

Raman Scattering Lab Supplement

Ben Nachumi

September 12, 2011

1 What is Raman Scattering?

Raman scattering is the inelastic scattering of light from vibrational or other excitations in a sample. Raman differs from, say, fluorescence in that Raman occurs over a very short timescale, and is a coherent process over virtual states. By contrast, in fluorescence, the intermediate excited state of the sample is “really there”. This means that while fluorescence emissions have a definite color, Raman scattered light is shifted from and tied to the incoming laser frequency. The *Raman shift*, measured from the incoming frequency, corresponds to how much energy has been deposited in the sample, and one sees peaks in the Raman spectrum corresponding to favored excitations of the sample.

In a Raman experiment, you will see excitations when the $\langle \mathbf{OUT} |$ (scattered) and $|\mathbf{IN}\rangle$ (incident) polarizations of the photon field match (transform like) the symmetry of the excitation in the sample. If you’re interested, I expand upon this mathematics in section 5, but for now, think of this fact as sort of like how, for a polarizer, you’ll see light on the other side when the incident electric field has a component in the transmission direction. So given a molecular or crystal structure, you can predict which excitations you ought to see, depending on how you select the polarization of the incoming and outgoing light.

Hence to see an understandable Raman spectrum, you need to pre-filter your incoming light (use a gas laser filtered for plasma lines), define the incoming polarization, select the outgoing polarization, and filter for color (energy). Using symmetry tables, you can predict which phonons (in solids) or vibrational/rotational modes (in liquids and gases) you are allowed by symmetry to see. You need a dark room, a quiet detector, and a rejection filter for your excitation, because as a higher-order process, Raman scattering is much weaker than the elastic scattering of the excitation. But with a decent spectrometer, it isn’t very difficult to see the effect.

A successful denumeration of the observed modes then tells you stuff about the bonds and the phonons. Further refinements of the technique can tell you about electron-phonon and electron-electron interactions. Add temperature dependence and you can often pick up evidence of structural and electronic phase transitions for much less than

the cost of a neutron scattering or X-ray spectrometer. From the larger philosophical perspective, Raman scattering presents an accessible exploration of the consequences of symmetry in quantum mechanics.

2 The Big Pieces

The spectrometer is a research-grade single-grating Czerny-Turner (slit-mirror-grating-mirror-out) spectrometer (JY HR-640) that fell upon some hard times. It's been fixed up so that the grating moves as it should. There are lasers of at least two colors for excitation, and enough lenses and mirrors to make a good transfer of laser light to your sample, and to focus the scattered light on the entrance slit of the spectrometer. For detection, there are, as of this writing, two usable low-dark-current CCDs. One is a liquid nitrogen-cooled camera (Photometrics CH250) and the other is a Peltier-cooled camera (SBIG ST-8300M).

3 Follow the Light Path!

3.1 Beam Transport: How Light gets to the Sample

The figure below shows the basic setup of the beam transport. A source of more-or-less monochromatic light is cleaned by bouncing it off a preliminary grating, then sent (via mirrors) over to the sample. Along the way it passes through one or more lenses and polarizers, so that it arrives at the sample in a custom-sized, custom-polarized spot. Note that the shorter the focal length and the larger the beam on the lens surface, the smaller the focused spot. Can you see why?

Now recall that you don't really want to see the laser light—you want the other light emitted when the laser light is shifted. This is a much weaker signal; and in fact **even a weak-looking, less than 1 mW direct laser beam quickly oversaturates, and can even damage, the detection CCD's**. You almost never want the specular beam, and you certainly don't want it at full strength. You want the laser to hit the sample such that the spectral (mirror-reflected) laser beam misses the collection lens. To hit the sample around the collection lens while the specular reflection escapes, you can bounce the laser beam off of one or more mirrors.

3.2 How Light Gets into the Spectrometer

Next, you want to be able to get the light from the sample into the spectrometer. The usual way to do this is with at least two lenses. The first lens collects light from the laser spot. The second lens refocuses the collected light on the entrance slit. For higher signal, you want to collect lots of light. So for a collection lens, you should use a nice




Figure 1: Beam transport. A high-power beam and some mist have been used to make the path more visible.




