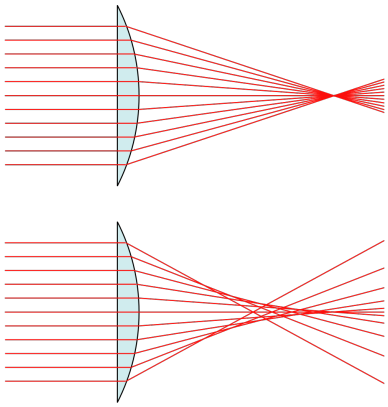

Figure 1.6: Spherical Aberration - the top lens exhibits no spherical aberration, the bottom lens exhibits severe spherical aberration. Figure from Wikipedia.

Chromatic aberration is something we will be (marginally) more interested in, and can be simply described as *'rays of different colours travelling through the same point of a lens focussing to different points'*, as illustrated Figure 1.7 (also from wiki). We will go further and split chromatic aberration into two kinds - *transverse chromatic aberration* (TCA), and *longitudinal chromatic aberration* (LCA). These cases are illustrated in Figure 1.7. Clearly TCA is where different colours are consistently focussed to different radial points on the detector[7]. LCA is where different colours have different focal points along the optical (X) axis.

[7] This is basically what we want to do with VMI in the first place!

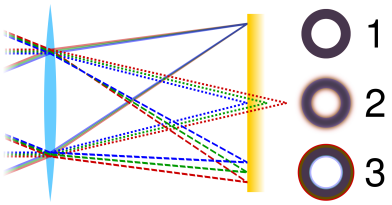

Figure 1.7: Chromatic Aberration. Image 1 exhibits no chromatic aberration; image 2 exhibits only longitudinal chromatic aberration; image 3 exhibits only transverse chromatic aberration. Figure from Wikipedia.

The analogy to electrostatic optics in this case is that the colour of a light ray is related to the photon energy, and so the particle energy is the analogue of colour in this case. LCA is most deletrious for VMI, as it means that different photon energies will be focussed at different points along the ToF axis (this corresponds to a curvature of the Fourier plane - see later). This means that a set of voltages that produce a well-focussed image of a 10 eV ion may not produce a well focussed image of a 20 eV ion. If a wide range of energies are desired to be focussed simultaneously, this aberration clearly is an issue. However, if only a single ion channel (with a single energy) is of interest, then it is always possible to find voltages that provide 'ideal' focussing for that channel. As we will see, many VMI spectrometers have LCA that would be considered 'bad' in most imaging systems, but we will also see that there are many other external factors that reduce our possible resolution, such that that caused by LCA is generally not the most pressing. LCA can be quantified using the chromatic aberration coefficent C_c, which is defined Equation 1.4.

$$\Delta X_0 = C_c \frac{\Delta E}{E_0} \tag{1.4}$$

Where ΔX_0 is the longitudinal distance along the optical (ToF) axis between the focal points for ions with an energy range ΔE. E_0 is the energy of an ion in the centre of this range. A plot of ΔX_0 against $\frac{\Delta E}{E_0}$ should therefore be linear, with a gradient C_c. This should make intuitive sense - it implies that for a given range of particle eneriges ΔE, the chromatic aberration is worse if ΔX_0 is larger.

I mentioned also, and it is clear from the figure, that TCA looks quite a lot like what we actually want VMI to do - to map different energies onto different parts of the detector. This links to all the stuff I implied earlier about the paraxial approximation breaking down - and this is documented in the literature REF EPPINK. In one sense, VMI requires that the paraxial approximation breaks down in order to map ions with different energies to different points on the detector - which would be considered severe TCA in conventional ion optics. This is one way to approach thinking about VMI, and helps to justify how almost every rule

of 'best practice' I have listed is broken by almost every VMI spectrometer in existence.

We will keep aberration control in mind when designing things, but also be aware that the system we are trying to design is not one that would be considered 'good' by a lot of imaging standards. We are not trying to focus a paraxial beam, so a lot of the aberrations and rules of thumb that apply to these systems are not appropriate for what we plan to do. Nevertheless, there is still value in understanding *what* 'rules' you're breaking, and *why* you break them.

1.3.3 Analogies between Electrostatic and Light Optics

We have already seen that $\sqrt{V}$ in electrostatic optics plays the role of refractive index n in light optics. Here we summarise some of the more useful analogies and key differences.

- Useful Analogies
  - $\sqrt{V} \rightarrow n$
  - The Lensmaker's Equation (Equation 1.5) from classical optics will prove interesting. Here I deliberately omit some detail which only applies to light lenses under certain circumstances, but the overall form of the equation is most relevant. f is the focal length of the lens, P is the *optical power* of the lens, n is refractive index, and R is the radius of curvature of the lens.

    $$\frac{1}{f} = P \propto (n) \times \left(\frac{1}{R}\right) \tag{1.5}$$

    A stronger lens (higher P) will have a short focal length. Clearly, a higher refractive index results in a stronger lens (i.e. flint glass vs crown glass, for the glass aficionados among you). Equally, a *smaller radius of curvature* - i.e. a more obviously curved lens - will produce a stronger lens. This is exactly as we have already seen in electrostatic optics - where a higher voltage applied to a lens resulted in stronger focussing action, and a smaller aperture (more 'bulging' field) resulted in a stronger focussing action. Overall, then, **higher lens voltage** $\rightarrow$ **stronger lens**, and **smaller lens aperture** $\rightarrow$ **stronger lens.**
  - **Colour of light** $\rightarrow$ **particle energy**. Colour of light is just the photon energy, so this makes sense.
- Key Differences
  - Generally, you cannot make a diverging electrostatic lens (at least when the lens is bounded by regions of constant potential). This means that

a simple achromatic doublet (as familiar in light optics) is not so simple in electrostatic optics.

- Aspheric lenses are much more difficult to produce, and generally require use of multipole geometries. For all of the work we will do here, the lenses will be of the simple cylindrical type - which have the advantage of being much easier to produce for workshops. Some VMI spectrometers have been built with interestingly shaped electrodes, which will be discussed in due course. That we cannot easily produce aspheric lenses puts limits on how effectively we can control aberrations. Generally in light optics aberrations are corrected using a wide array of aspheric lenses.
- The range of possible particle energies is absolutely massive in comparison to the range of photon energies generally used in light optics (particle energies from 0.1 eV to 100 eV are not uncommon, whereas the visible spectrum is only around 1 eV to 3 eV). This means electrostatic systems have to be designed to work for a wider range of acceptance energies, which fortunately is possible as it is possible to vary the voltage applied to an electrode over a *much* wider ranger than it is possible to vary the refractive index of a glass.
- As mentioned at the end of the previous section - we are not (in VMI) trying to focus a paraxial beam. So it is important to be aware that a lot of the concepts and metrics that are derived from the paraxial approximation *may* not be appropriate for VMI.

Having discussed some basic electrostatic optics, let us now try to discuss how VMI works with this analogy in mind, and try to develop an intuitive useful model we can use to guide our thinking later

1.3.4 Fourier Optics $\rightarrow$ Velocity Map Imaging

Fourier optics is an alternative way (compared to ray optics) to analyse and understand light optics using Fourier transforms. Ray optics (or *geometrical optics*) treats the light as simply 'lines' that propagate through a system and change their angles (wrt the optical axis) as they undergo refraction. Ray optics does not account for the wave nature of light, so cannot explain phenomena like diffraction or interference. Fourier optics accounts for the wave nature of light, and so can explain these phenomena. The drawback of Fourier optics is that it is not as intuitive as ray optics.

To see how Fourier optics works, it's worth considering what a Fourier transform really does, and what it physically means. To do this, let's first look at a

generic *Fourier series* of a function $f(t)$.

$$f(t) = \sum_{\omega} c_n(\omega) \cdot \exp(i\omega t) \tag{1.6}$$

From this expression, it is clear that the term $\exp(i\omega t)$ is simply a **plane wave**, which can be understood by writing the exponent in terms of sine and cosine. The plane wave this defines has an angular frequency ω, and therefore has an energy $\hbar\omega$. The $c_n(\omega)$ term can be simply regarded as an expansion coefficient, which varies depending on the exact ω in the sum. What we have effectively done, therefore, with this Fourier series, is *expand $f(t)$ as a sum of different plane waves of different frequencies*. Alternatively, you could think of it as expanding $f(t)$ in *different basis set*, the basis set of frequencies (in the same way you can write vectors as a sum of different components of unit vectors). So we are taking a function of time, and writing it as a superposition of waves with different frequencies - this should be familiar behaviour of a fourier series.

The Fourier *transform*, as opposed to *series*, is defined analogously but as an integral rather than a sum. For example:

$$f(t) = \int_{-\infty}^{\infty} c(\omega) \cdot \exp(i\omega t)\, \mathrm{d}\omega \tag{1.7}$$

An obvious question is what the form of the function $c(\omega)$ is. It turns out that we can write:

$$f(t) = \int_{-\infty}^{\infty} F(\omega) \cdot \exp(i\omega t)\, \mathrm{d}\omega \tag{1.8}$$

Where:

$$F(\omega) = \int_{-\infty}^{\infty} f(t) \cdot \exp(-i\omega t)\, \mathrm{d}t \tag{1.9}$$

These last two equations are the most common statements of the Fourier transform. Generally, Equation 1.8 is defined as the **inverse Fourier transform**, whereas Equation 1.9 is defined as the **Fourier transform**. The two functions f and F form a **Fourier transform pair**. Again, the key point for our discussion here is to note that essentially what this achieves is writing the function $f(t)$ in terms of it's corresponding frequency (ω) components, in the same way that one might write a vector as a sum of different components. Fourier optics, then, is just an approach to analysing the behaviour of light optics by treating the light as a combination of different plane waves with different frequencies - as opposed to geometrical optics (or 'ray optics'), where light is treated as a line that propagates in the same direction of a light wave. This is why geometrical optics cannot account for phenomena such as diffraction or interference - which arise due to different components of the light interacting in a way that depends on their relative phase.

