Lab 11A • 02/29/12

Review of NMR

Let's start with what exactly goes on during an NMR experiment – not so much the construction of the machine, but, physically, what's happening to the sample that generates this NMR signal that we're trying to analyze. What is happening to the sample that's causing it to respond in a way that generates this NMR signal? What physically is going on? What interaction? We know it's spectroscopy, so we know light's involved somehow – but how? Nuclei are, in the absence of a magnetic field, randomly oriented. The nuclei, for some reason, align either with the magnetic field or against it. This is for a kinda of specific case of nuclei that, fortunately, carbon and hydrogen happen to be one of these nuclei. What is it that's causing it to align? The magnetic field. What aspect of the nucleus? The spin. Nuclei have spin, just like electrons have spin. Unlike electrons, though, nuclei have the potential to have spins other than just 1/2. But it turns out that protium, 1H, and carbon-13, which are in organic molecules, they have spins of 1/2. Spins of 1/2 can only adopt one of two orientations: with or against the magnetic field. It makes for a relatively easy case to discuss. In more advanced NMR theory, there are other nuclei that can have multiple energy states, which makes it a little bit more complicated.

Why do we care about these two different states? What does that create? Why does the fact that nuclei can align either with or against the magnetic field cause them to absorb these radio waves of certain frequencies?

Many nuclei have spin, in the same way that electrons have spin. Specifically, carbon-13 and protium, 1H, have nuclei with spin of 1/2. The quantity is not just 1/2, it's 1/2h-bar; h-bar is, in turn, h divided by 2 pi. h is Planck's constant, which tells us the size of energy packets, and it shows up in a bunch of equations [and] in a bunch of physical constants – in this case, it's spin. We're going to use Planck's constant again when we start talking about photons. The reason that these systems of 1/2 spin are nice to have to discuss is because there's only two energetic possibilities. When nuclei of spin 1/2 are placed in the magnetic field, they'll either align with the magnetic field or against it.

We can draw a picture to represent this. Let's say that we start out with the case of no magnetic field, in which we'll have nuclei that have their spins randomly arranged. Then, we can talk about what happens when we have a magnetic field, in which some of the spins will be aligned with and some will be aligned against. Let's say that this is the direction of the magnetic field. Why have I written the ones that are aligned or antialigned, why did I write some lower and some higher? What am I setting up to do? Energy – one has more energy than the other. This is being accessed when we do an NMR experiment, because if you have these two different energy levels, there's a certain energy gap to it. Without the magnetic field, you have one average energy level; with the magnetic field, with spin 1/2, that gets split into two. That corresponds to an energy gap, which we can relate that, using Planck's constant, to light – specifically , to a certain frequency of light. With that correct frequency of light, you can promote a nucleus that's in the lower energy level up to the higher energy level. The two possible orientations of the nucleus create an energy gap, which can be crossed if a photon with the correct energy is absorbed. That photon corresponds to a specific frequency of light.

It turns out, at room temperature, not all of the spins are going to be in the lowest-energy state. In fact, that same Ea^(1/RT) expression, that exponential expression, there's another version very similar to it that we could use to say what fraction of the nuclei are likely to be in that lower state versus a higher state, at a given temperature. What we do, then, is we take the sample, irradiate it with a range of frequencies. Any of those frequencies of light that correspond to photons that match these energy gaps, that light gets absorbed; the rest of the light just disappears after it's been released. As soon as you turn the switch off, as soon as you stop irradiating that sample, the energy that was absorbed gets released, because even though there may be some nuclei that even at room temperature will be in the higher energy stated, not all of them will be. If you add energy, all of them go up to that higher energy state; some of them fall back down again. That falling releases energy, and that's what we measure. So we zap it with all of the different frequencies, but we only get a few of them back. That's the technique of Fourier transform NMR. This is what's going on in an NMR experiment: we have these energy gaps that we're looking for.

