Raman Scattering Lab Status

Ben Nachumi

September 8, 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 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. For visible light, the momentum difference between the incident and scattered light corresponds to very small wavevectors (of excitations in solids). So in solids it's a good approximation to treat the momentum transfer in Raman scattering as zero.

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. 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 Rayleigh scattering. Otherwise, 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 Apparatus

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 lest 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 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 How the Spectrometer Works

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

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

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 follws 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 $\lambda_c = 5320\mathring{A}$, for a 1200 groove/mm grating (use $\gamma = 7.67936°$).

3.2 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 $x\lambda_\nu = 2\epsilon \sin(\Delta+\nu/2)\cos(\Delta-\nu/2)$.

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 $\lambda_c = 5320\mathring{A}$ (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.

3.3 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 $\lambda_c = 5000\AA$, is about 1000 cm^{-1}, which is more than enough to see many phonons in a single exposure.

3.4 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 at angle Θ the angular width of the xth 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 be 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 corressponds to about half of an 11cm×11cm, 1200 groove/mm grating). Translate this number into a number of opixels, assuming that the width of one pixel is 24μm.

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 24μm, and the intrinsic spot size is 1μm, and the slit itself is 200μm, then the slit spot size will be about about 8 pixels (for $F = 640$mm and $\lambda = 5000\AA$, check to see that you can ignore diffaction 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 Status

As of August, 2011, I've familiarized myself with the software enough to start taking measurements. The spectrometer has been aligned with both the 1200 groove/mm and the 2400 groove/mm gratings, using lines from a mercury vapor lamp and from the argon laser, and I've begun to calibrate and look for Raman spectra in some well-known systems like crystalline silicon. These procedures are pretty close to what the undergraduates will have to do to get started.

For an example of how calibration of the spectrometer enhances the lab experience, consider my own most recent 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. The following thinking about this data is meant to represent the habits of imaginative reckoning and rechecking that we are trying to develop.

`CAL1C_text.JPG`

Figure 1: First multiline calibration image.

> With the grating set at a nominal central wavelength of 5018 $\mathring{A}$, the first picture shows three fairly tightly constrained spots of light from the laser. The lower spot is 4880 $\mathring{A}$, and the upper spot is 5145 $\mathring{A}$ (the students could know this by actually looking at the diffraction spots on a piece of white paper light held inside the spectrometer). From the known wavelengths of these lines, their positions on the CCD chip, and the size of the chip, one

`CAL2C_text.JPG`

Figure 2: 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.

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 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.

That is how modern spectroscopy experiments start. With this apparatus, and some guidance, the students will have the opportunity to do it for themselves.

5 What we still need for Undergraduate Raman

So all of the really difficult stuff (spectrometer, gratings, detectors, lasers, software) is out of the way. Here's what's still required for a good lab experience:

- I still have to track down notch filters or nice sharp long-pass filters for the laser lines so the Raman signal won't be swamped by Rayleigh scattering. We'd like to use lower-power diode lasers, and unfortunately the filters we have around are suited to the lines of the argon laser, which are too different for them to be of use. Good edge filters are expensive–on the order of \$ 1000 each, and good notches are about \$ 1700. To establish a Raman shift requires at least two separate excitations, so we'd need \$ 2000 for edge filters. At least one notch filter would enable the students to compare Stokes/anti-Stokes intensities and so take the phonon temperature of their samples.
- To look at selection rules, you need to reliably control the laser polarization. Because of the upper limits on our laser power, we need to do this without throwing

too much power away. There are several alternatives here, with educational benefit inversely proportional to price. Newport sells a polarization rotator for $ 3150. That would be easiest to use, because it works simply for any wavelength. Alternatively, we could get at least two wave plates, for $ 407 each from Thorlabs. Then again, we could make rotating polarization using what's at hand one of the experimental tasks, for the price of prisms (maybe $ 750).

- Delicate things need sturdy, steady table space. An optical table would be best, though new ones can be expensive. I have seen several suitable second-hand ones for around $ 2000. Alternatively, we might get away with laying a steel plate over a regular tabletop, and using magnetic mounts. But these are bulky and tend to stick together at inopportune moments, preventing some experimental arrangements involving short-focal-length lenses.

- Raman signals tend to be faint, and lasers need to be screened. So a ceiling-mounted blackout curtain, sized for the experiment, is essential if the experiment is to share a room with others. Newport and their ilk sell really expensive laser curtains for like $ 5000. We could probably make our own for less than $ 500 (ten feet of ceiling track and 130 square feet of heavy suede, held up with some strong magnets). I requested a quote from a maker of hospital curtains (Kentek) but haven't heard back yet.

- To see Raman, one must focus the image of the laser spot on the spectrometer entrance slit, while sending the specular reflection from the sample off-axis. It's good to be able to see the laser spot when you are trying to do this. I've modified an internal flip mirror in the spectrometer so it can send an image of the slit through a port in the side. It's safer for the students to have this image go to a cheap camera and thence to a viewing screen (so they can't accidentally flash themselves right in the eye while adjusting things). A tiny security camera and the associated mounting/focusing tubes will run about $ 300. As for the screen, I picked a working color security monitor out of Columbia's bin of discarded electronics about two months ago.

- And finally, we'd need a new computer (approximately $ 1000 from the likes of Dell) to dedicate to the newer astro camera (which uses a newer USB than the current machine supports). It would also be a good place to do data analysis.