The Nuclear Overhauser Effect

Eugene E. Kwan

March 24, 2011.

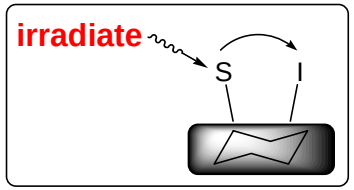

Scope of Lecture

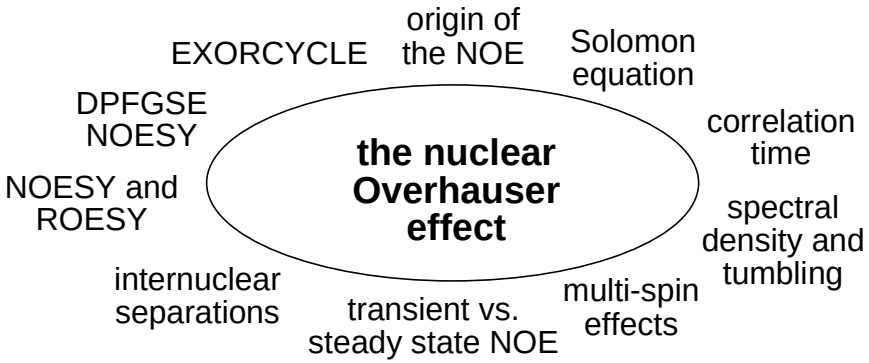

Key References

1. Understanding NMR Spectroscopy (2nd Ed.) Keeler, J. Wiley, 2010. (Chapter 9)

2. NMR Spectroscopy Explained: Simplified Theory... Jacobsen, N.E. Wiley, 2007. (Chapter 8)

3. The ABCs of FT-NMR Roberts, J.D. University Science Books, 2000. (Chapter 6)

4. High-Resolution NMR Techniques in Organic Chemistry (2nd Ed.) Claridge, T.D.W. Elsevier, 2009. (**Chapter 8**)

5. The Nuclear Overhauser Effect in Structural... (2nd Ed.) Neuhaus, D.; Williamson, M.P. Wiley-VCH, 2000.

Key Questions

(1) Where does the NOE come from?

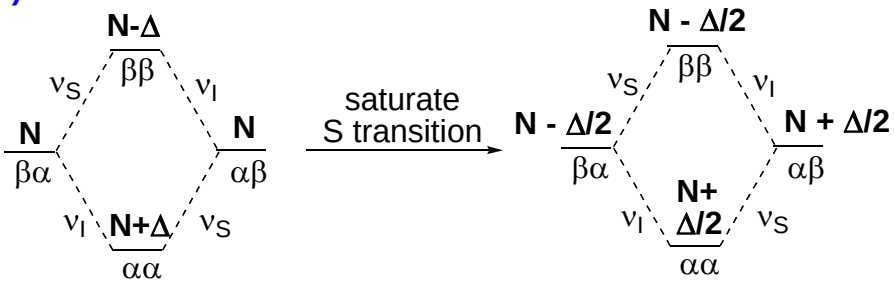

(2) Why does the NOE have a sign and magnitude? Why does the molecular weight of the compound matter?

$$n_I\{S\} = \frac{\gamma_I}{\gamma_S}\left[\frac{W_2 - W_0}{W_0 + 2W_1^I + W_2}\right] \qquad J(\omega) = \overline{B_{local}^2}\,\frac{2\tau_c}{1+\omega^2\tau_c^2}$$

(3) What are the best ways to get NOE correlations?

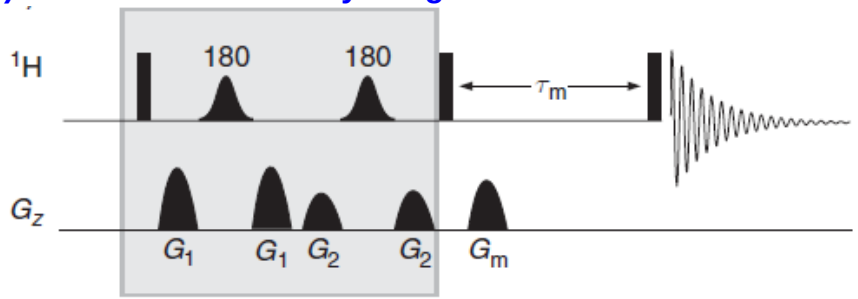

DPFGSE selective excitation

(4) What do NOE correlations tell us?

Many of the figures in this lecture are taken from references 1,2 and 4. The treatment mostly follows that of reference 4.

Origin of the NOE

Q: Where does the NOE come from?

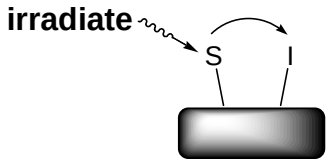

Consider two nuclei, I and S, which share a dipolar (through-space) coupling. These dipolar couplings depend on the relative orientation of I and S. In solution, molecular tumbling averages these couplings, so they do not appear in typical NMR spectra. However, magnetization can still be transfered between them.

We now consider the effect of perturbing the equilibrium populations of spin S in some way on the intensity of the signal from spin I. The NOE is defined as:

$$n_I\{S\} = \frac{I - I_0}{I_0} \times 100\%$$

For simplicity, let us assume that I and S do not share a scalar (*J*) coupling. The simplest NOE experiment is the **steady state** experiment. One *selectively* saturates spin S, and then applies a 90° pulse to observe the effect this has on the spin population on spin I:

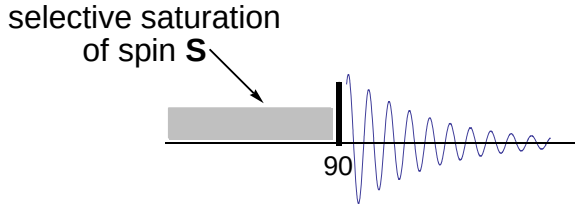

Remember, saturation means that the population of the α and β energy levels on spin S are equalized. In this case, the amount of time the saturation is applied for is the mixing time for the experiment; in a general sense, longer mixing times allow more time for magnetization to be transfered.

At equilibrium, the population *differences* between the energy levels is determined by the Boltzmann distribution. Let's call the Δ the population difference for spin I: roughly speaking, the number of excess nuclei in the lower energy α state. Because chemical shifts are much smaller than the Larmor frequency itself, Δ is also the population difference for spin S (assuming, as we are here, that I and S are both protons).

For 4N nuclei, the energy diagram is:

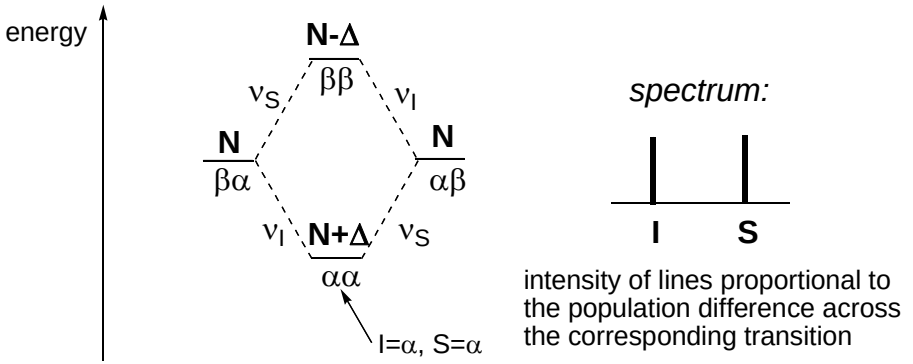

(Of course, there is no spectrum at equilibrium; one only sees this spectrum if an observation pulse is applied.) We see that the population differences across both the I and S transitions are both D. Now, upon saturation of S, we have:

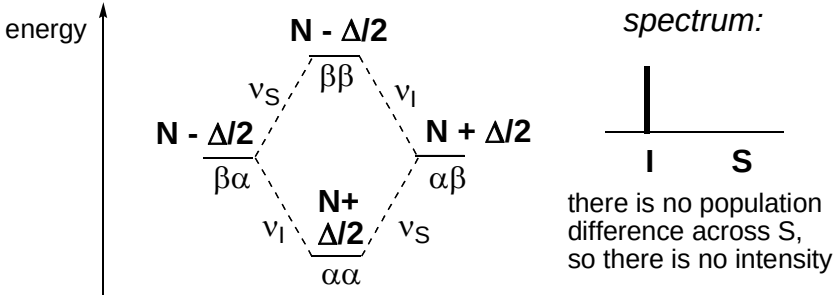

To get these numbers, note that the total number of nuclei is conserved. The population difference across I must remain Δ, but the population across S must be zero. Energy has been applied to the system, so now we must examine how relaxation sends it back to the surroundings.

Origin of the NOE

So far, we have:

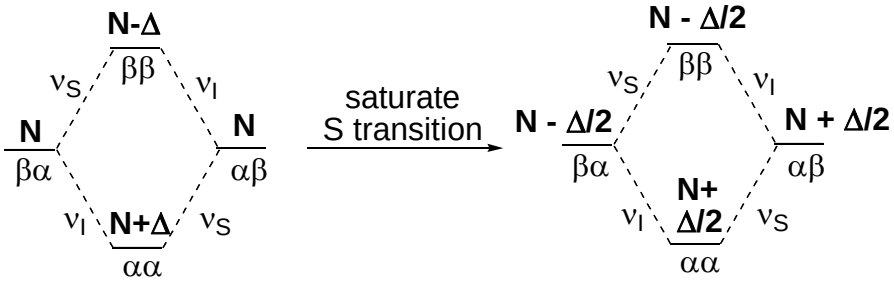

Let us examine the populations of αα and ββ. What were the populations at equilibrium, and what are they now?

	equilibrium	after saturation	
ββ	N-Δ	N - Δ/2	too much
αα	N+Δ	N + Δ/2	too little

(Subtracting a little less than before means the quantity is a little bigger than before.) Similarly,

	equilibrium	after saturation	
βα	N	N - Δ/2	too little
αβ	N	N + Δ/2	too much

Thus, relaxation must remove spins from ββ and put them into αα and remove spins from αβ and put them into βα.

We must now consider the relaxation processes themselves. In a two-spin system like this, there are six possible transitions. Every transition is labeled by the net change in spin quantum number, ΔM. Single-quantum transitions (ΔM=1) are "allowed" and give rise to observable NMR signals, while zero- (ΔM=0) and double-quantum (ΔM=2) transitions are forbidden and do not give any signal. Paradoxically, zero- and double-quantum transitions still occur; they're just invisible.

On the energy diagram, we have:

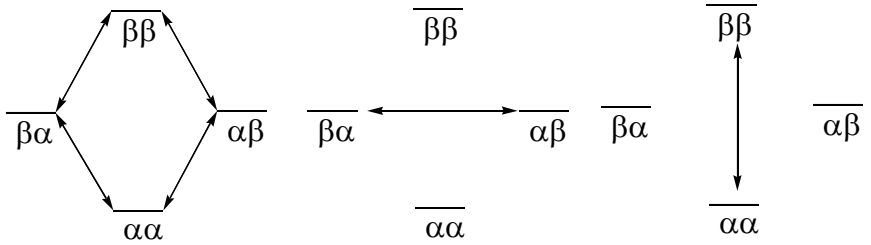

single-quantum transitions (W_1) **zero-quantum transitions (W_0)** **double quantum transitions (W_2)**

The W labels are the "transition probabilities" which can be interpreted as the rate of the transitions. The frequencies of the W_1 transitions correspond to the Larmor frequencies of the two nuclei. The frequency of the W_0 transition is the *difference* in the chemical shifts of the two nuclei and is therefore quite small. The frequency of the W_2 transition is the *sum* of the chemical shifts of the two nuclei and is therefore large and approximately twice the Larmor frequency.

The W_1 transitions will play no role in the NOE. The W_1^I transition will not cause a change in the poplations, since the population difference across βα/αα and ββ/αβ is already Δ, its equilibrium value. The W_I^S transition will simply reduce the saturation of the S transition, reducing the magnitude of the NOE, but not changing its direction.

Now, **if W_2 relaxation is dominant, then a positive NOE will be observed.** Conversely, if **W_0 relaxation is dominant, then a negative NOE will be observed.**

How can one have a "negative" NOE? In the "lightning bolt" analogy, a positive NOE means that "zapping" S will give a more intense signal on I. With a *negative* NOE, irradiating S will give a *weaker* signal on I. As it turns out, the sign of the NOE is determined by molecular tumbling rates, and therefore molecular weight.