Figure 2: A simple input train.

fast (large diameter/focal length, or low f-number) lens, and you want the sample as close to the lens as possible. A simple way to do this and make a real image is to put the sample at the focus of the fast collection lens. In ray optics, this configuration would turn the approximate point of the laser spot into a parallel beam. From this parallel beam, the refocusing lens produces an image at its focus—so the refocusing lens should have its focus right on the entrance slits. You want the focused spot to be as small as possible so that, while collecting full signal, the slit can be made as narrow as possible. Hence you also want the refocusing lens to have a short focal length. So as not to throw collected light away, you also want the refocusing lens to be as big as the collection lens.

3.3 How to Get a Good Image on the Slit

Go through these steps to put a focused image on the entrance slits.

- Turn on the little focusing camera and the monitor, and swing around the flip mirror. This mirror sends light from the slits through an internal lens that makes a real image of the slit on the camera. Now open the spectrometer shutter by issuing the appropriate software command. Check that the focus camera is really imaging the front (vertical) slits—these determine the spectral resolution. The slits are focused if you see their sharp edges move around on the monitor when you adjust the slit width.
- Next, open the slits and center a refocusing lens in front of the shutter. By moving the lens forward and back, try to focus the image of the most distant object you can see in the monitor. The focal point of this lens is now just a bit in front of the slits (after all, the object isn't infinitely far away), but this gets you near where you eventually want to be.
- Now center the collection lens on the refocusing lens. Leave some space between them so you can insert a polarizer or a filter later. Move a needle or some other tiny pointy thing around until you see its focused image appear in the monitor along with the slits.
- The image of the needle is now falling on the slits, but you can improve things a bit by just touching refocusing lens back toward the slits, and finding the image of the needle as before. If you can't get a focus, that means that the refocusing lens is too close to the slits. Move it a bit forward from them and try again.
- Now turn on the laser and steer the unfocused beam so you see it hit the needle in the monitor. Finally, focus the beam into a spot on the needle.

Your sample goes where the needle is, at the focus of the collection lens. When you swap in your sample, you'll have to tweak the beam steering and the sample position to get everything just right. Eventually, you'll want to close down the slits while keeping the spot visible.

4 Inside the Spectrometer

To begin to understand how the spectrometer works, refer to the following figure. The




Figure 3: The light path inside the spectrometer.

basic idea is to make separate copies (one in each color present) of the illuminated entrance slit, and put them on a detector. Now the slit is in the focal plane of the concave mirror (**M1**), so light from the slit that hits **M1** gets sent parallel, to the grating (**G**). The grating splits the multicolored parallel beam into a fan of copies. Those copies that hit the second concave mirror (**M2**) are refocused onto the plane of the detector. So for each color present at the **M2**, one gets a separate spot on the detector. Which section of spectrum you see depends on how the grating is turned, and on the pitch (grooves per mm) of **G**. The more densely grooved is **G**, the more rapidly the spectrum fans out, while **M2** of course stays the same size. So for a fixed grating angle, you'll catch less of the spectrum (but have better resolution) with a finer grating.

4.1 The Central Wavelength

For all of this to be useful, you have to know how position on the detector corresponds to color. The task is to calculate this spectral range, from the dimensions of the spec-

trometer and from the formulas for diffraction from a grating. You'll also learn how this particular commercial spectrometer associates a central wavelength with a particular grating position.

First, consider the undercarriage of the spectrometer (note the associated figure). You'll see a big drive screw. Upon this screw rides a shuttle, which in turn moves a long




Figure 4: The mechanism.

arm. For reasons that should soon become obvious, this arm is called the *sine bar*. Upon the sine bar rides a small wheel, attached to another arm, which appears to turn on a hinge that goes up into the spectrometer. In fact, this hinge turns the grating. When the screw turns, the sine bar rides up or down parallel to its axis, which pushes the arm that turns the grating.

Now consider the following diagram, adapted from the one in the spectrometer manual. The lines l , m , and n are mutually parallel. l connects the center of the slit to the center of **M1**, m goes through the center of **G**, and the normal to the grating makes an angle (the grating angle) Δ with m . n passes through the center of the exit port. Also note for your later convenience the line r , which connects the centers of **M1** and **G**, and makes the angle γ with m .




Figure 5: Determining λ_c .

Now, the spectrometer is tuned—that is, **M1** and **M2** are positioned—so that when $\Delta = 0$, the specular reflection of light moving parallel to r from the grating of light from **M1** is ultimately sent out the center of the exit port (via **M2**) along n . The color of the light that follows this path is called the *central wavelength* λ_c . For a perfectly calibrated system, this wavelength would correspond to the readout of the odometer on the side of the spectrometer.