There is no reason why the mathematical machinery of the Fourier transform is limited to transformations between the time and frequency domains (although this is perhaps the most common use). Infact, the variables **momentum** and **position** form another Fourier transform pair[8]. There are lots of nice analogies between this and the Uncertainty Principle (note that time and energy (frequency $\times\hbar$) have an Uncertainty relation), but these are beyond the scope of what we do here. That position x and momentum p form a Fourier transform pair is interesting, as this can help us understand the action of a VMI spectrometer.

[8] In general, the units of the two variables in the pair must be reciprocal. This is true for position and momentum to within a factor of $\hbar$. If position is measured in m, then momentum is measured in kg m s^{-1}. Momentum divided by $\hbar$ will have units of m^{-1}. We are free to do this as the Fourier transform doesn't specify the scales of either of our variables - so we are fine to just divide out momentum by $\hbar$ in this way.

A useful starting point here is to consider diffraction by a prism. This disperses the frequency components of the light, and is effect performing a 1D Fourier transform of the incoming light (separating it out into different frequencies). This is illustrated in Figure 1.8.

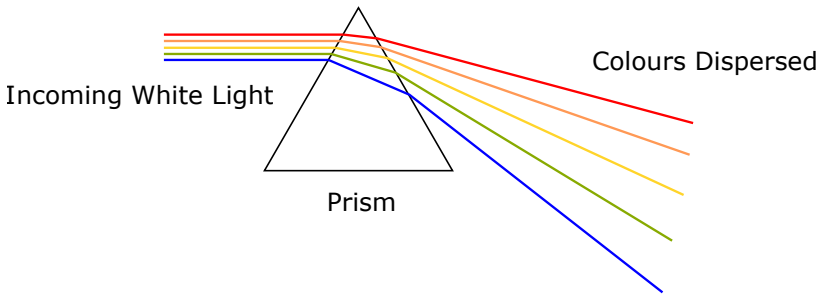

Figure 1.8: Diffraction by a prism. Incoming white light is separated into different frequencies (colours). This is performing a 1D Fourier transform on the incoming light, decomposing it into it's different frequency components.

We can extend this argument and consider diffraction through a slit, or small hole. This produces well known diffraction patterns, and the maxima and minima arise from constructive and destructive interference of plane waves with different phases. The plane waves all enter the slit at different angles θ, and the separation of the peaks in the diffraction pattern is proportional to $1/d$ where d is the size of the slit (i.e. smaller slit leads to stronger diffraction). $1/d$ will have units of inverse length, and is equal to $\sin\theta/\lambda$[9]. This should also intuitively make sense - a wider range of input angles leads to a larger separation of peaks, and shorter wavelength waves diffract more strongly (why lenses tend to focus blue light harder than red light). The quantity $1/d$ is called the **spatial frequency** - that is, how often the peaks repeat per unit distance. Think of it as analogous to normal frequency (sometimes called *temporal frequency* to differentiate it)-

[9] As $d\sin\theta = \lambda$ for the first maxima.

that is the reciprocal of time (units of cycles per unit time). Spatial frequency is the reciprocal of position (units of cycles per unit length). ADD FIGURE OR SOMETHING

Fine, but how does all of this relate to Fourier transforms? Well, it turns out that the diffraction pattern (in the far-field, so Fraunhofer diffraction) is 'simply' the Fourier transform of the input field across the diffracting object. That is, from the diffraction pattern it is possible to reconstruct an image of the diffracting object[10] - just as is commonly done in X-ray crystallography or image analysis.

[10] If you have all of the information (phase and power spectrum) from the Fourier transform available.

Let us now spend some time getting bogged down in units and dimensions. We have seen that **position** and **momentum** form a Fourier transform pair, but I just basically implied that **position** and **spatial frequency** are a Fourier transform pair. This implies that spatial frequency and momentum are related - which they are! Clearly they have the same units (to within a multiple of $\hbar$), so if we are mapping spatial frequency onto a screen, we're also mapping momentum onto a screen. In the case of Fraunhofer diffraction, the momentum is the momentum of the light, as defined by it's direction (θ) and energy ($\propto \lambda$) - infact you could define spatial frequency as $\omega_s = \frac{theta}{\lambda}$[11]. What is happening is that the different momenta in the incoming wave are being mapped onto a screen, this screen is called the **Fourier plane**, and we have effected a transformation from **position space** to **Fourier space**, and in this case, **Fourier space** is really **momentum space**. Or **k-space** if you're into physics.

[11] In the paraxial approximation where $\sin(\theta) \approx \theta$.

So, to take stock, we have seen that diffraction can effect a Fourier transform, turning position space into momentum space. This diffraction can be performed by a lens, as is illustrated in Figure 1.9. The lower panel of this figure shows that a light beam made up of different plane waves each with different angles to the optical axis (i.e. each with different spatial frequencies) is dispersed by the lens such that each plane wave focusses to a specific point in the focal plane. To make an analogy to light optics, beams with higher **divergence** will focus to larger spots when focussed by a lens - as the higher divergence $\rightarrow$ larger angles θ $\rightarrow$ larger spread in the x direction in the focal plane $\rightarrow$ larger focal spot. We are now in a position to think about VMI as Fourier transform process.

In VMI, we take ions with a certain spatial distribution in a source region (the 'source plane'), and map their momentum distribution onto a detector (the Fourier plane). This transformation is effected by a lens, and provides an intuitive way to understand VMI. Really, the VMI spectrometer is a Fourier transformer - mapping the ion momenta onto a screen, irrespective of where they were formed. This is akin to decomposing the initial source distribution into it's various momentum components, and mapping them onto a sphere. **The VMI spectrometer is performing a Fourier transform on the initial source distribution, to extract the momentum distribution.**

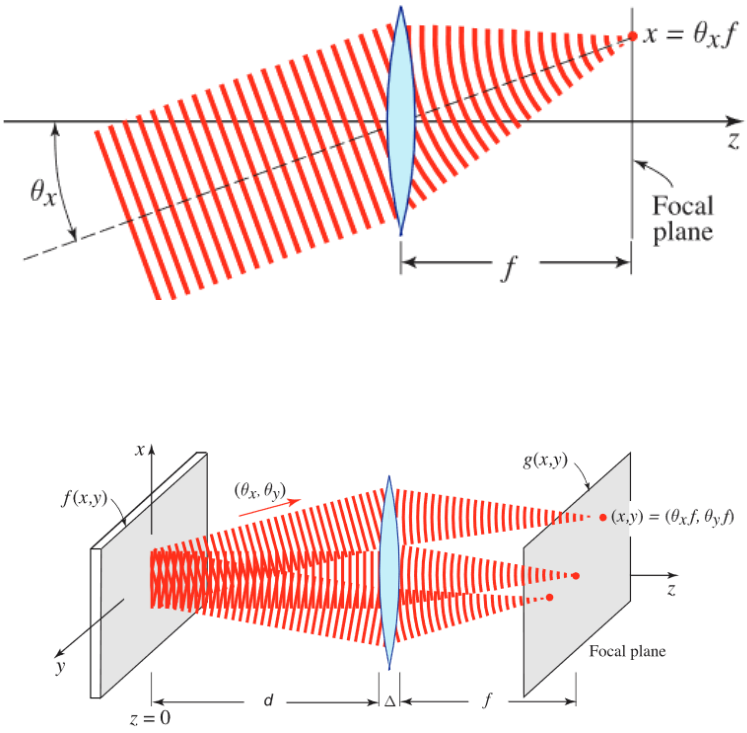

Figure 1.9: Top: A plane wave travelling through a lens at a given angle θ_x will be focussed onto a specific point x on the focal plane. Bottom: A light beam composed of many plane waves will result in many different spots on the focal plane - the lens decomposes the incoming beam into plane waves with different angles θ to the optical axis. These angles are essentially spatial frequencies (see above). Figure adapted from Saleh and Teich *'Fundamentals of Photonics'*.

To continue the analogy from above, we said that the momentum was proportional to $\sin\theta/\lambda$. For ions, θ is their emission angle away from the optical axis (if we consider a 1D case - as in the figure - for simplicity). And 'λ' in this case is inversely proportional to particle energy. Therefore, in our VMI image, we expect that the angle to the optical axis at which an ion is travelling to be mapped onto our detector[12], which it is. We also expect ions of higher energies to be mapped onto higher spatial frequencies - so to hit further out on the detector, which is also what is observed.

[12] Strictly it is mapped into a 3D velocity distribution, which is then mapped onto the detector.

