Lab 11B • 03/01/12

NMR review [aerodynamics of understanding]

What, physically, is going on, not with the NMR machine, but within the NMR sample that gives rise to the spectra that we acquire? The sample absorbs some frequencies of light, and then what happens? What causes that signal to be generated? Why is it that light gets absorbed? There's a magnetic field involved, and there's this thing called spin that some, not all, nuclei have. In this particular case, we're looking at carbon and hydrogen – specifically, protium, which is 1H, and then 13C, because they have spin of 1/2 – that's the same 1/2 as you'd find in an electron's spin. The reason that ends up being really convenient is that, with spin 1/2, you only have those two orientation possible; with other spins, there can be multiple energy states, which doesn't mean you can't do NMR, but it makes it a lot more tricky to interpret. By only having the up or down possibility, we have one energy gap that corresponds to a particular nucleus. That nucleus can then absorb light, but only at a particular frequency. If you can zap a sample with a range of frequencies, it'll absorb only the ones that correspond to those energy gaps, and then after the power's turned off, release that energy. We get that and interpret that as the spectra. Nuclei, such as 1H and 13C, have spin, and specifically of spin 1/2. Protium is the name when you want to refer specifically to 1H, the main form of hydrogen. It matters, because deuterium's a different nucleus, so it doesn't respond at the same frequencies that protium does. 1H and 13C have spin 1/2. The 1/2 designation's actually incomplete; it's actually 1/2 times h / 2pi, which is the same thing as one half of this symbol called h-bar. The reason 2 pi comes in is because it has to do with rotation, and 2 pi is the circumference of a circle if you're in radians. That h is Planck's constant, which tells you the size of energy packets. That's related, then, to spin, because spin is, indirectly, kind of a representation of a certain amount of energy.

Nuclei with spin 1/2 are randomly arranged when not in a magnetic field, but will either align with or antiparallel, against a magnetic field, once a magnetic field's present. To draw a rough picture of this, in the case where we have no magnetic field, there's nothing, really, to cause those spins to want to align one way or another. On average, it's going to be randomly oriented, and there's going to be, technically, therefore, only one energy level. But then, when we do have a magnetic field, we end up with the two possibilities: spins aligning with the direction of the field or aligned opposite direction. Those each have an energy associated with it, energy levels, and we can see, compared to the randomly oriented state, we have a split in energies that occurs. That causes an energy gap to form. The nuclei won't necessarily all be in the lowest energy level at room temperature, because thermal energy can, in a sense, cause nuclei to flip, so there'll be some portion of the nuclei at the lower state and some proportion at the higher state. That same e^(-Ea/RT) factor that shows up in the Arrhenius expression, that same form of equation shows up to tell you the proportion, the population of these two energy states is. The point of this is: if you were to add a photon, a packet of energy, that exactly matches that gap – because we do have an equation that relates the two: that is frequency of light that corresponds to a photon of a particular energy – then you can excite the nucleu[s] to be in that higher spin state, that higher-energy spin state. But of course, once you turn the light off, nothing is going to hold those nuclei up permanently, and they're start reverting back to their lower-energy state. That's the release of energy that we measure; that's how the spectrum gets generated. Only when light of the right energy gap hits it do we get a response back afterwards. If a photon with the correct energy, which corresponds to light of a particular frequency, is absorbed by a nucleus, the nucleus can flip from the low- to high-energy state. The energy released when the nucleus flips back to the low-energy state is what's measured by the NMR spectrometer.

If you only have one nucleus, it can't be two things at the same time, so if it's in the high energy state, it can't be the low energy; it's not like there's two spaces, it's one nucleus that's aligned one way or the other.If it's flipped this way, it's not that way.

There's three factors that can affect that energy gap. It matters because we're trying to measure those energy gaps. One of the factors is just what is the machine strength itself – the stronger the magnetic field of the machine, the larger the energy gap that you're going to generate. There's also the nucleus involved – whichever particular nucleus you have has its own sensitivity to the magnetic field. For example, hydrogen is almost exactly four times more sensitive than carbon, which means the frequencies that you scan hydrogen at are four times that of those frequencies that you scan carbon. There's one more general effect that can change that energy gap – that's the most important one, usually, because we're only going to be looking at carbon and hydrogen, so it's not like we're comparing multiple nuclei, and when we run an NMR experiment, we only do it on one machine. What's the main effect that's going to change these energy gaps, and therefore change which frequencies we observe that gives us that information that lets us get a molecular structure – the chemical environment, or we can also refer to it as chemical shift. Could you describe what happens to cause that energy gap to change?

