The Hadamard Chronicles

A Story of one Research Project -- Warts and All

Ray Freeman

Jesus College, Cambridge, UK

June 2003

INTRODUCTION

It is seldom that an entire research project is documented with any degree of completeness; few scientists keep comprehensive laboratory notes in this fast-moving age. Fortunately this particular collaboration is an exception because the two authors communicated almost entirely by electronic mail and met in person on only one occasion -- a scientific meeting in Prague in the Czech Republic. Although we did occasionally use the telephone, 95% of our discussions were by e-mail and a complete record exists. This provides an opportunity to follow the development of the original idea through to publication, with all the misconceptions, misunderstandings, mistakes, delays and omissions clearly set out for all to see. Notice how the activity started very slowly but reached a crescendo in July 2002.

The field of research is nuclear magnetic resonance (NMR) spectroscopy. My coworker is Dr Eriks Kupce, originally from Latvia, now working for Varian, Inc. in Oxford, UK. He performed the actual experiments while I sat in Cambridge and waited. Unfortunately his "day job" involves extensive travelling around the world to demonstrate Varian's latest NMR spectrometers; this leaves little time for research. Another problem is that we have different computers (PC and Mac) and there have been many difficulties transmitting figures.

Although only 70 miles apart, the two university cities (Oxford and Cambridge) have poor road and rail communications. We were saved by electronic mail which works just as well when Eriks is in Palo Alto, California as in Oxford, England. Since 1999 I have been formally retired and have no direct access to NMR equipment, but before retirement I worked with Eriks over a a ten year period, and have the greatest respect for his skills and ideas. We have published some 30 papers together.

For the last thirty years high-resolution NMR spectroscopy has employed the pulse/Fourier transform method, almost without exception. Signals are excited in the time domain and later transformed into the frequency domain. This research "Hadamard spectroscopy" proposes an alternative technique where the nuclear spins are excited directly in the frequency domain with a "comb" of frequencies. It involves an encoding scheme roughly equivalent to the operation of a scrambler telephone. Many conventional NMR ideas have to be radically revised in Hadamard spectroscopy. We fell into some of these traps ourselves.

I thought there might be some interest in this record of our electronic conversations, edited to remove the headings and addresses, with some of my retrospective comments (*in parentheses and italics*). The original intention was to reproduce our electronic mail without further editing, but this turned out to be too long and tedious. In the end I deleted material about problems with spectrometer access and absence from the laboratory on other "official" work. Eriks' replies are identified by the ">" sign and by smaller type. On several occasions there were multiple e-mail messages on the same day; these have been combined, with one line space between them.

Abbreviations:

FT:	Fourier transform
FID:	Free induction decay
COSY:	Correlation spectroscopy
2D:	Two-dimensional spectroscopy
HT:	Hadamard transform
TOCSY:	Total correlation spectroscopy
NOESY:	Nuclear Overhauser spectroscopy
DANTE:	Selective excitation sequence
P-box:	Pandora's box, a Varian software program

23 Nov 01
Dear Eriks,

FREQUENCY-DOMAIN HADAMARD SPECTROSCOPY [23 November 2001]

Last night I had an idea that might prove useful, and you are the one person who could make it work.
You will probably remember our papers on the "Hadamard trick" a few years ago.
Multiple selective excitation at N different sites with phases alternated according to a Hadamard matrix, followed by decomposition according to columns of the same matrix gave all N separate subspectra but with a square root of N advantage in signal-to-noise. See "Spin Choreography" page 181.
Chem. Phys. Lett. 215, 341 (1993)
Magn. Reson. Chem. 32, 134 (1994)

Toshi later used these techniques:
Nishida et al, Magn. Reson. Chem. 33, 596 (1995)
Nishida et al, Magn. Reson. Chem. 34, 377 (1996)

My idea is to extend this to "polychromatic" excitation with soft pulses spaced regularly across the high-resolution spectrum. This would give us the multiplex advantage. Note that, unlike conventional pulse-FT spectroscopy, any region can be selected at will. This is because it is a direct frequency-domain experiment, rather like Wes Anderson's "Prayer Wheel" method (never implemented).
At first, one wonders whether adjacent soft pulses would interfere ("Close encounters") but I think the interference would cancel owing to the +/– coding of the pulses. This could easily be checked.
The technique would be very similar to the one you used recently for the 4-bit quantum computer test where you let the computer set all the possible pulse frequencies automatically.
There are standard ways to generate large Hadamard matrices. Kaiser wrote a rather obscure paper on noise excitation with Hadamard matrices.
R. Kaiser, J. Magn. Reson. 15, 44 (1974).
What are the applications ? I believe there may be several.

(1) Everyone has been looking for an alternative to FT, without much success.
Some of the FT artifacts would be circumvented with frequency-domain excitation.
No aliasing problem ! No truncated FID's ! The dynamic range problem is moderated.

(2) Run only the interesting sections of an extensive HR spectrum. Multiple band-selective excitation. Cuts down the overall length of the Hadamard experiment (which could be quite long in the general case).

(3) If we run the high-resolution spectrum of a sample suspected of being a mixture of compounds, then run spectra of the individual components, the information could be used to decompose the "unknown" spectrum into subspectra of pure components. This might be a software program that identifies the relevant combinations of columns of the Hadamard matrix. Could be a clean separation, even in the presence overlap.

(4) Solvent suppression by leaving a gap in the frequency domain. The main drawback of FT is the dynamic range problem, but this is circumvented with frequency-domain excitation.

(5) Very wide-band excitation. No fall-off at the edges of a very extended spectrum. Phosphorus-31 spectra recorded in two separate sections, for example.