Origin of the NOE

To see how this works, consider the first case where **W_2 relaxation is dominant**. From before, we know that this will reduce the population of the $\beta\beta$ state and increase the population of the $\alpha\alpha$ state:

	equilibrium	**after saturation**	
$\beta\beta$	N-Δ	N - Δ/2	too much
$\alpha\alpha$	N+Δ	N + Δ/2	too little

W_2

Imagine that, a short time after saturation is stopped, W_2 has moved δ nuclei from the $\beta\beta$ state to the $\alpha\alpha$ state ($\delta << \Delta$):

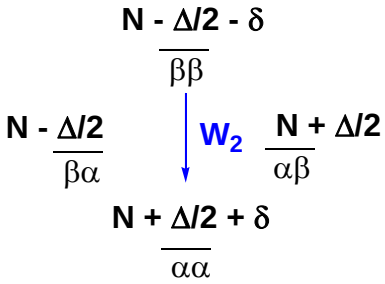

This has the simultaneous effect of increasing the population difference between $\beta\alpha$ and $\alpha\alpha$ from Δ to $\Delta+\delta$. You can see this by doing: N + Δ/2 + δ - (N - Δ/2) = $\Delta+\delta$. (The same goes for the other I transition.) Thus, the intensity of I increases and a positive NOE is observed.

The W_0 transition does the opposite:

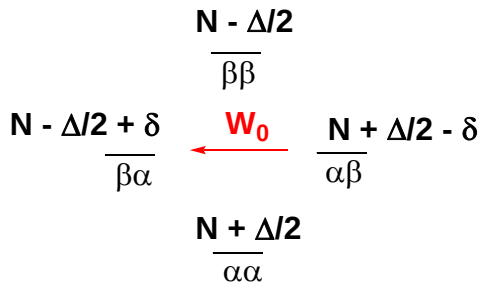

This *decreases* the population difference from Δ to Δ - δ. Thus, the signal intensity of I is reduced, and a negative NOE is observed.

Now, of course, if you wait long enough, W_I^S relaxation will eventually return everything to its equilibrium value and the NOE, whether positive or negative, will disappear. In summary:

W_2 relaxation - makes NOE positive
W_0 relaxation - makes NOE negative
W_1 relaxation - doesn't change sign, but reduces the magnitude of the NOE

In the steady-state NOE experiment, saturation can be thought of as essentially continuous. The relative rates of the various Ws changes as the experiment starts, and then reaches a steady state value. The **Solomon equation** describes the sign and magnitude of the observed NOE:

$$n_I\{S\} = \frac{\gamma_I}{\gamma_S}\left[\frac{W_2 - W_0}{W_0 + 2W_1^I + W_2}\right]$$

σ_{IS} is the "cross-relaxation" rate while ρ_{IS} is the "dipolar longitudinal relaxation" rate. Note that the ratio of the gyromagnetic ratios is included to account for the possibility of heteronuclear NOEs (e.g., broadband decoupling in 1D carbon-13 spectra).

Q: What determines the relative magnitudes of W_2, W_0, and W_1?

As I mentioned in earlier lectures, the principal mechanism of relaxation in organic molecules is dipole-dipole relaxation: the presence of fluctuating local magnetic fields induced by the tumbling of nuclei around each other. For relaxation to occur, the frequency of the tumbling must match the Larmor frequency of the nucleus.

Correlation Time

Q: What's the relationship between the rate of tumbling and the rate of relaxation?

In gases, molecules are essentially isolated. They spin around very quickly at rates near the Larmor frequency, which is what is required for relaxation. In a liquid, molecules can't tumble very far until they collide with another molecule. The collision of molecule A with B ends up rotating B a random amount which will generally be small, but could be large.

The **random walk** is typically used to describe this scenario. We imagine this is one dimension: the molecule starts at 0. Every Δt seconds, it is buffeted by a random amount. In the simulation shown below, 300 time steps are taken in total. With every step, the position of the molecule can increase or decrease by up to 0.04 (graph a) or 0.08 (graph b) units. The probability distribution is uniform--all the possible changes within the allowed range are equally likely. Every run is different, since this is a random process, but there are bigger changes for the runs in graph b.

graph a: maximum allowed change, 0.04

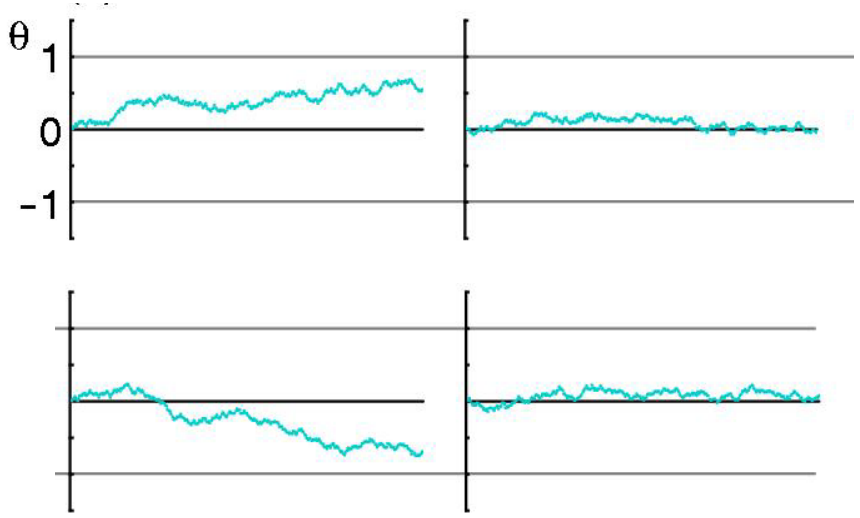

graph b: maximum allowed change, 0.08

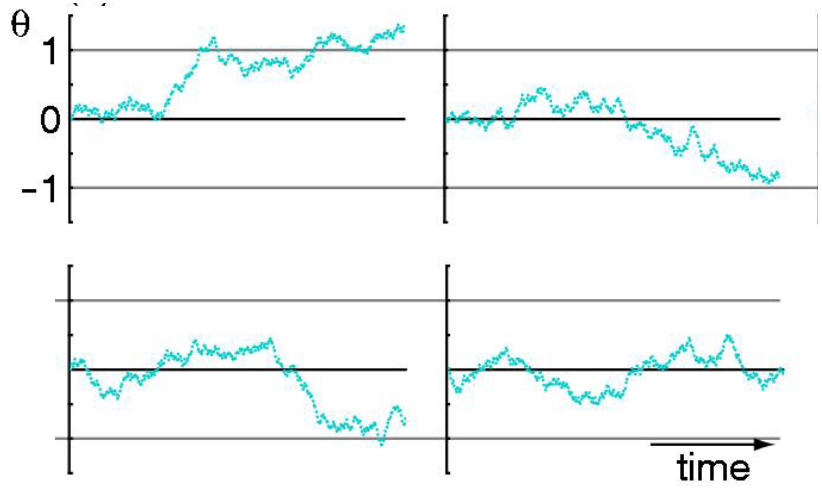

If we make the unit of change radians, then we can describe tumbling: **rotational diffusion**. The **correlation time** is the average amount of time a molecule takes to tumble through one radian. Bigger molecules tumble slower and have longer correlation times.

How does this relate to relaxation? The local field experienced by every spin will vary in time. How much does this local field fluctuate? The **autocorrelation** for a spin *i* is defined as:

$$A_i(t,\tau) = B^i_{local}(t)\,B^i_{local}(t+\tau)$$

For example, imagine that the local field for spin *i* is fluctuating such that every three seconds, the local field is the same, say, "10." Then, A(*t*,3) = 100 -- big. But if the field fluctuates in a more or less random way every three seconds, then if we average A(*t*,3) over all the times *t*, it will tend towards zero.

Overall, the ensemble average over all the *N* nuclei is:

$$G(t,\tau) = \frac{1}{N}\sum_i B^i_{local}(t)\,B^i_{local}(t+\tau)$$

Correlation Time

If this seems a bit too abstract, consider the picture given below (Keeler, figure 9.8). In (a), t = τ = 0. Of course, with a lag of 0, everything is perfectly correlated with itself. Therefore, the ensemble averaged autocorrelation is at a maximum here. The actual total value isn't shown, but you would get it by summing up all the values in the bottom left graph. Note that the position of the dots on the x-axis is meaningless--this just spreads out the values so we can look at them. In (b), t < τ_c, the correlation time. This means that the dots have not moved very much, and the autocorrelation is still big (but a little smaller than it was for τ=0). In (c) t >> τ_c, meaning that the dots have moved a lot. In effect, the dots don't remember where they were a long time ago. So the autocorrelation is much smaller, as evidenced by the many negative dots in the bottom right graph.

(a) t =0 **(b) t << τ_c** **(c) t >> τ_c**

local fields at time 0

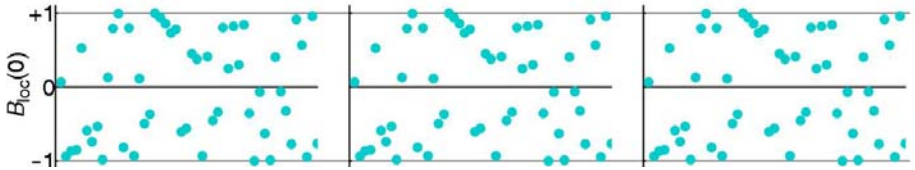

local fields at a later time τ

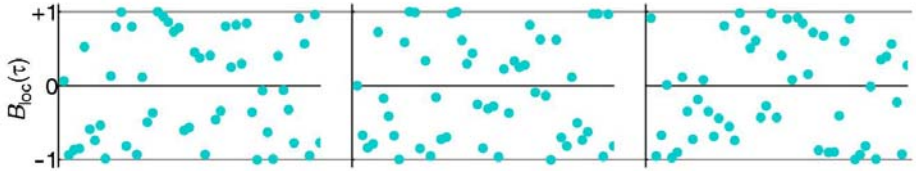

the autocorrelation

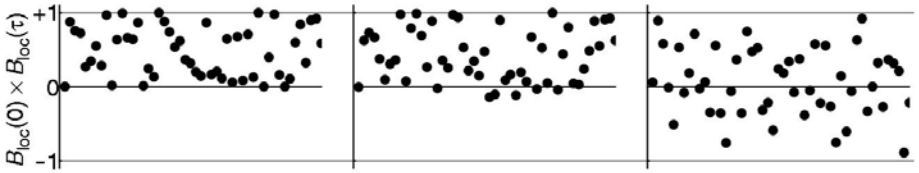

If all this talk of autocorrelation and periodicity gets you thinking that this is a lot like the Fourier transform, you're right! As it turns out, the **Weiner-Khinchin theorem** tells us that the autocorrelation and power spectrum are conjugates--the Fourier transform of one gives the other. (Recall that the Fourier transform gives a result that is, in general, complex. The power spectrum is the norm of the transform. This amounts to squaring the real and imaginary parts for every frequency and adding them up.)

$$G(\tau) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} g(\omega)\, g^*(\omega)\, e^{i\omega\tau}\, dt$$

g(ω) is the Fourier transform of the autocorrelation G(w). The asterisk denotes the complex conjugate. The power spectrum is $g(\omega)g^*(\omega)=|g(\omega)|^2$. (I have dropped the variable *t* from G to indicate that the autocorrelation is being averaged over all time. For an *ergodic* or *stationary* process, it doesn't matter what time interval you choose--the autocorrelation remains the same.)

What does this all mean? The power spectrum tells us *how much energy is available at every frequency*. So if we can determine the correlation function, we can determine the **spectral density**. The amount of energy available at the Larmor frequency will determine the size of W_1. Similarly, we can work out how big W_0 and W_2 are.

As it turns out, if you assume the molecules are hard spheres floating about in a solvent medium of some fixed viscosity, you find that (the bar means ensemble average):

$$G(\tau) = \overline{B_{local}^2} \exp\left(-|\tau| / \tau_c\right)$$

The spectral density/power spectrum is given by:

$$J(\omega) = \overline{B_{local}^2} \frac{2\tau_c}{1+\omega^2\tau_c^2}$$

Note that the total energy (integral of J over all frequencies) is always fixed (specifically, πB^2_{local}), regardless of τ_c.