[answer] For example, I'm a nucleus being scanned. If my neighbor next to me is really electronegative, it pulls electron density away from me, which means I experience the magnetic field, as a nucleus, to a greater extent – that's calle deshielding. When you deshield a nucleus, because that nucleus feels the magnetic field, in a greater sense, then that gap increases, which means you need a photon with greater energy, which means you have to have light of a higher frequency.

If you had something that wasn't electronegative, then you don't have that deshielding; you're shielded, which means you have to have a stronger magnetic field in order to get the same response, or if you were in the same magnetic field, that energy gap's going to shrink, which means a photon of less energy energy, which means light of a lower frequency. With that piece of information, chemical shift, often we can identify a functional group; it's useful information. If a nucleus has an electronegative neighbor, that neighbor will pull electron density away from the nucleus, so it will deshield that nucleus, exposing the nucleus more greatly to the magnetic field. This causes a larger energy gap, which means a higher frequency of light would be needed.

here is another use for the term chemical shift, which is a numeric quantity which we represent with the symbol lower-case delta. What [is] the definition of chemical shift, in this sense? It is a difference of frequencies: the observed frequency minus machine frequency, divided by the machine frequency, times one million; [since] this is in units of parts per million. Why do we care about this quantity; why care about chemical shift like this, why not just use frequency? Each machine's got a different magnetic field, which means you'd register a different frequency on each machine, so it wouldn't be easy to compare from one machine to another; it'd be worthless, in fact. But by the fact that the machine field strength is directly proportional to both the observed and the machine frequencies – because the same magnetic fields operating in both cases – the fact that we're taking frequency and dividing it by frequency, which buried in that, that means we're taking the magnetic field and dividing it by the magnetic field, we end up with a machine-independent quantity. That's why chemical shift is useful.

What's the machine frequency? The machine frequency is the frequency we get when we measure a specific standard, which for carbon and hydrogen, there's a silicon-based standard, which is what? TMS, which [in this case] does not mean trimethylsiyl – we have this TMS protecting group – TMS also stands for tetra, which means four, tetramethylsilane, which is the compound that's used for calibration. For protium and 13C, the machine frequency is determined by using the reference standard TMS, which stands for tetramethylsilane. Recall that, in actuality, when we are acquiring these spectra, we don't really get the end product at first, what we look at to interpret, because what comes out of the sample is just a set of waves, different frequencies of light all stacked on top of each other. We get this big overlapping wave that decreases over time as the nuclei relax; that's got all of these different frequencies buried in it. There's this technique we can use called a Fourier transform, which is able to pick out of that combined wave each of the individual frequencies and how much, in terms of intensity, each of those frequencies is present. Once you do the Fourier transform, then we're got our spectrum.

I want to draw an idealized picture of a spectrum so we can discuss it a bit. The y-axis is intensity, how much signal is generated. The x-axis is chemical shift, but recall that chemical shift increases from right to left. That's because there's a historical reason behind it. In older forms of NMR which sometimes are still performed today – there are still applications for it – in older systems, you might adjust the magnetic field instead of adjusting the frequency that you're trying to observe at. If you were adjusting the magnetic field, then what you were doing was scanning consistently at just one frequency. If you had something that was deshielded, which means it's more exposed to the nucleus, it means you would need to decrease the magnetic field in order to get that energy gap to be correct to be scanned at that one frequency. If you were a shielded nucleus, that means you'd have to have a stronger magnetic field to bring that energy gap up to the point where it would match that particular absorbance. So, magnetic field strength, in that old style of performing NMR, was written increasing from left to right, but because the effect is opposite – with the shielding you needed a bigger magnetic field, with a deshielding you needed a smaller magnetic field – if we're looking at it in terms of frequency, it's opposite. If we fix the magnetic field, keep it constant, then when we have shielding, we get the smaller frequency, and when we have deshielding, we have the larger frequency. That's why chemical shift is plotted in the opposite direction, because there's an inverse correspondence. [That's where these terms upfield/downfield come from]