(6) Non-uniform excitation. Detect carbon-13 satellites at high gain with the parent peaks at low gain.

(7) Template excitation. Look for a specified component selectively. Could be useful in monitoring the time evolution of a selected component; integrate the chosen signal but ignore the rest of the spectrum.

(8) Homonuclear decoupling. Continuous excitation of one chosen narrow band (the irradiated spins) without having to detect this region.

(9) Variable resolution across the spectrum. The irradiation frequencies need not be regularly spaced.

(10) Lends itself to data compression, when we want to compare an experimental spectrum with a library of standard spectra (combinatorial chemistry) and where too much detail is counterproductive.

(11) Pseudo COSY. [J. Magn. Reson. 75, 540 (1987)]. This has the great advantage that the diagonal peaks are almost completely suppressed. Hadamard coding would restore the multiplex advantage.

DISADVANTAGES

(1) The main drawback of the proposed Hadamard method is the requirement for N scans where N is the number of excitation points in the frequency domain.

(2) Second drawback is that the number of excitation frequencies has to fit an existing Hadamard matrix. We cannot use just any old number.

This seems to be an "enabling" technique that could lead to applications not yet dreamed about.

Note that you could test the idea with quite a small Hadamard matrix.

Can you do it ???

Eriks' reply:
>I actually thought about it some time ago and I even have all the
>algorithms ready for any number of lines. We just have to do it. The
>problem is that at the moment I am almost completely consumed by the 900
>while my 500 is being re-installed and will probably become available in
>about a week or so. Can you suggest an experiment in the meantime?

26 Nov 01
Dear Eriks,

I guessed as much, because you invented almost all of the tools required to do this Hadamard spectroscopy. I think it could be a hot property, so why not sneak a few hours for your own work. In the long run it should be a great asset for Varian; you could offer the software package (rather like the P-box) and Bruker would struggle even to try to keep up.

For the first experiment I would try a band-selective excitation of a complicated multiplet. Carbon-13 spectra of camphor or endo-borneol have two interpenetrating quartets near 2 ppm, for example. Then the two multiplets could be displayed separately.

The big question is the potential interference between adjacent irradiation frequencies. You addressed this problem in the "Close encounters" paper but I think a simpler solution would be interleaving -- two separate measurements (one immediately after the other) with the pulse frequencies first at odd-numbered frequency locations then at the even locations. Recombine into a Hadamard matrix afterwards. Although this would seem to halve the multiplex advantage, in fact much of the loss is recovered because while the "odd" regions are relaxing, we hit the "even" regions, then the "odd" regions again, and so on. So the "odd" and "even" acquisitions have to be interleaved in time.

Eventually we could try more complex interleaving: 0 - 2 - 1 - 3 for example.

Hadamard spectroscopy requires N measurements where N is the number of individual frequency components (N approximately = spectral range divided by linewidth). So the method

only competes with FT if the inherent signal-to-noise is low enough to require extensive time-averaging. Otherwise FT gets the result more quickly. This requirement is easier to satisfy for carbon-13 and spectra of restricted width. But (multiple) band selected excitation or non-uniform sampling are key advantages of Hadamard spectroscopy anyway. What we would be offering is an alternative to FT for situations where FT does not work too well.
One application would be (a) sparse sampling of an unknown spectrum, followed by (b) fine sampling of restricted "interesting" regions.
My guess is that any desired Hadamard matrix could be generated by a recursive algorithm starting from

$$\begin{matrix} + & + \\ + & - \end{matrix}$$

but they should be in the textbooks anyway.
I think this could be a really important new technique. If someone has been reading our papers carefully, all the elements are there to come up with this idea. We need a quick Communication to get a foot in the door. At least do a "quick and dirty" run on something. The question of interference could be critical.
We must keep in close contact while I am in California (12 December onwards). I don't think Dominique has e-mail at the house yet, but both Dominique and Louise have facilities at work.

4 Dec 01
Dear Eriks,
Best of luck for your Hadamard experiments. If it works this experiment has a nice feel about it and you have all the tricks already worked out.

12 Dec 01
Dear Ray,
>There is still a bit of delay with my 500, which will only be at field end
>of this week. Then it will probably take a week or so for it be mapped,
>shimmed etc., so we will have to do the work in cyberspace once again.
>My plan is to run a systematic frequency space mapping using the Hadamard
>excitation, maybe with 128 or so frequency components. Then we can
>demonstrate the advantages for the case when the frequencies are known and
>the frequency space can be sampled in a non-linear fashion. A simple HMQC
>type experiment should do.

19 Feb 02
Dear Eriks,
I suppose you could start on assembling all the programs needed for Frequency Domain Hadamard Spectroscopy. If it eventually works, we might just make the ENC meeting ??????????????. I see that Anatole Abragam will be there.

22 Jan 02
Dear Ray,
>Nice to see you are back. In the meantime the 500 magnet has been brought
>back to field. It is a magnet that has been rejected by a customer, so we
>get it (and all the trouble with it) from Oxford for free. After it is
>shimmed and reconnected, I can start doing some experiments. The only thing
>is, I am travelling for the next two weeks, but we will certainly do it
>right after.

11 Mar 02
Dear Eriks,
That is great news. In case you mislaid ny original notes, here they are again:
(not duplicated here)
The square root of two advantage is just the standard multiplex advantage (detecting all the information in every scan) and is fundamental to the Hadamard experiment.

11 Mar 02
Dear Ray,
>I guess this is the explanation for the sqrt(2) increase observed by
>Krishnamurthy, not the signal narrowing, as he explained it. Also, I wonder
>whether there is a connection with the sensitivity enhancement technique by

>Rance et al, which also gives sqrt(2) improvement.