Correlation Time

What does this all mean? The plot below shows the amount of energy available at every frequency for molecules that are tumbling slowly (large MW) and quickly (small MW). (I set the field strength to 1 arbitrarily):

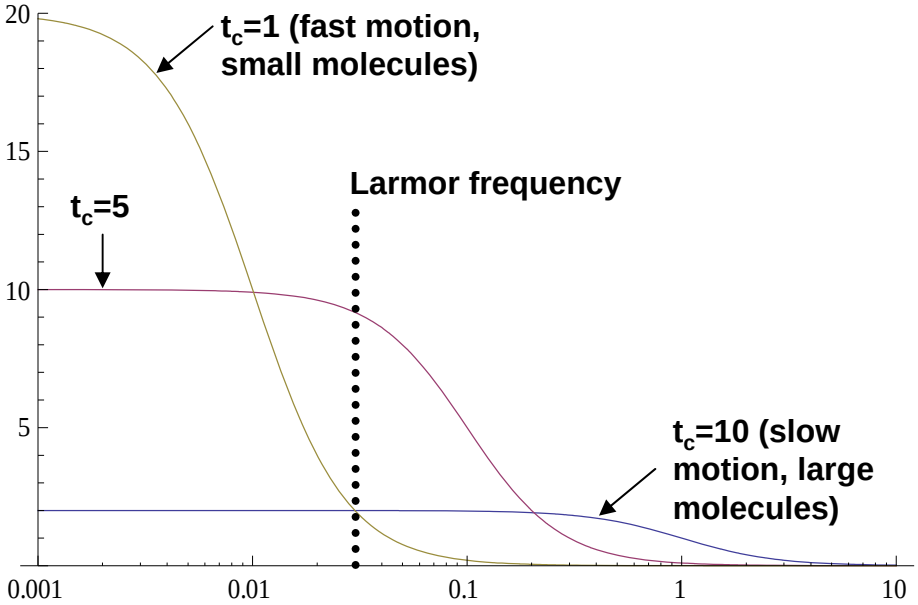

Since the field strength is normalized to 1 for each plot, the area under each curve (the total energy available) is the same. The vertical line represents the Larmor frequency. So this means that intermediate-sized molecules have the greatest amount of energy available for relaxation at the Larmor frequency, and therefore, the shortest relaxation times and fastest W_1 rates. Similarly, we can see that large molecules have a relatively small amount of energy at the Larmor frequency and relax relatively slowly.

Expressions have been derived that relate the rate constants to the correlation time and internuclear distance. Note that each cross-relaxation term depends on r^{-6}, which is why the NOE is useful. The magnitude of the NOE depends very strongly on distance, so we can use it figure out which protons are near each other.

$$W_I^1 \propto \gamma_I^2\gamma_S^2\left[\frac{3\tau_c}{r^6\left(1+\omega_I^2\tau_c^2\right)}\right]$$

$$W_0 \propto \gamma_I^2\gamma_S^2\left[\frac{2\tau_c}{r^6\left(1+\left(\omega_I-\omega_S\right)^2\tau_c^2\right)}\right]$$

$$W_2 \propto \gamma_I^2\gamma_S^2\left[\frac{12\tau_c}{r^6\left(1+\left(\omega_I+\omega_S\right)^2\tau_c^2\right)}\right]$$

The parctical consequence of this is that for any particular spectrometer frequency ω_0, there are certain correlation times that will give weak or **no observable NOE**. Here is a chart showing the theoretical NOE in a proton-proton system. For small molecules, it is +50%; for large molecules, it is -100%. Generally speaking, this "crossover region" occurs for molecules that weigh 750-2000 Da, although the specific number depends on the exact molecule.

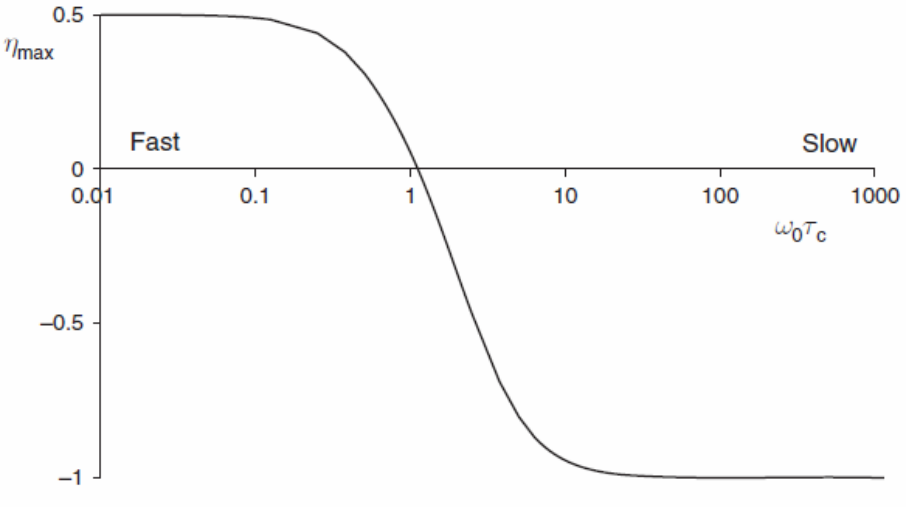

Correlation Time

Since the crossover region depends on spectrometer frequency, it is possible to change the spectrometer frequency to increase the magnitude of the NOE. This graph shows how the theoretical NOE depends on the spectrometer frequency (MHz):

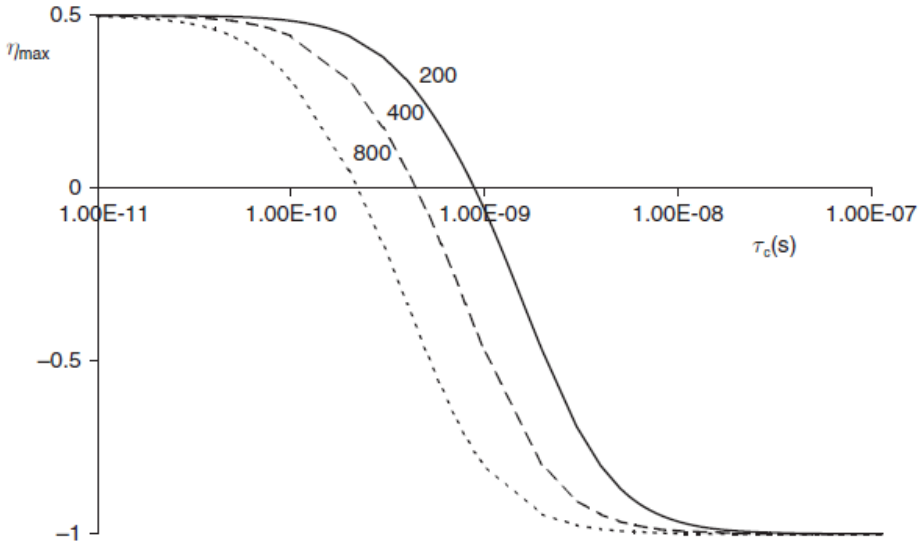

If $\omega_0\tau_c << 1$, which is true for small, rapidly tumbling molecules, then we are in the **extreme narrowing limit.** Under these conditions, the equations given above simplify:

$$\eta_I\{S\} = \left[\frac{\left(\frac{12\tau_c}{r^{-6}}\right) - \left(\frac{2\tau_c}{r^{-6}}\right)}{\left(\frac{2\tau_c}{r^{-6}}\right) + 2\left(\frac{3\tau_c}{r^{-6}}\right) + \left(\frac{12\tau_c}{r^{-6}}\right)}\right] = \left[\frac{12-2}{2+6+12}\right] = \frac{1}{2}$$

This means that regardless of the internuclear separation of the two nuclei in this two spin system, the theoretical NOE is 50% and *independent* of distance! Clearly, this is unrealistic--the problem is that in real molecules, there are other sources of relaxation. However, it seems clear that **the magnitude of the steady state NOE does not relate to distance in any simple way.**

Multispin Systems

So what is the effect of other relaxation pathways? These introduce a new source of W_1 relaxation, which tend to decrease the magnitude of the steady state NOE. Here is a graph showing the expected NOE with (solid line) and without (dashed line) external relaxation:

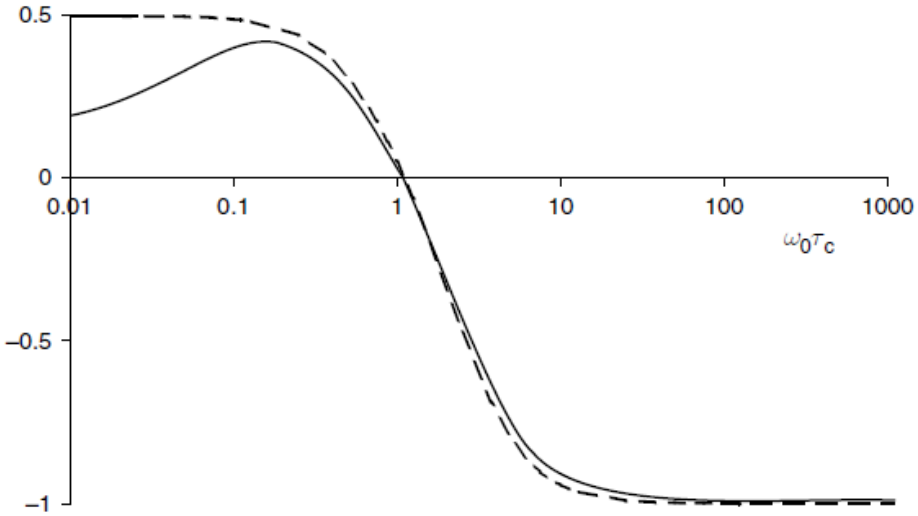

These external pathways could be dissolved oxygen (para-magnetic), paramagnetic ions, quadrupolar nuclei, the solvent (although deuteration means this is less important), etc. The graph shows that this is most serious for small molecules, where σ_{IS} is relatively small compread to ρ_I*.

However, we are ignoring a major source of undesired dipolar relaxation: other spins! Here, N is a neighboring spin:

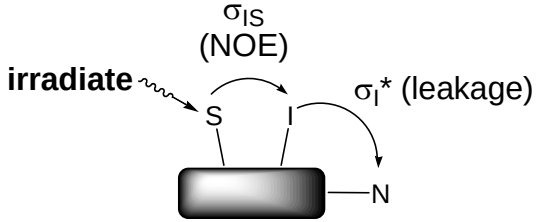

Multispin Systems

Here's where you can get a negative NOE, even if you have a small molecule with a short correlation time. Remember, saturation decreases the population difference across spin S, leading to an increase in the population difference across spin I, leading to a positive NOE. If a neighboring spin N is near spin I, then the increase in the population difference across N will lead, through the same reasoning, to a negative NOE on spin N.

Here, are some figures from Claridge to show you what's going on. Here, the numbers between the bars indicate the inter-nuclear distances in angstroms. If we irradiate B, then regardless of how close spin C is to B, we still get roughly a 50% NOE at A (although if C is closer, this is reduced a bit). This analysis assumes there's no dipolar coupling between A and C:

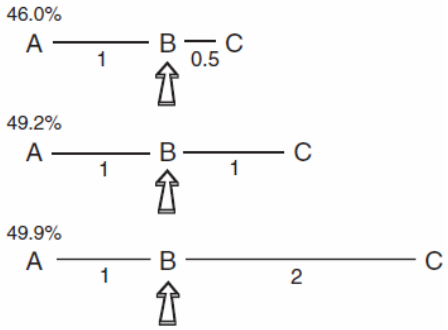

But when A is irradiated, you get a positive NOE at B, but a negative NOE at C:

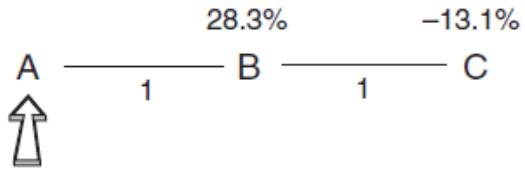

With a four spin system, things get more complicated:

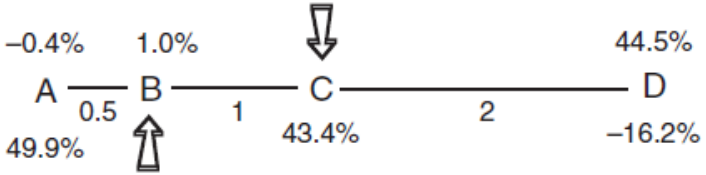

So negative NOEs can be a sign of these so-called relay effects. In a real molecule, there are a lot of spins, so over the steady state saturation period, one has a complicated series of "spin diffusions":

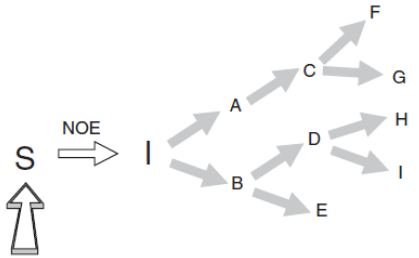

Negative NOEs can also arise from saturation transfer/chemical exchange. You can think of chemical exchange as a shuttle that can take saturated nuclei and put them at other sites, even though they haven't been selectively irradiated.