Let me draw an example spectrum. [rant about studying NMR] Let's look at this spectrum. Let's say that I have a formula of C2H5Br. Let's first figure out what are the key pieces of information that we can get from this spectrum. The chemical shift, that's the first thing we can notice. That tells us something about the functional group. If it's for hydrogen, if it's 10 or greater, it's almost automatically an aldehyde or a carboxylic acid; if it's between 6 1/2 and 8, it's probably a hydrogen on a benzene ring; if it's from 4 to 6ish, it's probably an alkene; from 2 to 4, an alcohol; if it's around 2, it could be next to a ketone (of course, with a ketone, you wouldn't have it on the carbonyl but next to it); and then below 2, it's probably some form of alkyl halide or just alkane. [Chart in text] You can use that chart to try to guess what functional group that you have. [chemical formula will be given] [You don't] go to that chart and looking at every single absorbance, cause if you don't have oxygen, you wouldn't look at the chart for anything that has oxygen in it, for example. Therefore, you can use those charts, to guess, structural patterns, motifs that you have. If you know, for example, that you have an oxygen and maybe you have something that shows up next to a ketone, then you know you've got a carbonyl. That way, you get little hints that lets you solve the structure. Chemical shift is one of the things that's important. That would be useful in identifying the form of functional group that you have.

What's another piece of information? We have splitting, and what's the last one? Integrals, which [are] proportional, for 1H NMR, to the number of hydrogens. Turns out it's not as useful for carbon NMR, largely because there's only 1/100th the amount of carbon versus hydrogen in an average sample, at least that responds because it's spin 1/2. For hydrogens, it's useful. Notice I say it's proportional, not equal. For example, if you had the compound cyclohexane, which at room temperature only has one type of hydrogen, and cyclcopentane, which again at room temperature only has one type of hydrogen – if you only have one hydrogen, you're only going to get one signal out of it.

So, if you have these two different samples, and if you didn't know their concentrations, and you just took the absorbances, you're going to get this number from the machine that's some form of intensity, but how can you relate that at all to how many hydrogens you have around? The answer is, you can't, unless you know the concentration of your sample and unless you put in a reference, like TMS, that you also know the concentration of. Then, you can do a quantitative measurement. [normally integrals will be given]

Then, there's splitting. Splitting comes from what? The fact that, if you're observing one nucleus, you've got another neighbor, or two or three, that are going to interact with the magnetic field that that nucleus is going to experience. For example, if you have a neighbor that is aligned the same way as the magnetic field is, then I'm still going to feel the general magnetic field, but also, because I'm so close to this one, I'm going to feel the magnetic field coming from over here. They're going to add to each other, which is going to increase my energy gap. At some point, maybe that neighbor's pointed the opposite way, instead. Although I've got the main magnetic field that I'm experiencing, now this one is subtracting from it, canceling it – just locally, but I'm local, so it's affecting me. Statistically, both of those possibilities are going to occur – not simultaneously on the same molecule, but we've got billions of molecules in the sample. So, at the same time, what used to have been a single signal will be split and generate two peaks, right next close to each other, which we call a doublet. If you have an additional equal neighbor, we'll have a triplet; three neighbors, a quartet. That's this n + 1 rule. Splitting is useful because the other information is talking about the position that's being scanned, but if you know something about what's being scanned and what's next to it, then you can start piecing together which is the neighbor of which, and that way you eventually end up with the structure of the molecule. Splitting gives information about neighbors.

Now that we know what the three main factors, the three pieces of information we're going to use are, let's go back to the spectrum and see if it makes sense that this would match a compound like C2H5Br. First, was this spectrum obtained in a chiral or an achiral solvent? Why does it matter? Because how many homotopic, enantiotopic, or diastereotopic protons does this compound have? What do those terms mean? Homotopic means that, if you were to replace one of the hydrogens at that position, you wouldn't generate multiple possible structures, there would be no stereocenter involved. The reason the potential to make a stereocenter even matters is because if you have something over here but pointing only one way that's electronegative, and if I had one hydrogen that's closer to it versus one hydrogen that's further away from it, on the same position, they're going to have different chemical shifts, they're going to have different chemical environments. If you don't have the potential to make a stereocenter, that means all of the hydrogens are going to be equivalent no matter what you do. Those homotopic protons, it doesn't matter what solvent that you're in, they'll never split. That would be what we have with the methyl group in this compound, because if you say methyl group, there's three of that hydrogen there, and if you replace one with some random group X, you still don't have a stereocenter, which means you wouldn't have any difference in the interaction of that positon with any neighbor it has, so we ignore it.

