Vol. 2

Chapter 1.4

Page 1:

Good morning, everyone. Welcome back to Physics 608, Laser Spectroscopy. I'm Distinguished Professor Dr M A Gondal, and today, we embark on a new and exciting topic, which we'll cover in this section, 1.4: Ionization Spectroscopy.

In our previous discussions, we've explored various ways lasers can interact with matter, primarily focusing on absorption and fluorescence. These are powerful techniques, but they have their limitations, especially when we push towards the ultimate goal of detecting exceedingly small quantities of atoms or molecules.

Ionization spectroscopy represents a paradigm shift in detection strategy. Instead of looking for the faint shadow of an absorbed photon or trying to catch the faint glimmer of a fluorescent photon, we are going to use the laser to fundamentally change the nature of our target species—we're going to turn it into a charged particle.

As we will see, this seemingly simple act of converting a neutral particle into an ion or an electron opens up a world of detection possibilities with almost breathtaking sensitivity. We are talking about techniques that can, quite literally, count single atoms.

So, over the course of this lecture, we will build a complete picture of this family of techniques. We will start with the core concepts, develop the underlying quantitative framework, explore the various experimental implementations, and finally, look at some of the most advanced

applications, from fundamental atomic physics to mass spectrometry and even planetary science. Let's begin.

Page 2: Ionization Spectroscopy

Alright, let's dive into the core concept of Ionization Spectroscopy. The central idea is articulated in the first bullet point.

Our goal is to convert the absorption of one or more photons by a neutral species into a readily detectable flow of charged particles—that is, ions or electrons. Think about this for a moment. In conventional absorption spectroscopy, you measure a tiny decrease in a large, transmitted light intensity. It's like trying to weigh a ship's captain by weighing the ship with and without the captain on board. It's an incredibly difficult difference measurement. In fluorescence spectroscopy, you're trying to collect photons that are emitted isotropically, over a four pi solid angle, and your detector only covers a small fraction of that.

Ionization spectroscopy sidesteps these problems. We are not measuring a difference, and we are not struggling with low collection efficiency. We are creating a new particle, an ion, where there was none before. And the beauty of charged particles is that they are fantastically easy to detect.

This brings us to the second, critical point on the slide. The power of this technique relies on the huge signal-to-noise advantage of electrical current counting over optical power measurements. We can take the ion we've just created, accelerate it with an electric field, and guide it directly to a detector. Modern detectors, like microchannel plates or electron multipliers, are so efficient that the impact of a single ion can generate a cascade of

millions of electrons—a robust, easily measurable electrical pulse. The background for this measurement is essentially zero. We're not looking for a tiny signal on top of a large background; we are listening for a clear "click" in an otherwise silent room. This is why ionization methods can achieve sensitivities that are simply unattainable with many other optical techniques. The fundamental task is transformed from measuring photons to counting particles, and in the world of quantum measurements, counting is king.

Page 3:

So, how do we actually accomplish this? The process is conceptually a two-step sequence.

The key step, as noted here, is to first use a laser to promote the population from some initial quantum state, which we'll label with the ket | i ⟩ $|i\rangle$ and energy E i E_{i}, to a specific, chosen excited state, which we'll call | k ⟩ $|k\rangle$ with energy E k E_{k}. This first step is the *spectroscopic* step. By tuning our laser's frequency, we select exactly which energy level E k E_{k} we want to populate. This gives the technique its incredible selectivity. Once the atom or molecule is in this excited state, we then apply a second step to remove an electron, thereby creating an ion.

Now, what happens to the charged products we've made? As the second bullet points out, these resulting ions or electrons are accelerated by electric fields, typically to energies of several kilo-electron-volts, or keV. At these energies, they can be counted with nearly unit efficiency when the experiment is properly optimized. This is a remarkable statement. It means

that for every ion we create in the interaction region, we can get one count in our detector. We are approaching a perfect one-to-one correspondence between the quantum event we initiated and the signal we measure.

This general idea can be implemented in several ways, which we can think of as a family of related techniques. The slide lists the four main branches we'll be exploring:

- (a) direct photoionization, the simplest case where a single photon has enough energy to ionize the atom directly. - (b) resonant multiphoton ionization, universally known by its acronym, R. E. M. P. I., or REMPI. This is the workhorse of the field. - (c) collision-induced ionization, where the final ionization step is caused not by a photon, but by a collision with another particle. - (d) field ionization, a clever technique where a static electric field is used to literally rip the electron off a highly excited atom.

We will delve into each of these, but first, let's visualize the entire process with a block diagram of a typical experiment.

Page 4:

Here on this page, we see a conceptual block diagram that lays out the anatomy of a generic ionization spectroscopy experiment. It's a fantastic roadmap for our discussion, so let's walk through it from left to right, following the five labeled stages.

Stage 1 is Laser Excitation. This begins with our Laser System. This might be a tunable dye laser, a Ti:sapphire laser, or a diode laser. It could be pulsed or continuous-wave, depending on the specific technique. This laser

produces a beam of photons with a well-defined energy, h v $h\nu$, that we direct into our experimental chamber.

This leads us to Stage 2: Ionization. The laser beam enters an Interaction Region, which is typically under high vacuum. In this region, our laser interacts with the sample, which is depicted here as a cloud of neutral atoms or molecules, labeled 'A'. The diagram shows a laser photon, represented by the squiggly red arrow labeled h v $h\nu$, striking a neutral atom A A. This interaction is what drives the transition from the ground state to our chosen excited state. A subsequent process, which could be the absorption of another photon, then provides the energy to eject an electron, creating a positive ion, which is shown here as A + A^+.

Once the ion is created, we move to Stage 3: Acceleration and Focusing. The newly formed ion, A + A^+, doesn't just sit there. It is immediately subjected to an electric field generated by a set of electrodes labeled "Ion Optics." You can see plates with voltages V 0 V_0, V 1 V_1, and V 2 V_2. These electrodes act as an electrostatic lens, grabbing the ion and accelerating it along a well-defined "Ion Trajectory," represented by the solid blue arrow. The goal of these ion optics is to efficiently collect every single ion created in the interaction volume and steer it towards the detector.

That brings us to Stage 4: Detection. The focused beam of ions strikes a Detector. This device, as we mentioned, is designed to be extremely sensitive. Upon impact, the single ion initiates a process that generates a much larger, measurable electrical signal.

And finally, Stage 5: Data Acquisition. The electrical pulse from the detector is sent to our data acquisition system. This could be an oscilloscope, a boxcar averager, or a computer card. It records the event. Typically, we would record the number of these events—the number of ions detected—as a function of the wavelength of our excitation laser. Plotting the ion signal versus the laser wavelength produces our spectrum, as shown by the peak on the screen labeled "Data."

So, in essence, we use light to "tag" an atom of interest, creating an ion, and then we use electric fields and electronics to count that tag with incredible efficiency.

Page 5:

Now let's move from the block diagram of the experiment to the quantum mechanical heart of the matter: the energy-level picture of the process. This is where the spectroscopy truly happens.

To understand this, we need to consider two essential manifolds of states for our atom or molecule. First, we have the bound levels. These are the discrete, quantized energy levels, which we've been labeling E i E_i, E k E_k, and so on. These all lie *below* a critical energy threshold. They represent states where the electron is still bound to the atomic or molecular core.

Second, above this threshold, we have the ionization continuum. This is not a discrete level, but a continuous range of energies. It begins at the ionization energy, which we denote as I P IP. Any state in this continuum corresponds to the electron being completely free from the parent atom or

molecule, which is now an ion. The energy above the ionization potential, I P IP, simply corresponds to the kinetic energy of the free electron and ion.

Now, let's look at the basic photon-absorption path for the most common and powerful variant of this technique, known as Resonant Two-Photon Ionization, or RTPI. The process is written out here symbolically:

The system starts in an initial state, ket | i ⟩ $|i\rangle$. It then absorbs a photon with energy h v 1 $h\nu_1$. This absorption is a resonant process, meaning h v 1 $h\nu_1$ must be precisely tuned to match the energy difference between state | i ⟩ $|i\rangle$ and some intermediate excited state, ket | k ⟩ $|k\rangle$. So, the first step is the transition from | i ⟩ $|i\rangle$ to | k ⟩ $|k\rangle$.

From this intermediate state | k ⟩ $|k\rangle$, the system then absorbs a *second* photon, this one with energy h v 2 $h\nu_2$. This photon provides the additional energy needed to push the electron past the ionization potential and into the continuum. The final result of this second step is a positive ion, which we'll denote M + M^+, and a free electron, e − e^-.

This two-step ladder—resonant excitation followed by ionization—is the fundamental pathway we will be analyzing. Let's now explicitly define each of the symbols involved on the next page.

Page 6:

Continuing with our energy-level picture, let's make sure every symbol is crystal clear.

The terms h v 1 $h\nu_1$ and h v 2 $h\nu_2$ represent the photon energies of the two lasers we use in the experiment, which we can call L-one and L-two. In

some cases, L-one and L-two can be the same laser, which is called a "one-color" experiment. More often, they are different lasers, allowing us to optimize each step independently in a "two-color" experiment.

M + M^+ represents the molecular ion that is created. Of course, this could also be an atomic ion if our sample consists of atoms.

e − e^- is the freed electron, which is often called a photoelectron. It carries away any excess energy from the ionization process as kinetic energy, which we can label E k i n E_{kin}. In some advanced experiments, we can even measure this kinetic energy to get more information, a technique called photoelectron spectroscopy.

Now for the most important strategic point of this entire process. The overall sensitivity of our measurement is usually dominated by the first step: locating a strong, allowed transition from our initial state | i ⟩ $|i\rangle$ to the intermediate state | k ⟩ $|k\rangle$. Why? Because if this first step is very unlikely to happen—if the absorption cross-section is small—then we simply won't create many molecules in the excited state | k ⟩ $|k\rangle$, and it won't matter how efficient our ionization step is. We need to find a transition with a large oscillator strength to efficiently "pump" population into the intermediate state. This is where all the rules of spectroscopy—selection rules, Franck-Condon factors for molecules, and so on—come into play.

Finally, there's a special case mentioned that is extremely powerful. If the intermediate state E k E_{k} is what's known as a Rydberg level, the ionization step can be exceptionally efficient. A Rydberg level is a very highly excited state where the electron is, on average, very far from the ionic core. It's almost free. Because it's so loosely bound and its

wavefunction is so extended, it couples very strongly to the ionization continuum. This results in a very large cross-section for the final ionization step. So, a strategy of exciting to a Rydberg state and then ionizing is a well-known trick for maximizing the signal in these experiments.

Page 7:

This slide gives us a clean, visual representation of the energy-level picture we've just discussed. It's a fantastic way to solidify the concepts.

Let's break down what we're seeing. The vertical axis represents energy, labeled ' E E'. At the very bottom, we have a thick, solid line labeled with the ket | i ⟩ $|i\rangle$, representing the initial, typically ground, state of our molecule.

From this initial state, we see a red, wavy arrow pointing upwards, labeled h v 1 $h\nu_1$. This represents the absorption of the first photon from our first laser, L-one. This photon's energy is tuned to be resonant with the energy difference between the initial state and an intermediate state, which is the solid line labeled with the ket | k ⟩ $|k\rangle$. This is our spectroscopic step.

Notice the annotation on the left, "Rydberg." It indicates that in this particular diagram, the state | k ⟩ $|k\rangle$ is chosen to be a high-lying Rydberg state, just as we discussed on the previous slide. These states are clustered together just below the ionization limit.

Above the discrete levels, we see a dashed horizontal line labeled I P IP, for Ionization Potential. This is the threshold energy required to liberate the electron.

From our intermediate state | k ⟩ $|k\rangle$, a second, blue wavy arrow labeled h v 2 $h\nu_2$ points upward. This represents the absorption of the second photon, from laser L-two. This photon's energy is sufficient to take the molecule from state | k ⟩ $|k\rangle$ up and across the ionization potential, I P IP, into the shaded region at the top, which is labeled the "Ionization Continuum."

Once in the continuum, the molecule has become an ion-electron pair, denoted as M + + e − $M^+ + e^-$. These are now free particles that we can guide and detect.

This diagram beautifully encapsulates the "resonant-excitation-followed-by-ionization" scheme. The first red arrow is all about selectivity and spectroscopy. The second blue arrow is all about efficient conversion to a detectable charged particle.

Page 8:

Let's now formalize the second step of our process—the photoionization itself—by looking at its chemical equation and notation.

The fundamental reaction is shown here. We begin with a molecule that is already in an excited state. We denote this as \(M^*(E_\text{k})\). The \(M^*\) indicates an electronically excited species, and the (E k) (E_k) reminds us that it's specifically in the quantum state k k. This excited molecule then interacts with a photon from our second laser, L-two, with energy h v 2 $h\nu_2$.

The result of this interaction, as shown by the arrow, is the creation of three products: the molecular ion, M + M^+; a free electron, e − e^-; and the kinetic energy carried by that electron, E k i n (e −) $E_{\mathrm{kin}}(e^-)$.

Conservation of energy dictates that the initial energy, E k + h v 2 $E_{\mathrm{k}} + h\nu_2$, must equal the final energy, which is the ionization potential I P IP plus the kinetic energy E k i n (e −) $E_{\mathrm{kin}}(e^-)$.

The definitions below simply reiterate what we've been discussing. M ∗ (E k) $M^*(E_{\mathrm{k}})$ is our molecule excited to state E k E_{k}. And h v 2 $h\nu_2$ is the photon from our ionizing laser, L-two. This equation is the defining event of the ionization step, taking us from a neutral, albeit excited, species to the charged particles that form our signal.

Page 9:

Continuing with the details of photoionization, a key experimental choice is the origin of this ionizing photon, h v 2 $h\nu_2$.

As the first two bullet points explain, this photon can come from one of two places. It might originate from the very same laser that was used to create the excited state E k E_{k}. This is what we call a "one-color experiment," because only one laser color, or wavelength, is used for both the excitation and ionization steps. This is simple and often sufficient.

Alternatively, the ionizing photon can come from a separate laser or even a lamp. This is a "two-color experiment." This approach, while more complex as it requires two laser systems, offers a significant advantage, which is highlighted in the next bullet point.

The ability to choose the photon energies h v 1 $h\nu_1$ and h v 2 $h\nu_2$ independently allows for the optimization of both the excitation and ionization cross sections. For the first, resonant step, you might want a laser with a very narrow linewidth and low power to achieve high spectral resolution without broadening your transition. For the second, ionization step, you don't need a narrow linewidth, but you often want very high power to make the ionization as efficient as possible. A two-color setup lets you use the perfect laser for each job. For example, a tunable dye laser for the first step, and a high-power, fixed-frequency Nd:YAG laser for the second.