>Because of the same thing - in phase encoded experiments and in the
>ordinary 2D experiments we detect only one component, Re or Im, at a time
>(in the indirect dimension). In the Hadamard experiment we detect both
"plus" and "minus" components simultaneously, similar to the sensitivity
>enhanced experiments where both the Re and Im parts are detected in one
>experiment.

11 Mar 02
Dear Eriks,
I'm not aware of the Rance paper -- more details ?

12 Mar 02
>This is his trick to detect simultaneously both the cos and sin modulated
>components followed by recombination (JMR, v. 91, pp. 429-436 (1991).
>A simple phase encoding gives the same sensitivity as 2D. Hadamard gives
>sqrt(2) times more!

Dear Eriks
No this is a different matter. Hadamard should have the same sensitivity as two-dimensional spectroscopy, but without some of its imperfections.
Go for it !

12 Mar 02
Dear Eriks,

Glad you have made a start. Remember it is easy to kill an idea (for the wrong reasons) at this stage. Many applications would be on a narrow spectral width anyway.

12 Mar 02
Dear Ray,
>What I mean is that in FT for a real signal the imaginary part is zero and
>vice versa, whereas for Hadamard technique (HT) there is always a signal,
>either positive or negative, which gives the advantage of square root of 2:

	Phase encoding		Hadamard	
	Re	Im	(Re)	(Im)
>FID 1	1	0	1	1
>FID 2	0	1	1	-1
>FT/HT	1	1	2	2
>S/N	1	1	1.4	1.4

I am doing simulations just to see how it all works. Interestingly, it works better with more components. The bad news is that the bandwidth is limited by the available RF amplitude (bandwidth = B1max), which for C-13 might be insufficient, especially at high fields. The other problem with simultaneous excitations is the Bloch-Siegert shifts, which distort the frequency domain.
A potential solution for both problems would be to use sequential 180 pulses, if we can ignore relaxation. It will be a bumpy ride...

19 June 02
Dear Eriks,

I am in the middle of a major rethink on the Hadamard trick and will send you some thoughts very soon. I believe we can avoid most of the interference effects and that a Fourier transform is NOT necessary, so truncation artifacts etc should disappear.

19 June 02
>I had similar thoughts. I look forward to see what they have in common...

19 June 02
Dear Eriks,

My idea for Hadamard spectroscopy would NOT involve Fourier transformation. However the processing is a little more complex than my first description. Essentially this is a multichannel excitation scheme with multichannel detectors. At the risk of seeming longwinded, I spell it out in excruciating detail:

Take a very simple example of a single resonance line defined by 4 digitization points ("channels") at frequencies 1,2,3,4, with amplitudes (say) 1, 9, 1, 0.
The "channels" are at finer frequency steps than the resonance linewidth.
We apply a composite sequence made up of four weak selective rf pulses.
[Maybe they have to have small flip angles, maybe not].
The selectivity should be high enough that only the chosen channel is mainly excited, with the nearest neighbours only slightly affected. This is NOT the wideband polychromatic excitation scheme that we wrote about in JMR A 108, 268 (1994), on a fine frequency scale the envelope is not of constant amplitude, although each channel receives the same level of excitation.
We make four measurements a, b, c, d with the relative pulse phases:

	1 2 3 4	
a	+ + + +	FID(a)
b	+ + – –	FID(b)
c	+ – + –	FID(c)
d	+ – – +	FID(d)

Here "+" means radiofrequency phase 0. "–" means 180. This is essential.
We record a free induction decay after each measurement: FID(a,b,c, or d)
Each column is associated with a known audiofrequency which we used to set up the soft pulse excitation stage. Call this F.
We combine the four free induction decays according to the pluses and minuses in that particular column, for example + + – – for column 2.
We multiply the combined free induction decay by F, and integrate the result. This is a single-channel Fourier transformation if you like. It measures only the component at or very near to F, all other signals vanish. This is why truncation does not really matter very much; the FID should however be long enough that contributions from neighbouring channels is small.
Phase shifts of spins in other channels cancel because of the +/– encoding.
Repeat this processing for all channels using the appropriate F values.
Each channel now shows a number (no phase information). They should now read 1, 9. 1, 0.
Note that the soft pulses MUST have phases 0 and 180 (not on and off).
Consequently we side-step the well-know drawbacks of the conventional FT method -- truncation artifacts, pulse breakthrough, uniform frequency-domain sampling etc, etc.
Bloch-Siegert shifts of nearest-neighbour resonances still occur, but should be small for rf pulses that are soft enough.
Let me know what you think.

Dear Ray,

>Our thoughts match as usually, with few exceptions (I think it is
>Polychromatic excitation, except we are now varying the phases of
>individual components). In a sense it is CW NMR with Hadamard transform (it
>would be interesting to see what is the analogue for FT and whether it
>would reduce itself to the conventional FT NMR, I guess it would!). You
>will be pleased to learn that I actually did some preliminary experiments
>just before I left. I did not have time to do a full Hadamard experiment,
>but I did get the experiment working with a single component. I used the
>doped MeJ sample (to have a reasonable relaxation rate and no J couplings)
>to lock on a single line and this worked quite well. So, all we have to do
>is to repeat this with a simple Hadamard matrix. Can you wait until the
>next weekend?

19 Jun 02
Dear Eriks,
You are right: it is very close to CW NMR but with the multiplex advantage of FT. I was wondering if there might be an operational mode where the signals were accumulated at the same time that the soft pulses were being applied in a repetitive sequence ????