This works because the grating inside the spectrometer is parallel to the arm that rides along the sine bar, and the sine bar moves along the main screw (see the diagram below), which makes the distance D along the screw proportional to $\sin \Delta$. Let θ be the angle between the normal to **G** and m , so that $\theta = \Delta - \gamma$. Finally, call the deviation between the ray approaching **G** and its outgoing partner ϕ . For the central wavelength, $\phi_c = \Delta + \gamma$. Call the spacing between neighboring grooves of the grating ϵ . The path length difference between two rays reflecting from neighboring grooves will be $\epsilon(\sin \phi + \sin \theta)$, and you'll see constructive interference along the central wavelength's path for

$$x\lambda_c = \epsilon(\sin \phi + \sin \theta)$$

(x is an integer), or (using a trig identity or two)

$$x\lambda_c = 2\epsilon \sin \Delta \cos \gamma.$$

Since γ is fixed, for first-order ($x = 1$) interference maxima, $\lambda_c \propto \sin \Delta \propto D$. The proportionality constant comes from the pitch of the carriage screw.

As an exercise, you can show that this same formula applies when the grating is turned so that the normal to **G** is on the other side of r and $x = -1$. As another, calculate Δ_{532} , the angle appropriate for $\lambda_c = 5320\text{\AA}$, for a 1200 groove/mm grating (use $\gamma = 7.67936^\circ$).

4.2 Beyond the Central Wavelength: Calculating Dispersion

If this were **BCCD** (**B**efore **CCD**'s—olden times), you'd be done. You would take your data by driving the grating through some angle, and recording the readout of a photomultiplier tube sitting on the center of the exit port. If times were even older, you'd be recording a broad spectrum in one go on a photographic plate—but this would be very difficult to calibrate for intensity. The CCD is an array of pixel detectors that acts like electronic film—the best of both older methods. However to understand the CCD spectrum, you need to modify the previous discussion for colors that are maximized off of the central path.

This isn't too difficult to do. Referring to the next diagram, note that I've added a new color λ_ν , coming off **G** at the angle ϕ_ν from **G**'s normal, with $\phi_\nu = \phi_c + \nu$. Going through essentially the same steps, you can show that $x\lambda_\nu = 2\epsilon \sin(\Delta + \nu/2) \cos(\Delta - \nu/2)$.

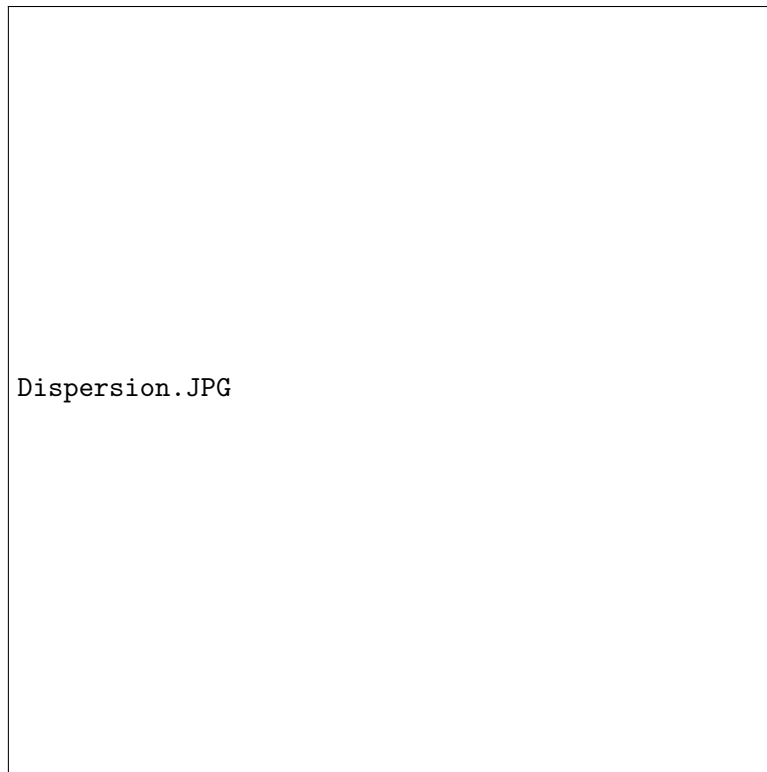


Figure 6: Dispersion geometry.