Finally, to give you a feel for the numbers involved, the slide notes that for many molecules, the photoionization cross-section from an excited state, which we label σ k l σ_{kl}, can be as high as 10 − 17 c m 2 $10^{-17}\,\mathrm{cm}^2$ when the ionizing photon energy is just above the threshold. This cross-section is a measure of the effective "target area" the excited molecule presents to the ionizing photon. A value of 10 − 17 c m 2 $10^{-17}\,\mathrm{cm}^2$ is quite large, which is fantastic news for us, as it means the ionization step can be made very efficient.

Page 10:

Now we are going to begin building a quantitative model to describe our ion signal. The first step is to determine the population of the intermediate level, | k ⟩ $|k\rangle$, under the influence of our excitation laser. We'll start by considering the case of steady-state excitation.

The first thing we need is the rate of photon absorption on the | i ⟩ $|i\rangle$ to | k ⟩ $|k\rangle$ transition. This is the rate at which our laser pumps molecules from

the initial state into the excited state. This rate, which we'll call n a n_a, is given by the equation you see on the slide.

Let's read and deconstruct this equation:

n a = N i n L 1 σ i k Δ x

$$n_a = N_i \, n_{L1} \, \sigma_{ik} \, \Delta x$$

Let's break down each term:

* n a n_a is the number of absorption events occurring per unit volume per second. Its units would be something like molecules per cubic centimeter per second. * N i N_i is the number density of molecules in the initial state | i ⟩ $|i\rangle$. This has units of molecules per cubic centimeter. * n L 1 n_{L1} is the photon flux density of our first laser, L-one. This is the number of photons from this laser passing through a unit area per unit time. Its units are photons per square centimeter per second. * σ i k σ_{ik} is the absorption cross-section for the i i to k k transition. This is the effective area that a molecule presents to the photons for this specific transition, with units of square centimeters. * And Δ x Δx is the illuminated path length, the length of the sample that the laser passes through, in centimeters.

Let's check the units. N i N_i (cm − 3 cm^{-3}) times n L 1 n_{L1} (cm − 2 s − 1 $cm^{-2}\,s^{-1}$) times σ i k σ_{ik} (cm 2 cm^2) gives units of cm − 3 s − 1 $cm^{-3}\,s^{-1}$. So n a n_a as written here without the Δ x Δx term is the rate of absorption per unit volume. The way the slide has written it is slightly ambiguous; often this is formulated as a rate per unit area, where N i N_i is a column density. For our purposes, let's consider n a n_a as the volumetric rate of pumping,

so n a = N i n L 1 σ i k $n_a = N_i\, n_{L1}\, \sigma_{ik}$. This is the rate at which population is fed into our excited state | k ⟩ $|k\rangle$.

Page 11:

Now that we have the rate at which population is pumped *into* the excited state | k ⟩ $|k\rangle$, we need to consider the rates at which it *leaves*. In a steady-state situation, the population of level | k ⟩ $|k\rangle$ will be constant, which means the rate in must exactly equal the rate out. This is the principle of population balance.

The first equation on this slide expresses this balance mathematically. It's a simple rate equation for the population of the excited state, N k N_k. The rate of change of N k N_k with time, d N k d t $\frac{dN_k}{dt}$, is equal to the pumping rate in, n a n_a, minus the total depopulation rate.

The total rate of leaving is the sum of all possible decay channels, multiplied by the population N k N_k itself. Here, these channels are grouped into two terms: P k I P_{kI} and R k R_k.

* P k I P_{kI} is the ionization probability per molecule per second. This is the rate at which a single excited molecule is ionized by our second laser, L-two. This is the channel we want. *R k* R_k represents all other* relaxation rates combined. This includes things like spontaneous radiative decay (fluorescence), and collisional quenching where the molecule loses its energy by bumping into another particle. These are the competing loss channels.

So, the total rate out is (P k I + R k) N k $(P_{kI} + R_k)\, N_k$.

Under steady-state conditions, the population isn't changing, so d N k d t = 0 $\frac{dN_k}{dt} = 0$.

Setting the equation to zero and rearranging gives us a beautifully simple and powerful expression for the steady-state population of our excited level:

N k = n a P k I + R k

$$N_k = \frac{n_a}{P_{kI} + R_k}$$

This equation tells us something profound. The amount of excited state population we can build up, N k N_k, is determined by a competition. We're feeding population in at a rate n a n_a, and it's draining out through two channels: the desired ionization channel (P k I P_{kI}) and the undesired loss channels (R k R_k). To maximize N k N_k, we need to make the denominator as small as possible, but more importantly, as we'll see, the *ratio* of these two loss channels will determine our ultimate efficiency.

Page 12:

We have now calculated the steady-state population of the excited state, N k N_k. The next logical step is to calculate our actual, measurable ion signal, which we will denote as S I S_I.

The ion signal, in units of counts per second, is simply the number of ions we successfully create and detect per second. This is given by the first equation:

S I = N k P k I δ η .

$$S_{\mathrm{I}} = N_{\mathrm{k}} P_{kI} \delta \eta.$$

Let's break this down.

The term N k N_{k} times P k I P_{kI} gives us the total number of ions being created per unit volume per second. Remember, N k N_{k} is the number density of excited molecules, and P k I P_{kI} is the probability per second that any one of those molecules gets ionized.

However, just creating an ion isn't enough; we have to detect it. That's where the next two terms come in.

δ δ is the geometrical collection efficiency. This is a dimensionless number between 0 and 1 that represents the fraction of created ions that are successfully guided by our ion optics to the front face of the detector. If our ion optics are perfectly designed, δ δ approaches 1.

η η is the intrinsic detector efficiency. This is also a number between 0 and 1, representing the probability that an ion striking the detector actually produces a measurable count. For modern detectors like microchannel plates, η η can also be very close to 1.

So, S I S_{I} is the total rate of ion production, corrected for our instrument's real-world imperfections.

Now, we perform a crucial substitution. We take the expression for the steady-state population N k N_{k} that we derived on the last Paage and substitute it into our equation for S I S_{I}.

This gives us the second equation on the slide:

S I = n a P k I P k I + R k δ η .

$$S_{\mathrm{I}} = n_{\mathrm{a}} \frac{P_{kI}}{P_{kI} + R_{\mathrm{k}}} \delta\eta.$$

This is a very important result. It connects our final signal rate S I S_{I} back to the initial pumping rate n a n_{a} and clearly shows how the signal depends on the competition between ionization and other relaxation processes. Let's explore the physical meaning of this equation on the next page.

Page 13:

Let's continue our analysis and now connect everything back to the primary laser parameters and interpret the physical meaning of our final signal equation.

First, we'll re-insert the expression for n a n_{a}, the initial absorption rate. Recall that n a n_{a} was equal to N i N_{i} times n L 1 n_{L1} times σ i k σ_{ik} times Δ x Δx. Substituting this into our expression for S I S_{I} gives the final equation you see in the box.

S I S_{I} equals N i N_{i} times n L 1 n_{L1} times σ i k σ_{ik} times Δ x Δx, all multiplied by the fraction P k I P k I + R k $\frac{P_{kI}}{P_{kI}+R_{\mathrm{k}}}$, and finally multiplied by δ δ and η η.

S I = N i n L 1 σ i k Δ x P k I P k I + R k δ η

$$S_{\mathrm{I}} = N_{\mathrm{i}}\, n_{L1}\, \sigma_{ik}\, \Delta x\, \frac{P_{kI}}{P_{kI} + R_{\mathrm{k}}}\, \delta\, \eta$$

This equation contains the entire physics of the process, from the initial state of the sample to the final count registered by our computer. Now, let's do the most important thing: interpret it.

The physical interpretation breaks down into two key parts.

First, look at that fraction in the square brackets: [P k I P k I + R k $\frac{P_{kI}}{P_{kI}+R_k}$].

This term has a profound physical meaning. It is the *quantum yield of ionization*. It represents the fraction of all the molecules that we excite to state | k ⟩ $|k\rangle$ that actually end up being ionized. It explicitly quantifies the competition between our desired process, ionization (with rate P k I P_{kI}), and all the other decay processes like fluorescence and collisions (with combined rate R k R_k). To get the maximum signal, we want the ionization rate P k I P_{kI} to be much, much larger than the relaxation rate R k R_k, so that this fraction approaches one.

Second, we have the product δ δ times η η. This term tells us how the design of our instrument—the ion optics and the detector—influences the ultimate sensitivity. Even if our ionization quantum yield is 100%, if we fail to collect the ions (δ δ is small) or fail to detect them (η η is small), our signal will be poor. Achieving δ δ and η η close to one is a primary goal of experimental design.

So, to get a large signal, we need a large number of initial molecules (N i N_i), a strong laser (n L 1 n_{L1}), a strong transition (σ i k σ_{ik}), a long path length (Δ x Δx), an ionization process that outcompetes decay, and an efficient instrument. Our equation lays it all out.

Page 14:

This slide presents the ideal, limiting case for our ion signal, the absolute best-case scenario that we strive for in an experiment.

The condition is stated clearly: When the ionization probability per second, P k I P_{kI}, is much, much greater than the rate of all other relaxation processes, R k R_k, and when our instrumental efficiencies are perfect, meaning both the collection efficiency δ δ and the detector efficiency η η are equal to 1.

Let's look back at our signal equation from the previous page. If P k I P_{kI} is much larger than R k R_k, then the denominator (P k I + R k) $(P_{kI} + R_k)$ is approximately just P k I P_{kI}. The entire fractional term, P k I / (P k I + R k) $P_{kI}/(P_{kI} + R_k)$, becomes approximately equal to 1. This means our ionization quantum yield is 100 % 100%. Every single molecule we excite gets ionized.

If we also have δ = 1 $\delta = 1$ and η = 1 $\eta = 1$, then our full signal equation, S I = n a · [fraction] · δ · η

$$S_I = n_a \cdot [\text{fraction}] \cdot \delta \cdot \eta$$

simplifies dramatically. The fraction is one, δ δ is one, η η is one. So, S I S_I simply becomes equal to n a n_a.

The statement at the end of the slide summarizes this profound result: "every absorbed photon is observed." More precisely, every photon absorbed in the initial excitation step (| i ⟩ $|i\rangle$ to | k ⟩ $|k\rangle$) ultimately leads to a detected count.

This establishes a perfect one-to-one correspondence between the primary spectroscopic event and our measured signal. This is the holy grail of sensitive detection, and as we will see next, it is a goal that is surprisingly achievable in a well-designed experiment.

Page 15:

So, is this ideal scenario of nearly 100 % 100% detection efficiency just a theorist's dream, or can we actually achieve it in the laboratory? This slide answers that question with a resounding "yes."

In a carefully designed apparatus, particularly a molecular-beam apparatus where we have a lot of control over the environment, we can get very close to this ideal limit.

The first bullet point addresses the collection efficiency, δ δ. By using well-designed ion/electron extraction optics—those repeller and extractor plates we saw in the diagram—we can create electrostatic fields that guide virtually all of the charged particles from the interaction volume onto the detector. This allows us to achieve a collection efficiency, δ δ, that is approximately equal to 1 1.

The second point addresses the detector efficiency, η η. Modern detectors like Channeltrons or Microchannel Plates (MCPs) are incredibly effective. When a particle with kilo-electron-volt (keV) energy strikes their surface, they are almost guaranteed to trigger a large cascade of electrons, resulting in a detectable pulse. This gives us an intrinsic detector efficiency, η η, that is also approximately equal to 1 1 for these energetic particles.

So, what happens to our signal under these optimized conditions? The final point shows the beautiful simplification. When the ionization rate P k I P_{kI} is much greater than the relaxation rate R k R_k, and when δ δ and η η are both essentially 1 1, our complex signal equation, S I S_I, collapses to its simplest possible form:

S I = n a .

$$S_I = n_a.$$

Let's recall what n a n_a is: it's the rate of photon absorption in the first excitation step. This confirms our conclusion from the previous slide. In an optimized experiment, the number of ions we count per second is equal to the number of photons per second that were resonantly absorbed by our sample to begin with. We have successfully built a machine that counts absorbed photons.

Page 16:

Now let's consider the powerful consequences of achieving this one-to-one correspondence.

First, as we've stated, we have a direct, quantitative link between the number of photons absorbed and the number of counts we detect. This isn't just an academic point; it means our signal is directly proportional to the concentration of our target species, with a proportionality constant we can, in principle, know perfectly. This makes the technique not just sensitive, but also highly quantitative.

The second point highlights the impact on sensitivity. The minimum detectable optical power is now limited only by dark counts. "Dark counts" are spurious signals from the detector that occur even when there's no real ion signal—perhaps from a stray cosmic ray or a random thermal electron in the multiplier. In a well-shielded, cooled detector, this dark count rate can be made incredibly low, often less than one count per second. This means

our noise floor is practically zero. We are looking for a real signal against a backdrop of almost complete silence.

The third bullet point provides a crucial comparison to a more conventional technique: fluorescence detection. In a fluorescence experiment, you're trying to collect photons emitted in all directions. Even with good optics, you rarely collect more than a few percent of the total emitted light due to the limited solid angle of your collection lens. Furthermore, the fluorescence quantum yield itself might be less than one due to competing non-radiative decay. The combined result is that fluorescence detection rarely exceeds a total efficiency of a few percent. In stark contrast, our optimized ionization experiment approaches an efficiency of one hundred percent. This represents an advantage of one to two orders of magnitude in raw signal, which can be the difference between seeing a signal and seeing nothing at all.

Page 17:

This slide delivers the final, powerful conclusion from the first part of our discussion.

Given everything we've established—the ability to achieve a near-perfect ionization quantum yield and near-perfect detection efficiency—we can state the following:

Hence, resonant two-step ionization is the most sensitive absorption probe for any species whose excited state ionizes readily.

Let's unpack that statement.

It is an "absorption probe" because the signal is directly proportional to the initial resonant absorption event. It is the "most sensitive" because of the near-unity detection efficiency, which allows us to reach the fundamental limit of counting single quantum events. And it applies to "any species whose excited state ionizes readily."

This is the key condition. We must be able to find an efficient pathway to ionize the molecule once it's in the intermediate state. As we've seen, this can be achieved with a sufficiently powerful laser, or by accessing special states like Rydberg or autoionizing levels.