>Sure there is. That's how I did the preliminary experiments. I have to run now....

It's a bit like stochastic excitation without the complications arising from excitation noise.
I don't believe quadrature detection would be required since there is no strong central transmitter frequency in Hadamard mode.
It's all so new, we should proceed by small simple steps (as you did with MeI).

24 Jun 02
Dear Eriks,
My revised notes on Hadamard spectroscopy are attached to this e-mail.

FREQUENCY-DOMAIN HADAMARD SPECTROSCOPY

In its simplest form, the idea is to use an array of soft radiofrequency pulses spaced regularly across the high-resolution spectrum. See "Polychromatic Selective Pulses" Ē. Kupce and R. Freeman, J. Magn. Reson. A 102, 122-126 (1993).
By exciting all the spins all the time, this gives the same multiplex advantage enjoyed by hard pulse / Fourier transform methods..
In the more general case, the spacing of frequencies could be arbitrary, or we could define regions of regularly-spaced frequencies alternating with "blank" regions where no NMR signals are expected. Note that, unlike conventional pulse-FT spectroscopy, any spectral region can be selected at will. This is because it is a direct frequency-domain experiment.
The pulse phases are determined (0 or180) by a Hadamard matrix. There are standard ways to generate large Hadamard matrices. Kaiser wrote a rather obscure paper on noise excitation with Hadamard matrices [R. Kaiser, J. Magn. Reson. 15, 44 (1974)].
Consider the very simple Hadamard matrix:

	1 2 3 4	
a	+ + + +	FID(a)
b	+ + − −	FID(b)
c	+ − + −	FID(c)
d	+ − − +	FID(d)

Each row of the matrix defines one experimental measurement where a free induction signal (FID) is recorded. Each cloumn defines a particular point in the frequency-domain excitation spectrum. Call this a "channel" of the multichannel excitation. Each channel is associated with a particular frequency "F" in the array of selective pulses. We can think of these as audiofrequencies that are used to modulate the main transmitter frequency.

In each channel the free induction decays are combined according to the pluses and minuses in that particular column, for example + + − − for column 2.

The combined free induction decay is multiplied by F, and the result is integrated. This can be thought of as a single-channel Fourier transformation. It measures only the NMR component at or very near to the audiofrequency F, all other signals vanish. This operation is repeated for all the channels, using the appropriate values for F.
Truncation of the free induction decay does not really matter much; unlike conventional Fourier spectroscopy, a step function at the end of the FID does not affect the shapes of all the resonance lines in the spectrum. However the FID should be long enough that the signal contributions from neighbouring channels is small.
Number of channels N

There will be an upper limit on the number of channels, set by the maximum radiofrequency power dissipation. It may help to use small flip-angle pulses.
The number of channels also determines the duration of the measurements, because N free induction decays must be accumulated. Weak NMR signals already require long-duration experiments by conventional Fourier transform methods but there the number of accumulations is arbitrary. N also determines the amount of data processing required.
At first sight it would seem that all frequency-domain Hadamard measurements require a large number of separate acquisitions N, where N is the number of frequency channels. That is because the Hadamard matrix is square. This may not in fact be necessary. We could consider using an array of 8x8 Hadamard matrices linked together to cover the entire spectral frequency range. Columns with identical +/– coding would then be separated by seven intervening channels with different encoding, and the two frequencies would be sufficiently far apart that signal separation at the processing stage would be quite satisfactory. The encoding would be repetitive. For the even simpler example of 4x4 matrices, only four accumulations would be required to cover the entire spectra range:

+ + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
+ + − − + + − − + + − − + + − − + + − − + + − − + + − − + + − − + + − − + + − −
+ − + − + − + − + − + − + − + − + − + − + − + − + − + − + − + − + − + − + − + −
+ − − + + − − + + − − + + − − + + − − + + − − + + − − + + − − + + − − + + − − +

In practice, we would use the largest possible Hadamard matrix compatible with the amount of multiscan averaging required to reach an acceptable signal-to-noise ratio. If it required Q square Hadamard matrices of size NxN tacked together, there would be N different composite free induction decays and each one would require processing with Q different audiofrequencies to give the NQ different signals.
Interference caused by the radiofrequency fields in adjacent channels.
For one adjacent channel there would be a Bloch-Siegert shift and a corresponding phase error. But for two interfering radiofrequency fields symmetrically placed on either side, these shifts and phase errors should be equal and opposite. Only at the edges of the excitation spectrum would there be an asymmetry and hence some residual Bloch-Siegert effects.
It might just be necessary to omit all columns encoded with all +. Then there would always be an equal number of free induction decays combinations where the adjacent channels had an equal number of + and – phases. Better cancellation ?

APPLICATIONS

(1) THIS METHOD DOES NOT INVOLVE FOURIER TRANSFORMATION
(Except in the sense of single-frequency Fourier analysis).
Everyone has been looking for an alternative to FT, without much success.
Some of the FT artifacts would be circumvented with frequency-domain excitation.
No aliasing problem ! No quadrature detection ! No problem with truncated FID's ! No zero-filling or apodisation required !

(2) The dynamic range problem is circumvented because a gap could be left in the excitation spectrum so that a strong solvent peak would not be excited. Alternatively the strong peak could be allowed to overflow without prejudice the rest of the spectrum, or simply processed at much lower vertical scale..

(3) Run only the interesting sections of an extensive high-resolution spectrum. Multiple band-selective excitation. Cuts down the overall length of the Hadamard experiment (which could be quite long in the general case). We could run a preliminary full spectrum and then go back to re-examine an interesting narrow region of the spectrum with very high definition.