Why are we looking for them? Because these energy gaps are really sensitive to three different effects. One of them is the machine itself – the bigger the magnet, the more the magnetic field affects the nucleus, the bigger the energy split's going to be. It also depends on what the atom is that we're searching for. Carbon versus hydrogen, hydrogen is about four times more responsive to a magnetic field than carbon is; that's in this little constant called the magnetogyric constant. There's one more effect that's the most important, because once we have a sample and we know that we're looking for carbon and hydrogen and we're just comparing different carbon hydrogens, and we're in the same machine, then there's one other thing that could affect that energy gap, which is what? Shielding or deshielding will affect this chemical shift.

Chemical shift is this idea: if I have an electronegative element that's next to me, I'm the nucleus being observed, but I've got a really electron-withdrawing neighbor, that neighbor is going to pull electron density away from me, which is going to expose my nucleus more to the external magnetic field.

The more that it's exposed to the magnetic field, the larger the energy gap gets, which the large frequency of light ends up getting associated with that. That's chemical shift. If a nucleus has an electron-withdrawing neighbor, that nucleus is more exposed to the magnetic field, which means it is deshielded. That causes a larger energy gap, which causes a higher frequency of light to be absorbed.

There is this numerical quantity, though, of chemical shift, which is written as this symbol delta. What is the mathematical definition of chemical shift? (observed frequency - machine frequency) / machine frequency * 1,000,000. Why do we care about chemical shift, why don't we just use frequency? We do use chemical shift to identify functional groups; in theory, we could use frequency to do that as well, but why is frequency a horrible way to go? Because every machine's got a different magnetic field, which means every machine's going to generate a different frequency. Buried on both the top and bottom of this equation, we have some proportionality to the magnetic field strength itself. Whatever would affect either the observe of the machine frequencies, they're both due to the same-strength magnetic field. By taking something involved with the magnetic field, dividing by the same type of quantity involving that magnetic field, the magnetic field's canceled out of the equation and we end up with a machine-independent quantity. That's why we care about chemical shift, cause it doesn't matter which machine we run it on, we're going to get the same values.

How do we figure out what the machine frequency is? We do need some kind of sample to test against. We do test against a reference sample; for carbon and hydrogen, what is that reference? TMS, which is not the same TMS as we had this quarter – this quarter we had trimethylsilyl [ethers], but this is tetramethylsilane, four methyl groups on that silicon. Why do we use TMS as the reference sample? It's shielded; it's less electronegative than other neighbors, which means the hydrogens and carbons attached to it don't have their electron density pulled away, which means they're not as exposed to the magnetic field, which means their energy gaps are smaller, which means the frequency you observe is also smaller. For 13C and 1H, TMS, tetramethylsilane, is used as the reference compound to determine the machine frequency.

Let's look at what a spectrum would look like. When we first do an NMR experiment, we don't get the spectrum like what we're going to interpret; we instead just get this wave that comes off, with all of the different frequencies being released by that sample being absorbed at once. This is why we call this the Fourier transform technique, Fourier transform NMR, because we have to do a mathematical operation to pull out how much of each of those frequencies is present. Once we do do that, we will end up with a graph, where the y-axis is intensity, and the x-axis is chemical shift. Remember that chemical shift increases from right to left; this is due to historical reasons, because in the past, more routinely if we were trying to do an NMR experiment, we'd have a machine that was really good at measuring this one frequency, and then we would vary the magnetic field. If you had one particular frequency you were trying to observe at, if you had a particularly shielded atom, then you'd have to increase its magnetic field in order to create the same energy gap to be able to be scanned. Or, if you had a very deshieled atom that's very exposed to the magnetic field, you'd have to turn the magnetic field down lower to lower that energy gap to match the frequency that you're scanning at. If we were to do this graph in terms of magnetic field, magnetic field increases to the right and decreases to the left, so we get upfield and downfield. But, modern NMR, FT-NMR, we don't change the field, we change what we're observing at – in fact, we observe everything at once. That's how we have this opposite relationship, because there, we have a fixed magnetic field, which means that if you did have something that's shielded, you'd end up with a smaller frequency; if you had something that was deshielded, you'd need a larger frequency, instead of you having to have, respectively, a larger magnetic field and smaller magnetic field.