Transient NOEs

Q: If steady state NOEs are not related to internuclear distance in any simple way, then what good is the NOE?

The key here is that, in the buildup to the steady state, NOEs between nuclei that are far apart take longer to develop than NOEs between nuclei that are closer together. This gives rise to the idea of *transient* NOE experiments. Here, what we do is selectively invert spin S, wait a mixing time τ_m, and then observe the longitudinal populations:

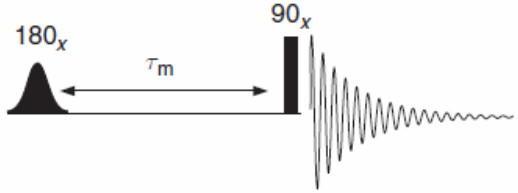

In the old days, people used steady state NOE experiments in the form of "NOE difference" experiments. Nowadays, people use transient NOE experiments in the form of 1D-NOESY and 2D-NOESY. I'll consider them in detail in a bit.

The key question here now: **how does the transient NOE intensity depend on internuclear separation?**

Transient NOEs

The clever thing is that the NOE starts growing linearly (the *initial rate approximation*), reaches a maximum, and then decreases back to equilibrium as the equilibrium population differences are restored:

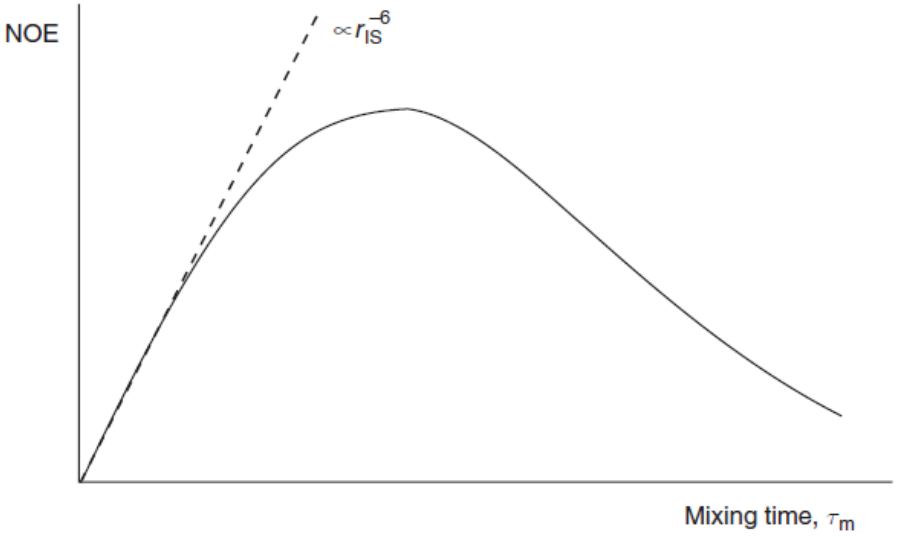

As it turns out (there's factor of 2 because the population of S is initially inverted, rather than equalized):

$$\frac{dI_z}{dt} = 2\sigma_{IS} S_z^0; \quad \sigma_{IS} \propto \gamma^4 \left\{ \frac{6}{1+4\omega_0^2\tau_c^2} - 1 \right\} \frac{\tau_c}{r_{IS}^6}$$

Note that the maximum homonuclear enhancement is now 38%, rather than 50% as in a steady state experiment. (For larger molecules, it is still -100%.) So if you want to measure internuclear separations, you can measure two transient NOEs in the initial linear growth region. If the internuclear distance is known in one case, then the distance for the other case can be inferred.

Here it is mathematically. In the linear growth region:

$$n_A\{B\} = k\sigma_{AB}\tau_m \propto r_{AB}^{-6}\tau_m$$

This means that we can consider the *ratio* of the NOEs for two pairs A-B and X-Y:

$$\frac{n_A\{B\}}{n_X\{Y\}} = \frac{r_{AB}^{-6}}{r_{XY}^{-6}}$$

Now, in practice, this is more complicated. Molecules often have more than one conformation, and the NOE will see an averaged structure which more heavily weights the conformers that have small internuclear separations (because of the heavy r^{-6} dependence). Typically, people use this for proteins, rather than small molecules, although in principle the ideas are the same. In a protein, one can measure a large number of NOEs to generate "distance constraints" on the various internuclear separations in the molecule. Computer modeling can then provide a picture of the solution state conformation of the molecule.

In conclusion, one can take NOEs and turn them into a quantitative measure of solution state distance by measuring a series of NOEs with different mixing times. The initial rates correspond to the distances:

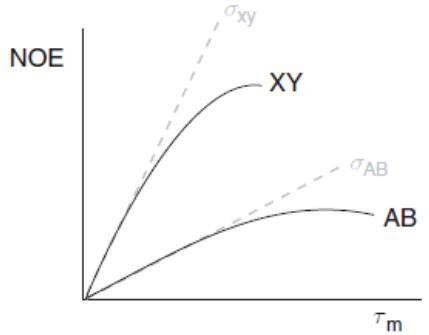

Q: What about molecules in the crossover region?

I already mentioned that you can adjust the spectrometer frequency, but in many cases, this is not a satisfactor solution. As it turns out, one can perform a different kind of NOE called the **rotating frame NOE (ROESY).**

Rotating Frame NOE: ROESY

The idea is that this experiment has a different kind of intrinsic cross-relaxation rate which is given by:

$$\sigma_{IS} \propto \gamma^4 \left\{ \frac{3}{1+\omega_0^2\tau_c^2} + 2 \right\} \frac{\tau_c}{r_{IS}^6}$$

For all realistic values of correlation time, this remains positive. Hence, the ROE is always positive. It has a maximum of +68% for large molecules, but it never crosses zero:

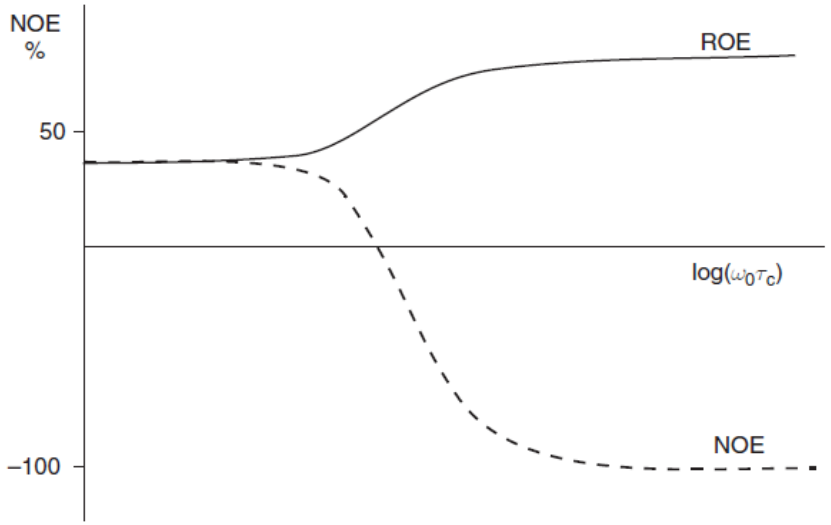

The growth rate of the NOE and ROE are similar for small molecules, but the ROE grows twice as fast as the NOE for large molecules.

How does it work? The pulse sequence involves selective inversion of the target peak, and then rotation of all resonances into the xy plane:

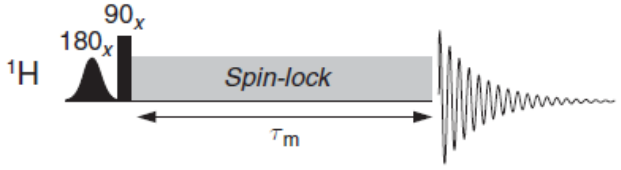

The experiment is analogous to the NOE. In the NOE, magnetization is transfered from an inverted resonance along -z to upright resonances along +z (thus, one doesn't have to worry about chemical shift offsets). Magnetization is under the influence of the large applied B_0 field.

In the ROE, magnetization is transfered from a source along -y to other resonances along +y. Magnetization is under the influence of the relatively small spin-locking field B_1, which serves to refocus chemical shifts continually. Because $B_1 << B_0$, the relevant "Larmor frequency" is much smaller, and therefore, there is a lot of energy at this smaller frequency, regardless of how large the molecules are.

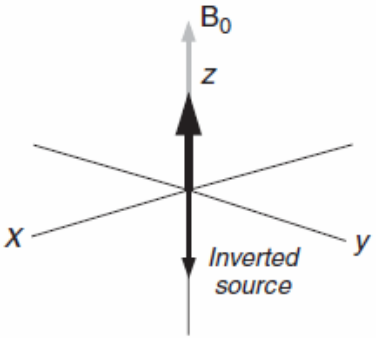

Transient NOE experiment

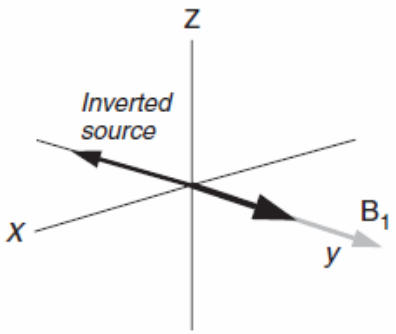

Rotating-frame NOE experiment

Artifacts and Peak Phase

Like many other NMR experiments, NOESY and ROESY spectra suffer from artifacts, which must be recognized if a spectrum is to be interpreted correctly. In general, whether 1D or 2D experiments are being run, experiments are presented in **phase-sensitive mode**. By convention, diagonal peaks are always phased up.

For NOESY:

diagonal: up
positive NOE: down **chemical exchange: up**
negative NOE: up **COSY: antiphase**

Artifacts and Peak Phase

As it turns out, the COSY and NOESY (and EXSY) pulse sequences are basically the same. In COSY, we are interested in magnetization transfer through bonds; in NOESY, we are interested in magnetization transfer through space. COSY correlations arise from zero-quantum coherence, which we will discuss in detail later in the course. These *cannot* be removed by gradients or phase-cycling, but some special methods like "z-filtration" have been developed to get rid of them. (I'll tell you about that later, too.) The point is that COSY crosspeaks are easy to identify because they have an "up-down" phase. They are particularly common when *J* is large; for example, between two trans-diaxial protons in a cyclohexane.