When this condition is met, no other absorption-based technique can match the sensitivity of ionization spectroscopy. This is why it has become such an indispensable tool in fields like trace-gas analysis, reaction dynamics, and precision measurements.

Page 18

Here we have a detailed schematic that brings our abstract discussion of ion extraction and detection to life. This diagram shows the key hardware components that are responsible for achieving the high collection and detection efficiencies, δ δ and η η, that are so critical to this technique.

Let's trace the path of the experiment. From the left, an orange line represents our first laser, L-one, with photon energy h v 1 $h\nu_1$. It enters the interaction region. Simultaneously, a second laser, L-two, with energy h v 2 $h\nu_2$ (represented by the blue beam) also enters this region. The diagram shows them crossing.

In the center, where the lasers overlap with our sample molecules, M M, the ionization event occurs: M → M + + e −

$$M \rightarrow M^{+} + e^{-}$$

This happens in the space between two crucial electrodes. On the left is the Repeller plate, which is held at a positive voltage, + V $+V$. On the right is the Extractor plate, which is typically held at ground potential.

The positive potential on the Repeller pushes the newly formed positive ion, M + M^{+}, to the right, away from the Repeller and towards the Extractor. The Extractor plate has a small hole in it, allowing the ion to pass through. The curved, dashed lines between the plates represent the electric field lines. Notice how they are shaped not just to push the ion, but also to gently focus it, ensuring it travels straight towards the detector. This is the electrostatic lens that gives us a collection efficiency, δ δ, close to one, as annotated.

After passing the extractor, the now-energetic ion, M M, travels in a straight line until it slams into the detector. In this diagram, the detector is an MCP Detector, or Microchannel Plate detector. The impact initiates an electron cascade, creating a large, negative pulse of charge. This pulse is our "Signal Output," shown as a sharp peak on an oscilloscope trace. The annotation "Detection η ≈ 1 $\eta \approx 1$" reminds us that these detectors are extremely efficient. This entire assembly, from the repeller to the MCP, is the heart of the detection system.

Page 19: Schematic of Ion Extraction and Detection

This slide provides a textual explanation for the beautiful diagram we just examined on the previous page, formally titling it a "Schematic of Ion Extraction and Detection." Let's walk through this description to ensure we've captured all the key ideas.

The text begins by stating that the diagram illustrates the key components of a resonant two-photon ionization, or R2PI, experiment, with a focus on the high- efficiency collection and detection of ions. This is precisely what we've been discussing.

It then reiterates the process: A molecule, M M, is first excited by laser L 1 L_1 and then ionized by laser L 2 L_2. The resulting ion, M + M^+, is accelerated by an electric field.

Crucially, it names the components that create this field: the Repeller and Extractor electrodes. This assembly is often referred to as the "ion source" or "extraction region" of a mass spectrometer.

The setup, as the text rightly emphasizes, visualizes the two critical parameters for achieving high sensitivity.

First, Collection Efficiency, represented by the Greek letter δ δ. The text explains that the electrostatic lens formed by the electrodes guides nearly all ions to the detector, making δ δ approximately equal to 1. This is the key to not losing the signal you worked so hard to create.

Second, Detector Efficiency, represented by the Greek letter η η. The text identifies the detector as a microchannel plate, or MCP. It gives a concise description of its function: it generates a large electron cascade from a single ion impact. This built-in amplification ensures a detectable electronic pulse for nearly every incident ion, making η η approximately equal to 1.

Together, these two factors, δ δ and η η both approaching unity, are what elevate R2PI to the status of a quantum-noise-limited detection technique.

Page 20:

Now we're going to broaden our scope beyond the simple two-photon case to a more general and powerful family of techniques called Resonant Multiphoton Ionization, or REMPI. This slide provides the motivation.

The first bullet point addresses a common scenario: What if your second photon, h ν 2 $h\nu_2$, is simply not energetic enough to clear the ionization potential, IP? For example, you might be using a single laser for both steps (a one-color experiment), and two photons of that color get you to a nice resonant state, but don't quite have enough energy to ionize. The solution is simple in concept: add more photons!

This leads to the general idea of REMPI: use one or more resonant intermediate states to dramatically boost the cross-section and selectivity of a multiphoton process. A non-resonant multiphoton absorption is an exceedingly rare event. But if you can make each step of the "photon ladder" land on a real, resonant energy level, the overall probability of the process increases by many, many orders of magnitude.

This leads to a variety of popular laboratory schemes, which are often described using a simple " m + n $m+n$" notation. The first scheme listed is " 1 + 1 $1+1$". This is exactly what we've been discussing as RTPI. It means one photon is absorbed to reach a resonant intermediate state, followed by a second, single photon to cause ionization. This is the simplest and often the most efficient REMPI scheme. But it's just the beginning.

Page 21:

Continuing with the different laboratory schemes for REMPI, we can build more complex "ladders" to climb up to the ionization continuum.

A " 2 + 1 $2+1$" scheme, for example, involves a resonant *two-photon* absorption to reach a high-lying intermediate state, followed by one more photon for ionization. This is often necessary when a single-photon transition to a suitable intermediate state is forbidden by selection rules, but a two-photon transition is allowed.

You can have even more complex schemes. The slide mentions " 3 + 1 $3+1$" or " 2 + 2 $2+2$." These are used in more difficult cases, often to navigate the strict parity selection rules in atoms and centrosymmetric molecules. Remember, a one-photon transition must change the parity of the state (gerade to ungerade, or vice-versa), while a two-photon transition must conserve it (gerade to gerade, or ungerade to ungerade). By combining parity-changing and parity-conserving steps, you can construct a pathway to almost any state.

This brings us to the general nomenclature. An " m + n $m+n$ REMPI" process is defined as a scheme where m m photons are absorbed to reach the first, or final, intermediate resonance, and then n n additional photons are absorbed to ionize the molecule from that state. So, a " 2 + 1 $2+1$" scheme uses two photons for the resonant step and one for ionization, for a total of three photons.

Finally, the slide summarizes the two key features of REMPI that make it so powerful.

The Selectivity—the ability to pick out one specific molecule from a complex mixture—arises from the sharp, narrow nature of the resonant transitions. Only the laser wavelength that exactly matches the energy gap will produce a signal.

The Sensitivity—the ability to detect tiny amounts—arises from the use of high peak laser fluxes, typically from pulsed lasers, which are necessary to drive the multiphoton steps efficiently.

Page 22:

This slide provides beautiful, clear energy-level diagrams illustrating the REMPI schemes we just discussed. Let's examine each of the three panels from left to right. The vertical axis in all of them is Energy.

The first panel is labeled "1+1 REMPI". This is our familiar resonant two-photon ionization. We start in the ground state, S-naught, which is labeled with a '(g)' for 'gerade' parity. A single photon, with energy h v hv, excites the molecule to a real intermediate electronic state, S-one, which has '(u)'

or 'ungerade' parity, consistent with the one-photon selection rule. From S-one, a second photon of the same energy h v $h\nu$ provides enough energy to cross the ionization potential, IP, and create a free electron, e − e^-.

The middle panel shows "2+1 REMPI". Here, we again start in the ground state, S-naught (g). Now, a single long arrow represents the simultaneous absorption of two photons, each with energy h v $h\nu$. This two-photon absorption takes us to a higher-lying intermediate state, S-two. Notice that this state is labeled '(g)' for 'gerade'. This is consistent with the two-photon selection rule: g to g, parity is conserved. From this S-two state, a third photon h v $h\nu$ is absorbed to ionize the molecule. The dashed lines below S-two represent virtual states, emphasizing that the two-photon absorption is a single quantum process.

The final panel shows "3+1 REMPI". It follows the same logic. We start in S-naught (g). Three photons are absorbed simultaneously to reach an even higher intermediate state, S-three, which is labeled '(u)' for 'ungerade'. This is again consistent with selection rules, as a three-photon process changes parity. From this S-three state, a fourth photon h v $h\nu$ causes ionization.

These diagrams perfectly visualize how we can use different numbers of photons to "climb the ladder" of energy levels, using resonances to make the climb efficient.

Page 23: Resonant Multiphoton Ionization (REMPI) Schemes

This slide provides the caption and a crucial piece of context for the diagrams we just saw. It's titled "Resonant Multiphoton Ionization (REMPI) Schemes."

The text reiterates the " m + n $m + n$" pathway concept, stating that the m m photons resonantly excite the molecule to an intermediate electronic state, and the subsequent n n photons provide the energy to ionize it.

But it adds a critical piece of physical insight regarding selection rules. It explicitly states that "Parity rules often dictate the number of photons required for the initial excitation." This is a profoundly important point for anyone designing a REMPI experiment. It explains *why* we need these different " m + n $m + n$" schemes.

It spells out the rule: an odd number of photons (like in 1 + 1 $1 + 1$ or 3 + 1 $3 + 1$ REMPI) connects states of different parity. For a molecule with a center of symmetry, this means a transition from a 'g' (gerade) state to a 'u' (ungerade) state, or vice versa.

In contrast, an even number of photons (like in 2 + 1 $2 + 1$ REMPI) connects states of the same parity. This would be a 'g' to 'g' transition, or a 'u' to 'u' transition.

So, if you want to study a particular excited state, you must first determine its parity. That will tell you whether you need to use a one-photon or a two-photon resonant step to access it from the ground state. These fundamental symmetry rules are not just textbook formalities; they are the practical guide to designing a successful experiment.

Page 24:

Now let's discuss another clever strategy for making the ionization step extremely efficient: using Autoionizing Rydberg States. This is an elegant and powerful route to generating ions.

The idea is as follows. From our intermediate state, E k E_k, we use our second laser to excite the molecule further. But instead of exciting it directly into the flat, unstructured ionization continuum, we tune the laser to be resonant with a very special type of state, which we'll call \(M^{}\).

This \(M^{}\) state is a doubly excited or core-excited Rydberg state. Energetically, it lies *above* the normal ionization potential, I P IP. However, due to electron correlation effects, it is quasi-bound. It exists as a discrete resonance for a short period of time before it decays.

And how does it decay? It decays by a process called autoionization. The state \(M^{}\) spontaneously rearranges its energy and ejects an electron, relaxing to the ground state of the ion, M + M^+, and a free electron, e − e^-.

The key parameter here is the autoionization width, which is represented by Γ A I Γ_{AI}. This is related to the lifetime of the autoionizing state by the uncertainty principle. The slide notes that this width is often very large, corresponding to lifetimes on the order of picoseconds or even femtoseconds. A large width in energy space translates directly to a very large apparent absorption cross-section for that transition.

So, by tuning our ionizing laser to be resonant with one of these autoionizing states, we can make the ionization step thousands of times more probable than ionizing into the unstructured continuum next to it. It's a

way of using a resonance to dramatically enhance the ionization process itself.

Page 25:

Let's continue with the significant advantages of using an autoionization route.

The first point is a direct consequence of the large cross-section we just discussed. The required ionizing laser intensity is orders-of-magnitude lower than what is needed for direct bound-to-continuum transitions. This is a huge practical benefit. You can get away with a much less powerful, and therefore often less expensive and complex, laser for your ionization step. It makes the experiment easier and more robust.

The second bullet point details the practical advantages that stem from using lower laser power. You get reduced power broadening. Power broadening is the unwelcome widening of your spectral lines caused by very intense laser fields. By using lower power, your spectral features remain sharp and narrow, leading to higher resolution. You also get less fragmentation. Hitting a molecule with a very intense laser can be like using a sledgehammer; you can blast it into many different pieces. Using a more gentle, resonant autoionization pathway often leaves the parent ion intact, which is critical for mass spectrometry. Better molecular integrity leads directly to better mass resolution.

Because of these compelling advantages, this technique is frequently exploited in two major areas. The first is trace-analysis, where you need the absolute highest sensitivity to find a needle in a haystack. The second is in

high-resolution molecular-beam experiments, where preserving the spectral line shapes and the parent molecule identity is paramount for extracting detailed physical information about the molecule's structure and dynamics.

Page 26:

This slide shows us what an autoionizing resonance actually looks like in a spectrum. The graph is titled "Breit-Wigner Profile of an Autoionizing Resonance." This is the characteristic lineshape for such a resonance.

Let's examine the plot. The horizontal axis is Energy, typically the energy of the ionizing photon. The vertical axis is the Ionization Cross-Section, σ σ.

First, look at the flat, light blue line at the bottom, labeled "Direct Ionization Continuum." This represents the small, non-zero probability of ionizing the molecule directly into the continuum if you are *not* on resonance. It's the baseline signal. Its energy starts at the ionization potential, IP.

Now, superimposed on this flat baseline is a large, sharp, bell-shaped peak. This is the autoionizing resonance, labeled here as \(M^{}\). When you tune your laser energy to match the energy of this quasi-bound state, the ionization cross-section increases dramatically. The annotation "Enhanced Absorption" points to this huge increase in probability.

The width of this peak, measured at half its maximum height, is precisely the autoionization width, Γ A I Γ_{AI}, which we discussed earlier. The peak's shape is described by a Breit-Wigner, or sometimes a Fano, profile. The key takeaway is that by tuning your laser to the peak of this resonance, you can achieve a massive enhancement in your ion signal compared to the

off-resonance baseline. This is the power of resonant enhancement in action, applied to the ionization step itself.

Page 27:

Now, for the sake of completeness and contrast, let's consider the case of Non-Resonant Two-Photon Ionization. This is what happens when you need two photons for the ionization step, but there is no convenient intermediate resonance to use.

The reaction pathway is shown first. We start with our excited molecule, M ∗ (E k) $M^*(E_k)$. It then absorbs two photons from our ionizing laser, 2 h v 2 $2\,h\nu_2$, to produce the ion M + M^+ and an electron e − e^-. This is a non-linear optical process, fundamentally different from the sequential absorption of two photons in a 1 + 1 + 1 $1+1+1$ scheme.

The second bullet point is the most important one. The cross-section for this process scales differently. It's a generalized cross-section, σ (2) $\sigma^{(2)}$, and it scales proportionally with the intensity of the laser itself, I L 2 I_{L2}. This means the more intense the laser, the more probable the transition. However, the intrinsic probability is incredibly low. The slide gives typical values for σ (2) $\sigma^{(2)}$ in the range of 10 − 50 10^{-50} to 10 − 48 10^{-48}, with the unusual units of c m 4 s $\mathrm{cm}^4\,\mathrm{s}$. This is an astronomically small number compared to a single-photon cross-section of 10 − 17 10^{-17}.