The creation of this Fourier plane has interesting connotations - firstly in terms of magnifying the VMI image. We can (using the transfer-matrix method) treat

this Fourier plane as a 'new' source plane, and then use electrostatic lenses to 'zoom' in or out of this image - more later. Secondly, there is possibly information to be gained by applying the mathematics of Fourier theory to the VMI image. This is not something that I am aware of in the literature. For instance, in an analogous way to crystal diffraction it may be possible to extract information about the source distribution from the velocity-mapped image.

VMI Fourier Planes

Typically in VMI, the position of the ideal VMI focus along the ToF direction depends on the energy of the ion being measured. For a given set of voltages, generally the VMI focal point[13] for a high energy ion lies further away from the detector than for a low energy ion. This results in a *curvature* of the Fourier plane. This is illustrated in Figure 1.10.

[13] Defined as the position along the ToF direction where ions produced with the same momenta all arrive at the same point (or where the circle of confusion is smallest).

Figure 1.10 (a) shows a SIMION workbench where ions with energies from 3 eV to 30 eV are flown with the same velocity from different source positions. Each colour corresponds to a unique ion energy. The ions were born either 1 mm above or below the central ToF axis, this is drastically worse than a realistic source distribution (as in reality the laser spot is around 50 µm in diameter), but illustrates the idea of the Fourier plane well. Figure 1.10 (b) shows a zoom of the region near the detector with black circles placed at the VMI focal points for each ion energy. There is a clear and dramatic curvature of the plane away from the detector at higher energies. Figure 1.10 (c) shows the same zoom as in (b), but now the source positions are only 0.1 mm above or below the central ToF axis. This clearly makes the curvature of the Fourier plane all but invisible - showing that under realistic experimental conditions this effect is essentially negligible.

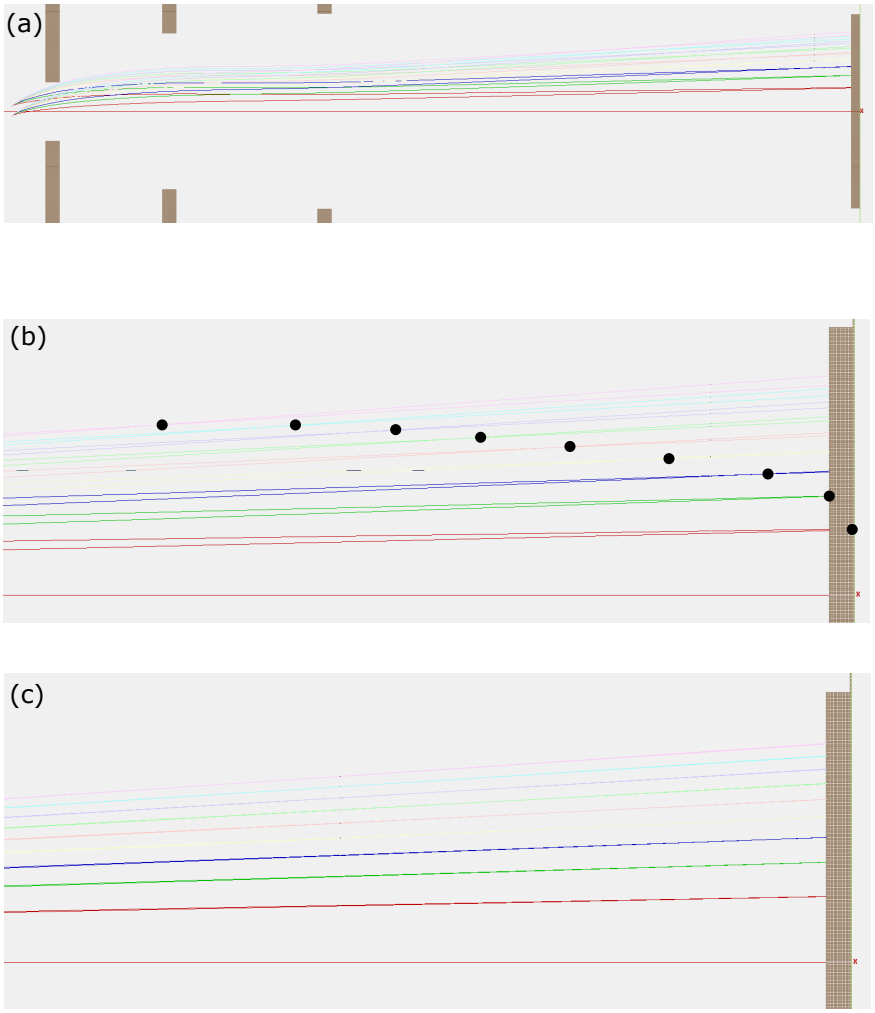

Figure 1.10: Illustrating the concept of a Fourier plane within VMI. (a) A workbench showing a test setup for visualising the Fourier plane. Ions are born 1 mm either side of the central ToF axis, with a range of energies (each colour is a unique energy). (b) A zoom of the region near the detector, points where isoenergetic ions from different source positions cross (VMI focal points) are marked with black circles. (c) The same zoom as in (b), but with a realistic source region size.

Chapter 2

Spectrometer Design

A VMI spectrometer for the new ultrafast imaging instrument in the MBB group needs to be designed and built. This spectrometer will initially be used for velocity-map imaging of the nascent ions from Coulomb explosion of neutral molecules, induced by femtosecond laser pulses. Ultimately it will also be used for imaging of ions from the explosion of larger molecules that are brought into the gas-phase via electrospray ionisation.

2.1 Requirements

The requirements for the spectrometer are fairly typical of many similar instruments. They are outlined below.

- Physical Constraints
  - Spectrometer must fit inside the planned interaction chamber (CF200 - 200 mm ID). As such the diameter of the plates cannot be larger than around 180 mm to avoid arcing to the chamber.
  - Spectrometer must achieve metrics (below) using a 40 mm effective area MCP detector.
- Experimental Capabilities
  1. Spectrometer must be able to resolve energy differences on the order of 0.1 eV clearly.[1]
  2. Spectrometer must be able to clearly resolve mass peaks separated by 1 Da over a reasonable mass range - in practice this requires a $\frac{m}{\Delta m}$ resolution of at least 500, and probably substantially more.
  3. Spectrometer must be capable of operating in a 'slice-imaging' mode, where the Newton sphere is stretched along the ToF axis to facilitate slice imaging.

[1] Michael should be aware that this request is meaningless without saying 0.1 eV *in* something else. 0.1 eV out of 100 eV is very different to 0.1 eV out of 1 eV.

4. Good velocity mapping must be maintained down to small image radii, as imaging of large, slow-moving ions is expected ultimately.

Addition of a zoom lens to enable the lower-energy portions of the image to be magnified is also a desired addition (more later).

2.2 Definition of Metrics

Simulation of the spectrometer was performed using SIMION. Several metrics were mentioned above as ways to characterise the performance of the spectrometer. They are defined in more detail below.

2.2.1 Energy Resolution

[2]Here we use the term 'velocity-mapping' not 'momentum-mapping' to be consistent with previous literature. Although maybe people just doing this for consistency is what led to the current shitshow of nomenclature.

The fundamental principle of good velocity-mapping[2] is that an ion produced with a unique momentum vector in the detector plane (here the YZ-plane) should be 'mapped' to a unique radius on the detector. That is, that there must be a near-perfect correlation between the initial momentum vector in the detector plane v_{DET} of an ion, and the radius r_{DET} it hits on the detector. As such, the parameter used to assess how good the velocity mapping is is the squared Pearson coefficient R^2 of these variables, to assess the goodness-of-fit of a plot of v_{DET} and r_{DET}. Typically an R^2 value greater than 0.99 was considered good.

This parameter gives a measure of how good the velocity-mapping is, but doesn't quantify the actual *energy resolution* - i.e. how small of change in ion energies is resolvable. The energy resolution is conventionally defined as $\frac{\Delta E}{E}$ where E is the ion energy and ΔE is the spread of particle energies. This formula can be easily linked to the experimental observable R (image radius), by noting that $E \propto R^2$, such that $\frac{\Delta E}{E} = \frac{2\Delta R}{R}$. Thus the velocity resolution can be calculated by measuring the radius that a particle with a particular initial energy hits, and measuring the spread of different radii that a particle with that energy hits.

There are many definitions of ΔR given in literature - some define it simply as the range of the distribution of radii that a given particle energy hits (i.e. $R_{max} - R_{min}$) (Kling2014ThickLens). Others define it as the FWHM of the peak in a plot of ion yield vs radius (Marchetti2015Thiophene) of a pancaked ion distribution, and others use the standard deviation of a Gaussian fit to the peak profile. Still others produce a simulated ion image which is then subjected to an Abel transform to extract the 3D momentum distribution from the 2D image, from which an energy spectrum is plotted and the ΔE is obtained directly as the FWHM of a Gaussian fit to the spectrum (Rading). Conceivably, one could also simulate an ion image and then take the central time-slice in SIMION and find the width of this peak.

Of these methods, as far as I can tell the best method is to simulate an ion image, perform an inverse Abel transform and then fit the peaks in the resulting radial distribution. Initial attempts demonstrated to me that not accounting for the projection of the 3D Newton sphere onto the detector meant that finding an accurate peak width was difficult and likely to cause an inaccurate value for the resolution (especially as generally experimental images are inverted anyway). Finding the central slice in SIMION is straightforward in principle, but requires that you are able to pick out the central arrival time from the distribution of ions along the ToF axis. This is, in itself, trivial - but doing this in a computationally robust way would require that you deconvolute the projection of the ellipsoid Newton 'sphere' onto the detector from the arrival time distribution, which is non-trivial (at least for a mediocre mathematician like me).