What about the other hydrogens? If we were to replace one of these, if you look at the Hb hydrogens, if we replace one of those versus the other, what will we generate? A stereocenter, but specifically in terms of the molecules that we would make, what kind of molecules would we make? Chiral, but ... what would be the relationship between the molecules that would be generated? They would be enantiomers, so we therefore call these enantiotopic protons. Why do those matter? Because, in achiral solvents, they actually don't matter, but in chiral solvents, then you're in that particular situation: if you have a stereocenter, where you have a hydrogen pointing one way and some other functional group the other way, and then you have a chiral solvent that has a particular arrangement to it as well, then you're going to have differing interactions with different sides of the molecule. So in chiral solvents, Hbs, those would be different from each other. But, in an achiral solvent, they would be the same. The number of different hydrogens is the number of different NMR signals. We have two NMR signals that match these two hydrogens, so it means this was acquired in an achiral solvent. There is one more relationship – diastereotopic protons, that if you replace one position and you can form diastereomers, then those are called diastereotopic, which are always potentially different from each other.

Leaving that little bit of history behind, let's again look at this and see if this matches up. I have two hydrogens and three hydrogens; I have an integration of 2 and an integration of 3, so we might think that the 2 corresponds to Hb, the ones that are in red, and the 3 would correspond to Ha, the ones I've written in blue. Does that make sense? Look at the groups of peaks there, these multiplets get formed because of splitting. The Ha hydrogens, the blue ones, are split into a triplet. If you look, they only have two neighbors, so that matches – two neighbors would make a triplet. In turn, the red hydrogens, the ones that integrate to 2, they have three neighbors, which is why that's split into a quartet, so that does match. Next, we're going to go from the spectral information to a structure; let's prepare for that by taking a structure and seeing if we can predict what its spectrum is going to look like.

Here's the example molecule I want to use. In achiral solvents – 99% of the time we're only going to be discussing things in achiral solvents, so normally you're not going to have to worry about [enantiotopic] protons. How many different types of hydrogens would there be on this molecule, then? Four, cause if you don't worry about the possibility of making stereocenters, we have one hydrogen where the chlorine is – I'll call that Ha – we have two hydrogens on the position next to it, which would be the same – I'll refer to one of them as Hb – then next position over, same situation, two hydrogens which would be equivalent – which I'll call Hc – and then we have the methyl group on the other side of the ketone – which I'll call Hd.

There are no hydrogens on the carbonyl. Why does that matter? You can't split if you don't have neighbors, so notice that the red hydrogens have no neighbors, so they're going to have no splitting. Let's just look at that first set of hydrogens. What I'd like to do is show you how you'd express the spectral information that would come from this. Three piece of information we have – chemical shift, splitting, and integration. That's the order that, generally, the information's listed in. Where could we get the chemical shift for the hydrogen like that? [table] It turns out that, for a methyl group on a ketone, it generally shows up around 2.0. There'll be no splitting, which is super-important, cause if you have a singlet, the only way that that's possible is by having something next to you that doesn't have any hydrogens on it. If you can clearly identify you have possibly a methyl group – if it integrates to have three hydrogens – three hydrogens with a singlet is almost always a methyl group attached to a carbonyl, attached to an oxygen, attached to a quaternary carbon (one that has no hydrogens on it). Knowing you have that blocking occurring often lets you solve, very quickly, the structure of a compound, cause it limits the types of structures that are possible.