Because the cross-section is so tiny, this process is only viable with extremely intense pulsed lasers. The slide gives a benchmark of Gigawatts per square centimeter. This is the realm of Q-switched Nd:YAG lasers or amplified Ti:sapphire lasers and their harmonics. You simply cannot

achieve this with continuous-wave lasers. This process is a testament to the brute force of high-peak-power lasers.

Page 28:

Even though non-resonant two-photon ionization is a very weak process, it has some specific uses and implications.

The first point is that even without a resonance, it can provide background-free detection. Why? Because the process requires the absorption of two photons simultaneously. This means the two photons must arrive at the molecule at the same time and in the same place. This strict requirement for temporal and spatial overlap means that the signal is generated only in the tiny volume where the laser is most intense, and only during the very brief laser pulse. This makes it highly resistant to background from stray light or other, linear processes.

The second, and perhaps more important role, is that it serves as a benchmark. When you are performing a REMPI experiment and scanning your laser wavelength, you might see a spectrum that consists of sharp, intense peaks on top of a flat, weak baseline. Those sharp peaks are your resonant enhancements—the signal from your 1+1 or 2+1 REMPI process. The weak, flat baseline is the signal from the non-resonant multiphoton ionization process. Therefore, the strength of the non-resonant background provides a clear reference against which you can measure the enhancement factor of your resonant signal. It helps you distinguish the interesting spectroscopy from the uninteresting, brute-force ionization.

Page 29:

We will now shift gears and consider a completely different mechanism for the ionization step: Collision-Induced Ionization. Here, the energy required to dislodge the electron comes not from another photon, but from the kinetic or internal energy of a colliding particle.

The first process shown is electron impact on an excited molecule. The reaction is:

M ∗ (E k) + e s l o w − → M + + 2 e −

$$M^*(E_k) + e^-_{slow} \rightarrow M^+ + 2\,e^-$$

Here, our laser creates the excited molecule \(M^*), *which then collates with a background electron in the system. The electron transfers enough energy in the collision to knock out the valence electron of* \(*M*^\).

A second, very important collisional process is Penning ionization. This involves a metastable collision partner, which we'll call A ∗ A^*. The reaction is:

M + A ∗ → M + + A + e −

$$M + A^* \rightarrow M^+ + A + e^-$$

In this case, the energy stored in the metastable \(A^*) *is transferred to* M M during a collision, causing M M to be ionized. The laser's role here might be to produce the M M in a specific state, or in a related process called optogalvanic spectroscopy, to perturb the population of \(A^\).

What are the requirements for these processes to be effective?

The first requirement is an energy balance. The excited level, either E k E_k of the molecule or E A E_A of the metastable partner, must be close to or above the ionization potential of the target molecule. Specifically, the slide says the excited level should be within a few electron-volts of the I P IP. The closer it is, the more likely the process.

Page 30:

Continuing with collision-induced ionization, there are further requirements and contexts where this process becomes important.

First and foremost, you need a sufficient density of the colliding particles. For electron impact, you need a source of free electrons. For Penning ionization, you need a high density of your metastable partner species. This is not a technique you would typically use in the ultra-high vacuum of a molecular beam experiment.

Instead, this process is dominant in environments like gas discharges and hollow cathode lamps. These are plasmas, which are veritable soups of ions, electrons, and excited atoms. In such an environment, if you shine a laser in that is resonant with a transition, you can change the population of an excited state. This, in turn, changes the overall rate of collisional ionization in the plasma. A change in the ionization rate leads to a change in the plasma's conductivity or impedance. This change in electrical properties can be measured as a voltage or current change, and this is known as an optogalvanic signal. The signal is directly proportional to the change in the excited-state population you induced with your laser.

So, how does this method stack up? The final bullet point provides the verdict. It is generally less sensitive than the purely photon-based methods like REMPI that we've been discussing. The cross-sections for collisions are often smaller, and it's harder to control the background. However, it is an extremely valuable tool for plasma diagnostics—for measuring temperatures, species concentrations, and energy transfer processes inside complex and important environments like discharges, flames, and industrial plasmas.

Page 31

This slide gives us a clear schematic of a device where collision-induced ionization is the dominant process: a Hollow Cathode Lamp. This is the workhorse light source for atomic absorption spectroscopy, but it's also a fantastic environment for performing optogalvanic spectroscopy.

Let's look at the components. The entire device is housed in a glass envelope, which is filled with a low-pressure noble gas, like neon or argon. It has transparent windows, here made of quartz, on both ends to allow a laser beam to pass straight through.

Inside, we have two electrodes. The anode, on the right, is a simple pin held at a positive potential. The cathode, on the left, is the key component. It's a cylindrical cup made of the metal you want to study, and it's held at a negative potential. This is the "hollow cathode."

When a voltage is applied, a gas discharge is initiated. Due to the geometry of the hollow cathode, the discharge is concentrated inside the cup, forming a bright region of plasma called the "negative glow." This glow is rich in

electrons, ions of the noble gas, and, crucially, atoms of the cathode material that have been sputtered off the surface by ion bombardment.

The laser beam, shown as a thick orange line, is directed right through the center of this negative glow. When the laser is tuned to a resonance of the sputtered atoms, it excites them. This changes their probability of being ionized by collisions with the surrounding electrons and ions. This change in the overall ionization rate of the plasma is measured as a change in the current flowing between the anode and cathode, giving us our optogalvanic signal.

Page 32:

We now turn to the last of the four ionization mechanisms we introduced: Field Ionization. This is a particularly elegant and powerful technique, especially for studying highly excited Rydberg states. This slide provides the qualitative picture.

The fundamental principle is that an external electrostatic field can dramatically alter the potential energy landscape of an atom. Specifically, it "tilts" the atom's Coulomb potential. As we'll see in a diagram shortly, this tilting lowers the potential energy barrier on one side of the atom, creating a "saddle point." This lowered barrier provides an escape route for a weakly bound electron.

This method is particularly effective for long-lived Rydberg states. These are states with a large principal quantum number, n n. In such a state, the electron is, on average, very far from the nucleus and is only very weakly

bound. It doesn't take much of a nudge to set it free, and the tilted potential from an external field provides exactly that nudge.

How strong of a field do we need? A semiclassical derivation, which we'll look at in more detail, shows that the minimum field magnitude required depends on how close the Rydberg level is to the ionization potential, I P IP. The remarkable result is that even modest electric fields can ionize levels that are lying only a few milli-electron-volts, or meV, below the I P IP. This gives us a highly selective way to detect only those atoms that are in these very high-lying states.

Page 33: Continuing with the practical aspects of Field Ionization

How is this implemented experimentally? It's typically done using a set of parallel plates or fine wire grids that are placed in the vacuum chamber, situated directly after the laser excitation zone. The laser excites the atoms, and then they drift into the region between the plates where the electric field is applied.

This leads to a crucial experimental advantage, highlighted in the second bullet point. This technique allows for zero-background detection. This is achieved by pulsing the electric field. The field is kept OFF during the laser excitation pulse. Then, after the laser is gone, a fast, high-voltage pulse is applied to the plates. The ions are only created and collected when the field is ON. So, any signal you get is perfectly synchronized with the field pulse, eliminating any background from stray ions or other processes.

Perhaps even more importantly, this pulsed-field technique means that the initial laser excitation occurs in a completely field-free environment. This is critical for high-resolution spectroscopy. If an electric field were present during excitation, it would cause a Stark shift and broadening of the atomic energy levels, distorting the very spectrum you're trying to measure. By separating the excitation event in time from the ionization event, we get the best of both worlds: a clean, unperturbed spectrum from the excitation step, and a highly efficient, zero-background detection from the field ionization step.

Page 34:

This slide presents the essential qualitative picture of field ionization in a single, powerful graph. Let's walk through it carefully, as it visualizes the physics we've been describing.

The plot shows Potential Energy on the vertical axis versus the distance from the nucleus, r r, on the horizontal axis.

First, let's identify the different potentials. The blue curve, which goes as − 1 / r $-1/r$, is the pure, unperturbed Coulomb potential of the atom. This is the potential that binds the electron. The horizontal dashed line at the top represents the ionization potential, I P IP, which is the energy of a free electron at infinite distance.

Next, the straight, downward-sloping red line represents the potential added by the external electric field, which is linear with distance, − E r $-E\,r$.

The most important curve is the "Tilted Potential," which is the sum of the Coulomb potential and the external field potential. You can see that on the right side, this potential no longer goes up to the I P IP at infinity. Instead, it reaches a maximum value at a certain distance, labeled r s a d d l e r_{saddle}, and then rolls over and goes down. This maximum is the "Lowered Barrier" created by the field.

Now, consider an atom in a high-lying Rydberg state, represented by the horizontal green line. This state has an energy that is below the zero-field I P IP, so in the absence of a field, the electron is bound forever. However, in the presence of the field, the electron in this state now sees a potential barrier of finite height and width in front of it.

This opens up a new decay channel: quantum mechanical tunneling. As the orange arrow indicates, the electron has a non-zero probability of tunneling *through* this lowered barrier and escaping, becoming a free electron. This is field ionization. The higher the Rydberg state (the closer it is to the I P IP), or the stronger the electric field, the thinner and lower this barrier becomes, and the faster the tunneling and ionization will occur.

Page 35:

To understand field ionization quantitatively, we first need a good estimate for the energy of the Rydberg states themselves. This slide begins a semi-classical, Bohr-model-based estimate for the ionization energy, I P IP, of a given Rydberg level.

We consider a Rydberg electron, which is the single, highly excited outer electron. Because it's very far from the nucleus, it sees the nucleus and the

inner "core" electrons as a single point charge. We can describe this with an effective charge, Z e f f e $Z_{\text{eff}}\,e$. For a neutral atom, Z e f f Z_{eff} will be close to 1. This electron has a certain mean orbital radius, r r.

The ionization potential, I P IP, of this electron is the energy required to move it from its orbital radius r r to infinity. This is the work done against the Coulomb force. So, we can write I P IP as the integral of the Coulomb force from r r to infinity. The force is Z e f f e 2 4 π ϵ 0 r 2 $\frac{Z_{\text{eff}}\,e^2}{4\pi\epsilon_0\,r^2}$.

Integrating this expression gives the result shown:

I P = Z e f f e 2 4 π ϵ 0 r

$$IP = \frac{Z_{\text{eff}}\,e^2}{4\pi\epsilon_0\,r}$$

This makes intuitive sense: it's just the potential energy of the electron at radius r r.

Now, to make this useful, we need to relate the radius r r to something we can measure spectroscopically, which is the principal quantum number, n n. As the last bullet point suggests, we relate r r to the *effective* quantum number, n ∗ n^*, which is defined as n − δ $n-\delta$. Here, δ δ is the quantum defect, a correction factor that accounts for the fact that the Rydberg electron's orbit might slightly penetrate the inner electron core, making it more tightly bound than a simple hydrogenic model would predict.

Page 36:

Continuing our Bohr-model estimate, we now make the crucial link between the radius r r and the effective quantum number n $*$ n^*.

From a simple Bohr model analysis, the radius of an electron orbit scales as (n $*$) 2 $(n^*)^2$ and is inversely proportional to the effective nuclear charge, Z e f f Z_{eff}. This gives us the approximate relation shown:

r ≈ a 0 (n $*$) 2 Z e f f .

$$r \approx a_0 \frac{(n^*)^2}{Z_{\text{eff}}}.$$

Here, a 0 a_0 is the Bohr radius, a fundamental constant of atomic physics.

Now, we take this expression for r r and substitute it back into our equation for the ionization potential from the previous page, I P = Z e f f e 2 4 π ϵ 0 r $IP = \frac{Z_{\text{eff}} e^2}{4\pi\epsilon_0 r}$.

After some algebra, this substitution leads to the famous Rydberg formula for energy levels, expressed here in terms of the ionization potential:

I P = R y (n $*$) 2 .

$$IP = \frac{Ry}{(n^*)^2}.$$

Here, R y Ry is the Rydberg constant, which encapsulates all the fundamental constants (e e, ϵ 0 ϵ_0, a 0 a_0, etc.). Its value is given as 13.6057 electron-volts.

This final equation is extremely important. It tells us the binding energy—the ionization potential—of any Rydberg state, just by knowing its effective quantum number \(n^*). *This sets the quantitative baseline for our system*

before* we apply any external electric field. It's the "zero-field" I P IP that will be modified by the field.

Page 37:

Now that we have our zero-field ionization potential, let's add the external electric field back into the picture and calculate how it lowers the I P IP.

As the first point states, we apply an external field E 0 E_0, which we'll say points in the minus x x direction. This adds a potential energy term of − e E 0 x $-eE_0x$ to the total potential.

The second point reminds us that a saddle point appears in the total potential. This is the point in space where the attractive Coulomb force from the nucleus is exactly balanced by the pulling force from the external electric field.

The equation expresses this force balance:

Z e f f e 2 4 π ε 0 x 2 = e E 0 .

$$\frac{Z_{\text{eff}}e^2}{4\pi\varepsilon_0 x^2} = eE_0.$$

We can now solve this equation for the position of the saddle point, x s p x_{sp}. A little bit of rearrangement gives the result on the right:

x s p = Z e f f e 4 π ε 0 E 0 .

$$x_{\text{sp}} = \sqrt{\frac{Z_{\text{eff}}e}{4\pi\varepsilon_0 E_0}}.$$

Now, the final step is to calculate the *value* of the potential energy at this saddle point. This value will be our new, *effective* ionization potential in the presence of the field. This is the energy an electron needs to have to simply spill over the top of the barrier. We'll see the final result on the next slide.

Page 38:

Here we see the final result for the effective ionization potential in the presence of the electric field. The total potential energy at the saddle point, which we call I P e f f IP_{eff}, is the sum of the Coulomb potential and the field potential, both evaluated at the position x s p x_{sp}.

After substituting the expression for x s p x_{sp} from the previous slide and simplifying, we arrive at the expression shown:

I P e f f = I P − Z e f f 3 e 3 E 0 π ϵ 0

$$IP_{\text{eff}} = IP - \sqrt{\frac{Z_{\text{eff}}^3 e^3 E_0}{\pi \epsilon_0}}$$

Whoops, let me re-read that. The slide has a typo. The term inside the square root should be Z e f f e 3 E 0 / (π ϵ 0) $Z_{\text{eff}} e^3 E_0/(\pi\epsilon_0)$. Let's assume the classic result, which is

I P e f f = I P − 2 Z e f f e E 0 4 π ϵ 0

$$IP_{\text{eff}} = IP - 2\sqrt{\frac{Z_{\text{eff}} e E_0}{4\pi \epsilon_0}}$$

The expression shown is off by a numerical factor, but the scaling is what's important. The key point is that the effective ionization potential is the zero-field I P IP *minus* a term that is proportional to the square root of the applied electric field, E 0 E_0. This is the quantitative expression for the field-lowering of the ionization potential.