The method used here was to simulate an ion image, perform an inverse Abel transform using PyAbel, and then integrate over the angular coordinates to produce a distribution of ion yield against radius for the inverted image. Then the widths and locations of the peaks in this distribution can be reliably determined using a peak picking algorithm. Fitting each peak to a Gaussian is also possible, but it seemed to me unnecessary to add this extra possible source of error when there are sufficient statistics to determine the peak widths directly (as the data are simulated anyway).

2.2.2 Mass Resolution

Mass resolution is conventionally defined as $\frac{m}{\Delta m}$, where m is the mass of the particle and Δm is the uncertainty with which m is determined. Higher values of this ratio correspond to better mass resolution (as Δm is smaller). This is equivalent to $\frac{t}{2\Delta t}$ as $m \propto t^2$. Thus, the mass resolution can be measured by taking an arrival time distribution and finding the mean t and standard deviation Δt - approximating the distribution as a Gaussian is normally appropriate.

While this is generally how mass resolution is defined in more traditional imaging mass spectrometry experiments, there is a difficulty in applying this definition to a VMI experiment. The arrival times of the ions in VMI do not immediately follow a straightforward Gaussian distribution. In reality, the Gaussian is convoluted with the projection (along the time axis) of the ellipsoid 'Newton Sphere' onto the detector. Deconvolution of this ellipsoid projection in a rigorous way should be possible, but lots of time has not been spent on this here[3]. Mostly, we will find a setup which gives good velocity resolution and then confirm empirically through SIMION that the mass resolution is within the required range. The bottom line is that we need to be able to have isotopic mass resolution - but as almost every VMI spectrometer I've seen has this, I don't think it is a critical parameter to work on, and it can always be improved by adding (cheap) flight tube such that the ions diverge in time more before they hit the detector.

[3] Somebody better at maths than me can probably do this in about 10 seconds.

2.3 'Experimental' Setup

All simulations were performed in SIMION 8.1.1.32. Analysis was performed either using a Lua workbench program or Python.

2.3.1 Workbench

The backbone of the workbench was a grounded tube of internal diameter 200 mm, reflecting the CF200 chamber. The ion optics were contained within this tube, and a 40 mm MCP detector was mounted in one of two places. Either mounted at the end of the interaction chamber (to simulate having no additional flight tube), or mounted at the end of a 500 mm flight tube. Having this flexibility in the flight tube greatly extends the usable range of particle energies that can be imaged, as will be seen later. The lengths of the flight tube was arbitrarily chosen to be a length that will be easy to machine in a workshop. It is not critical, however, as any slight (mm) deviations from the SIMION geometry can be easily corrected by using slightly different voltages.

The coordinate system has already been defined. The ToF axis is the X-axis, the molecular beam also propagates along the X-axis, and the laser propagates along the Y-axis, such that the XZ-plane is the polarisation plane. The YZ-plane is the detection plane, and it is velocity within this plane that will be measured.

The workbench geometry was defined in a .gem file, which provides a flexible and repeatable way to define the electrode geometries. SIMION creates potential arrays ('PA') built up of a number of 'grid points', and each point stores geometry and potential information. By applying different geometry and potential information to every grid point, the full workbench is built up. The scaling between real space (millimetres) and grid points (grid units) is (by default) set to $1\,\mathrm{mm\,gu^{-1}}$. However, even with electrode surface enhancement (see manual) turned on, this does not allow sub-mm or fractional mm spacings to be defined. Therefore, throughout these simulations we instead use $10\,\mathrm{mm\,gu^{-1}}$, for increased precision (at the cost of a larger and slower simulation).

2.3.2 Particles

SIMION requires that the ions to be flown in the simulation are defined in a .fly2 file before running. As there are several metrics that need to be assessed, it is necessary to define multiple sets of ions to rigorously explore the capabilities of a spectrometer. As such, I define multiple sets of ions to determine the velocity resolution, mass resolution, and to assess how it might behave under typical experimental conditions. In this section I detail the different sets of initial ions, but there are some parameters common to all the ions.

2.3.3 Common Parameters

Ionisation Source

All ions are assumed to be created instantaneously and at the same time (reflecting the femtosecond timescale ionisation). The shape of the region that the ions are produced in is modelled by a 3D Gaussian function, reflecting the intersection of the laser focus and molecular beam. This Gaussian is located at the center of the interaction chamber (internal workbench coordinates $(175, 0, 0)$), reflecting the location of the windows.

The laser is assumed to be an 800 nm Gaussian beam of initial diameter $(1/e^2)$ 6 mm that is focussed by a 300 mm lens to produce a focal spot size of 25 µm×25 µm ($\omega_0 \times \omega_0$). This is fairly typical of Coulomb explosion imaging experiments. The beam waist ω_0 is simply the point at which the transverse intensity in the the xy-plane (polarisation plane) drops to $1/e^2$ it's maximum value (assuming a circular beam profile), and is measured at the point at which the beam is smallest. This is the 2σ width, so half of the ω_0 value defines the standard deviation of the 3D Gaussian in the xy-plane.

The width of the 3D Gaussian along the propagation axis (x) of the laser beam can be estimated by taking into account the variation in spot size along the propagation direction of the beam, and the width of the molecular beam along this axis. For a Gaussian beam with beam waist ω the intensity varies along x as:

$$\omega(x) = \omega_0 \sqrt{1 + \left(\frac{x - x_0}{x_R}\right)^2} \tag{2.1}$$

Where x_0 is the location of the focus, and x_R is the *Rayleigh Length*, given by the location where the spot size is $\sqrt{2}$ larger than at the focus[4], such that:

[4] i.e. $\sqrt{2}$ larger than ω_0.

$$x_R = \frac{\pi \omega(x_0)^2}{\lambda} \tag{2.2}$$

Where λ is the wavelength of the light. The beam parameters given create a Rayleigh length of x_R =2.4 mm. This would give a standard deviation of 1.2 mm for the 3D Gaussian if it was interacting with a uniform region of gas density.

However, the finite width of the molecular beam means that not the entire Rayleigh length is able to produce ions. The gas pulse has passed through a 1 mm skimmer, and is assumed to be collimated such that the gas density in the x-direction is described by a Gaussian function with $(1/e^2)$ width (2σ) of 1 mm. To account for possible non-uniform collimation, the standard deviation of the Gaussian describing the region of ion production in the x-direction was set to 0.75 mm. The parameters describing the 3D Gaussian describing the region of ion production are summarised below.

Coordinate	Mean (mm)	Standard Deviation (mm)
x	175	0.75
y	0	0.0125
z	0	0.0125

As will be seen in due course, the closer the ionisation source is to being a point source, the better the velocity mapping performance. One of the main advantages of 'fancier' designs incorporating aspheric electrodes (or similar) is that they allow a much larger ionisation source to produce good spectrometer performance. The ionisation source chosen here was chosen to be realistic, but errs on the side of being large - so we are hopefully getting a realistic 'minimum' performance from simulations (or at least, the real interaction volume should be smaller than this - and could be easily made smaller than this by focussing the beam harder).

Ions

SIMION allows the definition of the ions flown in each simulation in a .fly2 file. This all relevant parameters (positions, velocity vectors, energies, masses, charges) for all flown ions to be set in a robust and flexible way.

For general testing, a set of 1000 ions with a mass of 28 Da and charge of 1 e was used (this being a pretty 'average' ion that falls in the middle of the general usable range). This number of ions is sufficient to determine the R^2 coefficient for a given geometry. For more refined tests a variety of ion masses were flown - detailed in the relevant results section. In general when testing for velocity resolutions a large number of ions need to be flown to perform a reliable Abel inversion - normally 20000.

2.4 Geometry Optimisation

In designing a VMI spectrometer, there is a very large parameter space that can be explored. The parameters that can be tuned include the geometries of each electrode (generally the outer and aperture diameters for conventional 'disc' shaped electrodes - but also other geometrical parameters for differently shaped electrodes); the spacings between neighbouring electrodes; the voltages applied to each electrode; and the overall geometry of the spectrometer. Many of these parameters are interdependent - by analogy with classical optics it is clear that moving a lens has a similar effect to changing the strength of the lens - so thoroughly exploring this parameter space is challenging.

2.4.1 Genetic Algorithms - yay or nay?

An initial way to try to get a feel for this parameter space is to employ a genetic algorithm. This algorithm creates a wide number of different spectrometer designs and tests them all against a certain metric - the metrics are then compared and the 'fittest' designs are then used to generate new designs, in a manner akin to biological selection. This process is iterated until a fittest design is produced. The optimised metric in the simulations was the goodness of fit between the initial velocity vector and the final position on the detector, as discussed previously. In tests 1000 ions were flown per geometry to be optimised, which is sufficient to establish the R^2 parameter and see how good the geometry is.

I haven't attempted to plot and document the result of every single GA run, as this would be very dense (and some the data failed to save due to some of SIMION's many quirks). However, I want to summarise the results of my testing here and document some useful/interesting results. This is mostly based on many days/weeks of running things to get a feel for how it works - I wasn't attempting to rigorously document things at this point.