Here's a NOESY spectrum without any removal of zero-quantum artifacts:

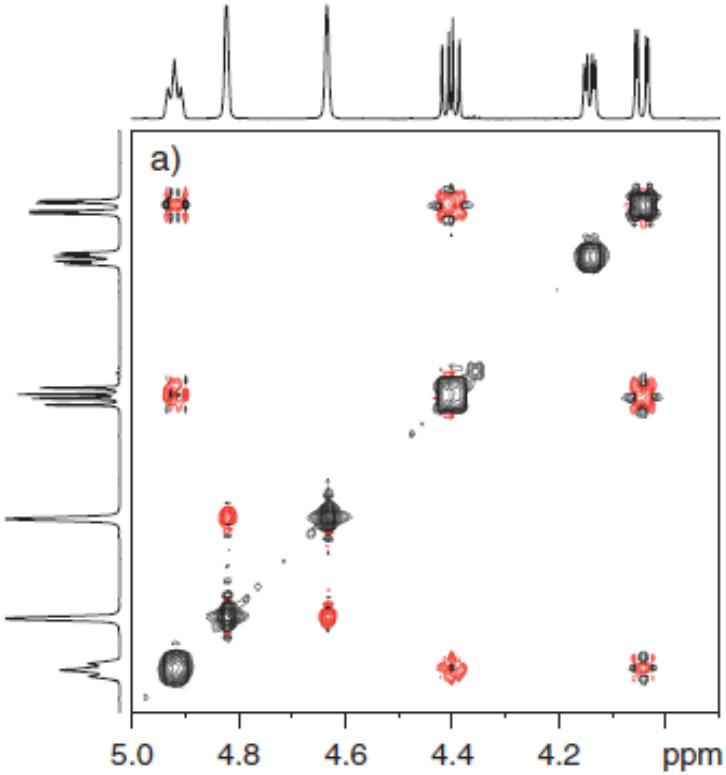

Here, a sophisticated method called "swept-frequency inversion" was used to remove the antiphase COSY signals:

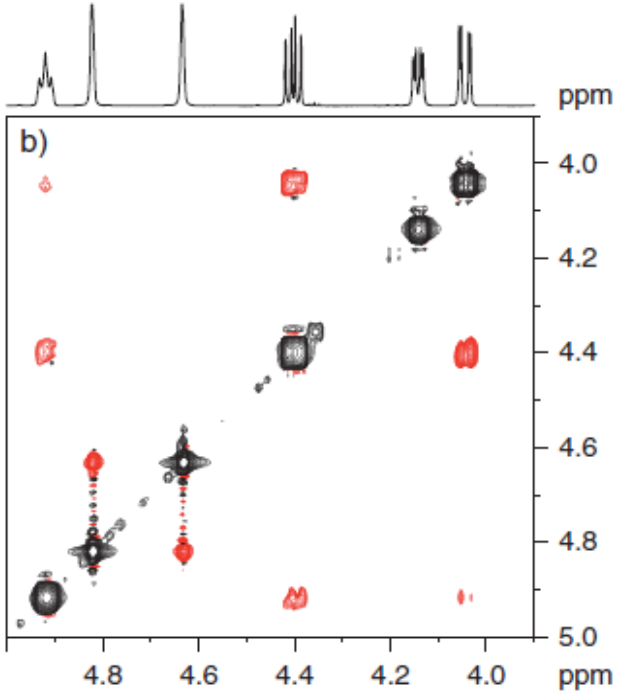

ROESY spectra aren't immune to artifacts, either. Instead of COSY artifacts, they have TOCSY artifacts. Clever people like AJ Shaka (*JACS* **1992** *114* 2157) have figured out that the "transverse ROESY" experiment, or t-ROESY can get rid of most of these annoying TOCSY crosspeaks. The price is that there is some sensitivity loss (none for small molecules, 2x for intermediate molecules, and 4x for very large molecules). It works by alternating the phase of the spin-locking train:

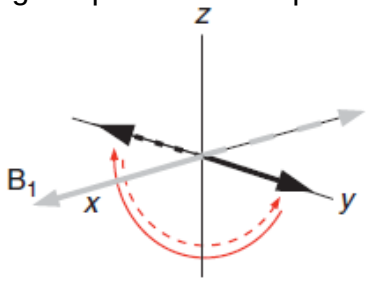

Artifacts and Peak Phase

The fact that you can have TOCSY transfer means that you can have potential confusion because of sequential TOCSY/ROE or ROE/TOCSY transfers:

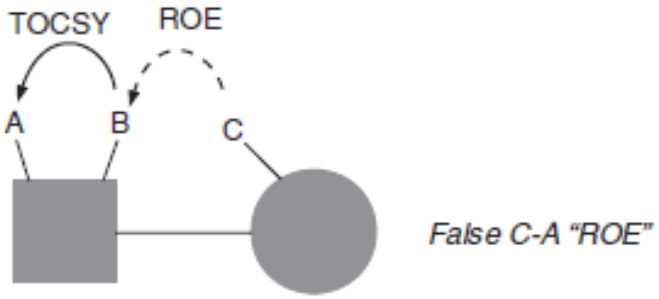

False C-A "ROE"

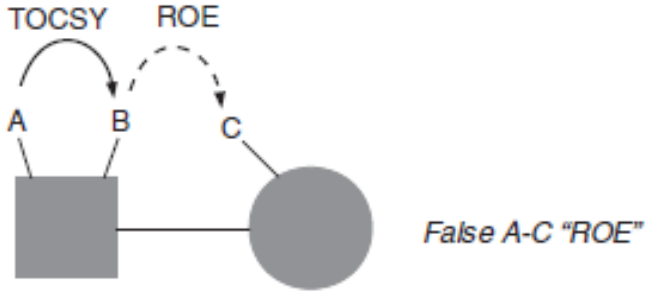

False A-C "ROE"

Thus, everybody uses t-ROESY in the crossover region, even if there's some loss in sensitivity. For small or very large molecules, NOESY is preferred. Finally, I should mention that chemical exchange peaks are also positive in ROESY spectra, so this can be a useful way to identify OH protons.

The bottom line: if your molecule has a MW between 750 and 2000 Da, and NOESY gives weak correlations, use t-ROESY instead.

You don't have to believe just my words. Here are some spectra from Professor Reynolds showing that for small molecules, NOESY and ROESY look quite similar, but for large molecules, ROESY performs a lot better. These first spectra were taken with at+d1=1.8 s, ni=256 with linear prediction to 1024, and are presented at the same scale in phase-sensitive mode.

NOESY (a small molecule)

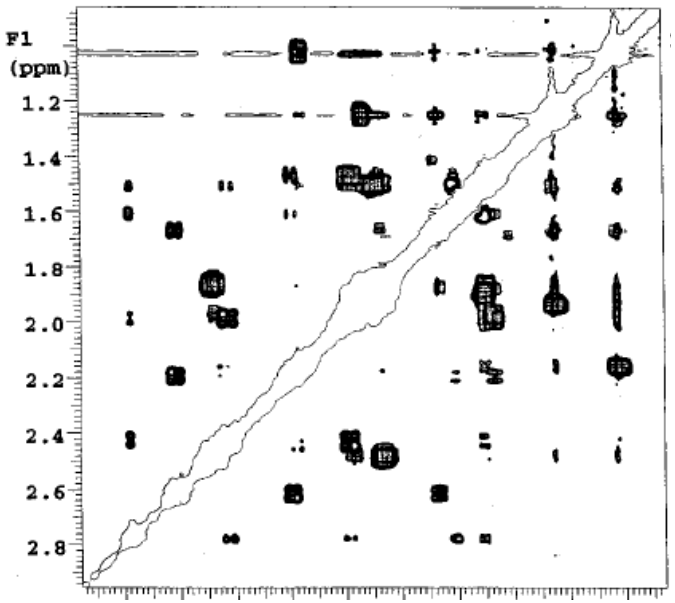

ROESY (the same molecule and conditions)

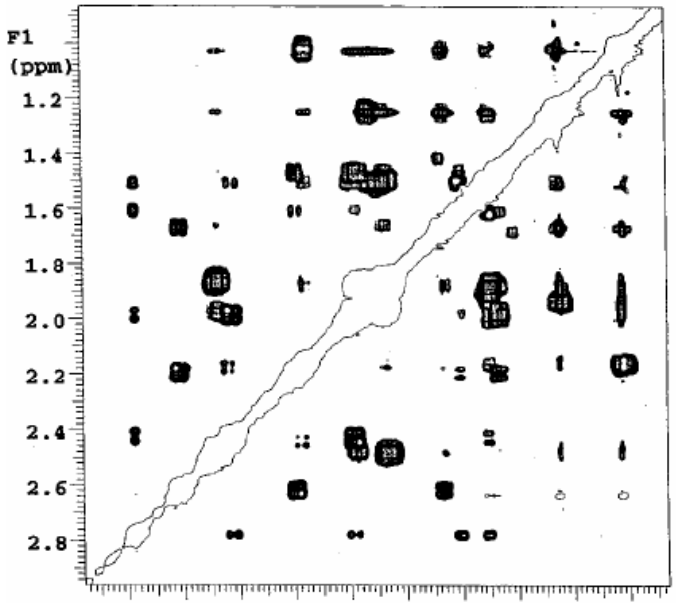

NOESY (an intermediate-sized, crossover region molecule)

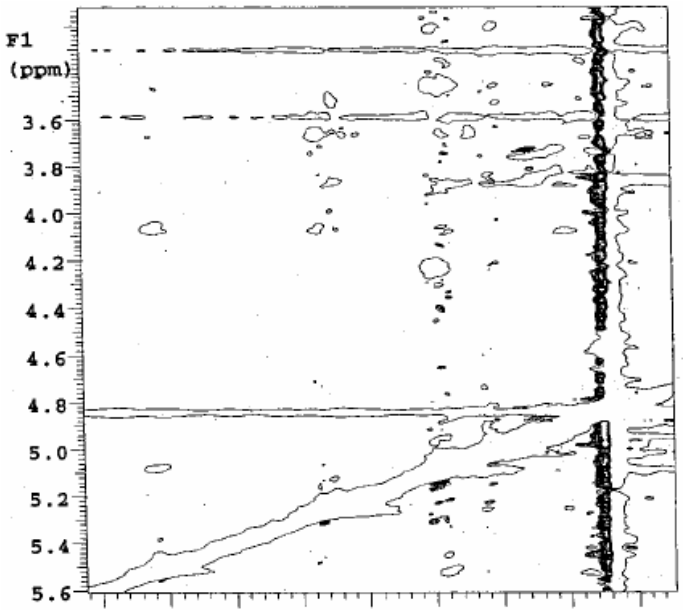

ROESY (the same molecule and conditions)

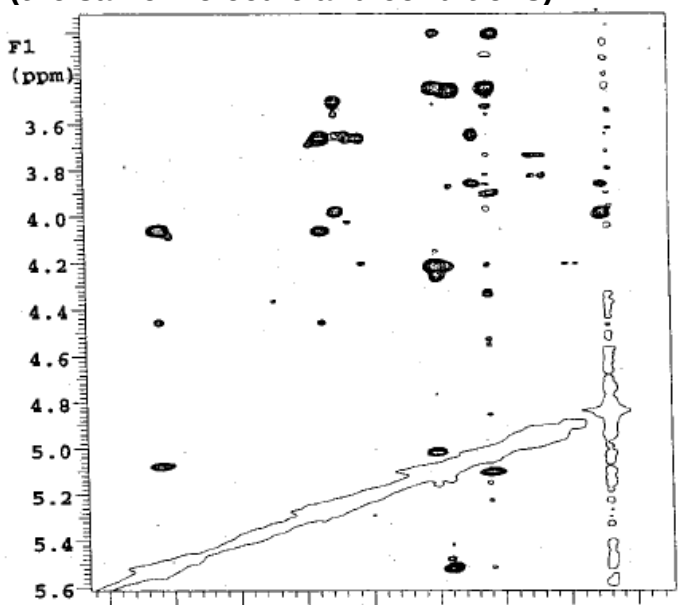

1D NOE Experiments

For historical reasons, I'm going to talk a little bit about NOE difference experiments. The idea is to take pairs of scans, one where the selective presaturation is applied on resonance, and one where it isn't. The presaturation is not just turned off in the control experiment so as to have conditions that are as similar as possible. The more non-NOE related differences there are, the more noise. The off-resonance presaturation is typically applied at the edge of the spectrum.

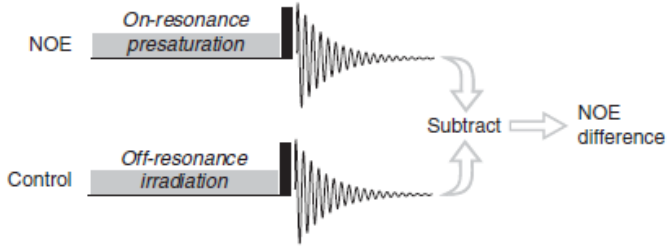

Here are a few spectra (top: NOE diff., bottom: regular 1D):

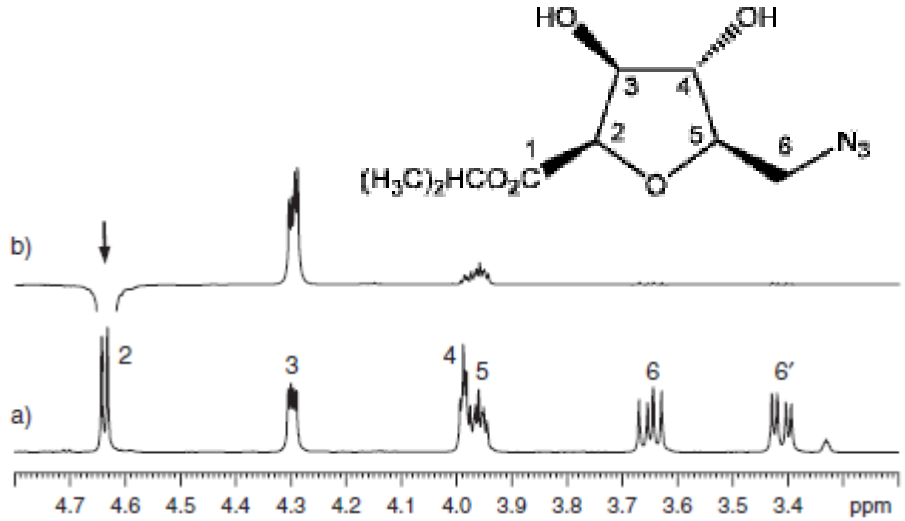

Here, it seems NOE was used to assign the 2,5-cis geometry in this five-membered ring. I don't know the specifics of this case, but in general, using NOE to assign the relative configuraiton in five-membered rings is dangerous since they are so conformationally flexible.