Now this extra angular deviation just carries through the rest of the system. Coming off **M2**, λ_ν will be deflected from the central path of λ_c by the angle ν . Call the focal length of **M2** F . In the focal plane of **M2**, at the CCD, λ_ν will make a bright spot a distance $y = F \tan \nu$ from the spot made by λ_c .

Using the dimensions of the CCD and the known Δ_c , you can determine the range of wavelengths that you can see in one shot (the *wavelength coverage*)

$$\lambda_{Max} - \lambda_{Min} = 2\epsilon \sin \nu_{Max} \cos(\Delta_c + \nu_{Max}),$$

as well as the linear dispersion along the CCD in units of (wave)length/length

$$\frac{d\lambda_\nu}{dy} \approx \frac{1}{F} \epsilon (\cos(\Delta_c + \nu/2) \cos(\gamma + \nu/2) - \sin(\Delta_c + \nu/2) \sin(\gamma + \nu/2)).$$

As an exercise, find the wavelength coverage that you expect for $\lambda_c = 5320\text{\AA}$ (use 24.4 mm for the width of the CCD chip and 640mm for F). Dividing the range by the width of the CCD, get the average linear dispersion for this setting. Then generate a plot of the true linear dispersion as a function of position (from -12.2 mm to +12.2 mm) across the CCD. On the same axes, plot the % deviation from the average linear dispersion.

4.3 Frequency Units

Because frequency, rather than wavelength, is the fundamental quantity of interest, you will usually convert your spectra from “Intensity vs. Wavelength” to “Intensity vs. Frequency”. The peaks in this latter plot will correspond to the energies of the excitations in the sample. In spectroscopy, the usual unit of frequency is the *inverse centimeter* (cm^{-1}), which is just the usual frequency (in s^{-1}) divided by the speed of light in cm/s . A frequency difference of 1 cm^{-1} corresponds to 30 GHz.//

You must be more careful to properly convert dispersion from wavelength/mm to cm^{-1}/mm . Since $\lambda = 1/f$, $\delta f = -\delta\lambda/\lambda^2$, and a constant linear dispersion $d\lambda/dx$ in wavelength units becomes, via the chain rule,

$$\frac{df}{dx} = \frac{-1}{\lambda^2} \frac{d\lambda}{dx}$$

when converted to frequency units. The spectral coverage of the spectrometer with the 1200 groove/mm grating installed, with $\lambda_c = 5000\text{\AA}$, is about 1000 cm^{-1} , which is more than enough to see many phonons in a single exposure.

4.4 Resolution

Colloquially speaking, the *resolution* is the degree to which you can distinguish bright spots on your CCD. The spectral resolution is the corresponding minimal difference in λ . For a working definition of what is “barely distinguishable” is, the common choice is

the *Rayleigh criterion*: two peaks are barely distinguishable if the maximum of one falls on the minimum of the other. To apply it, you need to know the width of the bright spot—the distance along the CCD from maximum intensity to zero. For a given grating of line spacing ϵ , with N lines illuminated, with a bright beam going off at angle Θ the angular width of the x th maximum is

$$\delta(\sin \Theta) = \frac{\lambda}{N\epsilon}.$$

This just describes the change in angle necessary to get a phase shift of π over half the width of the grating, so that reflections in lines cancel in pairs (as in the standard analysis for diffraction from a single slit). On the other hand, from the diffraction maximum condition $\sin \Theta = x\lambda$ (x integer), the next distinguishable peak, $\delta\lambda$ away, will deviate from Θ by $\delta(\sin \Theta) = x\delta\lambda$. Setting these two $\delta(\sin \Theta)$'s equal gives the *grating resolution*

$$\frac{\delta\lambda}{\lambda} = \frac{1}{Nx}.$$

Using the dispersion to convert this information to a width on the CCD is straightforward, and you should find the *intrinsic spot size* for $x = 1$ (first-order diffraction) and $N \approx 5 \times 10^4$ (this N corresponds to about half of an 11cm \times 11cm, 1200 groove/mm grating). Translate this number into a number of opixels, assuming that the width of one pixel is $24\mu\text{m}$.