Let me write some fake spectrum up; I'm going to write something on here that we would need in real life. What three types of information can I pull off of this graph? What three things do I care about when analyzing compounds am I going to need off of this graph? In order to solve the problem, I'm going to need to know how many different types of protons there are; let's tackle that later. I could use this graph to solve the structure to determine how many different kinds of protons I have, but how am I going to get there? What numerical pieces of information do I need? What do we care about the fact that one of these clumps is located to the left of the other? [peek and peak, pick] What does that tell us? One is more deshielded, which means what? Different functional groups. [functional group reference with approximate chemical shifts] For example, any time for hydrogen that you've got a chemical shift greater than 10, it's got to be an aldehyde or carboxylic acid; between 6.5 and 8, it's a benzene ring most likely; between 4 and 6, it's probably an alkene; between 2 to 4, probably an alcohol; in the 2s or 3s, an alkyl halide; below that, just alkyl groups. One of the pieces of information that comes from this spectrum that we need is chemical shift, which can be used to determine functional group.

What about those numbers I wrote down, what are those? Integrals, the area of the curve. The intensity is related to how many hydrogens there are. Integrals are really only used for proton NMR. [Integrals] are proportional to the number of hydrogens. I say proportional because, depending on the compound and the machine, there may not be a fixed relationship, at least from compound to compound to compound, as far as what kind of absorbance you get versus how many hydrogens are really there. Another way to put it is, if I put cyclopentane into the machine and cyclohexane, at room temperature they each only have one type of hydrogen, so you'd just get one signal. Well, if you had one signal and one signal, how could use a single signal in order to determine which compound is which, especially if their chemical shifts end up being so close to each other that they might as well be identical? The answer is, you can't, unless you do a quantitative experiment, where you know the molarity of the compound, and you put in a reference sample, like TMS, and you know its molarity as well.

Then, you could calculate a direct correspondance between the signal and the number of hydrogens that are there. But, I am normally going to give you those integral values. The two and the three up there refer to two hydrogens and three hydrogens, respectively.

There's one more piece of information, though, which is what? Splitting, which provides information about the number of neighbors. This is what makes proton NMR so terribly useful, because you don't just know information about each individual atom, you know information about their neighbors, automatically, due to splitting. If you know my configuration and what's going on with my neighbors, and you put all of the different pieces of information together to try to overlap them, connect them, that's how you're eventually going to end up with a full molecular structure. That's what we have to learn now after this review – how do we take this information and work our way back to a structure. Before we can do that, though, is we need to take an example molecule and see if we can possibly guess what the spectrum would look like.

Before we do that, even, let's take the top spectrum I already wrote and see if we could agree that a compound like bromoethane would match that spectrum. How many different types of hydrogens are there on that molecule? That's the wrong answer. The right answer is: am I in a chiral or achiral solvent? Remember, it matters. I'll ask the question again: how many hydrogens are on there? Your response is ... well, what kind of solvent do you want it to be: chiral or achiral? Achiral? How many hydrogens are there? Two is the right answer to that question, because in a methyl group, it doesn't matter which hydrogen that we looked at, they'd all end up being equivalent, no matter what their neighboring environment would be. These are what are known as homotopic protons, ones that if you replace any one of the protons with some random group X would not generate a stereocenter. The reason characterizing it that way matters is because by ability or the potential to generate a stereocenter, that fact matches the fact that that position would have different interactions with the neighbors around it. In this case, if we say replacing any one hydrogen doesn't make a difference, then no matter what you do, no matter what solvent you're in, those three hydrogens are always going to be equivalent. What about these two hydrogens? If I replaced one of those two hydrogens with some random group X, then what kind of molecules could I potentially generate? Enantiomers. Replacing those, we could potentially make enantiomers, which in achiral solvents doesn't matter, but in chiral solvents, it does, so in a chiral solvent, there are three different hydrogens there, because the two blue ones would become inequivalent.