1D NOE Experiments

The results don't look as good when there are some instabilities in the spectrum between scans. Here, true NOEs are marked with Ns and difference artifacts are marked with asterisks:

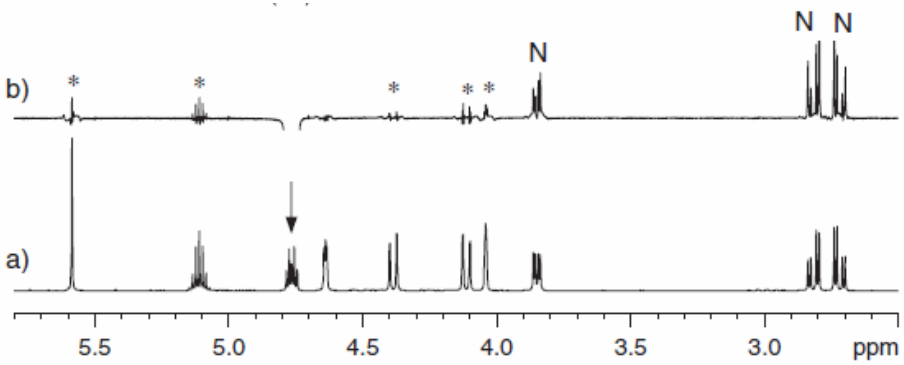

Short-term instabilities are removed by signal averaging; long-term instabilities are removed by interleaving the control and NOE scans. "Steady state" scans are also required at the beginning, since the repetition delay is usually much less than $5T_1$.

There are a lot of details about how to optimize the pre-saturation period, off-resonance effects, saturation transfer, etc. but I won't get into them. The bottom line is NOE difference experiments take a long time to do because you need to supress the non-NOE peaks, and even then, they might not be supressed very well.

Is there a better way? Why, yes: the 1D-NOESY sequence. The basic form is shown below:

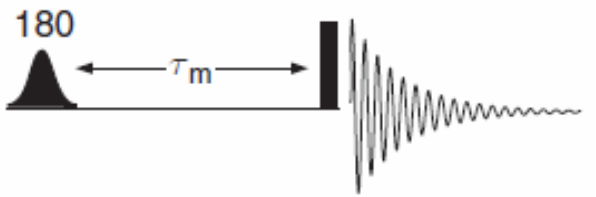

Note that this is a transient experiment, as opposed to the steady state NOE difference experiments I just mentioned. Now, it turns out this is not the experiment people use. Instead, they use the "DPFGSE-1D-NOESY" sequence:

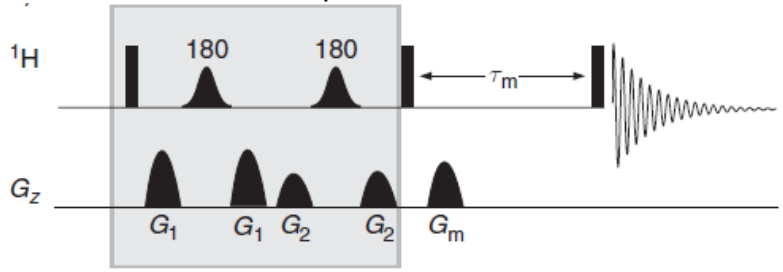

DPFGSE selective excitation

This stands for "double pulsed field gradient spin echo." Aside from being a mouthful, this is by far the most robust experiment to date, and gives absolutely beautiful results in a much shorter time period. From Professor Reynolds again:

NOE difference **DPFGSE-NOESY**

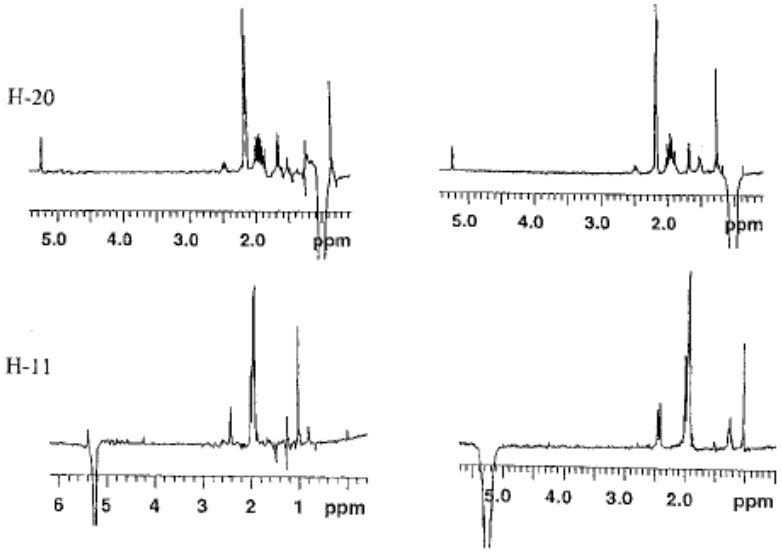

Clearly, the DPFGSE-NOESY is much cleaner (total acquisition times were held constant). So how does it work? We must understand what pulsed field gradients do...

Pulsed Field Gradients

Developed in the 1990s, PFGs revolutionized NMR. The idea is to apply a magnetic field on top of the B_0 field that is *not* uniform, but instead varies over the sample. Typically, when we talk about PFGs, we mean z-gradients:

$$B_g(z) = B_0 + G_z z$$

This says that the field now depends linearly on the z-position in the tube. Graphically:

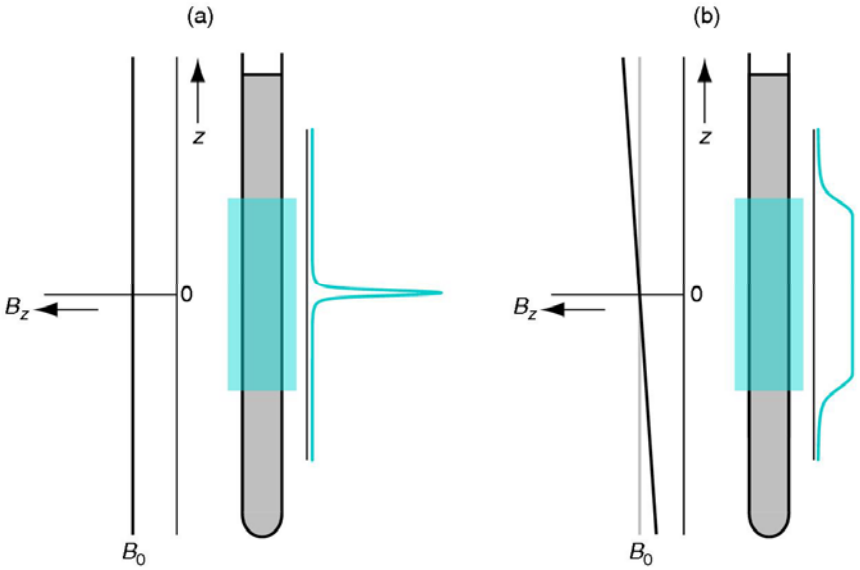

From our prior discussions of the spin echo, you should realize that this is a tremendous amount of magnetic field inhomogeneity, and consequently, T_2* dephasing is extreemly fast and no signal is observed. Were we to look at the chemical shifts of the same nucleus across different z-slices of the tube, we would find they vary linearly with the gradient.

The picture, then, is of a **coherence helix**. Imagine applying a 90 degree pulse, followed by a gradient. The vectors dephase in a helical way, depending on their z-position in the tube:

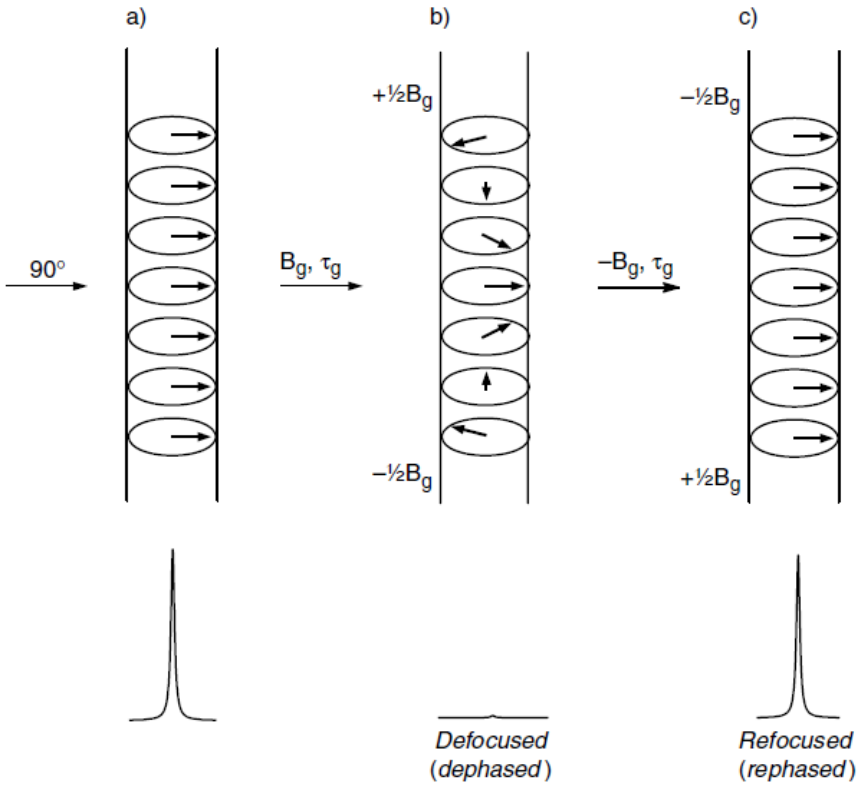

If we then apply a negative gradient of the same magnitude for the same period of time, the coherence helix gets unwound, and therefore, the peak comes back. This is a **gradient echo**.

Q: How can gradients be used for selective excitation?

Why is selective excitation advantageous? Remember, in an NOE difference experiment, the peaks that are not related to NOE enhancment still appear, but are cancelled on alternate scans. The desired signal is the difference between two signals of essentially maximum height. So this is getting the difference between two big numbers (105 - 100 = 5). In 1D-NOESY, only the irradiated peak and any NOE crosspeaks appear, so every scan essentially collects useful information. This means we see the "5" signal directly, giving better signal to noise.

Soft Pulses and PFGs

Recall from our discussion of the Fourier transform that peaks that are narrow in the time domain correspond to broad peaks in the frequency domain--this is a manifestation of the uncertainty principle. The visual representation is as follows.

Suppose I give you two waves and I ask you tell me what frequency ranges they cover:

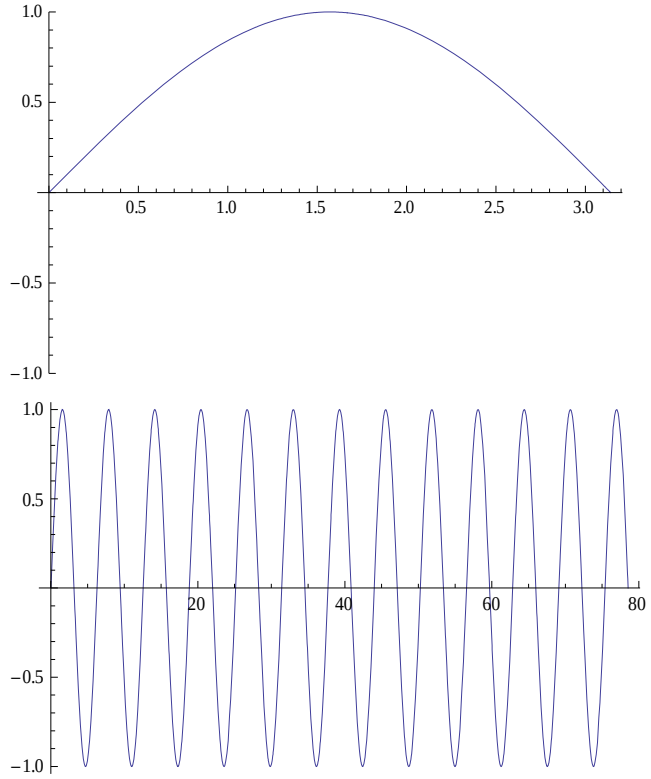

In the first one, you observe the pulse for quite a long time, and you can with pretty good certainty that it has a frequency of 2π rad/s. But in the second one, you only see one crest, so you're not really as sure what frequency it is--its representation in the frequency domain is braoder to reflect this uncertainty.