Let's go to the green hydrogens, Hc, next. It turns out that, when you have carbon as a neighbor, there ends up being a slight electron withdrawing effect that occurs, so the chemical shift for the green hydrogens is going to be a little bit higher than the red ones, and, based on the tables, I'm going to guestimate that it shows up at about chemical shift 2.3. What would be the splitting that would occur? Triplet, because one neighbor would cause the one absorbance to be split into two, the next next neighbor would cause each of those to split, but since it the same neighbors, when they split a second time, from those two peaks you get one peak but with double the intensity. You also get a peak on either side of it – one splits into two then splits into three, a triplet. There's only two hydrogens there, so it should integrate to two. Let's skip over Hb, the blue ones, because that's going to be the trickiest one; let's go to Ha. Ha is right next to, or alpha to, the chlorine position. There's only one hydrogen there; it's going to have a lot of electron withdrawing around it. Based on the tables, this has a chemical shift somewhere around 5.3. It's got two neighbors, just like the green hydrogens did, so it'll also be a triplet. It's only going to integrate to one because there's only one hydrogen at that position.

Now let's look at the blue hydrogens, and let's write a splitting tree – an explanation for the splitting we would observe. If Hb were all by itself, had no neighbors what soever, you would only generate one signal. But, if I had one of the Hc's that interact with it, then it splits – that neighbor, Hc, might be up or down, which then would affect the magnetic field experienced by Hb. That distance between those two peaks has what's called a coupling constant. I've labeled it Jbc because it represents the extent of interaction between B and C. If we throw one more neighbor in, then each of these two peaks that we just made is going themselves split, but because it's being split by the same neighbor, Hc, the coupling constant is identical again. If you have Hb and it equally splits this first peak – and that means half of it's that way and half of it's that way – when it splits again, half of it's going to come back, which means it will be in the same position that Hb was in to begin with. That'll happen from both directions, so that's why we end up with something like this. Notice that, in terms of intensity, we now have a pattern of 1, 2, 1. If we he had another neighbor that was exactly the same, we would then have a pattern of 1, 3, 3, 1. If we had another neighbor, 1, 4, 6, 4, 1. This is Pascal's triangle.

[Previously,] we said that if you know how many neighbors you have, add one to that, and that'll tell you how many peaks you end up with. We didn't talk about the fact that what if you had two different kinds of neighbors. Look at Hc, the one in green, and Ha, the one in purple – two totally different chemical environments: one that's next to a ketone, one that's next to two halogens. Their coupling constants are very likely going to be different. Let's imagine if the coupling constant between A & B is much smaller than the coupling constant between B & C. Then, when we throw Ha at it, again each one of the peaks we had is going to split, but now they're no coalescence, there's no overlap that occurs, because the coupling constants are different. So, you start out with a singlet. With one neighbor, you make a doublet; with one more equivalent neighbor, you make a triplet. But now, with this different neighbor, each part of that triplet has made a doublet. Depending on exactly which one is more prevalent, we could call this a triplet of doublets, or we could call it a doublet of triplets. When you use a symbol like that, that specifically means you have two different neighbors. That's useful information. If I know that I have just a triplet, then I know that maybe I'm up against a ketone like this, and I can only have splitting on one side. If know I have a doublet of triplets, that automatically means I have neighbors at least on two sides of me – possibly three if it's a symmetric molecule. Going back to the blue hydrogen, it'll have a chemical shift of roughly 2.5. It'll be a doublet of triplets that integrates to two.

How do we work the problem in reverse? You start with that and the chemical formula.

_____

Many nuclei have spin. For example, 1H (protium) and 13C have spin 1/2 (technically 1/2 * 1/2pi * h = 1/2 h-bar). Nuclei with spin are randomly arranged in the absence of a magnetic field; in the presence of a magnetic field, nuclei with 1/2 can align with the magnetic field or against it. This causes an energy gap to form. If a photon with the correct energy (which corresponds to light of a particular frequency) i absorbed by the nucleus, the nucleus can flip from the low- to the high-energy spin state. The energy released when the nucleus returns to the lower state is measured by an NMR spectrometer.

If a nucleus has an electromagnetic neighbor, that neighbor will pull electron density away from the nucleus (deshield), exposing the nucleus more greatly to the magnetic field. This causes a larger energy gap, which means a higher frequency of light would be needed.

For 1H and 13C, the machine frequency is determined by using the reference standard TMS (tetramethylsilane).

C2H5Br
1) Chemical shift – functional group
2) Integrals – proportional to the # of Hs
3) Splitting – gives information about neighbors

delta 2.0, s, 3H; 2.3, t, 2H; 2.5, dt, 2H; 5.3, t, 1H

_____

Structures – Identical to those from lab 11A (02/29/12)