Does this match the spectrum? Yes, because if I circle this lefthand peak and look at that integration of two, that matches up with the two blue protons we have. Those two blue protons have three neighbors, which having those three neighbors causes the peak to be split into a quartet, which are the four lines that are clumped together like that. Similarly, if we were to look at the three red hydrogens and match them up with this peak that integrates to three, we'll see that the three red hydrogens only have two neighbors; two neighbors cause a triplet, which is what we've got, so that spectrum matches the compound.

Let's do the same thing, but now take a compound and predict what the spectrum's going to look like. Same question I'll ask you: how many different types of hydrogens are there – in an achiral solvent? Four? Yes, four different types of hydrogens, in an achiral solvent. How about a chiral solvent? Six? That's right, because we have two enantiotopic positions – if you replace one of the blues or one of the greens with some random group, you'd make a different set of enantiomers, so both blues and greens are enantiotopic, which means they would show up differently in chiral solvents, but 99% of the time, we're going to be working in achiral solvents, so normally when I ask how many [kinds of] hydrogens, you can assume I mean achiral solvents. In this case, four. What would be the splitting for the red hydrogens? What causes splitting in the first place? Splitting is caused by my neighbors. This is a why it's a useful technique, because it's not telling me information about a particular atom; it's telling what's next to that atom. If I've got one neighbor over here, maybe that neighbor's spin is aligned with the magnetic field. If it is, from my perspective, I'm experiencing both the real magnetic field and this guy's magnetic field over here, and since they're aligned with each other, they add to each other, which means my energy gap is bigger, which means my frequency is bigger which means my chemical shift is larger. Just as likely, I could have the neighboring spin pointed down, which I'm experiencing the main magnetic field by this one is subtracting from it, and I'm close enough to it that I experience that. That causes a smaller energy gap, which means a smaller frequency, which means a smaller chemical shift. Since, over the course of an NMR experiment, given the number of molecules around, you end up with pretty much equal probability that that neighbor could be up versus down, which means one absorbance is effectively going to be split into two when it acts with one neighbor.

Let me come back to that methyl group and point out how important it is to realize the role of the carbonyl. What if I just told you I have a spectrum that has a singlet in it – what does that mean? A singlet means only one kind of hydrogen, if all I've got is just a singlet. But let's say I have multiple peaks, and I see that one of the peaks is a singlet. The only way you can generate a singlet is if you fail to have neighbors next door. If you look at the carbonyl, there are no hydrogens here on that carbonyl, so there's nothing for the red hydrogens to interact with, so there's no splitting. The carbonyl blocks the splitting. That's an important way to analyze a compound. Let's say we had something like acetone: when we see a singlet that integrates to 3, with some practice, you'll be able to jump automatically to: methyl ester or methyl ketone, or something like that. Carbonyl has no hydrogens on it, so it blocks splitting of the methyl group.

The way that I report spectral information will be like this: I'll say what it's chemical shift is. It turns out for a methyl group next to a carbonyl, it's usually somewhere around 2.0 for its chemical shift. Normally, then, next you would list its splitting: s stands for singlet, the fact that there are no neighbors.

Then, we would list the integration, which in this case, that peak would correspond to three hydrogens. What about the green hydrogens? They do have neighbors: how many? How many neighbors do the green hydrogens have that they're going to interact with? Just two, the blue ones. Again, the carbonyl prevents interaction with the red hydrogens on the other side of that carbonyl, so the green ones are only going to interact with the blues. If one neighbor causes a peak to split into two, then another neighbor causes each of those two peaks to split again. Since we're talking about the same hydrogen, what we're going to find happens is in that split, yes, you generate two more signals, but you also generate one signal that comes from the other two signals that were split, so you end up with a total of three signals, so you get a triplet. Those green protons, with a chemical shift value of approximately 2.3, we'd have a triplet that'll integrate to 2.