Typical pulses are rectangular, "hard" pulses that excite a range of frequencies. As it turns out, the rectangular pulse excites a range of frequencies like $\sin(\omega)/\omega$--the "sinc" function:

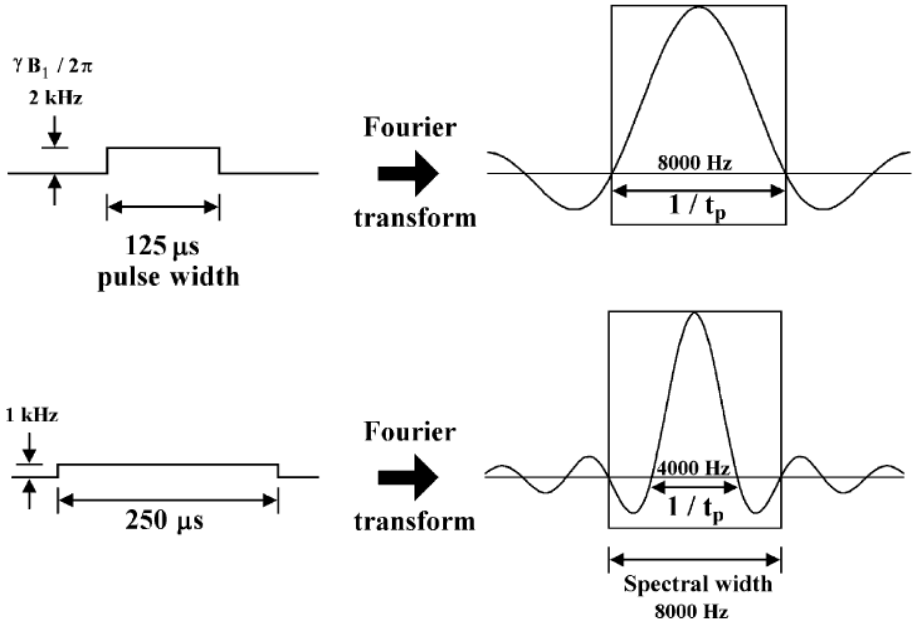

Notice that the shorter the rectangular pulse, the wider the sinc. So one way to get a selective excitation is to use a time-domain Gaussian (I'll discuss this in more detail in the next lecture):

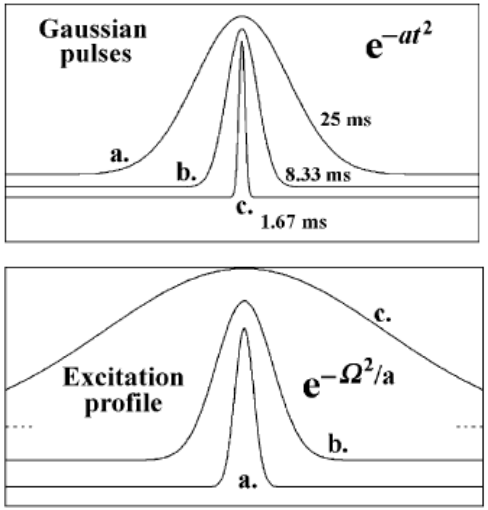

Soft Pulses and PFGs

Now, this is pretty good, but it's not perfect. Here's a spectrum:

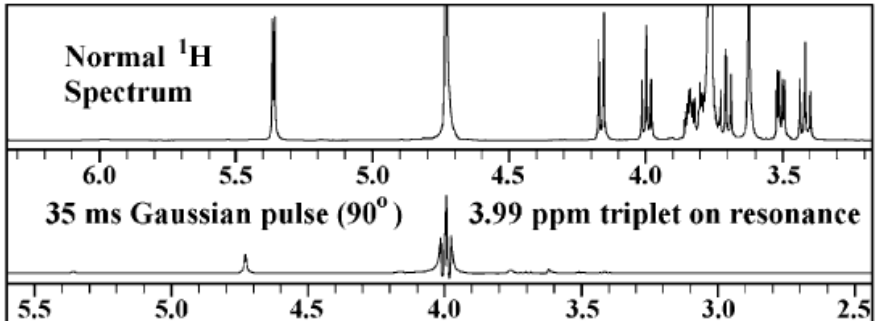

In the case of 1D-NOESY, this is just not good enough. A more poignant demonstration is to show the "excitation profile." To make one, you take a sample with one resonance, and move the center of the spectral window around (a standard Gaussian pulse excites maximally at the center of the spectral window). Typically, people use HDO for this purpose. Here are the excitation profiles for some Gaussians:

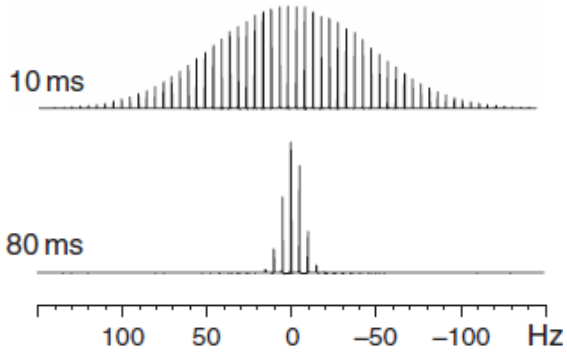

We need a way to get rid of the undesired signals. Remember, signals arise from phase coherence in the xy-plane, which can be destroyed by a gradient. The "PFGSE" procedure is:

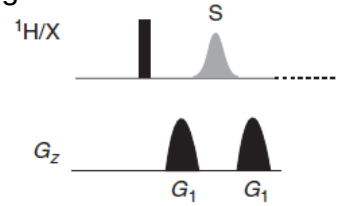

Consider the effect of this sequence on any peaks that are *not* inverted by the selective pulse in the middle. Just consider one resonance frequency. The first 90° pulse puts the vector into the xy plane. Then, a gradient is applied, creating a helix. The xy phase of every vector depends on its position in the tube, so a coherence helix is created. Applying the same gradient again simply twists the helix around twice. The net result is no signal: the phases of the vectors in the x-y plane are distributed in a uniform way; signal requires phase coherence.

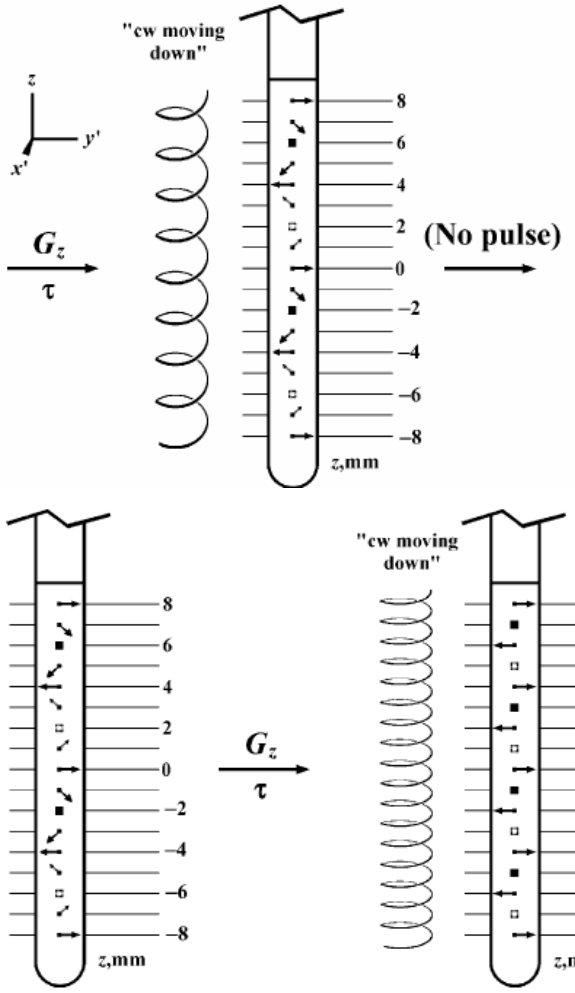

Soft Pulses and PFGs

Now, let's look at what happens to a resonance that *does* get inverted. The 180° pulse reverses the handedness of the helix such that application of the second gradient (identical to the first) causes the reversed helix to unwind. Thus, at the end, coherence returns in a gradient echo:

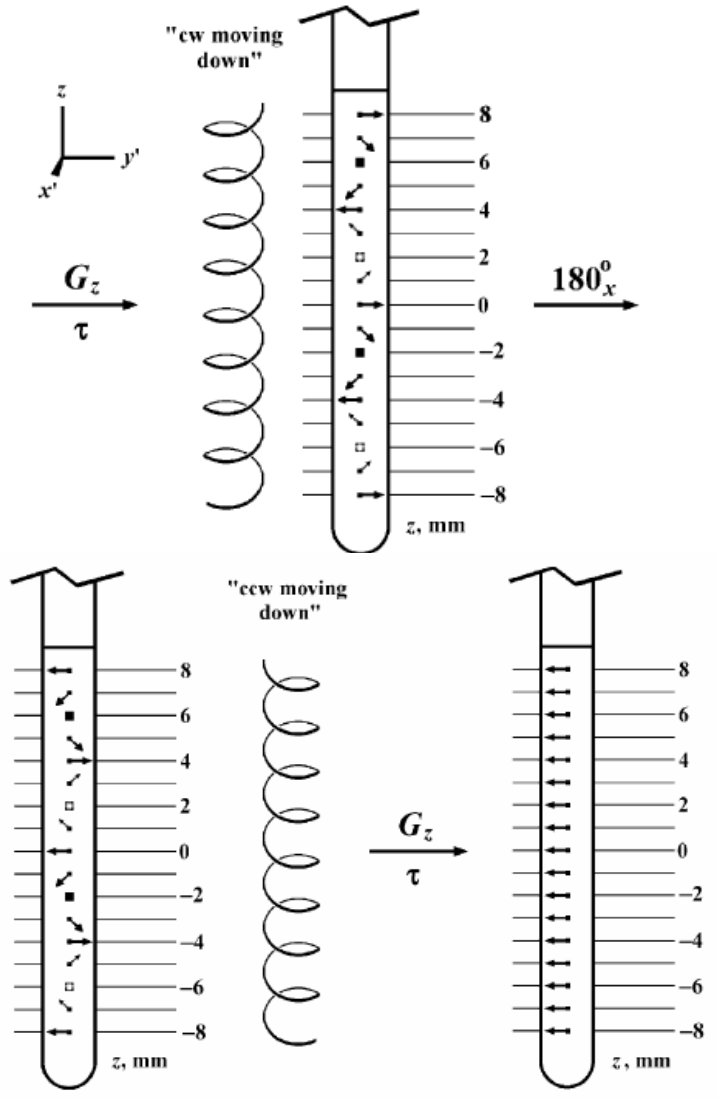

This strategy, known as **excitation sculpting** was developed by A.J. Shaka. Now, as it turns out, the *phase* properties of the selective excitation are sometimes problematic (see Lecture 13). The application of a second PFGSE will cancel these errors. The gradient strength is set to a different number, so as to avoid accidentally refocusing any of the previous dephased signals. This is the DPFGSE-1D-NOESY sequence:

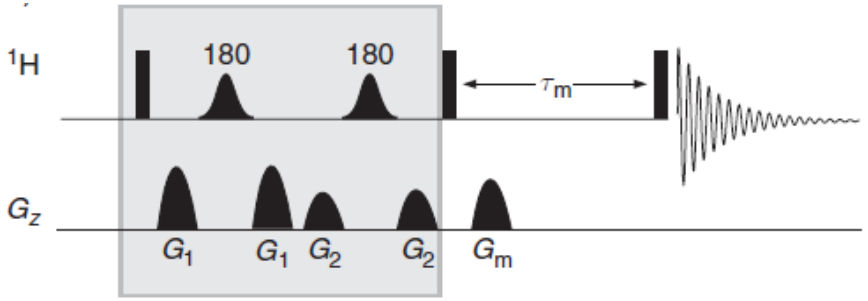

DPFGSE selective excitation

The results are very good: notice that the very strong signals are more or less removed.