(4) If we run the high-resolution spectrum of a sample suspected of being a mixture of compounds, then run spectra of the individual components, the information could be used to decompose the "unknown" spectrum into subspectra of pure components. A software program

could be written to identify the relevant combinations of columns of the Hadamard matrix. Could be a clean separation, even in the presence overlap.

(5) Very wide-band excitation. No fall-off at the edges of a very extended spectrum. Phosphorus-31 spectra recorded in two separate sections, for example. Possible limitation due to total radiofrequency power dissipation.

(6) Non-uniform excitation. Detect carbon-13 satellites at high gain with the parent peaks at low gain.

(7) Template excitation. Look for a specified component selectively. Could be useful in monitoring the time evolution of a selected component; integrate the chosen signal but ignore the rest of the spectrum.

(8) Homonuclear decoupling. Continuous excitation of one chosen narrow band (the irradiated spins) without detecting this region. Simply use higher rf and all "+" in these columns of the Hadamard matrix.

(9) Variable resolution across the spectrum. The irradiation frequencies need not be regularly spaced.

(10) Lends itself to data compression, when we want to compare an experimental spectrum with a library of standard spectra (combinatorial chemistry) and where too much detail is counterproductive.

(11) Pseudo COSY [J. Magn. Reson. 75, 540 (1987)]. This has the great advantage that the diagonal peaks are almost completely suppressed. Hadamard coding would restore the multiplex advantage that was lost in the original implementation.

(12) Editing of two-dimensional spectra, where only selected correlations, chemical exchange, or cross-relaxation effects are recorded. Tailor the excitation spectrum in the indirect dimension.

DISADVANTAGES

(1) The main drawback of the proposed Hadamard method is the requirement for N scans where N is the number of excitation points in the frequency domain. But see the modification suggested above.

(2) Second drawback is that the number of excitation frequencies may have to be quite large. May reach the radiofrequency power dissipation limit.

(3) Square Hadamard matrices ? We probably don't need to use all the N columns of the Hadamard excitation matrix, but we have to use all the N rows in the acquisition stage for proper selection.

This seems to be an "enabling" technique that could lead to applications not yet dreamed about.

2 Jul 02
Dear Ray,
>Below is what I have so far. It is a 256 point Hadamard spectrum of MeI
>with 1Hz per point resolution. I will try a more realistic sample now.

>I am sending you three pictures with separate e-mails. The first is an
>overlay of three individual traces from the MeI spectrum.

>below is the second spectrum. It is the build-up signal for the left
>component of the MeJ spectrum.

>Attached is the Me-salicylate spectrum. The spectral width in the Hadamard
>dimension is 512 Hz, with 128 points the resolution is 4 Hz per point.

>Spectrum # 3 is not very good. There is an axial peak (everything is done
>with a single transient) and a solvent (CDCl3) peak. Attached is a
>1-ethyl-indanone spectrum, the aromatic region, with Hadamard 256 matrix (1
>Hz/point). It looks to me that I will have to write a piece of software for
>a proper processing. The standard software is only good for a quick glance
>at the result. Perhaps we should try vinyl-phosphate (which we used for
>quantum computing paper). It gives a fairly compact and also fairly complex
>spectrum.

2 July 02
Dear Eriks,

Strangely enough, we wrote a paper on methyl salicylate in 1963, Spectra on a Varian 60 MHz spectrometer. The coupling constants were:
J(AR) ortho 7.20 Hz
J(AX) ortho 7.98 Hz
J(MR) ortho 8.37 Hz
J(AM) meta 1.15 Hz
J(RX) meta 1.78 Hz
J(MX) para 0.43 Hz
where the ring protons are labelled A,M,R,X from high-field to low-field.
There is a long-range coupling of 0.40 Hz between proton R and the OH.
Freeman, Bhacca and Reilly, J. Chem. Phys. 38, 293 (1963).

Received both picture and have printed them. Great ! The lineshapes look a lot better now. Absolute value still ??? The build up to a steady state is very obvious. Which makes me wonder if we shouldn't make the first rf pulse stronger than the rest so as to reach a steady state earlier ? Of use a dummy set of "prepulses" .But it all looks very encouraging !

I cannot match it up with the methyl salicylate spectrum we ran at 60 MHz almost 40 years ago. Are there impurities ?

3 July 02
Dear Eriks,

I liked the indanone spectrum. I assume you are still using absolute value mode. We may have to worry about line shapes and line widths. But it is really encouraging to get real spectra so soon. FT took years before we were confident of the results. I remember Howard and I had nightmares trying to get the first carbon FT spectra.

I have been thinking about the shape of the time-domain signal of methyl iodide that you sent yesterday, where the steady state builds up on a rising exponential. I don't think the shape matters too much (not in the way it does in conventional FT spectroscopy). In a sense we measure only how much oscillatory component there is -- which translates into just one ordinate in the frequency domain. If there had been a constant level throughout the time-domain, it would give a delta function in the frequency domain (of course) and that would be quite OK. However, if the initial point is zero (or near zero) that means that the frequency-domain must have zero (or near zero) integral which translates into negative sidelobes. In absolute-value mode these sidelobes become very long tails on each side of the line.

One simple trick would be to read the time-domain signal backwards, after Hadamard decoding but before the one-point Fourier transformation.

5 July 02
Dear Eriks,

I have been thinking about the linewidths. When you examine the time-domain signal "MeJ_fid.tif" it seems to be a pretty pure "monochromatic" sinewave with no other frequency components present. So I wonder if the observed line broadening is not due to the shape of the envelope of the soft pulses. If this envelope is rectangular there will be off-resonance excitation. Perhaps it needs to be rounded off a bit ?