- Trying to get the GA to optimise electrode spacings, apertures, and voltages is a **very** computationally intensive task. Having four electrodes (from the Townsend design) means that there are four apertures, five spacings, and four voltages to iterate through (keeping all plates the same thickness). Optimising these with millimeter (not even sub-millimeter) precision creates an enormous parameter space that takes a very long time to iterate through. Refining a PA with appropriate grid unit precision[5] can take up to 60 s - when added to the time for ions to fly and data to be analysed/written, each iteration can take between two and 5 minutes. To thoroughly explore the parameter space would require many hundreds, if not thousands, of iterations to be run - and the computational time quickly becomes impractical.

[5] Normally 10 grid units per mm.

- However, an obvious rebuttal to the previous point is: '*If you can just let it run, why not just leave it and let it work it out?*'. My testing has led me to the conclusion that the results of the GA are **not particularly reliable**. That is, you can run the same parameter set (with a very large number of iterations), and produce different (good) results every time. This is, I think, due to the following reasons:

  - The #1 reason: **many different geometries can produce good velocity mapping if appropriate voltages are set**. That is, if you allow the voltages *and* the geometry to vary (and voltages are always variable in a lab), then many geometries can give good spectrometer performance. **There is no one perfect geometry.** This intuitively makes sense (see below).

  - Thinking about ion optics and the effect of electrode apertures/voltages on the lensing effects (from the first chapter) - clearly a lot of these effects are interdependent. For instance, you could make an aperture smaller and make an electrode a stronger lens - or you could simply increase the voltage and keep the aperture the same size. Both of these could have an equivalent effect on the field lines. Optimising them together therefore seems inefficient.
  - The parameter space is sufficiently large that even a very large simulation fails to adequately cover it.

- To illustrate this somewhat, here is some anecdotal[6] evidence of things that the GA has optimised:
  - Starting with the DC slicing geometry, running the algorithm twice with a wide parameter space and a large number of iterations gave two different geometries, each of which were subtly different from the starting geometry. The same performance could have been achieved by taking the starting geometry and just optimising voltages.
  - Running the GA to determine the best possible size and shape for a spherical extractor electrode concluded that the best possible shape was just an annular disc (not spherical). Running it again produced the result that a slightly curved electrode produced essentially the same performance.
  - Allowing more flexibility in electrode spacings resulted in the 'best' VMI occurring when two of the plates overlapped.
- Clearly some of, or all of, the above problems could result from user error. Clearly letting the GA produce an unphysical geometry is poor programming on my part, however I think it serves to illustrate my main point which is that **you should not just run the GA and think it will optimise everything for you.**
- Perhaps doing a more refined simulation where many parameters (not just the R^2 value) are optimised simultaneously would lead to better results. But this seems like even more of a 'ram it all in and wait' type of approach, so I am sceptical.

[6] Feel free to run them yourself if you want the actual numbers.

With all this in mind, I would say that I **do not think that the genetic algorithm is an appropriate tool for this kind of simulation**. I think a better approach is to be guided more by the principles of electrostatic optics, and to try and optimise things 'by hand' as best you can. The fact that many geometries produce good performance means that the GA may well find you a good geometry - but it will stop you learning about how these things *actually work* if you accept what it tells you unquestioningly. My personal preference is to optimise geometry

manually and then run a simple simplex optimiser (there is an inbuilt SIMION library for this) to optimise the voltages on the plates for a given geometry. This is a much more 'hands on' approach and I think it gives a much better feel for how these electrostatic optics really work.

This is not to say that the GA is useless - genetic algorthims are widely used in many areas of the natural sciences. Within spectrometer design, I can imagine it being useful in design of a very non-standard spectrometer, but *only if used sparingly to optimise small parts of the geometry*. Trying to ram everything into it and optimise everything at once is, I think, a quite lazy approach to this kind of problem. Overall then, in answer to the section heading - I would say a strong 'nay', at least to this specific problem.

2.4.2 General Observations Regarding Geometry

It is also worth stating here about the link between spectrometer geometry and overall performance. Extensive simulation suggests to me several things[7]

[7] Again this was mostly all learnt during my 'playing around with it' phase and isn't rigorously documented.

- It is almost always possible to make any geometry produce good velocity-mapping for a single ion species with a single energy, by choosing appropriate voltages for all electrodes.
- Geometry of the spectrometer has more influence on the quality of velocity-mapping over a *wider range of particle energies* than the absolute energy resolution possible. That is, a 'good' spectrometer geometry will produce good VMI over a wider range of particle energies (using a constant set of voltages) than a 'poor' spectrometer geometry. Therefore, one can expect to get better VMI over the entire image using a better spectrometer geometry, as the curvature of the Fourier plane is lower. This means that a wider range of initial particle energies are accurately mapped (LCA is lower).
- This being said, it seems that almost any array of annular disc electrodes can produce VMI that is as good as that obtained in a 'standard' spectrometer (e.g. Eppink + Parker), provided the voltages are correctly chosen such that the Fourier plane of a certain ion coincides with the detector. More complex designs involving curved, aspheric, or conical electrodes (Marchetti/Wrede) can reduce the curvature of the Fourier plane and produce good VMI across a wider portion of the image, or can mean that a larger interaction region can be used to produce good VMI (in general if the interaction region is larger, then the VMI mapping is worse).
- However, design of these electrodes can be complex - and in most cases is performed with professional assistance from companies that specialise in electrostatic lenses. In addition to this, for the experiments planned in the MBB group, the limiting factor will **not** be the spectrometer performance.

These high-resolution spectrometers (see, for example, the Velocitas 'Double Prime' spectrometer), find most applicability in narrowband experiments where exceptional energy resolution is needed - velocity resolution of around 1% is attainable, and at this point is limited by the properties of the molecular beam. In contrast, the typical designs I have optimised tend to have velocity resolution of $<5\%$, and using a non-ideal[8] molecular beam, and a broadband laser source mean that it is unlikely we will be limited by this.

[8] Not an Even-Lavie Valve.

- Reducing the curvature of the Fourier plane is obviously a good thing to do, but again I think the benefit gained by doing it would be pretty marginal. In the end slight asymmetry in the manufacturing, and the finite pixel resolution of the imaging camera will mean that the difference in focussing caused by the curvature will not really be visible in a realistic experiment.

The above considerations lead me to think that the best thing to do is to use an existing design (such as the DC slicing design), and then see if it meets the requirements we have set. Whilst this probably isn't producing the theoretically best spectrometer anyone could possibly design, it will most likely be more than sufficient for our needs, and will certainly have a very good resolution over a reasonable range (if not the largest possible range).

2.4.3 Getting SIMION To Do What You Want

As I would not wish the 16 weeks of rage-inducing SIMION learning which made me continually want to rend down my facial features in hot quicklime rather than continue being employed on anyone, I will try to document some tips and tricks that I've learnt here.

- SIMION is very logical, but also hates you.
- Spending time learning how to use Lua for user programming is the best use of time in the beginning. You can make SIMION produce a working spectrometer very quickly (I did this in one day), but you need to produce that spectrometer programatically if you want to do serious optimisation.
- Learn to use .gem files rather than the 'modify' screen to design potential arrays. While it seems a bit daunting (and the documentation is a bit hit and miss), it is much easier ultimately to be able to change spacings and things in a program rather than by counting pixels on a screen.
- SIMION's inbuilt data recording works, but I found recording all the data via a user program to be much more intuitive and easy. You can then relatively easily interface the Lua script with python which is **much** better than Lua for data analysis and plotting. My preferred workflow was: Fly ions $\rightarrow$ write data to a temporary file $\rightarrow$ load file into Python for processing and plotting.

- A Lua script you write for a user program is not executed line-by-line in the way you might expect. The functions in each segment are executed in the order defined by the SIMION Flow Diagram. This caused confusion to me initially and other students, as not every SIMION variable is accessible in every function - so it is worth getting the manual (for some reason paper copies only..) to see the scope that each variable has.

- For some reason on my MBB computer, SIMION wouldn't let me save and reload a workbench (giving an error about a corrupted PA). I never fixed this and nor could David Manura on the help forum - a workaround was to put your working SIMION directory in the MBB Q drive.

- The SIMION Users Group (on the website under 'community/support') is really helpful, and David Manura (developer) actually answers questions actively and helpfully. Use it. Simple questions also don't seem to get sarcastic responses from insecure physicists that you see on stackexchange.

- If you rename a workbench to the same name as an existing .lua file, then SIMION sometimes overwrites that .lua file with the .lua file that was associated with the previous workbench. This can be annoying if it happens to you.

- While the SIMION manual says that turning the trajectory quality down to zero is normally fine and helps speed, I found that it creates problems with VMI simulation. Leaving it at the default of 3 seems to work fine.

Chapter 3

Simulation Results

Here results from various simulations are summarised. Approximately 12 weeks were spent just becoming conversant in Lua and discovering SIMION's many foibles, so lots of initial work was subsequently found to be either invalid or not done in the best way possible. It is really only in the previous week or two before writing this document that I feel like I understand what I'm doing and can get SIMION to do what I want (and even then, not really). Hopefully eventually this document will serve as a template for the next person who is tasked with doing a big simulation like this...

Initial designs were based on the DC slicing spectrometer made by Townsend, Minitti, and Suits. Running this geometry into the genetic algorithm and allowing it to optimise things produces a geometry which is very similar to the initial geometry - which confirms my hypothesis that many geometries can give good performance, by just changing the applied voltages. So, for simplicity we stick with the geometry defined in REF. REF is unique in that it actually gives dimensions for the VMI spectrometer - which is not the case for many other papers! Particularly, the two more 'exotic' designs by Wrede REF and Marchetti REF give few dimensions in the literature - presumably as (in the Marchetti case) the spectrometer is sold commercially by Velocitas.