When the pixel size on the CCD is larger than the intrinsic spot size, the effective smallest possible spot size is just the pixel size. However the width of the entrance slit, which is almost always larger than the size of the pixel, usually controls the experimental resolution. If the pixel size is $24\mu\text{m}$, and the intrinsic spot size is $1\mu\text{m}$, and the slit itself is $200\mu\text{m}$, then the slit spot size will be about about 8 pixels (for $F = 640\text{mm}$ and $\lambda = 5000\text{\AA}$, check to see that you can ignore diffraction from the slit). Hence for a large range of slit values, you can predict your experimental resolution pretty well by multiplying the dispersion by the slit width.

4.5 CCD Wavelength Calibration

The two figures below are representative exposures of the CCD to a multiline source (in this case, an argon laser, but a mercury vapor lamp would do as well), the second with the same source, but less well-filtered than the first, and with a longer exposure time. Here is what I was thinking about this data:

With the grating set at a nominal central wavelength of 5018 \AA , the first picture shows three fairly tightly constrained spots of light from the laser. The lower spot is 4880 \AA , and the upper spot is 5145 \AA (I confirmed this by actually looking at the diffraction spots on a piece of white paper held in front of the exit port with the CCD out).

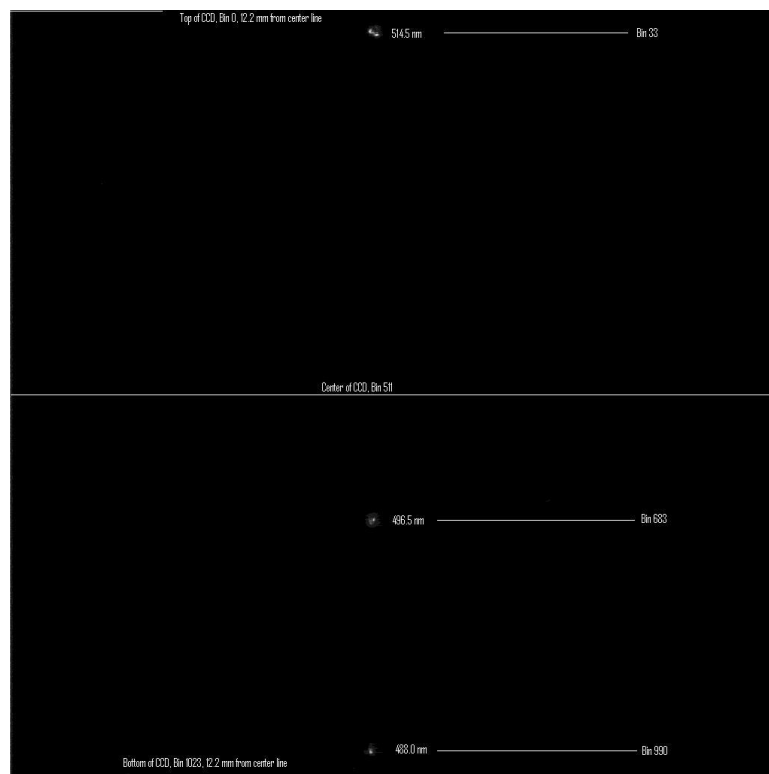


Figure 7: First multiline calibration image.



Figure 8: Second multiline calibration image. The beam is less filtered, so more lines are present. The lines from the first figure have saturated and leaked into adjoining pixels because of the longer exposure time.

From the known wavelengths of these lines, their positions on the CCD chip, and the size of the chip, one calculates the measured dispersion for the spectrum. In this case, the calculated linear dispersion is a touch low (11.62 \AA/mm), compared to what one expects to see by doing the standard calculation and using the dimensions of the chip and the nominal focal length of the spectrometer (11.78 \AA/mm). Yet the experimental dispersion accurately predicts the position of the third expected line (4965 \AA), and when the collimation irises on the input train are opened to allow more light in and the exposure time is lengthened, the same dispersion predicts the position of a new spot that's expected to contaminate the input (just about 5017 \AA , second picture). So the measured dispersion is believable.

An examination of the formula for the expected dispersion reveals the source of the discrepancy: the dispersion ought to vary inversely with the focal length of the setup, and the ratio of the measured dispersion to the manufacturer's quoted dispersion is about $63/64$. For a 640 mm focal length, the CCD chip is about a centimeter too close, and should be placed a bit farther back from the last mirror. This would allow the diffraction fan to spread out just a bit more, leading to the nominal dispersion—but would the focus (and hence the resolution) suffer? In fact, the focus as it stands is not as perfect as it could be: the 5145 \AA spot in the first picture is decidedly blurrier than the 4880 \AA spot. In addition to moving the CCD, some tilting of the last focusing mirror may be necessary.