8 Jul 02
Dear Ray,
>I am working on software issues to allow essentially unlimited number of
>data points in the Hadamard matrix and also to do a proper Hadamard
>transform, so that we do not have to use absolute value projections. I got
>carried away by some urgent Varian requests last week, as they are in the
>process of releasing new software, but that is over now and I am back to
>our experiments. I will send you new pictures as soon as they appear on the
>screen.

15 Jul 02
Dear Ray,
>Sorry, it took so long. Finally we do have a proper processing. The
>indanone spectrum has a linear phase gradient across the spectrum (as you
>would expect from polychromatic pulses), which is phase-able, Fig 2. The
>absolute value spectrum looks very similar to what we saw from the
>projection of the 2D spectrum. There are some peculiarities in the
>baseline. I will check whether Bloch-Siegert compensation or other tricks
>makes it any better.

>Attached is the first trace of the Hadamard experiment (all pluses) before
>and after the phase correction. Let me see whether the phase correction can
>be done within the pulse.
>The S/N at this stage will not be great because we are working with small
>flip angles (ca 10 degrees). To make it better we could either try to
>increase the flip angle or we could start with XY magnetization, very much
>like in quantum computing experiments.

15 July 02
Dear Eriks,

Well done ! The absorption-mode lines look much better than absolute value. I am not sure I see yet why there is a phase gradient; it's not for the same reason as conventional FT.
The baseline stuff seems to arise because the early points of the "FID" are weak, the very first may even be zero. This means that the integrals in the frequency-domain are smaller than they should be -- hence the negative lobes on each side of the line. We shall need to run a conventional FT spectrum of 1-ethyl indanone for comparison of signal-to-noise and resolution. How about a solvent suppression experiment with a notch at the water frequency ? Could you write a resume of the data processing that you have used ? Looks like we are in business.

I was wondering whether this phase gradient is an artifact of the phase incrementation scheme used in polychromatic pulses. Perhaps the incrementation should be "anchored" at the end rather than at the beginning, so all components end up with the same phase ?????

Dear Ray,
>It is "anchored" at the end of the pulse. I believe the phase gradient
>comes from Bloch-Siegert effects. It is zero in the middle and increases on
>the edges, just as you would expect from BS effects. I am now trying to see
>whether we can account for that in the shaped pulse.

>The S/N at this stage will not be great because we are working with small
>flip angles (ca 10 degrees). To make it better we could either try to
>increase the flip angle or we could start with XY magnetization, very much
>like in quantum computing experiments.

15 Jul 02
Dear Eriks,

I would be surprised if it were the Bloch-Siegert effect. The artifacts at the edges of your last indanone spectrum might be Bloch-Siegert though.

15 Jul 02
Dear Ray,
>It seems that phase correction works !
>Attached is a simple pc-9 pulse excitation (Mx) and phase profiles - the
>original and corrected.

>The phase gradient was 11.7 degrees per isochromat (100 Hz). This required
>increase of B1 from 400 Hz to 700 Hz (for a pulse duration of 5 ms).

15 July 02
Dear Eriks,

About the phase gradient -- "anchoring" the phase at the end of the sequence would be correct if we were acquiring a free induction decay starting AFTER the excitation. In your present scheme, you acquire data throughout the excitation sequence, so what is the correct "anchor" ? Perhaps it should be half way along the excitation sequence ?

>I suppose this is because the signals in our "channels" have slightly
>different phases due to the phase gradient in the PC-pulses.

16 July 02
Dear Eriks,

Your solution to the phase gradient problem is an elegant one. I noticed that it improves the amplitude profile as well, and it looks as if the edge artifacts are now from the Bloch-Siegert effect. If you don't make this correction, I think the phase should be anchored at the mid-point, so that dispersion signals from the first half cancel those from the second half. But we would lose some intensity. What about reading the "FID" backwards to avoid the rising signal at the beginning ? *(Here I fell into the trap of thinking in terms of Fourier transforms).*

17 July 02
Dear Ray,
>I tried to reverse the FID, but there was no effect. I guess that is not
>unexpected. Since we are only doing one point FT, it does not matter
>whether the FID is going up or down. The phase gradient was slightly
>higher. I still think that is hardware related and we should be able to
>adjust it. I am travelling for the rest of this week. I shall be back on Friday.

17 July 02
Dear Eriks,

It seems to me that the way you generate frequencies should not introduce any phase gradient, and if you use the same method to get each frequency for one-point Fourier transformation, there should also be no phase gradient. Which raises the question - what were those two simulations showing ?

>That was the situation at the end of a PC-pulse, not really related to the
>situation during our experiment, nevertheless an interesting result. I will
>do more experiments when I get back. We need to get rid of the phase gradient!

We know that if the channels alternate in phase (+/–) this would create a large phase gradient, but you are decoding the Hadamard modulation before you do the one-point Fourier transform, so there should be no (+/–) alternation at that time.
It seems more likely that it is the result of acquiring between rf pulses.
There's always some mystery with new techniques.

(We were looking for a fundamental reason for the phase gradients, when in fact it later turned out to be the result of a programming error, so all these speculations were pointless)

19 July 02
Dear Eriks,

Keep up the good work. I have written a second draft of the text which is an attachment to this e-mail. A few misconceptions have been corrected.