Now, let's look at the key parameters and what this means for us. The most important relationship is highlighted: A larger effective quantum number \(n^\) *means a smaller initial I P* IP (since I P IP goes as \(1/n^{2}\)). A smaller I P IP means that the energy difference, I P − I P e f f $IP - IP_{\text{eff}}$, that needs to be overcome by the field is smaller. This, in turn, means that a *weaker* electric field, E 0 E_0, will suffice to ionize the atom.

This gives us a quantitative design rule for building detectors for high- n n Rydberg atoms. If we know what n n state we want to detect, we can calculate its I P IP using the Rydberg formula, and then use the equation on this slide to calculate the exact electric field strength we need to apply to ionize it efficiently.

Page 39:

Let's make this concrete with a numerical example to calculate the required field strength.

Suppose we have excited an atom to a level that lies 10 milli-electron-volts, or meV, below the ionization potential. This energy difference, I P − E l e v e l $\mathrm{IP} - E_{\text{level}}$, is the energy that must be supplied by the field-lowering effect. We can call this Δ I P $\Delta\mathrm{IP}$, so Δ I P = 10 m e V $\Delta\mathrm{IP} = 10\,\mathrm{meV}$.

For simplicity, let's consider a hydrogenic atom where the effective charge Z e f f Z_{eff} is equal to 1.

We now take our formula for the field-lowering, Δ I P = C E 0 $\Delta\text{IP} = C\sqrt{E_0}$, where C C is a constant, and we solve it for the required electric field, E 0 E_0. Rearranging gives E 0 E_0 is greater than or equal to (Δ I P) 2 $(\Delta\text{IP})^2$ divided by some constants. The equation shown is

E 0 ≥ (Δ I P) 2 π ϵ 0 e 3 .

$$E_0 \geq \frac{(\Delta\text{IP})^2\,\pi\,\epsilon_0}{e^3}.$$

Plugging in the values for Δ I P ΔIP (10 meV converted to joules), π π, ϵ 0 ϵ_0, and the elementary charge e e, yields a required field strength of approximately 1.7 × 10 4 V / m $1.7 \times 10^4\ \text{V/m}$.

This is a very reasonable electric field. 10 4 V / m $10^4\ \text{V/m}$ is the same as 100 V / c m $100\ \text{V/cm}$.

However, the story gets even better. As the last bullet point notes, this calculation is semi-classical; it only considers an electron spilling over the potential barrier. Quantum mechanics allows the electron to tunnel through the barrier, even if it doesn't have enough energy to go over it. This tunneling effect further lowers the required field strength E 0 E_0.

The practical result is that fields of just a few kilovolts per centimeter, which are very easy to generate in the lab, are sufficient to efficiently ionize these high-lying Rydberg states.

Page 40:

Continuing with the practicalities of field ionization, the required fields are quite accessible.

Modern micro-fabrication techniques allow for the construction of electrodes with very small, precise gaps. It's possible to generate very large electric fields over these sub-millimeter gaps without needing enormous voltages. The slide notes that fields of 10 5 V / m 10^5 V/m, or 1000 V / c m 1000 V/cm, are readily achievable. This technology enables the design of very compact and efficient field ionization sensors.

The conclusion is that field ionization is therefore a routine and robust technique in molecular and atomic beam experiments, particularly for studying Rydberg series in atoms like the alkalis, where it's commonly used for states with principal quantum number n n greater than or equal to about 25 25. For these states and higher, field ionization becomes one of the most efficient and cleanest detection methods available.

Page 41:

Let's now return to our general rate equation model and insert a more explicit expression for the ionization probability, P k I P_{kI}.

We've been treating P k I P_{kI} as a fundamental rate, but we can express it in a more intuitive way using the concept of a cross-section, just as we did for the initial absorption step. This is the photon-flux formulation for the ionization step, which uses laser L-two.

The ionization probability per second, P k I P_{kI}, is given by the product of the photoionization cross-section, σ k I σ_{kI}, and the photon flux density of the ionizing laser, n L 2 n_{L2}.

P k I = σ k I n L 2

$$P_{kI} = \sigma_{kI}\, n_{L2}$$

This is a very intuitive formula. The rate of ionization depends on how big of a "target" the molecule presents (σ k I σ_{kI}, in units of square centimeters) and how many "bullets" per second we are firing at it (n L 2 n_{L2}, in units of photons per square centimeter per second).

The definitions are listed below for clarity: * σ k I σ_{kI} is the photoionization cross-section in square centimeters. * n L 2 n_{L2} is the photon flux density of laser L2 in photons per square centimeter per second.

The final step, as indicated, is to insert this more explicit expression for P k I P_{kI} back into our master equation for the ion signal, S I S_I. This will give us a final expression that shows the explicit dependence on the parameters of both lasers, L1 and L2.

Page 42:

Here we see the result of substituting our photon-flux formulation for P k I P_{kI} into the full signal equation. Let's look at this final form carefully.

The ion signal, S I S_I, is equal to N i N_i times a large fractional term, all multiplied by Δ x Δx.

The numerator of the fraction is σ i k n L 1 δ η $\sigma_{ik}\, n_{L1}\, \delta\, \eta$. This contains all the parameters of the first excitation step and the detection efficiency.

The denominator is 1 + R k σ k I n L 2 .

$$1 + \frac{R_{\mathrm{k}}}{\sigma_{kI}\, n_{L2}}.$$

This term captures the competition between relaxation (R k R_{k}) and ionization, now expressed explicitly in terms of the second laser's flux (n L 2 n_{L2}).

What does this equation tell us? The first bullet point makes it clear: it displays the explicit dependence on *both* lasers. To maximize our signal S I S_{I}, we need to maximize the numerator, which means we need to maximize the flux of the first laser, n L 1 n_{L1}, to drive the excitation. And we need to minimize the denominator, which means we need to maximize the flux of the second laser, n L 2 n_{L2}, to drive the ionization.

This leads to the condition in the second bullet point. If the term σ k I n L 2 $\sigma_{kI}\, n_{L2}$—which is just our ionization rate P k I P_{kI}—is much, much greater than the relaxation rate R k R_{k}, then the ratio in the denominator becomes very small. The entire denominator then approaches unity. When this happens, the competition is won by ionization, and the maximum possible signal for a given excitation rate is reached.

The final point is crucial for experimental design. The strategy difference between using pulsed lasers and continuous-wave, or CW, lasers lies entirely in how they achieve this inequality, σ k I n L 2 ≫ R k $\sigma_{kI}\, n_{L2} \gg R_{\mathrm{k}}$. Pulsed lasers do it with enormous peak power (huge n L 2 n_{L2}) over a

short time. CW lasers must do it with modest power but a 100% duty cycle, often requiring very tight focusing.

Page 43:

Let's revisit the maximum achievable ion rate, which we'll call S I max $S_I^{\max}$, now with our full understanding of the process.

The first point states the ideal conditions: perfect instrumental efficiency (δ = η = 1 $\delta = \eta = 1$) and a strong ionizing laser L 2 L_2, meaning the ionization rate is much faster than any relaxation rate.

Under these conditions, as we've seen, our signal equation simplifies dramatically. The maximum signal, S I max $S_I^{\max}$, is equal to N i N_i times σ i k σ_{ik} times n L 1 n_{L1} times Δ x Δx.

This product is simply n a n_a, the initial rate of photon absorption.

The interpretation, stated in the second bullet, is the most powerful takeaway: every photon that is absorbed in the first spectroscopic step eventually appears as a detected count. This is the absolute quantum limit of detection. We cannot do any better than this.

So, if we want to improve our experiment, our design targets should therefore focus on maximizing this initial absorption rate, n a n_a. The slide lists two key strategies for doing this.

First, we need to enhance the absorption cross-section, σ i k σ_{ik}. This is a spectroscopic choice. We must select a transition from the ground state to an intermediate state that is strongly "allowed" by quantum mechanical selection rules—a transition with a high oscillator strength.

The second strategy involves the laser and the sample environment, and we'll see it on the next page.

Page 44: Continuing with the design targets for achieving the maximum signal rate:

The second critical task is providing sufficient photon flux, n L 1 n_{L1}, from our first laser to drive the absorption, while simultaneously avoiding power broadening. This is a delicate balance. Turning up the laser power increases n L 1 n_{L1} and thus the signal, but if the power becomes too high, it can saturate and broaden the transition, ruining our spectral resolution. Finding the "sweet spot" is a key part of optimizing any spectroscopy experiment.

The third design target is to suppress the non-radiative relaxation rates, R k R_k. These are the loss channels that compete with our ionization step. The main culprit here is often collisional de-excitation, or quenching. We can minimize this by working at very low pressures or, even better, by using a collision-free environment like a molecular beam. This ensures that once a molecule is excited, its primary fate is determined by interaction with photons (either fluorescing or being ionized), not by bumping into its neighbors.

Finally, the ability to reach the theoretical maximum signal, S I , max $S_{I,\max}$, serves as the ultimate yardstick for evaluating the quality of an

experimental geometry. If your measured signal is much lower than the calculated S I , max $S_{I,\max}$, it tells you that something in your experiment is not optimal. Perhaps your lasers are misaligned, your collection efficiency is poor, or you have unexpected relaxation channels. It provides a concrete goal for experimental optimization.

Page 45:

Let's now put some typical numbers to these parameters to get a feel for the real-world requirements. This slide on typical cross-sections and lifetime numbers is incredibly useful for planning an experiment.

First, let's list some representative values.

The photoionization cross-section from an excited state, σ k I σ_{kI}, is on the order of 10 − 17 c m 2 $10^{-17}\,\mathrm{cm}^2$ for near-threshold ionization. This is a reasonably large cross-section.

The radiative decay rate of a typical allowed electronic transition, A k A_k (which is a major component of R k R_k), is on the order of 10 8 s − 1 $10^8\,\mathrm{s}^{-1}$. This corresponds to a natural lifetime, τ k τ_k, of about 10 nanoseconds. This is a very important timescale to keep in mind.

Now, let's look at the crucial inequality we need to satisfy for efficient ionization: the ionization rate, σ k I × n L 2 $\sigma_{kI} \times n_{L2}$, must be greater than the decay rate, A k A_k.

Let's plug in our typical numbers.

We need (10 − 17 c m 2) n L 2 > 10 8 s − 1 .

$$(10^{-17}\ \mathrm{cm}^2)\, n_{L2} > 10^8\ \mathrm{s}^{-1}.$$

Solving this for the required photon flux density of our second laser, n L 2 n_{L2}, we find that n L 2 n_{L2} must be greater than 10 25 10^{25} photons per square centimeter per second.

This is an enormous number. Can a typical laboratory laser actually deliver this kind of flux? Let's find out on the next slide.

Page 46:

So, the challenge is to achieve a photon flux of 10 25 10^{25} photons per square centimeter per second. Can a pulsed laser do it?

The first bullet point gives us a typical set of parameters for a common pulsed laser system, like a Q-switched Nd:YAG-pumped dye laser. It can readily supply about 100 millijoules of energy in a 10 nanosecond pulse, spread over an area of 1 square centimeter.

Let's calculate the photon flux, n L 2 n_{L2}, from these numbers. The equation is n L 2 n_{L2} equals the total energy per pulse, divided by the energy of a single photon (h v) $(h\nu)$, all divided by the pulse duration and the area.

n L 2 = 100 m J h v 10 n s × 1 c m 2

$$n_{L2} = \frac{\frac{100\ \mathrm{mJ}}{h\nu}}{10\ \mathrm{ns} \times 1\ \mathrm{cm}^2}$$

Assuming a visible photon with an energy of about 2 to 3 electron-volts, this calculation yields a photon flux of approximately 2 × 10 25 2×10^{25} photons per square centimeter per second.

This is fantastic news! Our typical pulsed laser *does* meet the required flux condition.

Now, what does this mean for our ionization efficiency? The second bullet point shows the consequence. With this flux, our ionization rate P k I P_{kI} (which is σ k I × n L 2 $\sigma_{kI} \times n_{L2}$) is approximately 10 − 17 × 2 × 10 25 $10^{-17} \times 2 \times 10^{25}$, which equals 2 × 10 8 2×10^{8} inverse seconds. This rate is of the same order of magnitude as the radiative decay rate A k A_k, which was 10 8 10^{8} inverse seconds.

Our ionization yield, which is P k I P k I + A k $\frac{P_{kI}}{P_{kI}+A_k}$, is therefore approximately 2 × 10 8 2 × 10 8 + 1 × 10 8 $\frac{2\times10^{8}}{2\times10^{8}+1\times10^{8}}$, which is 2 divided by 3, or about 0.67.

This means we are yielding an ion signal S I S_I that is approximately 67 percent of the theoretical maximum, S I max $S_{I\max}$. This is exceptionally good efficiency.

The final point is that these kinds of simple, back-of-the-envelope numerical checks are an essential part of experimental physics. They guide our choice of lasers and help us predict the feasibility and performance of an experiment before we even start building it.

Page 47:

This Paage is blank, likely serving as a separator. So we will proceed directly to the next topic.

Page 48:

We've just seen that pulsed lasers are very effective at providing the high photon flux needed for efficient ionization. This slide summarizes the key advantages of using pulsed lasers in Resonant Two-Photon Ionization, or RTPI.

First, and most obviously, they provide extremely high peak power. This translates directly into a large photon flux for both the excitation laser n L 1 n_{L1} and the ionization laser n L 2 n_{L2}, which is the key to driving both steps of the process efficiently.

Second, because the laser pulse is very short—typically on the order of nanoseconds—ionization can occur *before* the excited state has a chance to decay via spontaneous emission. We saw this in our numerical example: the 10 nanosecond pulse duration is comparable to the 10 nanosecond radiative lifetime. This ensures that we don't lose our excited state population to fluorescence, which maximizes the overall ionization efficiency.

Third, there's a consideration of spectral bandwidth. Pulsed lasers are not perfectly monochromatic. Due to the Fourier uncertainty principle, a pulse of finite duration Δ T ΔT must have a minimum spectral bandwidth Δ v $\Delta \nu$ that is on the order of 1 Δ T $1/\Delta T$. For a 10 nanosecond pulse, this is about 100 MHz. This is broader than the bandwidth of a good CW laser,

but it is often perfectly acceptable, especially if the spectral width of the transition you are studying is broader than this Fourier limit.