5 Origin and Use of Raman Selection Rules

This is a slightly more advanced introduction to the symmetry aspects of Raman scattering. For the sake of definiteness, this treatment assumes that the scattering occurs from phonons in a crystal, although the general results apply to any sort of excitation in any kind of sample.

The term “selection rule” is shorthand for the set of conditions under which a given transition matrix element is nonzero. The term “Raman scattering experiment” might well be shorthand for the determination of the selection rules for light-scattering from a sample, in the event that the incident and scattered light have different frequencies.

5.1 Some Simple Examples of Selection Rules. The Polarizer Example

Perhaps the simplest selection rule in quantum mechanics is the example of parity. The parity operator \mathbf{P} flips spatial coordinates $(x, y, z) \rightarrow (-x, -y, -z)$. In an introductory quantum mechanics class, one often sees the matrix element of an operator like the electric polarization $\mathbf{q}\vec{R}$, between two parity-labeled states, $\langle a' | \dots | a \rangle$. The question is: which

pairs of states are connected, and which are not?

The answer is that, since the polarization operator changes sign under the parity operation, and the entire matrix element is really just some integral, and integrals can't care whether we number our axes left-to-right or right-to-left, then the sandwich of states $\langle a' | \dots | a \rangle$ must also change sign. Hence the electric polarization operator can only connect states of opposite parity. This is the “parity selection rule”.

Here's another familiar example: the absorption of light by a polarizer. The polarizer seems to split the light into two parts, one with the electric oscillation \vec{E} perpendicular to, the other parallel to, the direction of the little absorbing wires in the polarizer. Current density \vec{J} is induced along the wires. In the usual derivation, we say that the power absorbed is proportional to $\vec{J} \cdot \vec{E}$, and this dot product is zero for $\vec{E} \perp \vec{J}$.

Now suppose that the little wires point in, say, the $+x$ direction, and \vec{E} points wherever it wants. If absorption occurs, how should \vec{J} go? Imagine that we flip the whole experiment around: nothing should change, because the electrical interaction can't depend on our choice of coordinates. Next, just flip the polarizer, but keep \vec{E} the same. If \vec{J} also flipped, this would introduce a phase (-1) in the dot product. But the whole experiment looks the same! *Nothing should have changed.* If \vec{J} also flipped when we turned over the polarizer, then the absorption would have to be zero. The absorption can be nonzero only when \vec{J} can follow \vec{E} , and it must be zero when \vec{J} can not.

5.2 Definition of the Raman Tensor

Consider the following matrix element as a representative Raman scattering event:

$$\langle \text{OUT} | \mathcal{S} | \text{IN} \rangle = \langle n' \gamma' | \mathcal{S} | n \gamma \rangle.$$

This matrix element corresponds to the absorption of a photon by the sample, an internal transition of the phonon field, and the reemission of a (different frequency) photon. The scattering operator \mathcal{S} represents the effect of the sample. By examining the way the different parts of the matrix element transform under the symmetry operations of the total system (ideal lattice plus phonons plus light), we may determine the selection rules for Raman scattering.

For example, the light part of the matrix element must be invariant with respect to general three-dimensional rotations (because Maxwell's equations are invariant under general rotations and translations). \mathcal{S} has the symmetry of the sample, that is, it is unchanged under the rotations, reflections, etc. that leave the idealized lattice looking microscopically the same. Thus \mathcal{S} always has a phase of +1 under the symmetry operations.