22 July 02
Dear Eriks,

Hope Varian will give you time to try a few more experiments; it looks very promising.
I thought about the phase gradient again. Is it just the delay caused by interleaving pulses and acquisition ? That would seem to give a phase gradient of the right magnitude. If so it should

be possible to introduce a (frequency-dependent) phase shift by delaying the start of the reference frequency F by the same (fixed) time interval (before multiplying the signal by F in the single-point FT operation). Alternatively the phase gradient could be calculated and used as the conventional phase correction. If the first of these suggestions works, we could claim it as an advantage over the usual Fourier transform method -- no phase corrections needed !

22 Jul 02
Dear Ray,
>I am back and pulsing! I looked through the second version of the
>manuscript. On page 7, paragraph 1 - perhaps we should emphasize that there
>is no free induction decay in the Hadamard spectroscopy.
>Looking into phase shifts now.

(Telephone conversation established that the phase gradients were an artifact of a programming error)

>Attached is the first trace (all pluses) of the indanone with zero phase
>correction and the HT spectrum. For some reason the later did require a
>little bit of the first order phase correction, but we can ignore that for
>the time being. The baseline now is much better.

23 July 02
Dear Eriks,

New indanone spectra look good. Is the extreme low-field line the aliased solvent peak you mentioned ? Strictly there should be no aliasing in HTS; it must be a bit of the solvent peak tail excited in the wrong channel.

>It looks that way.

Keep pulsing !

>I am going for 512 or maybe even 1024 matrix. That requires a bit of extra
>programming. Will send you a picture as soon as I get one.

Big matrices are good. I look forward to your results. You already anticipated one application of Hadamard spectroscopy in your 1992 "Fast-relaxing nuclei" paper.
Attached is a second draft of the proposed text.

26 Jul 02
Dear Ray,
>I ran a 2048 Hadamard experiment last night and I am processing it "as we
>speak". The technology is now in place, but the experiment takes about 8
>hours to run, mainly because of the 15 s delay between individual
>transients. That is time required at this stage to calculate each
>individual row of the Hadamard matrix. If you think about it, it is 2048 x
>2048 or more than 4 million selective excitations in one experiment. I
>would guess it is a new record in the Guiness book !

>Attached is the spectrum. There is a little glitch at the base of the
>signal near 2 ppm. Since the signal originates from a closely spaced AB
>system, I would guess that this has something to do with the Tsetse effect.
>The sample is strychnine. Let me know what do you think.

26 July 02
Dear Eriks,

2K is enormous. Well done ! Any dynamic range difficulties ?

>I saw some digitization errors due to the dynamic range problem in some
>shape-files. However, I would think that this merely adds to the excitation
>noise and for our experiment would not matter too much.

Anyway you are surely in the Guiness Book of Records already . . .
I look forward to seeing some pictures,

26 July 02
Dear Eriks,

That is amazing ! I can rewrite the text taking 2K by 2K as a "typical" Hadamard matrix size. (I mentioned my comments on the telephone) Let me have some experimental numbers. Keep pulsing !

26 Jul 02
Dear Ray,
>sw = 2000 Hz
>at = 1s
>pulse duration for the Hadamard pulse = 1s
>spacing between the frequency components in Hadamard pulses - 1 Hz
>Flip angle for each rectangular pulse component in the Hadamard matrix - 10
>degrees
>Number of Hadamard components 2048
>Time between transients - ca 15 sec (this is also time required to
>calculate one raw in the Hadamard matrix
>Total experiment time 8h 50 min (could be reduced to 53 mins using 1.5 sec
>repetition time).
>Please let me know if there is anything I have missed.

>I have now set up experiment to exclude the solvent signal. Will run it
>over the weekend. I will think about the other tricks too.

>Just a few remarks.
>1) I am using a 500 MHz spectrometer;
>2) I think the last statement is not quite correct. We are applying
>constant amplitude RF fields, but with different phases (either positive or
>negative) simultaneously at 2048 channels. So the total flip-angle at the
>end of the pulse in each channel is either +10 or -10 degrees.

26 July 02
Dear Eriks,

A quotation from that book on "Transform Techniques in Chemistry".
"There exist cyclic codes intimately related to orthogonal binary codes. These codes can be folded into a two-dimensional array, much as the one-dimensional lines of printed words are folded into a 2D array of printing."
We could look at the Hadamard matrix and see whether it could be reordered into a long linear array that would reproduce the entire matrix simply by shift operations.
As you say, we don't need this degree of sophistication at the moment.

Thanks for the numbers. I was thinking that frequency-domain excitation could be useful in 2D spectroscopy, for example COSY. Suppose we run a quick conventional 1D spectrum, extract the lines of interest and use these frequencies for simultaneous excitation as 90 degree pulses. After normal evolution and a hard 90, we would have all the information in the conventional COSY spectrum but as 1D traces. I guess you thought of this already.

(The extension to two-dimensional spectroscopt (COSY) turned out to be important but I was wrong to assume that we would need an evolution period)

Here is a rough first draft of an experimental section (attachment)

It would be fun to use all your 2048 channels on a restricted region of strychnine (say 3.7 to 4.4 ppm) with a much longer acquisition time (10 secs ?). See what kind if definition/resolution you could get.

29 July 02
Dear Eriks,

How many pulses do you use in the DANTE sequence ? Presumably this has to be appreciably larger than half the spectral width so that the first DANTE sidebands are well out of the NMR spectrum. Could this be the origin of the little negative glitch ?
IF, for example the number of DANTE pulses were 2K, you would have applied 8 billion pulses by the end of the experiment !
Sorry about the mistake.

>I used the vector addition scheme rather than DANTE. This has the advantage
>of giving no (or rather less) sidebands. So, the number is still 4 million.
>Sorry, if that sounds like a disappointment...

But I thought you had a sequence of pulses intersperced with one-point signal acquisitions ("explicit acquisition") essentially a DANTE sequence ????