Before we embark on a discussion of the simulations and results, now is probably a time to restate our initial requirements for the spectrometer.

- Physical Constraints
  - Spectrometer must fit inside the planned interaction chamber (CF200 - 200 mm ID). As such the diameter of the plates cannot be larger than around 180 mm to avoid arcing to the chamber.
  - Spectrometer must achieve metrics (below) using a 40 mm effective area MCP detector.

- Experimental Capabilities

    1. We will rephrase the 0.1 eV resolution requirement to actually be that we need to be at around 3% $\mathrm{d}E/E$ - absurdly high end spectrometers quote performance of 'approaching 1%', so this is probably fine.
    2. Spectrometer must be able to clearly resolve mass peaks separated by 1 Da over a reasonable mass range - in practice this requires a $\frac{m}{\Delta m}$ resolution of at least 500, and probably substantially more.
    3. Spectrometer must be capable of operating in a 'slice-imaging' mode, where the Newton sphere is stretched along the ToF axis to facilitate slice imaging.
    4. Good velocity mapping must be maintained down to small image radii, as imaging of large, slow-moving ions is expected ultimately.

Both of the physical constraints are easily satisfied, so will not be discussed further beyond saying that a larger detector clearly increases the range of energies that can be imaged using one set of voltages - but beyond this the benefit is not clear to me. The other constraints (numbered 1 through 4) will be considered in turn for each design.

3.1 DC Slicing Design

The initial design chosen was that described by Townsend *et al* for Direct Current (DC) Slice Imaging. This design is an extension of the original Eppink and Parker VMI designs to allow the extracted Newton sphere to be stretched along the ToF axis, allowing a fast imaging camera to image multiple 'slices' of the Newton sphere[1]. Slicing in this way is beneficial as it extracts more information and can reduce the need for inversion algorithms to extract the central slice. This design doesn't used pulsed electric fields to perform the slicing, hence *direct current* slice imaging. This design is widely copied (or reproduced with superficial modifications, such as on the MB imaging rig) labs worldwide so seems a reasonable place to start as slice imaging is desired. A schematic of the design (taken from the original paper) is reproduced in Figure 3.1.

[1] Which is indubitably not a sphere after this process, but hey.

Having spent upwards of three months attempting to find improvements to this design (documented vaguely in the previous chapter), it became clear that this design gives as good performance as almost any other design when appropriate voltages are chosen. With this in mind it is not surprising that almost every lab doing this kind of science uses a variation on this sort of design[2]. One change we will make is to make the plates slightly thicker (to 3 mm) for ease of machining, and because thicker plates are better according to the theoretical electrostatic optics. It doesn't actually make any difference to the VMI performance, however.

[2] And those that don't use designs that give improvements to problems that we aren't really affected by.

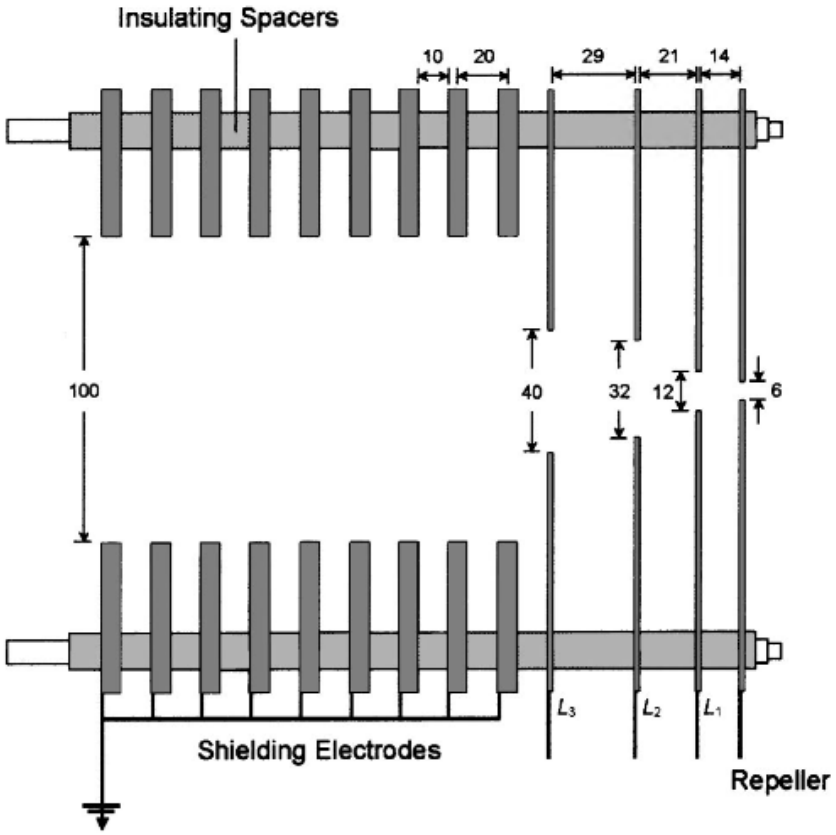

Figure 3.1: Schematic of the DC slicing spectrometer design. Reproduced from REF

3.1.1 Our Design

Our design is essentially a copy of the design in Figure 3.1, but with a removable flight tube so that the detector can either be placed directly on the interaction chamber ($\pm$ the length of an adapter flange), or can be placed an arbitrary distance away using a cheap CF nipple extender. The flight tube in simulation was chosen to be 500 mm long, but the length can be arbitrarily selected. This means that a much wider range of ion energies can be focussed onto the 40 mm screen - an enhancement over instruments with fixed ToF tubes (like the MB imaging rig), even with a much smaller detector. The tube could even be extended to be much longer if very high mass resolution was desired. The flight tube also provides space for an additional optional zoom lens, if desired. Figure 3.2 and Figure 3.3 show the two options (no tube and flight tube).

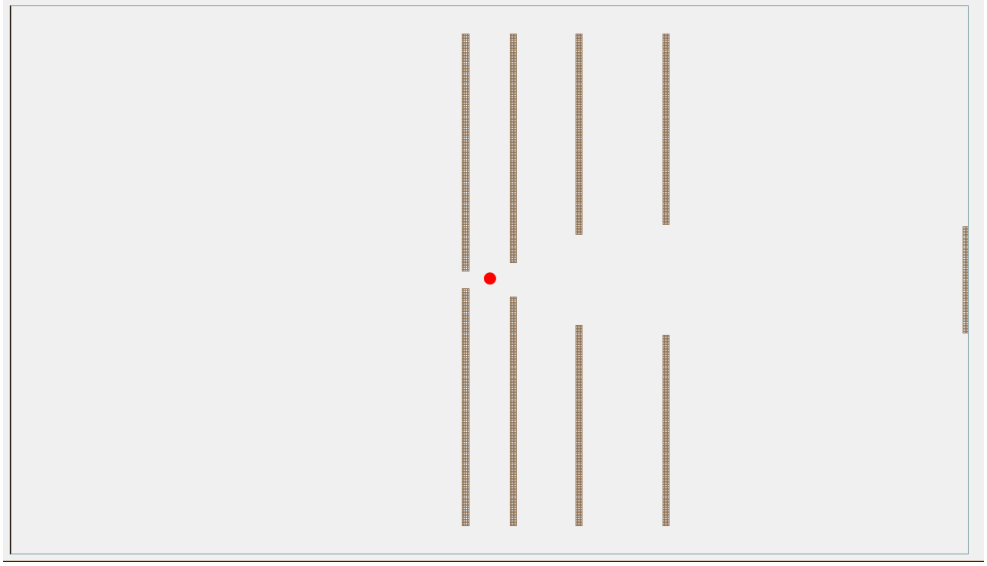

Figure 3.2: Our spectrometer design with no flight tube. The interaction region (marked with a red circle) is aligned with the chamber windows. From left to right, the electrodes are the repeller, extractor, third electrode, and a grounded plate. The detector is on the far right.

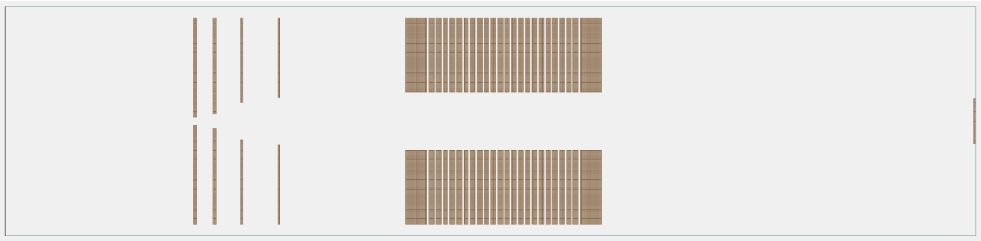

Figure 3.3: The design with a flight tube and optional zoom lens fitted. The geometry of the interaction region is identical to that in Figure 3.2. The zoom lens operation is described in more detail in the following chapter.