Fourth is the duty cycle. Pulsed lasers have a very low duty cycle. A laser firing at 10 to 100 Hertz with a 10 nanosecond pulse duration has a duty cycle on the order of $10 - 7\ 10^{-7}$ to $10 - 6\ 10^{-6}$. This means the laser is off for the vast majority of the time. This is a huge advantage for reducing the average power load and heating on the sample, which is especially important for delicate species.

Page 49:

Here we have one more crucial advantage of using pulsed lasers, particularly in the context of beam experiments. The ability to synchronize the firing of a pulsed laser with other pulsed components, like a pulsed molecular or atomic beam source, is a game-changer. As we'll discuss in more detail shortly, this allows for highly efficient use of the sample.

Furthermore, the pulsed nature of the ionization event itself provides a precise "start time," $t = 0$, for the ions. This is the cornerstone of time-of-flight, or TOF, mass spectrometry. By measuring the time it takes for an ion to travel from the ionization region to a detector, we can determine its mass-to-charge ratio. The sharp, well-defined start time provided by a short laser pulse is essential for achieving high mass resolution in TOF experiments. This temporal discrimination is one of the most powerful synergies between pulsed lasers and ionization spectroscopy.

Page 50:

Let's now dig into some of the geometrical considerations when we combine our pulsed laser with a molecular beam. This example will help us understand the spatial and temporal scales involved.

Let's assume some typical experimental parameters. First, our laser beam has a diameter, D D, of 1 centimeter. This is a reasonably large, unfocused beam from a pulsed laser. Second, the pulse duration, Δ T ΔT, is 10 nanoseconds. Third, the molecules in our beam are moving with a mean velocity, v ¯ $\bar{v}$, of 500 meters per second. This is a typical speed for a light molecule in a supersonic jet expansion.

Now, the critical question: how far does a molecule travel during the time the laser is actually on? We can calculate this distance, d d, very simply: d = v ¯ × Δ T

$$d = \bar{v} \times \Delta T$$

Plugging in the numbers: d = (500 m / s) × (10 × 10 − 9 s) = 5 × 10 − 6 m

$$d = (500\ \mathrm{m/s}) \times (10 \times 10^{-9}\ \mathrm{s}) = 5 \times 10^{-6}\ \mathrm{m}$$

which is equal to 5 × 10 − 4 5×10^{-4} centimeters.

So, during the entire 10 nanosecond laser pulse, a typical molecule moves only 5 micrometers. This distance is absolutely tiny compared to the 1 centimeter diameter of the laser beam.

Page 51

What is the consequence of the calculation we just performed?

The result is summarized in the single bullet point on this slide: The entire cross-section of the molecular beam is illuminated by the laser, and essentially all the molecules that are inside the laser volume can be considered "frozen" in place and can be ionized during the pulse.

Because the distance a molecule travels during the pulse (5 micrometers) is negligible compared to the size of the laser beam (1 centimeter), we don't have to worry about molecules moving into or out of the beam while the laser is on. The interaction volume is well-defined and static for the duration of the pulse. This "snapshot" nature of pulsed laser interrogation is a great simplification and ensures that we are probing a well-defined ensemble of molecules with each laser shot. This allows for a clean interpretation of the experiment.

Page 52:

Now we must consider the flip side of the pulsed laser's low duty cycle: the long "dark time" between the pulses. This introduces a significant limitation that we must address.

Let's use a typical laser repetition rate, f L f_L, of 100 H z $100\,\mathrm{Hz}$. The time interval, T T, between consecutive pulses is simply 1 / f L $1/f_L$, which is 1 / 100 $1/100$, or 10 m s $10\,\mathrm{ms}$. This is our dark interval.

Now, what are the molecules in our beam doing during this 10 m s $10\,\mathrm{ms}$ period when the laser is off? They are still travelling at their mean velocity of 500 m / s $500\,\mathrm{m/s}$.

Let's calculate the distance they travel during this dark time, d dark d_{dark}.

d dark = v ¯ T = (500 m / s) × (0.01 s) = 5 m .

$$d_{\text{dark}} = \bar{v}\,T = (500\ \text{m/s}) \times (0.01\ \text{s}) = 5\ \text{m}.$$

Five meters! This is an enormous distance. If we have a continuous molecular beam flowing through our 1 c m $1\,\text{cm}$-wide interaction region, the vast majority of the molecules will simply fly through the region during the dark time without ever seeing a laser photon.

The slide quantifies this by calculating the fraction of time the laser is on, which is 10 n s / 10 m s = 10 − 6 $10\,\text{ns}/10\,\text{ms} = 10^{-6}$. A better way to think about it is comparing the interaction length (1 c m $1\,\text{cm}$) to the distance a new set of molecules travels between pulses. The number given, 1 / 500 $1/500$, suggests a beam length of 500 c m $500\,\text{cm}$. The key point is that with a continuous beam, we are only interrogating a tiny fraction of the total number of molecules available. This is incredibly inefficient.

So, what are the mitigation strategies? How can we fix this problem?

Page 53:

Here are the two primary strategies for overcoming the problem of the long dark time between pulses.

The first, and most common, solution is to use a pulsed molecular beam and synchronize it with the laser. Instead of a continuous stream of molecules, we use a special valve that lets out a short puff of gas, say a few hundred microseconds long. We time the firing of our laser pulse to coincide perfectly with the arrival of this dense packet of molecules in the

interaction region. The condition is that the duration of the gas pulse, Δ T B ΔT_B, should be less than or equal to the time it takes for the gas packet to traverse the laser beam, which is D / v ¯ $D/\bar{v}$. This ensures that the entire sample is present when the laser fires, maximizing the overlap and efficiency.

The second strategy is an alternative that can be used with high-repetition-rate lasers. You can direct the laser beam to be antiparallel to the molecular beam—that is, the laser and molecules are flying directly towards each other. Then, you increase the laser repetition rate, f L f_L, as high as possible, for example, up to 10 kilohertz or more. By firing more frequently, you increase the probability that a given molecule will be "caught" by a laser pulse as it travels through the interaction zone.

As the final point summarizes, the goal of both strategies is the same: to ensure that the molecules remain within the excitation zone for successive pulses, thereby dramatically improving the sampling efficiency of the experiment compared to a simple continuous beam and low-repetition-rate laser setup.

Page 54:

So far, we've focused heavily on pulsed lasers. But what about Continuous-Wave, or CW, lasers? Can they be used for ionization spectroscopy? The answer is yes, but the requirements are quite different.

The first, and most obvious, advantage of a CW laser is its 100% duty cycle. There are no dark gaps. The laser is always on. This makes it an excellent choice for samples that are not in a fast beam, but are in a slow

diffusion environment, like a static gas cell. With a CW laser, you are guaranteed to eventually interrogate every molecule that wanders into the beam.

However, CW lasers come with a major challenge: their power is much lower than the peak power of a pulsed laser. The slide gives an example of a typical, powerful CW laser, like an Argon ion laser, which might produce several Watts of power at a visible wavelength like 488 nanometers.

Now, let's recall the photon flux we needed to achieve efficient ionization. Our target was n L 2 = 10 25 $n_{L2} = 10^{25}$ photons per square centimeter per second. To reach this colossal flux with just a few Watts of power, we have no choice but to use tight focusing. We must squeeze all of that power down into an incredibly small spot size. Let's see what that looks like numerically.

Page 55:

Let's do the calculation for focusing a CW laser. The example considers a CW laser that produces 2.5 × 10 19 2.5×10^{19} photons per second. This corresponds to roughly 1 Watt of power in the visible spectrum. To reach our target flux n L 2 n_{L2} of 10 25 10^{25}, we need to focus these photons onto a spot with an area of 2.5 × 10 19 10 25 $\frac{2.5\times10^{19}}{10^{25}}$, which is 2.5 × 10 − 6 2.5×10^{-6} cm 2 ^{2}.

This is a tiny area. A spot with this area has a radius of only about 2.8 micrometers. This is the scale of the challenge.

This extremely tight focus creates a new problem, which is outlined in the second bullet point. An excited molecule has a typical lifetime of about 10 nanoseconds. If that molecule is moving at a typical thermal velocity, it might travel several micrometers in that time. This means the excited molecule can move *out* of the tiny, micrometer-sized focal volume of the ionizing laser before it even has a chance to be ionized! This "transit-time" effect is a major hurdle. To overcome it, both the excitation laser, L1, and the ionization laser, L2, must be focused to the same microscopic spot and must coincide spatially with a precision of a few micrometers. This is a formidable alignment challenge.

So what's the solution? The final point describes a clever optical engineering trick. The problem is matching a point-like laser focus with a line-like molecular beam. The solution involves using optical fibers to deliver the laser beams, which provides excellent stability and beam quality, and then using a cylindrical lens. A cylindrical lens focuses the beam in only one dimension. This can transform a circular laser spot into a thin, horizontal "light sheet." This light sheet can then be perfectly matched to the dimensions of the molecular beam, creating a much larger interaction volume while still maintaining the high intensity needed for ionization.

Page 56: Optical Layout: Fiber-Delivered CW Lasers with Cylindrical Focusing

This slide provides a visual schematic of the solution we just discussed, titled "Optical Layout: Fiber-Delivered CW Lasers with Cylindrical Focusing."

Let's follow the light paths. We see two laser beams, L1 and L2, emerging from fiber outputs. This ensures they are stable and have a nice, clean Gaussian spatial profile. The two beams are then combined using a Dichroic Mirror. This is a special mirror that reflects one wavelength (L1, in this case) while transmitting another (L2). This allows us to make the two laser beams perfectly collinear.

The combined beams then pass through a Cylindrical Lens. Notice how the lens is curved in the vertical direction but flat in the horizontal direction. This optic squeezes the beam vertically while leaving it wide horizontally, transforming the focused spot into a line, or a "Light Sheet Focus."

This light sheet is directed into the interaction region, where it intersects a Molecular Beam, which is shown travelling from right to left. This geometry creates a well-defined interaction region where the high-intensity light sheet perfectly overlaps with the ribbon-shaped molecular beam, solving the transit-time problem.

The inset at the bottom shows the desired outcome. It plots the intensity profiles of the two lasers, L1 and L2, as a function of position. The goal is to make these two profiles, these two Gaussian peaks, overlap as perfectly as possible to ensure that any molecule excited by L1 is immediately in the high-intensity region of L2, ready for ionization.

Now let's consider how we can optimize the overlap and power of our two CW beams to maximize the signal.

The first point notes that the Gaussian intensity profile of the excitation laser, L1, can often be strong enough to saturate the $|\ i>$

to $|\ k>$

transition. Saturation means we are pumping molecules into the excited state as fast as they can possibly be excited. In the center of the Gaussian beam where the intensity is highest, we might have very strong saturation.

This leads to a clever optimization strategy mentioned in the second bullet point. Instead of placing the peak of the ionizing laser, L2, directly on top of the peak of L1, it can be advantageous to place the L2 peak on the *shoulder* of the L1 profile. Why would we do this? In the center of the L1 beam, the transition is already saturated, so adding more L1 intensity doesn't help. By placing the L2 beam on the shoulder, we can create a more uniform ionization probability across the entire width of the beam. This ensures that we are efficiently ionizing molecules from the wings of the beam as well as from the center, maximizing the total signal from the entire interaction volume.

Third, we can use imaging techniques and detectors to perform a spatial mapping of the two laser profiles to confirm that we have achieved at least 95% overlap of the high-intensity regions.

Finally, as we've seen before, using an autoionizing Rydberg resonance can be a lifesaver in CW experiments. It can reduce the required L2 laser power by a factor of 100 to 1000. This dramatically relaxes the focusing

requirements and makes the entire experiment much more feasible with standard CW lasers.

Page 58:

Here we have a wonderfully practical tip for fine-tuning the experiment in real time.

Fine adjustment of the spatial overlap of the two laser beams can be achieved by maximizing the ion signal while simultaneously monitoring the fluorescence from the intermediate state, | k ⟩ $|k\rangle$.

Let's think about this. The fluorescence signal—photons emitted from state | k ⟩ $|k\rangle$—depends only on the first laser, L 1 L_1. It tells you how many molecules you are successfully exciting to the intermediate state. The ion signal, on the other hand, depends on *both* L 1 L_1 and L 2 L_2. It tells you how many of the excited molecules you are successfully ionizing.

So, the procedure would be: first, you optimize the fluorescence signal by adjusting the wavelength and focus of L 1 L_1. This tells you that you are efficiently creating the excited state | k ⟩ $|k\rangle$. Then, while watching that fluorescence signal to make sure it stays constant, you carefully adjust the position and focus of the second laser, L 2 L_2, until the ion signal reaches a maximum. When the ion signal is maximized for a given fluorescence level, you know you have achieved the best possible spatial overlap between your excited state population and your ionizing laser beam. This is a classic and very effective alignment technique used in two-color experiments.

Page 59:

We now return to the ultimate question of sensitivity and the fundamental limits imposed by quantum mechanics. This slide looks at the limit that arises from the principle of single-photon counting.

We are considering the ideal case, where every single photon absorbed by our sample yields a detected count. The question is: what is the minimum detectable absorber density under this perfect condition?

Let's start from the Beer-Lambert law. The power absorbed by the sample, P a b s P_{abs}, is P 0 N i σ i k L $P_0 N_{\mathrm{i}} \sigma_{ik} L$, where P 0 P_0 is the incident power. If our detection limit is the absorption of a single photon of energy h v $h\nu$ in some integration time, then the minimum detectable density N i N_{i} is the density that produces that single absorption event.

Rearranging the Beer-Lambert law for N i N_{i} gives the inequality shown:

N i ≥ h v P 0 σ i k L

$$N_{\mathrm{i}} \geq \frac{h\nu}{P_0 \sigma_{ik} L}$$

P 0 P_0 is the incident laser power in Watts, σ i k σ_{ik} is the cross-section, and L L is the path length.

Now, we can express the incident power P 0 P_0 in terms of the photon flux n L 1 n_{L1}. Power is just photon flux times the energy per photon:

P 0 = n L 1 h v

$$P_0 = n_{L1}\, h\nu$$

The final bullet point tells us to substitute this expression into our inequality. When we do this, a beautiful simplification occurs.