Suppose now that we rotate the sample through the symmetry operations of the sample's group, while keeping the rest of the apparatus (laser polarizations, lab room) fixed. Clearly, the result of the experiment should not change: the sample should behave the same way when considered in any equivalent orientation. In this case we are effectively rotating the “partial matrix element”

$$\langle n' | \mathcal{S} | n \rangle$$

against a fixed background

$$\langle \gamma' | \dots | \gamma \rangle.$$

Unless $|n'\rangle$ and $|n\rangle$ are the same, that is, unless

$$\langle n' | n \rangle = 1,$$

the partial matrix element will pick up a nonzero phase corresponding to the relative phase of the different phonon states under the transformation. For example, a phonon consisting of just an upward atomic displacement will acquire a phase of -1 under a rotation of π about an axis perpendicular to the displacement:

$$\Uparrow = -1 \times \Downarrow.$$

Now make an assumption: if we rotate the entire world of the experiment, crystal plus apparatus, corresponding to the entire matrix element, then there can be no additional phase. If the phases of the partial matrix element and the photon fields differ, then by symmetry, the entire matrix element must be zero. The transformation mismatch of the $|\gamma'\rangle$ and $|\gamma\rangle$ states of the vector potential must exactly compensate for the phase which we get by rotating the crystal alone, e.g., the change in phase due to the phonon states. In other words, the product of the phonon states must transform “in the same way” as the product of the photon states. But photon states can be described by the vector potential, and vector potential states transform like vectors, so the operator connecting them must transform like a rank-two tensor.

Hence we find: the partial matrix element representing the transition of the phonon fields is a rank-two tensor. We'll call this the *partial Raman tensor for the phonon*, $\mathcal{R}(ph)$. The three possible values (x, y or z) of the two indices (out/row, in/column) of $\mathcal{R}(ph)$ correspond to polarization directions of the incoming and outgoing light. What a Raman experiment measures is the total Raman tensor, \mathcal{R} , which is a sum of the $\mathcal{R}(ph)$'s and similar things for other kinds of modes that can live in a sample. In a typical experiment, one controls the polarization of the incident light, and selects different polarizations of the scattered light. In this manner the Raman experiment probes different components of \mathcal{R} .

5.3 Classification of Phonons: Allowed and Forbidden

Now consider the actions of the sample's symmetry group on $\mathcal{R}(ph)$. The Raman tensor is a product of phonon states along with an object that can't pick up a phase. We know that this whole sandwiched product is a 3×3 tensor.

Furthermore, because $\langle n|n \rangle = 1$, the states that are different by a single phonon $\langle n+1|n \rangle$ must transform like a single phonon. This means that under all symmetry operations, $\mathcal{R}(ph)$ must behave the same way as the displacement pattern of the single phonon.

Suppose you are interested in a particular phonon. From the sample structure, you have guessed at the displacement pattern of the atoms involved in the phonon. You can use this information to predict which incoming/outgoing combination of light polarizations could create one of these. Just put the displacement pattern through the sample's symmetry operations, and then perform the same set of transformations on the elements of an arbitrary 3×3 tensor. Whichever of the 9 tensor elements transform the same way as the phonon, these correspond to experimental conditions in which the phonon could be observed. This is an *allowed* phonon. On the other hand, for a given polarization condition, if the relative phases add up to zero, then the phonon should be *forbidden*, because (according to the symmetry of your sample) you won't be able to see it.

For example, consider a flat (x, y) plane square molecule with symmetry C_4 :

$$\begin{array}{cc} \uparrow & \rightarrow \\ \leftarrow & \downarrow \end{array}.$$

Assume the excitation comes in along the z axis, and the scattered light comes out in the opposite direction, so we can represent our \mathcal{R} with a 2×2 matrix and not worry about polarizations in z . Let's also denote the experimental polarization conditions by $[\vec{i}\vec{n}, \vec{o}\vec{u}\vec{t}]$.

The square molecule is invariant under the identity, and turns by $\pi/2$, π , and $3\pi/2$ about a z axis through its center. Under the same operations, the polarization combination $[\vec{x}, \vec{y}]$ rotates to $[\vec{x}, \vec{y}]$, $[\vec{y}, -\vec{x}]$, $[-\vec{x}, -\vec{y}]$, $[-\vec{x}, -\vec{y}]$, and $[\vec{x}, -\vec{y}]$. Hence the phonon states that you should see in $[\vec{x}, \vec{y}]$ and $[\vec{y}, \vec{x}]$ must pick up phases of 1, -1, 1, -1. A phonon visible in these conditions should have the same transformation properties, so the corresponding $\mathcal{R}(ph)$ for this phonon will be antisymmetric under the exchange $x \iff y$, with diagonal elements =0, that is,

$$\mathcal{R}(ph) = \begin{pmatrix} 0 & B \\ -B & 0 \end{pmatrix}.$$

It's easy to find a displacement pattern which fits the bill: a diamond stretch/squish distortion.