3.2 Actual Results

Now we show some actual results. For each of the geometries (tube and no tube), I tested the operation at both high and low energies over a wide range of particle energies. I then simulated an ion image, inverted it using PyAbel[3], and calculated a radial distribution to find the energy resolution. I chose a range of ion energies that would (roughly) fill the detector at both a high (9 kV) and low (2 kV) repeller voltages. These voltages were chosen as 9 kV is about as high as a standard PSU can go without arcing/needing specialist kit; and below 2 kV the ions probably wouldn't have enough kinetic energy to hit the detector hard enough to produce

[3]Using a Cartesian Basex algorithm

a detectable signal.

I will present results in subsections and comment more generally on things after they have all been presented.

3.2.1 No Flight Tube

Simulation results without a flight tube are presented here. Figure 3.4 shows the results taken at low extraction voltages. Figure 3.5 shows results taken at high extraction voltages.

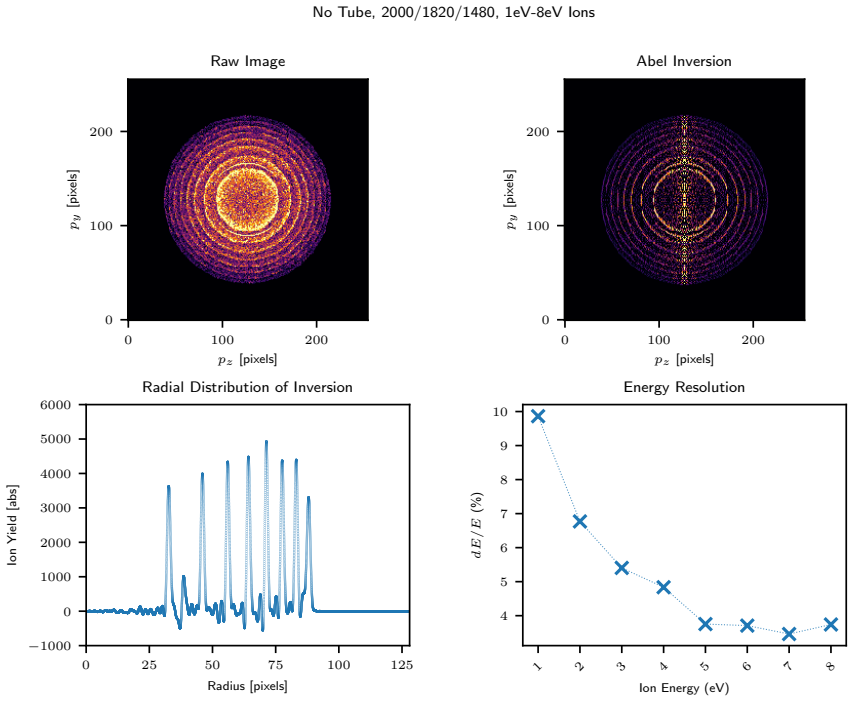

Figure 3.4: Clockwise from top left: Simulated ion image; Inverted image; calculated energy resolution; radial distribution. Ion energies ranging from 1 eV to 8 eV were used. Repeller/Extractor/Third Electrode voltages given in title (in volts).

Without the flight tube, the range of usable particle energies is probably around 3 eV to 40 eV. This is very wide, but without the flight tube the mass resolution is relatively poor[4] - isotopic resolution is probably difficult depending on the gate width. This is probably most useful for very high energy ions.

[4] According to the very imperfect simulation

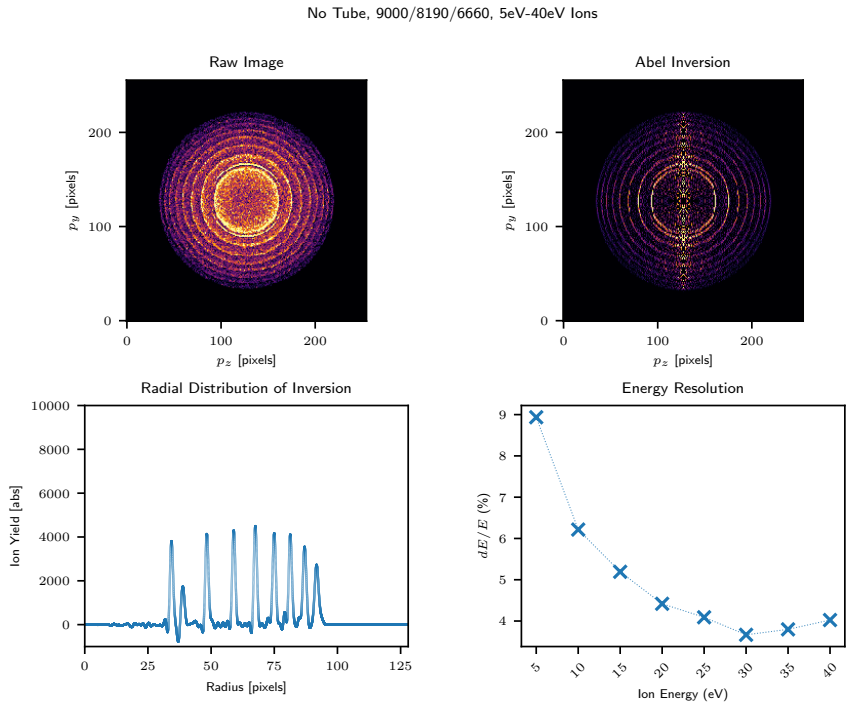

Figure 3.5: Clockwise from top left: Simulated ion image; Inverted image; calculated energy resolution; radial distribution. Ion energies ranging from 5 eV to 40 eV were used. Repeller/Extractor/Third Electrode voltages given in title (in volts).

3.2.2 With Flight Tube

Simulation results with a flight tube are presented here. Figure 3.6 shows the results taken at low extraction voltages. Figure 3.7 shows results taken at high extraction voltages. The zoom lens was present but grounded for these simulations - so effectively is not present.

With the flight tube, the range of usable particle energies is probably around 0.01 eV to 4.5 eV. This is not as wide as without the tube, but the resolution at the low energy end of things is better (possibly useful if you're using massive electrosprayed ions). The mass resolution is much better with the flight tube (as expected), and ions of mass 28 and 29 are easily resolvable in a simulated ToF spectrum.

3.2.3 General Comments

I think the performance shown here is pretty decent and more than sufficient for what we need. Most interesting perhaps is that the energy resolution decreases

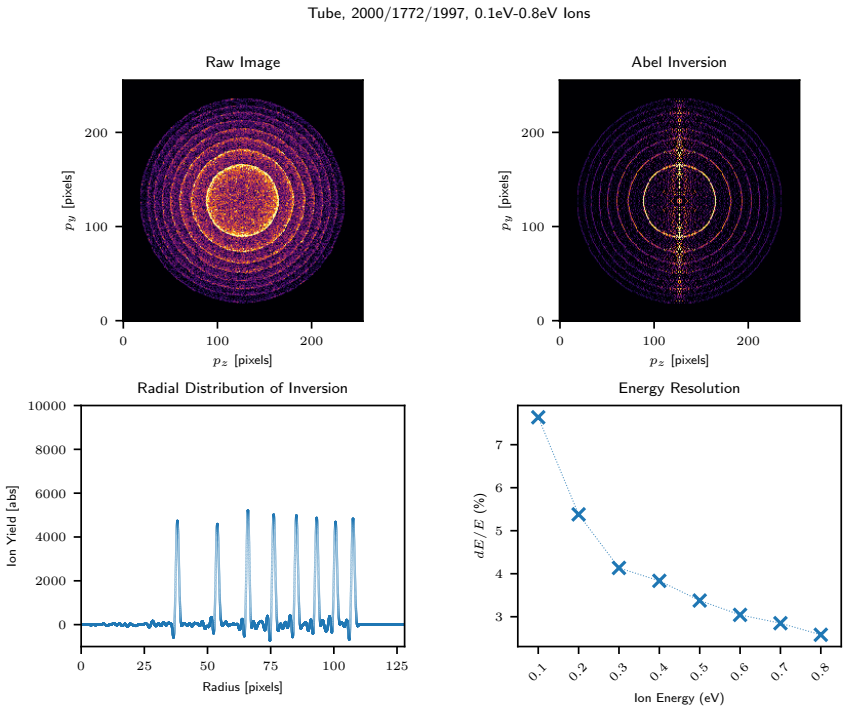

Figure 3.6: Clockwise from top left: Simulated ion image; Inverted image; calculated energy resolution; radial distribution. Ion energies ranging from 0.1 eV to 0.8 eV were used. Repeller/Extractor/Third Electrode voltages given in title (in volts).

dramatically when going towards the center of the detector - this is because there are fewer camera pixels per image segment as the image radius decreases, so the detected resolution is lower. The implication of this is that using a camera with a higher resolution (here using 256×256 pixels to simulate PImMS2 or TimePix3) will increase the attainable energy resolution. This is absolutely the case. For a 256×256 image of some 1 eV ions with energy resolution of 4.4%, increasing the sensor size to 512×512 maintaining the same VMI conditions gives 2.8% resolution, and increasing to 1024×1024 gives around 1.3% resolution.