Now let's do a splitting tree. Let's say that I'm doing it with the blue hydrogens [letter designations]. If I look at proton Hb, it has two Hcs next to it. The first Hc, in the way that I described earlier, is going to cause a higher chemical shift and a lower chemical shift peak. How far those peaks are split from each other is called a coupling constant. Each interaction has its own coupling constant. This is when we just have interaction with one hydrogen; let's throw a second hydrogen in. As I was saying a moment ago, that will cause the formation of a total of three different peaks. In the middle, the reason there's two lines is because that middle peak will be twice the intensity of the two outside ones, cause it has two unique ways that it was generated. The reason we have those two overlapping exactly is cause we still have exactly the same coupling constant. Remember that we had a way of looking at spin states and using spin states to rationalize: one, why do we get the different numbers of possibilities; and[, two,] why do we get these what are called degeneracies – the fact that multiple peaks overlap with each other.

Now, what about Ha? That's a third neighbor, isn't it? How many total neighbors to Hb do we have? Three. You might remember this (n + 1) rule, which says that if n is the number of neighbors you have, then n + 1 is the number of peaks that get generated, so if I have three neighbors, that means I should make a quartet. But that's only if the coupling constants with all of the neighbors are the same. Even if you have different functional groups, you can have different neighbors that their environments are so similar enough that their coupling constants end up being roughly the same, which means the n + 1 rule works. But what if we say that the coupling constant from Hb to Hc is much greater than the coupling constant between Ha and Hb, it's one neighbor to the other side? That one neighbor's still going to split it, but not as much. Notice that what we had as a triplet, each on of those portions of the triplet becomes a doublet; another way that we sometimes say this is that it's a doublet of triplets. If I'm writing out this information, I'll say chemical shift 2.5; it's a doublet of triplets that integrates to 2. The reason why knowing it's a doublet of triplets is useful is because we know that specifically means we have one neighbor that's one type and we have two other neighbors that are of a different type. Multiple splitting. The last peak we didn't touch on at all would be that Ha. It would have a chemical shift of 5.3; it would be a triplet, and would integrate to 1 hydrogen.

We're [now] going to do this in reverse. We're going to take these pieces of information – the chemical shift, splitting, and integral – and we're going to go backwards to a [compound].

Before we do any interpretation, the spectrum looks something like this: a sine wave, a superposition of multiple sine waves with all of the different frequencies that are getting absorbed [and then released again]. Fourier transform is an integral where if the one function you're comparing against exactly matches the function you're comparing it to, if the test function matches the function you're trying to analyze, you get an integral of one; otherwise, if they don't match, you get an integral of zero, effectively. Related to that is the Dirac delta function. This is intensity vesus time; Fourier transform converts that into inverse time, so it converts the domain into the inverse domain. Time, if you take one over time, that's frequency, so you're going from time to one over time. Essentially what happens is: take a sine wave with a certain frequency, cointegrate it with that function, and out will pop out how much of that one sine wave is located in that conglomeration of sine waves. You do that for every single frequency, and you can find out how much of each frequency of light is present. This technique wasn't possible until we had computers, so that's why we have non-Fourier transform NMR, where you slowly scan at different frequencies, or you slowly adjust your magnetic field and keep plowing away at the same frequency.

_____

Nuclei have (usually) spin in much the same way that electrons have spin. For example, 1H (protium) and 13C have spin 1/2 h-bar. Nuclei with spin 1/2, when placed in a magnetic field, will align either parallel with the field or opposite the field. The two possible orientations of the nucleus create an energy gap, which can be crossed if a photon with the correct energy is absorbe. That photon corresponds to a specific frequency of light. If a nucleus has an electron-withdrawing neighbor, that nucleus is more exposed to the magnetic field (deshielded), which causes a larger delta E, which causes a high frequency of light to be absorbed.
For 13C & 1H, TMS (tetramethylsilane) is used as the reference compound to determine the machine frequency.

Valuable info:
1) Chemical shift – used to determine functional group
2) Integrals – proportional to # of H
3) Splitting – provides information about the # of neighbors
_____

Structures

2/29/12 lab • 1

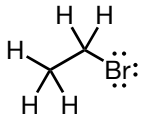

2/29/12 lab • 2

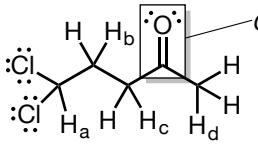

C=O has no H on it, so it blocks splitting of the –CH3