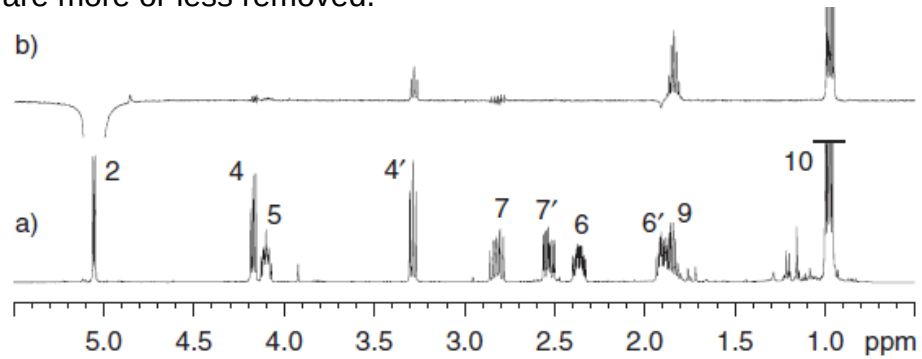

Now, of course, nothing is perfect. You can see that there are still a few "wiggles" from incomplete suppression. However, this is still a huge improvement from NOE difference.

Bottom line: 1D-NOESY is awesome--use it.

If you do this experiment in real life, you will find that you get much better results if the number of scans is a multiple of 4. Why is that?

EXORCYCLE

The usual answer to questions like that is that interrupting a phase cycle will produce incomplete cancellation. Since phase cycles need blocks of four scans to work, not using a multiple of four for the total number of scans will result in an incomplete block at the end of the experiment. This gives rise to all sorts of "weirdness" in the spectrum.

In spin-echo sequences like DPFGSE-1D-NOESY, these artifacts are referred to as "ghosts" or "phantoms." These can arise from imperfections in the 180° refocusing pulse or more insidious effects from the cooperativity of errors in both the 90° and 180° pulses. The solution is to use the EXORCYCLE phase cycle, which is the analog of CYCLOPS for spin echo sequences.

The strategy is to increment the phase of the 180° pulse and the receiver phase by 90° in each step:

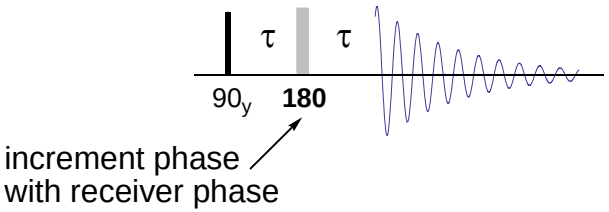

The phase cycle is:

180°: x, y, -x, -y; **receiver:** x, -x, x, -x

This gives alternately positive and negative echoes (from the perspective of someone at +x). Let's take a look at the vector diagrams. For this, I will consider two vectors of different chemical shift and place the rotating frame reference frequency at the average of their offsets. Let us assume there is no *J*-coupling. The big dot below indicates the position of the receiver (i.e., its phase).

For step 1 of the phase cycle, we have 180°: x, receiver: x. This gives a positive echo:

M_z, then $90°_y$ gives:

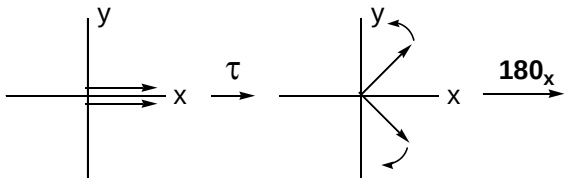

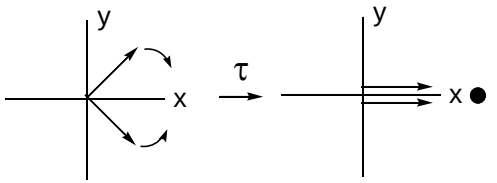

receiver and vectors are both at +x, so the echo is positive

In step 2, you get a negative echo, but since the receiver phase has been incremented as well, the signal is positive again:

M_z, then $90°_y$ gives:

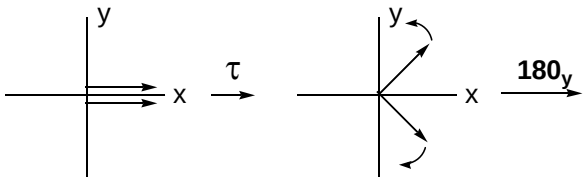

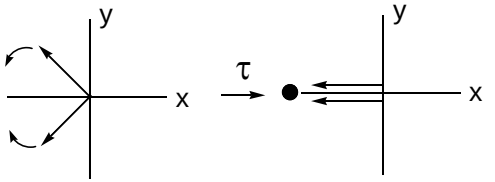

echo is negative but signal is positive

Steps 3 and 4 mirror 1 and 2, respectively. Now, any unwanted signals will be alternately inverted by the 180° pulse, so on adding up the data from the four scans, they will cancel. The desired signals will add up constructively.

Pulse imperfections are not the only source of problems that require some phase cycling. To see this, we have to consider the disposition of the other spins in the DPFGSE-1D-NOESY sequence.

Relaxation Artifacts

Here's Figure 8.27 from Jacobsen:

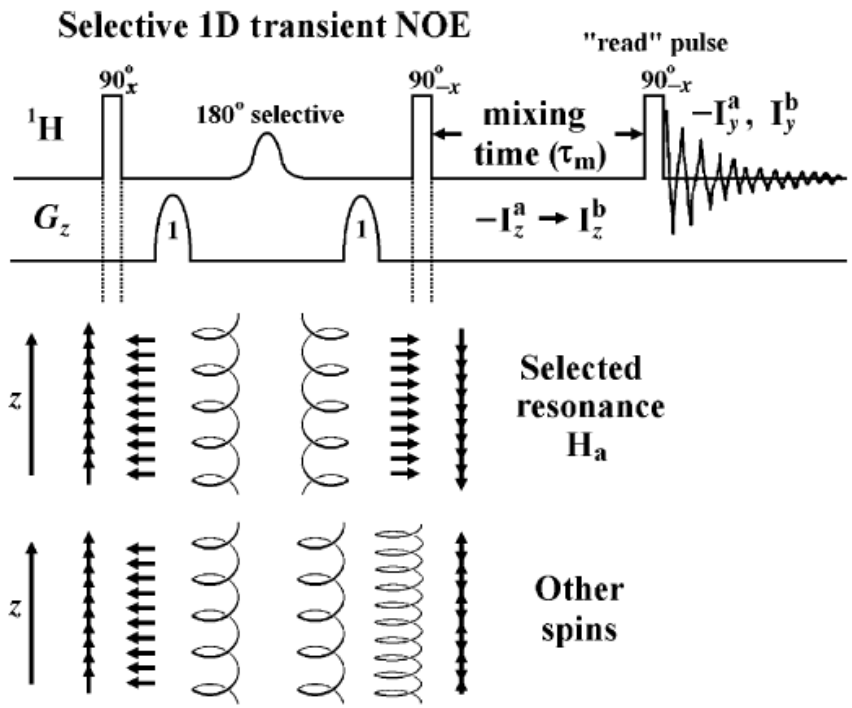

(To make life simpler here, let's just imagine there's only one spin echo.) At the end of the spin echo, the selectively refocused vectors lie along the x-axis. The subsequent 90° pulse moves the "source" nuclei to the -z axis, where they can begin to transfer coherence through the NOE and eventually come back to equilibrium.at +z.

The trouble is with the other, non-refocused spins. At the start of the mixing period, they are alternately disposed with +z and -z components. An immediate 90° read pulse will show that these have been fully dephased. However, after a finite mixing time, the -z components will begin to recover towards the +z direction, whereas the +z components, already being at equilibrium, will not. Thus, the unwanted signals will start to "grow back."

Additionally, since on average half of the undesired peaks are inverted to some degree (even though this is made invisible by the lack of phase coherence), they will start to develop NOEs as well. Thus, we can have a whole host of NOEs from the undesired signals, even though they are apparently invisible!

The solution is to alternate the phase of the middle 90° pulse--the one after the PFGSE. For the desired peak, this means that alternate scans will be control experiments. The phase will alternate between -z (which does give NOEs) and +z (which doesn't). For the undesired peaks, this means that the helix just gets flipped over, so that for both pulse phases, the same artifacts appear. This basically means that every other scan (those where the desired peaks are along +z) is a control experiment which produces only artifacts. The other half of the scans produce artifacts and NOE. By subtracting the former from the latter, we can get rid of the artifacts.

So in some sense, PFGSE-NOESY experiments are also difference experiments. However, the dynamic range of the signals being subtracted is much smaller than in the NOE difference experiment. So the experiment is still much better

Applications of the NOE

Here are some of the common applications of the NOE in organic chemistry.

Olefin Geometry:

Aromatic Substitution Patterns:

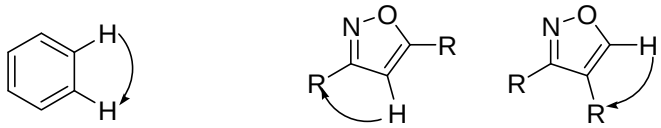

Applications of the NOE

1,3-Relationships in Cyclohexanes:

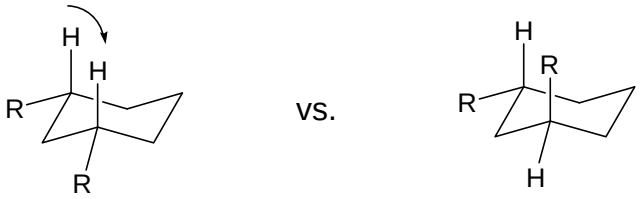

Resonance Assignment:

You can use NOESY as a complementary tool to things like COSY or HMBC. However, this may be dangerous because you don't yet know what all the peaks represent.

Endo vs. Exo (Diels-Alder):

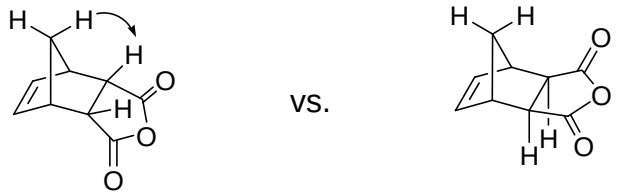

i, *i*+1 Assignments in Peptides:

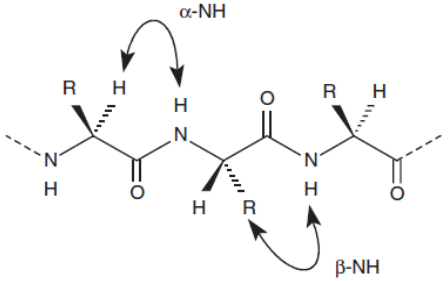

Solution State Conformations:

The presence of different NOEs can help narrow down the conformation of a molecule. In *J*-based analysis, NOEs can help distinguish between two relative configurations that otherwise give the same couplings.

In protein work, NOEs are used to create constaints on intra-molecular distances. These constraints can then be used in molecular mechanics optimizations to provide a picture of the solution state conformation of the protein:

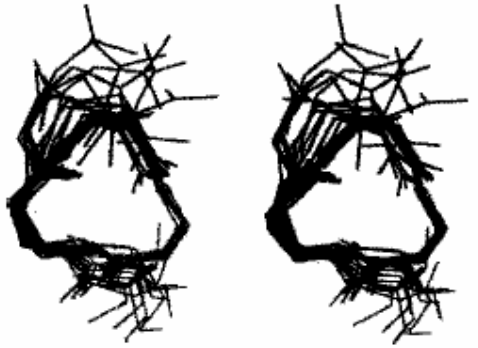

What This Means in the Lab

(1) Use 1D-NOESY if you have a specific correlation you're looking for. Use 2D-NOESY if you want to look at all the correlations at once. Pay attention to the mixing time (**mix**) and recycle delay (**at**+**d1**). **mix** should be about T_1; **at** + **d1** should be about $2T_1$. (These are for natural products.)

(2) Switch from NOESY to t-ROESY if your compound is in the crossover region (750-2000 Da).

(3) COSY artifacts are common in NOESY spectra and have an anti-phase appearance. Exchange peaks are negative.

(4) Use DPFGSE-1D-NOESY, not NOE difference, since this gives much better results.

(5) For either steady state or transient NOEs, the *percentage* NOE is, in the absence of calibration, not quantitatively related to distance. For most work, it suffices to call a crosspeak

(6) NOESY peaks are *not* symmetric because of spin diffusion. That means the crosspeak from A to B is not necessarily the same intensity as a crosspeak from B to A. (!)