Otherwise this all seems reasonably within the ranges we want. The mass resolution without the tube is a bit lacking (not quantified for reasons given earlier, but just empirically tested), but is better with the tube. It could also be improved by turning off the third electrode so that 'crush' imaging is effected - this will obviously improve the mass resolution as the Newton spheres are more compressed. Attempts at refining the geometry (and using other non-standard geometries using more interesting electrode shapes) didn't improve on this performance (under our experimental constraints), so I think this is good as a place to start and build

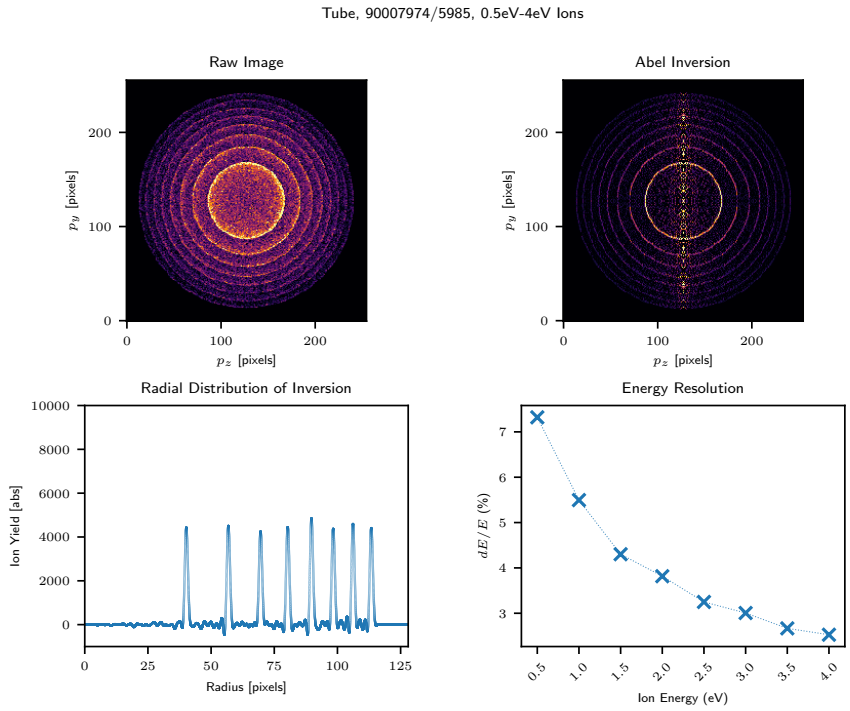

Figure 3.7: Clockwise from top left: Simulated ion image; Inverted image; calculated energy resolution; radial distribution. Ion energies ranging from 0.5 eV to 4 eV were used. Repeller/Extractor/Third Electrode voltages given in title (in volts).

(and it's tried and tested). The ion optic stack will be designed to be modular so it can easily be modified if desired down the line.

Chapter 4

Zoom Lenses

4.1 Electrostatic Zoom Lenses

It is instructive to consider the highest and lowest energy particles that a spectrometer can image. The highest energy particles will be considered later, but the lowest energy particles are more complex and will be considered here. A low energy particle will hit the detector at low radii. While it is possible to choose voltages such that the inherent VMI at this position is good, the ion signal is detected by a camera, and there are fewer pixels per unit angle at low radii than at high radii. This limits the measurable velocity resolution as was seen in the previous chapter, so it can be desirable to 'magnify' the image such that the lower energy particles actually hit a larger radius on the detector.

This can, in most cases, be achieved by simply lowering the repeller and extractor voltages. However, if they are lowered to lower than $\sim$2 kV, then the ions may not have sufficient energy to produce a signal when they strike the detector - so in this case it is not possible to magnify the image purely by lowering these voltages. This is where an external **zoom lens** is desirable, which can magnify the image without lowering the repeller voltage and reducing detection efficiency.

The design process of making a zoom lens was guided by the familiar action of a camera lens, and has been discussed widely in literature on electrostatic lenses - mostly in the context of electron beam focussing (REFS). A 'true' zoom lens (called a 'parfocal lens' in light optics) is one which can magnify an image whilst keeping other parameters - object distance, image distance, and focus - constant. The equivalent of this lens in VMI would be a lens which keeps the object distance, image distance, and image energy constant while the magnification is varied. In general, for a lens consisting of n fixed elements, then $n-2$ parameters can be kept constant by varying the properties of each element. Thus, a five-element lens would (theoretically) be needed to produce this behaviour. Alternatively, if the elements in the lens are movable (as is generally the case in light optics), then

$n-1$ parameters can be kept constant, so a four element lens would be needed for the zooming action. There are numerous examples of electrostatic analogues to these lenses in the literature (REFS) - including interesting examples with pseudo-movable lenses that are controlled via computers.

Attempts were made to build up these movable zoom lenses in SIMION, and they do function as intended (although not as nicely as indicated in previous work, as VMI tends to break the paraxial approximation and aberrations can get severe). However, the electronics needed to control such lenses would quickly get complex - and as is shown below, similar action can be achieved using a different, simpler, method.

4.1.1 Zoom Lenses in VMI

There are examples of zoom lenses being applied to VMI spectrometers, most notably by Vrakking (REF) and in the excellent paper on transfer matrix analysis by Harb (REF). The work by Harb inspired deeper thought about the analogy between VMI and Fourier optics - and here is where we return to the idea of creating a Fourier plane somewhere within the spectrometer, and then using an additional lens to spatially map this Fourier plane onto the detector with a different magnification. This idea is illustrated in Figure 4.1.

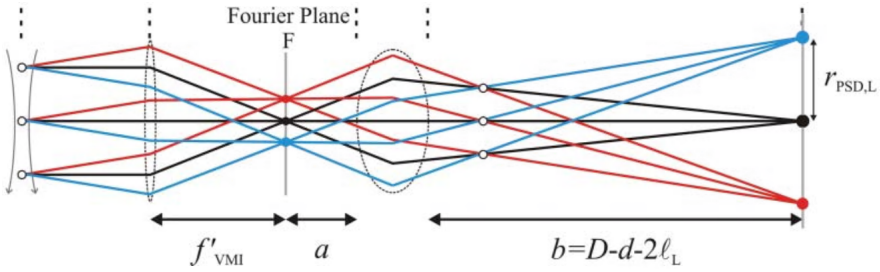

Figure 4.1: Schematic of the idea of magnifying a Fourier plane inside the spectrometer (F) onto a detector (PSD) further away. Taken from Harb *et al.*

These zoom lenses do not produce the 'true zooming' action where the magnification can be simply varied whilst everything remains constant, but achieve the same thing through slightly more steps - and this is perfectly acceptable for VMI purposes. Initial testing of a simple design where an einzel lens is placed between the interaction region and the detector show that magnifications of around a factor of 5 are possible, which is substantial. However, one (very minor) drawback of the design posited by Harb and Vrakking is that they require the initial focussing (as set by the extractor/repeller/possible third lens) to be set such that a Fourier plane is established not at the detector, but within the spectrometer. Then the magnification is performed that maps this Fourier plane onto a second Fourier plane that coincides with the detector. This means that by turning off the lens,

the VMI is not returned to how it was without the magnifying lens, and requires that the voltages in the interaction region be adjusted to get good focussing back.

My approach which circumvents this problem[1] draws on experience of using camera lenses. Generally speaking when a camera lens is zoomed, the image does not remain in focus - and the focussing needs to be corrected following the zoom. SIMION simulations suggest that a similar approach to designing VMI zoom lenses can be fruitful. By leaving the repeller/extractor region set at voltages which produce good VMI at the detector, a magnifying einzel lens can be turned on - this will produce some magnification), but the new Fourier plane is generally nowhere near the detector. By adjusting a 'focussing' electrode placed in the interaction region by a small amount, it is possible to restore the final Fourier plane to the detector plane. Then, if the magnification is no longer desired, the magnifying lens and focussing lens can be turned off - and the initial VMI condition is restored. I believe that this is a modest improvement on existing designs, largely due to ease of use and the intuitive nature of the zooming-focussing behaviour.

[1] More of a very minor inconvenience than an actual problem.

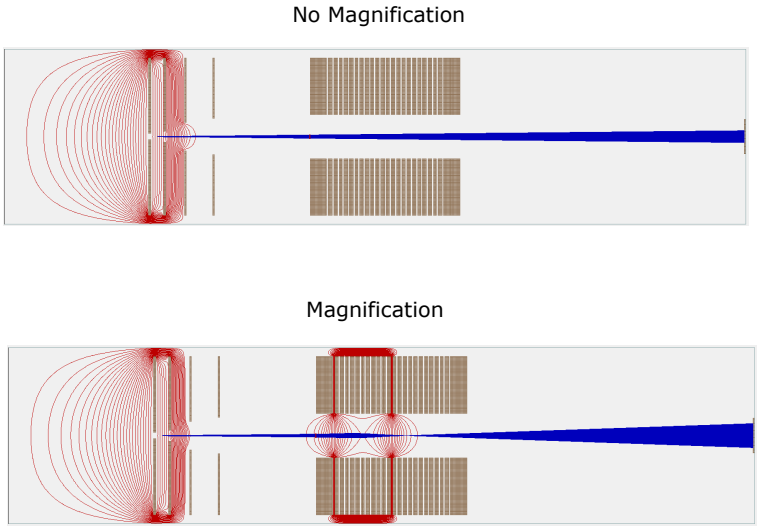

Figure 4.2: Example of how the magnifying zoom lens works. Top: Initial VMI of low energy ions that hit the middle of the detector. Bottom: The magnifying lens turned on, and the image filling a visibly larger area on the detector. Note the change in equipotential lines in the source region, as the third electrode has been adjusted to maintain VMI focussing under magnification.