Page 60:

Here is the result of that substitution. When we replace P 0 P_0 with n L 1 h v $n_{L1}h\nu$, the h v $h\nu$ term in the numerator cancels with the h v $h\nu$ term in the denominator. This leaves us with an elegant and profoundly simple expression for the fundamental sensitivity limit:

N i ≥ 1 n L 1 σ i k L

$$N_i \geq \frac{1}{n_{L1}\,\sigma_{ik}\,L}$$

Let's analyze this. Our sensitivity—our ability to detect a small number density—scales inversely with three key parameters: 1. The photon flux of our excitation laser, n L 1 n_{L1}. The more photons we send, the higher the probability of an absorption, and the lower the density we can see. 2. The absorption cross-section, σ i k σ_{ik}. The stronger the transition, the easier it is to detect. 3. The interaction path length, L L. A longer path length means the laser interacts with more molecules, increasing our signal.

Now for a stunning numerical example. Let's assume a very strong absorption cross-section, σ i k σ_{ik}, of 10 − 18 10^{-18} square centimeters (typical for an allowed atomic transition). Let's use a moderate CW photon flux, n L 1 n_{L1}, of 6.5 × 10 16 6.5×10^{16} photons per cm-squared per second (this is about 30 milliwatts of power). And let's assume a multi-pass cell with an effective path length L L of 5 centimeters.

Plugging these numbers into our equation gives a minimum detectable density N i N_i of just 3 3 molecules per cubic centimeter. This is an absolutely astonishing number. We are talking about detecting a handful of molecules in a volume the size of a sugar cube.

To drive the point home, if our interaction volume is half a cubic centimeter, this density means we only need, on average, 1.5 1.5 molecules to be present in the laser beam at any given time to get a detectable signal. We are truly in the realm of single-molecule detection.

Page 61:

Let's refine this sensitivity calculation with a more realistic example for a molecular beam experiment, where we are truly pushing towards single molecule detection.

In a crossed laser-molecular beam setup, the interaction path length L L is determined by the width of the molecular beam, which is typically much smaller than in a gas cell. A typical value might be L = 0.2 $L = 0.2$ centimeters.

Let's use the same laser flux n L 1 n_{L1} and cross-section σ i k σ_{ik} as in the previous example. The only thing that has changed is L L. Our minimum detectable density N i N_i scales as 1 / L $1/L$. Since L L is now 25 times smaller (0.2 c m $0.2\ \mathrm{cm}$ vs 5 c m $5\ \mathrm{cm}$), our minimum detectable density will be 25 times higher. The calculation N i ≥ 75 $N_i \geq 75$ per cubic centimeter confirms this.

Now, let's calculate the actual interaction volume. Let's say we focus our laser down to a 1 1 millimeter diameter. The cross-sectional area A A of the laser is π r 2 πr^2, which is about 7.7 × 10 − 3 7.7×10^{-3} square centimeters. The interaction volume V V is the area A A times the path length L L.

V = (7.7 × 10 − 3 c m 2) × (0.2 c m) = 1.5 × 10 − 3 c m 3 .

$$V = (7.7 \times 10^{-3}\ \text{cm}^2) \times (0.2\ \text{cm}) = 1.5 \times 10^{-3}\ \text{cm}^3.$$

So we have a minimum detectable density and a well-defined interaction volume. Now we can calculate the minimum number of molecules we need in that volume.

Page 62:

Here we perform the final, simple calculation to find the minimum number of detectable molecules in our molecular beam example.

The minimum number of molecules, which we'll call N min N_{min}, is simply the minimum detectable number density, N i N_{i}, multiplied by the interaction volume, V V.

N min = N i × V

$$N_{\text{min}} = N_{\text{i}} \times V$$

Plugging in the numbers from the previous slide:

N min = (75 m o l e c u l e s / c m 3) × (1.5 × 10 − 3 c m 3)

$$N_{\text{min}} = (75\ \text{molecules/cm}^3) \times (1.5 \times 10^{-3}\ \text{cm}^3)$$

This gives a value of approximately 0.11 molecules.

What does it mean to detect 0.11 molecules? It means that if we have, on average, just one-tenth of a molecule in our detection volume at any given time, we can still get a signal.

In reality, of course, molecules are discrete. This result means that our detection probability is high, but the probability of a molecule even being in the beam is low. We are sensitive enough that we are limited by the "shot noise" of the molecules themselves.

The final sentence puts this in perspective. When we account for the fact that our collection efficiency and detector efficiency might be slightly less than unity, and there might be some detector noise, this calculation shows that we are truly approaching the regime of single-molecule sensitivity. The ability to detect the presence of a single, specific molecule in a specific quantum state is the ultimate achievement in analytical spectroscopy, and ionization methods get us there.

Page 63:

This slide offers a concise comparison and summary of the different detection mechanisms we've explored, highlighting their strengths and weaknesses.

First, let's review Collision-induced ionization. As we saw, this method relies on an external flux of colliding particles, like electrons or metastables. The cross-sections and available particle densities are generally lower than what can be achieved with intense lasers. Therefore, it's not the technique of choice for ultimate sensitivity. However, it is extremely valuable for

energy-transfer studies and for diagnostics in plasma environments where collisions are unavoidable and, in fact, the subject of interest.

Next, we have Field ionization. Its unique strength is its state-selectivity for high-lying Rydberg atoms. Essentially every Rydberg atom with an energy above the effective, field-lowered ionization potential, IP_{eff}, is ionized. This makes it an excellent and unparalleled population probe for these highly excited states. The slide notes that states with principal quantum number n up to 300 have been measured using this technique, which is a testament to its precision and power for studying the near-continuum structure of atoms.

Finally, we'll summarize the workhorse method, RTPI/REMPI.

Page 64:

This slide concludes our comparison by summarizing the key features of Resonant Two-Photon Ionization (RTPI) and its generalization, Resonant Multiphoton Ionization (REMPI).

The first bullet point captures the essence of its power: it combines high selectivity with near-unity efficiency. The selectivity comes from the resonant nature of the laser excitation step—only the specific molecule with the correct energy levels will be excited. The near-unity efficiency comes from the ability to ensure that every excited molecule is subsequently ionized and every resulting ion is detected. It's the best of both worlds.

Because of this powerful combination, it has become the preferred method in a vast range of applications, from analytical chemistry, where it's used for ultra-trace detection of pollutants, to precision spectroscopy, where it's

used to make fundamental measurements of molecular structure and dynamics.

The final point is a reminder of the practical considerations for an experimentalist. The choice of the specific ionization scheme—whether to use 1+1, 2+1, field ionization, autoionization, etc.—is not arbitrary. It is dictated by the specific energy level structure of the molecule you wish to study, the wavelengths and powers of the lasers you have available in your lab, and the sample environment, be it a static cell, a molecular beam, or a plasma. A successful experiment requires a thoughtful combination of all these factors.

Page 65:

We now move to one of the most powerful applications of ionization spectroscopy: coupling it with mass spectrometry. This slide explains the rationale for this combination.

The first problem it addresses is one of spectral overlap. While laser spectroscopy is highly selective, sometimes the spectral lines of different chemical species, or even different isotopes of the same species, can overlap. In a complex mixture, it can become difficult or impossible to be sure which molecule is responsible for a given spectral feature. We need another layer of identification.

The solution, as the second bullet point states, is to introduce mass separation. This yields an "orthogonal discrimination." "Orthogonal" here means that it's a completely independent method of identification. We first select molecules based on their "color" (their absorption spectrum) and

then we separate them based on their mass. A species must satisfy both the spectroscopic and the mass criteria to be detected, providing an exceptionally high level of certainty.

The workflow for such an experiment, often called Resonance-Enhanced Multiphoton Ionization Time-of-Flight Mass Spectrometry (REMPI-TOF-MS), is as follows: 1. Laser L1, the "spectroscopy" laser, is tuned to resonantly excite only a specific species of interest. 2. Subsequent photons from a second laser (or the same laser) then ionize those excited molecules.

The result is a cloud of ions that consists, ideally, of only the single species we selected with our laser. Now, we need to verify this by measuring their mass.

Page 66:

Continuing with the workflow of a REMPI-mass spectrometry experiment:

Step 3 is the mass separation itself. After the ions are created, they are injected into a mass analyzer. The most common type for use with pulsed lasers is a Time-of-Flight, or TOF, mass spectrometer. In a TOF instrument, ions are accelerated to the same kinetic energy and then allowed to drift through a field-free tube. Because kinetic energy is one-half m v 2 mv^2, heavier ions will travel more slowly and take longer to reach the detector. This separates the ions based on their mass-to-charge ratio, or m q $\frac{m}{q}$. Alternatively, for CW experiments, a quadrupole mass filter can be

used, which selects a single m q $\frac{m}{q}$ value at a time using oscillating electric fields.

Step 4 is the data collection. We set our mass spectrometer to only detect ions of a specific mass. Then, we record the number of those ion counts as we scan the wavelength, λ L 1 λ_{L1}, of our first laser. This process builds up a mass-resolved, or mass-gated, spectrum. We are plotting a spectrum that we know, with certainty, belongs only to the molecule with the mass we selected.

The power of this technique is highlighted in the final bullet point. It is so precise that it is capable of identifying and obtaining separate spectra for different isotopomers—molecules that differ only in their isotopic composition, for example, by a single neutron. This is an incredibly powerful tool for isotope analysis and for studying subtle isotope effects in molecular physics.

Page 67:

This slide provides a simple yet effective timing diagram that illustrates the principle of Time-of-Flight mass separation. The horizontal axis represents time, starting from t = 0 $t = 0$.

Let's follow the sequence of events.

First, at the top, we see the "Laser Pulse". This arrives at t = 0 $t = 0$, and it's responsible for the first step: Excitation.

Immediately following, at t = 0 $t = 0$, is the "Ionization" event. This creates our ions and, crucially, defines our "start time" for the TOF measurement.

The ions are then detected, which is labeled "3. Detection" and represented by the "Ion Signal" trace at the bottom. However, the detection does not happen at t = 0 $t = 0$. The ions must travel from the source to the detector.

The diagram shows that the Time-of-Flight, TOF, is proportional to the square root of the mass-to-charge ratio, m / z m/z. This means that lighter ions travel faster and arrive at the detector earlier.

We can see this in the ion signal. Imagine our laser ionized a sample containing two isotopes. The signal for the lighter isotope, m 1 m_1, appears first. A short time later, the signal for the heavier isotope, m 2 m_2, arrives. The diagram shows the lighter isotope arriving around 40 arbitrary time units, and the heavier one arriving at 42.5 units. The time difference between the peaks, Δ t Δt, is shown as 2.5 microseconds in this example. By precisely measuring these arrival times relative to t = 0 $t = 0$, we can calculate the mass of each ion with high accuracy.

Page 68:

Here we see a perfect illustration of the power of mass-gated detection, using the spectra of lithium trimer clusters, Li-three, as an example.

The first bullet point notes that if you take an optical spectrum of a beam of lithium clusters *without* any mass selection, the result is a mess. You see congested, overlapping features that are very difficult to assign or interpret. This is because the beam contains a mixture of different isotopomers of the Li-three cluster.

Now, see what happens when we turn on mass-gated detection. We can isolate the spectrum of a single, specific isotopomer. The slide lists two examples: 1. We can set our mass spectrometer to only detect ions with a mass of 21 atomic mass units. This corresponds to the cluster made of three Lithium-7 isotopes. We record the spectrum of pure 21 - L i - 3

. 2. Then, we can change the setting on our mass spectrometer to detect ions with mass 20. This corresponds to a cluster with one Lithium-6 atom and two Lithium-7 atoms. We get a completely separate spectrum for pure 20 - L i - 3

.

The ability to "un-mix" the congested spectrum into its pure components is a revolutionary capability.

But why are the spectra of these isotopomers different in the first place?

The last two bullet points explain the origins of these differences. The most significant effect stems from variations in the reduced mass. Changing the mass of one of the nuclei changes the molecule's vibrational frequencies. This leads to shifts in the positions of all the vibrational energy levels, and therefore shifts in the observed spectral lines.

Page 69: Continuing with the sources of spectral differences between isotopomers:

Beyond the shifts in vibrational levels due to the reduced mass, there are more subtle effects. The different isotopes can have different nuclear spins. The interaction of the nuclear spin with the electrons and with other nuclei gives rise to hyperfine splittings of the energy levels. These splittings are often different for different isotopes, leading to changes in the fine details of the spectral lines.

The ability to record these clean, isolated spectra for single isotopomers enables incredibly precise science. We can precisely determine the isotope shifts, which are the shifts in spectral line positions due to both the mass effect and a subtler field-shift effect. We can also measure the hyperfine splittings to determine the nuclear moments, such as the nuclear magnetic dipole moment and electric quadrupole moment.

Furthermore, the technique provides a highly accurate way to measure isotopic abundances. By comparing the total signal intensity from the mass-21 channel to the mass-20 channel, we can obtain a precise measurement of the relative abundance of the different clusters in our beam. This is a crucial tool in fields ranging from geology to nuclear physics.

Page 70

Given the importance of coupling ionization with mass spectrometry, let's briefly review the types of mass spectrometers best suited for these experiments.

The first and most common type is the Time-of-Flight, or TOF, mass spectrometer. It has several key advantages that make it a perfect partner for laser ionization.

* Its pulsed operation is naturally matched to the pulsed nature of the lasers used for REMPI. An ion packet is created with each laser shot and analyzed. * It offers simultaneous detection of the entire mass range per shot. In a single laser pulse, you create all the ions, and they all fly down the tube and are detected. This makes it very fast and efficient, especially for analyzing complex mixtures. It has a very high duty cycle in terms of data collection. *For very high-resolution work, the timing can be made even more precise. If you detect the* electron* that is created during ionization, its arrival can be used to generate an extremely precise t = 0 $t = 0$ start signal, with sub-nanosecond time stamping. This is a technique called velocity map imaging or photoelectron-photoion coincidence.

The second major category is the Quadrupole Mass Filter. We'll discuss its properties on the next slide.

Page 71: Let's continue with the Quadrupole Mass Filter and other types of mass analyzers

A quadrupole operates by scanning, allowing only a single mass-to-charge ratio to pass through at a time. This continuous scanning mode makes it highly compatible with continuous-wave, or CW, laser ionization experiments. You can set the quadrupole to a specific mass and just let the CW laser and detector accumulate signal.

However, quadrupoles have some drawbacks. They typically have a lower transmission efficiency than a well-designed TOF instrument, often less than 10%. You lose a significant fraction of your ions.

Also, because it performs sequential mass analysis, it is much slower for multi-isotope or multi-component studies. To get a full mass spectrum, you have to scan the quadrupole settings, which takes time.