You can also start with a displacement pattern and predict polarization conditions. A displacement pattern that preserves the square would be like the displacement patterns mentioned above, but with all displacements in phase. The matching polarization conditions are $[\vec{x}, \vec{x}]$ and $[\vec{y}, \vec{y}]$, which each have a phase of +1 under all rotations. The partial Raman tensor for this phonon should look like

$$\mathcal{R}(ph) = \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix}.$$

5.4 Enumerating the Modes in a crystal: “Diagonal Ions” (requires some group theory)

The phonon modes in a crystal can be classified according to the irreducible representations of the crystal space group. The actual procedure for enumerating the modes hinges on the orthogonality of the characters of irreducible representations. The character tables of the space group for important points in the Brillouin zone are available in books and online. In addition, you need to identify the symmetry elements (e. g. rotational axes, mirror planes, etc.) in the unit cell. For this purpose, you might use the excellent diagrams in *The Analytical Expression of the Results of the Theory of Space Groups*, by Ralph W. G. Wyckoff (Carnegie Institute of Washington, 1930).

For visible light, Raman phonons live near the center of the Brillouin zone (zero wavevector) because the wavelength of visible light is large compared to the lattice spacing. We can therefore concentrate on the subgroup of the crystal space group corresponding to zero wavevector. This is just the point group of the unit cell, with symmetry elements distributed according to the particular space group (e. g. , for D_{4h} , it could be D_{4h}^{14}).

For this subgroup, the representation of all zero-wavevector phonons is the one which simultaneously depicts the vector displacements of all of the atoms in the unit cell—it is DN dimensional, if N is the number of atoms in the unit cell and D is the dimension of the unit cell. Under a symmetry operation of the subgroup, two tiers of transformations occur.

Firstly, sets of equivalent ions are shuffled among themselves. If, as a basis of the DN dimensional representation, we use the displacement coordinates of each ion, then for a given element of the group, only those ions whose identities are unchanged in the transformation can contribute to the character. For crystals, these are ions that reside on the rotational axes, reflection planes, etc. The others will be by definition off-diagonal, and so cannot contribute to the trace.

For each one of these special “diagonal” ions, the contribution to the trace is just the contribution of a vector displacement of unit length under the given transformation. The sum of the vector-trace contributions for the diagonal ions gives the trace of the

DN -dimensional representation for the transformation; the collection of these traces for the entire group gives the character of this representation for the group.

Once the character of the DN dimensional representation is known, decompose it into irreducible pieces by using the orthogonality of characters of irreducible representations. Each allowed phonon will appear once, and so the number of times $m(\alpha)$ a given irreducible representation α comes up gives the number of allowed phonons of that symmetry species. If d_α stands for the dimension of representation α , the sum rule

$$\sum_{\alpha} d_{\alpha} m(\alpha) = DN$$

must of course hold. This is a good check of your arithmetic.

For an example, consider again the fictitious square molecule of C_4 symmetry. The character table of C_4 is

| | e | C_4 | C_2 | C_4^3 |
|----------------|-----|-------|-------|---------|
| A | 1 | 1 | 1 | 1 |
| B | 1 | -1 | 1 | -1 |
| E ₁ | 1 | i | -1 | $-i$ |
| E ₂ | 1 | $-i$ | -1 | i |

Furthermore, the only operation that leaves atoms in place is the identity e , which leaves all four. In two dimensions, the vector character of the identity is obviously 2, and therefore the number of atoms \times the vector character, $N_A \chi^v = 8$, and we should subtract off 4 for the center of mass degrees of freedom. So using the orthogonality relation

$$\sum_i \chi_i^{\alpha*} \chi_i^{\alpha'} = g \delta_{\alpha, \alpha'},$$

to decompose the $2N$ -dimensional representation (here $g=4$), compute

$$4m(\alpha) = \sum_i (N_A(i) \chi_i^v - 2\chi_i^v) \chi_i^{\alpha} = 4$$

for each representation α , or one internal vibrational mode per irreducible representation. These match up with the square-preserving dilation mode (A), diamond stretch/squish (B), and two degenerate rectangular x and y stretches (E₁ and E₂). For the E doublet, the phase factors of i suggest that under $x \rightarrow y$, $\mathcal{R}_E \rightarrow \pm i \mathcal{R}_E$. One might best isolate these modes by looking with circularly polarized ($\vec{x} \pm i\vec{y}$) light.