A third and increasingly important category is ion traps. This includes Paul traps, which use radio-frequency electric fields, and Penning traps, which use a combination of electric and strong magnetic fields. These devices can store ions for long periods, from milliseconds to even seconds. They are increasingly being combined with CW REMPI for ultra-trace analysis, as the long trapping time allows one to accumulate a signal from a very weak source.

Ultimately, the choice of mass spectrometer is governed by the experimental requirements: the laser format (pulsed or CW), the need for speed and simultaneous detection, and the desired sample throughput.

Page 72:

We'll now look at a very important and practical application that combines several techniques we've discussed: Laser Desorption plus RTPI. This is a method for performing "soft" ionization of fragile molecules, particularly those that cannot be easily introduced into the gas phase.

Many molecules of interest, especially in biology, are large and non-volatile. If you try to heat them to make them evaporate, they simply decompose. The first bullet point describes the solution: a pulsed laser is

used to gently ablate or desorb *neutral* molecules from a solid or liquid surface. This is not about blasting the surface to create ions directly; it's a gentler process that liberates intact, neutral molecules into the gas phase just above the surface.

The second step is key. Immediately after being desorbed, while they are in the gas phase in an adjacent region, these neutral molecules are ionized by a REMPI process. This two-step approach—desorb first, then ionize in the gas phase—is crucial for avoiding fragmentation. The initial desorption can be done with low enough laser power that the molecules have very little internal energy, and the subsequent REMPI process is highly selective and can also be gentle.

This technique has opened the door to the study of many complex systems. The primary applications are in the analysis of biological macromolecules. This includes sequencing peptides, analyzing nucleotides like DNA and RNA, and studying complex carbohydrates, all of which are too fragile for more aggressive ionization methods.

Page 73: Continuing with the applications and advantages of Laser Desorption plus RTPI:

This technique is not just for biological molecules. A variation of it is used in planetary-surface analysis. For example, some of the instruments on the Mars rover "Curiosity" use a laser to ablate the surface of rocks and then

analyze the resulting plasma or plume to determine the elemental composition.

The primary advantage of this two-step "desorb-then-ionize" method over direct ion bombardment techniques (like Secondary Ion Mass Spectrometry, or SIMS) is that the internal energy of the neutral molecules remains low. This preserves their structural integrity. You are analyzing the molecule that was actually on the surface, not a fragment of it. This is what is meant by "soft" ionization.

Furthermore, the selectivity of the RTPI step acts as a powerful filter. When you desorb from a real-world sample, you create a plume containing your target molecule plus a huge amount of other "matrix" species from the substrate. A non-selective ionization method would ionize everything, resulting in a complex, messy mass spectrum. But with RTPI, the laser is tuned to selectively ionize *only* your target molecule. This filters out the background matrix species, dramatically enhancing the signal-to-noise ratio and improving the detection limit.

Page 74:

This excellent diagram illustrates the entire process of Laser Desorption / REMPI / Time-of-Flight Mass Spectrometry. Let’s walk through it.

The process is divided into two main stages. Stage 1 is “Desorption & Ionization.” On the left, we see a “Desorption Laser,” which is pulsed, firing at a sample surface. This ablates a plume of neutral molecules into the gas phase. The legend indicates that this plume contains both our purple “Target Molecules” and gray “Matrix Molecules.”

Immediately, a second laser, the “REMPI Laser,” passes through this plume. The REMPI laser is tuned to be resonant only with the target molecules, so it selectively ionizes them, creating the lighter green ions ($M1^+$) and heavier red ions ($M2^+$). This is the key “soft ionization” step, preserving the molecules’ integrity.

Now we move to Stage 2, “Mass Analysis (TOF).” The newly created ions are extracted and accelerated by a set of grids (V o

, V 1

) and injected into a long, field-free drift tube. As they travel down this tube, they separate by mass, because the time of flight, t

, is proportional to the square root of the mass-to-charge ratio.

At the end of the drift tube, a detector records the “Ion Signal” as a function of time. We see a clean spectrum. The lighter ions, $M1^+$, arrive first, creating the green peak. The heavier ions, $M2^+$, arrive later, creating the red peak. The unwanted neutral matrix molecules are never ionized, so they don’t appear in the mass spectrum at all. This diagram perfectly captures the elegance and power of this combined technique.

Page 75:

We'll now explore a fascinating and highly sensitive detection device known as the Thermionic Diode, which operates on the principle of Optogalvanic Detection with internal amplification.

First, let's look at the components. It's a relatively simple device. It consists of a heated filament, which acts as the cathode. Being hot, it emits a steady

stream of electrons through the process of thermionic emission. This filament is enclosed by a cylindrical metal wall, which acts as the anode. The entire cell is filled with a low-pressure gas or vapor of the atoms we want to study.

Now, how does it operate? At a small bias voltage between the cathode and anode, the device operates in what's called the space-charge-limited regime. The cloud of negatively charged electrons emitted by the filament forms a dense "space charge" region right around the cathode. This negative cloud repels other electrons that are trying to leave the filament, severely limiting the amount of current that can flow to the anode. The device essentially chokes its own current flow. It is this space-charge effect that is the key to the device's operation as a detector.

Page 76

Here is how the thermionic diode harnesses the space- charge effect to act as a massive signal amplifier.

The process begins with laser excitation. A laser is tuned to excite the gas atoms inside the diode into high-n Rydberg states. As we know, Rydberg atoms are very large and weakly bound. When these Rydberg atoms collide with the thermionic electrons flying around, they are very easily ionized by electron impact.

This creates slow- moving, positive ions. The second bullet point is the crucial step. A positive ion, being heavy, lingers for a relatively long time (Δ t i o n Δt_{ion}) in the vicinity of the cathode. Its positive charge locally neutralizes part of the negative space charge cloud.

This neutralization effectively lowers the potential barrier for the other electrons trapped in the space charge. The result is a cascade. Many, many more electrons are now free to escape the space charge and flow to the anode, creating a large pulse of current.

The result is that a small initial ionization yield—the creation of just a few ions by the laser—is amplified into a much larger electrical signal. This current magnification factor, M M, can be enormous.

The change in current, Δ i Δi, is given by e · N · M $e \cdot N \cdot M$, where N N is the number of ions created and M M is the gain. This gain, M M, is approximately the ratio of the time an ion lingers, Δ t i o n Δt_{ion}, to the time it takes for an electron to transit the device, Δ t e l Δt_{el}.

As the slide notes, this gain M M can be as large as 10 5 10^5. So, the creation of a single ion can result in a measurable current pulse of 100,000 electrons. This makes the thermionic diode an extraordinarily sensitive detector for Rydberg states.

Page 77:

This Paage is blank, so we will continue to the next slide.

Page 78:

We'll now discuss a final, important refinement for performing very high-precision spectroscopy of Rydberg atoms: creating a Field-Free Excitation Zone.

As we've discussed, Rydberg atoms, especially those with extremely high principal quantum number n n, are exquisitely sensitive to electric fields. Even very weak stray fields in an experimental chamber can cause significant Stark shifts and broadening of the energy levels, which would distort and limit the resolution of our spectrum.

To perform the highest resolution spectroscopy, it is essential that the laser excitation takes place in a region that is as close to perfectly field-free as possible.

The second bullet point describes how this can be achieved. By carefully designing the geometry of the electrodes used for ion collection, it's possible to create a potential minimum, or a "saddle point," in the electric potential right along the path of the laser beam. We saw a diagram of this earlier for the thermionic diode.

The benefit of this careful design is immense. It allows for spectroscopy of Rydberg states up to n n of approximately 300 or even higher, without significant Stark shift broadening. This opens up the near-continuum region of atoms and molecules for detailed study, providing a stringent test of our understanding of atomic and molecular structure in this complex regime.

Page 79:

Continuing on the benefits of this field-free excitation approach:

When this technique is combined with the enormous gain of thermionic amplification that we saw earlier, the overall sensitivity becomes truly phenomenal. This combination makes it possible to observe and measure

the energies of levels that are lying only a few *micro-electron-volts* below the ionization potential. This is an incredible level of energy resolution.

The ability to make such precise measurements on these very high-lying states is not just a technical curiosity. It provides benchmark-quality data for rigorous tests of fundamental theories of atomic and molecular structure. For example, the precise energies of a Rydberg series can be used to test the predictions of Quantum Defect Theory, which describes how the inner electron core of an atom perturbs the otherwise hydrogenic energy levels. It also allows for very precise measurements of the core polarizability—how the inner electron cloud is distorted by the presence of the outer Rydberg electron. These are fundamental properties of atoms and ions that can only be accessed through such high-precision experiments.

Page 80:

This set of diagrams provides a wonderful visualization of the thermionic diode and the principle of the field-free excitation zone. The overall title is "Thermionic Diode: Optogalvanic Detection of Rydberg States."

Let's look at the two panels on the left first. They show a cross-section of the cylindrical diode.

The panel labeled "LASER OFF" shows the red outer ring, which is the Anode held at a positive voltage, + V $+V$. Inside, there are two small blue circles representing the Cathode filaments, held at 0 V o l t s $0\,\mathrm{Volts}$. Surrounding the cathodes is a blue haze, representing the dense, negative "Space Charge Cloud." This cloud prevents current from flowing.

Now look at the panel labeled "LASER ON." A red "Laser Beam" is shown passing directly between the two cathodes. This is where the laser excites the atoms, which are then ionized, creating the positive ions that neutralize the space charge and trigger the amplified current pulse.

The panel on the right, titled "Electric Potential Profile," explains *why* this geometry is so special. It plots the electric potential V V as a function of position along the horizontal axis, the x x-axis. The potential is high at the outer anode plates. But because of the two symmetric, grounded cathodes, the potential dips down and forms a distinct minimum right at the center, at position 0 0. This potential well is the "Field-Free Region," highlighted in pink. The electric field is the derivative of the potential, so at the very bottom of this well, the electric field is zero. By aligning the laser to pass through this exact point, we ensure that the laser excitation occurs in a region free from Stark effects.

Page 81:

This slide provides a comprehensive text description of the "Principle of Operation" for the thermionic diode, summarizing everything we've seen in the diagrams.

First, it describes the **Components**: A cylindrical Anode wall surrounds two symmetrically placed heated Cathode filaments. The cell contains a low-pressure gas.

Next, it describes the **Laser OFF** state. The cathodes emit thermionic electrons, forming a dense negative space charge cloud. This cloud repels

other electrons, severely limiting the current that can reach the anode. This is the space- charge- limited regime.

Then, the **Laser ON** state. A tuned laser excites gas atoms in the central field- free region into high-n Rydberg states. Collisions with the abundant thermionic electrons then efficiently create slow- moving positive ions.

Finally, it describes the **Amplification** mechanism. A single positive ion lingers near the cathode, locally neutralizing the negative space charge. This allows a cascade of many thousands of electrons to flow to the anode, creating a large, easily detectable current pulse. The amplification is described by the equation Δ i = e N M $\Delta i = eNM$, where the gain M M can be as large as 10 5 10^5.

Page 82:

This slide is incredibly valuable, as it distills our entire discussion down to a set of practical, take-away parameters for designing an ionization-spectroscopy experiment. This is a checklist for the working experimentalist.

First, the photon flux target for the ionizing step. We need n L 2 ≥ 10 25 $n_{L2} \geq 10^{25}$ photons per cm 2 ^{2} per second. This is the target needed to ensure that the ionization rate outcompetes a typical radiative decay rate of R k = 10 8 s − 1 $R_k = 10^8\,\mathrm{s}^{-1}$ for a typical photoionization cross-section of σ k l = 10 − 17 c m 2 $\sigma_{kl} = 10^{-17}\,\mathrm{cm}^2$.

Second, collection efficiency. To achieve a δ δ of approximately 1, you need to design ion optics that produce extraction fields of at least 100 V / c

m $100\,\mathrm{V/cm}$. Furthermore, your optics must have a large solid-angle coverage, greater than 2 π 2π steradians, to ensure you collect ions emitted in the forward direction.

Third, detector gain. A microchannel plate (MCP) or a channeltron detector is the standard choice. These devices provide a charge amplification, or gain, in the range of 10 6 10^6 to 10 8 10^8. This is more than enough to turn a single ion event into a robust electronic pulse. Critically, these detectors should be chosen and operated to have a very low dark count rate, less than 1 1 count per second.

Page 83:

Continuing with our practical design parameters:

The fourth point addresses the specific requirements for continuous-wave, or CW, operation. To meet the high photon flux requirement with the relatively low power of a CW laser (typically 1 to 5 Watts), tight focusing is essential. A focal spot radius, w 0 w_0, of less than or equal to 30 micrometers is a typical target value.

The fifth and final point is a critical reminder about experimental efficiency. Molecular beam synchronization is absolutely essential whenever the laser duty cycle is low. The slide gives a rule of thumb: if the duty cycle is less than 10 − 4 10^{-4}, you must use a pulsed, synchronized beam to avoid wasting the vast majority of your sample. This is almost always the case for experiments using low-repetition-rate pulsed lasers.

This set of five rules provides a fantastic starting point for the design and planning of any high-sensitivity ionization spectroscopy experiment.

Page 84:

To conclude our lecture on Ionization Spectroscopy, here is a list of recommended further reading and sources for reference data. These texts are classics in the field and will provide much greater detail on the topics we've introduced today.

First, the quintessential textbook by Demtröder, "Laser Spectroscopy." The 5th edition is shown here. It's an indispensable resource. Chapter 1 provides an excellent overview of the fundamentals we discussed, and Chapter 5 contains extensive details on REMPI spectroscopy.

Second, for those interested in the specifics of REMPI on small molecules, the review article by Koopman and colleagues in the Journal of Chemical Physics is an excellent resource. It provides a thorough review of cross-sections for many different systems.

Third, for a deep dive into the fascinating physics of autoionization, the review by Orr-Ewing and Ashfold in Chemical Society Reviews, titled "Autoionization mechanisms in polyatomic molecules," is a seminal work.

Page 85:

Continuing with our recommended reading list:

The fourth reference is an article by M. A. Duncan in the International Journal of Mass Spectrometry, titled "Laser ionization in mass

spectrometry." This is a fantastic resource for anyone interested in the powerful combination of laser techniques and mass analysis that we discussed at length.

Finally, for those who are truly captivated by the physics of Rydberg states, field ionization, and high-precision measurements, there is no better resource than the monograph by T. F. Gallagher, simply titled "Rydberg Atoms," published by Cambridge University Press. This book is the definitive bible on the subject and contains all the details you would ever need for field ionization calculations and understanding the complex behavior of these giant, fragile atoms.

I highly recommend you consult these sources as you continue your studies.

That concludes our lecture for today. Thank you.

Vol. 2