>We do this after the pulse is made. This requires adjustment of the RF
>amplitude for the duty-cycle, which is small - only 5% - to reduce the
>sensitivity losses due to the receiver gating.

29 Jun 02
Dear Ray,
>I am thinking about a COSY experiment. That should be an interesting one

Attached is the solvent suppression experiment. I excluded +-20 Hz (40 points) around the water signal from the excitation pulse. I wonder whether we should also exclude it from detection ?

Just to let you know that Cosy seems to be working just fine. I will be in London on Tuesday. Back in action on Wednesday.

>I did the solvent suppression experiment over the weekend. It seems to be
>working OK. Will send you some pictures soon.

I can't wait !

(At this stage we were already thinking about a two-dimensional version of the Hadamard method, but we decided that this would be best kept for a second paper. In retrospect this was the more important aspect and led to several subsequent papers on multidimensional Hadamard spectroscopy, speeding up conventional methods by orders of magnitude.)

29 July 02
Dear Eriks,

It looks good, but I agree, a bit of nulling in the acquisition stage would be nice.

I have revamped the draft quite a bit and wondered if you could take a close look at it. I am concerned that the "average" reader might find it a bit difficult to follow since we are breaking so many established customs. Would there be a point in running strychnine with (say) only 16 channels (0.25 ppm steps) to show that we can use the Hadamard method for direct data reduction (combinatorial chemistry) ? It would greatly speed up pattern recognition routines when searching a library. I was wondering whether NOESY might be more interesting than COSY ???????

30 July 02
Dear Eriks,

Good news about COSY. I look forward to seeing the results.
I found another mistake in my text. Aliasing does happen after all. Your pulse repetition rate (acquisition rate) has to be high enough to avoid aliasing the channel frequencies.
HOWEVER: There is another mode where we don't need the reference frequency for each channel. All we really need is the integral of the signal, relying entirely on the Hadamard coding to separate the channels. Take the absolute magnitude of the oscillating NMR signal and

integrate. This is a RECTIFIER but not a SYNCHRONOUS DETECTOR. No single-point Fourier analysis !
Then if a frequency is aliased it does not matter, so you could use a slower pulse repetition rate (and a lower duty cycle) if necessary.
The beauty of this method is that the HADAMARD separation is paramount -- there is no question that the ++++++++++++++++++ scan could be used alone. I was always a bit afraid of the possible criticism that the Hadamard trick is not strictly necessary.
The only caveat I can see is that the method of tacking together small Hadamard matrices would still need the reference frequency method, otherwise there would be no discrimination between similar channels in different blocks.
Could you send details of your choice of pulse repetition rate and overall duration of the polychromatic pulses you used ?
The duration determines the frequency selectivity, but it might be possible to operate with low selectivity (short duration) and still get away with it because of the Hadamard encoding ????
Long durations would improve signal-to-noise, but since we do so many scans this is hardly necessary. How do we choose the pulse duration ???? Let me know what you think.

31 July 02
Dear Eriks,
Looking forward to seeing your results.
I have been having second thoughts about my idea for tacking several (N by N) Hadamard matrices together. It may be nonsense. If there were M blocks we would need to separate M composite signals by a kind of Fourier transform process, and this would have to be repeated N times. I would appreciate your comments on practical feasibility.
Otherwise we would need a separate "layer" of encoding, perhaps using TRUE phase-encoding with angles not equal to 0 or 180. The only other way out (that I can see) is to perform these M blocks sequentially, and that sacrifices sensitivity per unit time.
This may be the only real drawback of the Hadamard method.
On the other hand, I think the "rectifier" mode is worth trying for its simplicity compared with the "synchronous detector" scheme.

31 July 02
Dear Eriks,
At the risk of boring you, I quote below a couple of paragraphs from the latest draft:

The NMR signal is acquired one data point at a time in the intervals between radiofrequency pulses (11); no free induction decay is involved. At this stage the response is a composite signal comprising all NMR frequencies within the excitation band, and in each scan these signal components have been phase-encoded in a different way. Decoding involves combining the results of all 2048 scans according to the appropriate column of the Hadamard matrix. In the simplified example of an 8 by 8 matrix shown above, the response from the second frequency channel can be selected by combining the composite signals from the first four scans with inverted versions of the signals from the last four scans according to the second column of this matrix. All other responses cancel, leaving only one NMR frequency.

After this decoding stage the signal in any given channel is monochromatic, with a uniform envelope except for a few early points where the spin system is still approaching its steady state. There are two ways to process this signal. The first is equivalent to employing a synchronous detector. Since each excitation channel is associated with its characteristic frequency F, this can now be used as a reference frequency to demodulate the signal in that channel. This involves multiplication of the NMR response by F, followed by integration -- equivalent to a single-point Fourier analysis. This first method is subject to aliasing if the data acquisition rate (equal to the pulse repetition rate) is not fast enough in comparison with the highest frequency component.

The second method is equivalent to the use of a rectifier rather than a synchronous detector. It does not require the reference frequency F, but simply integrates the absolute magnitude of the oscillating NMR response. Discrimination of responses from different channels then relies entirely on the Hadamard encoding scheme. An important advantage of this method is that aliasing of the NMR frequency can be ignored -- the integral is unchanged by down-conversion of the frequency. Slower pulse repetition rates can therefore be contemplated.

31 July 02
Dear Eriks,
You were right, Ping used:

hard 90 (+X), soft 180 (+Y) ACQ(+)
hard 90 (+X), soft 180 (+X) ACQ(+)

Then to convert into a COSY experiment just add a hard 90 pulse.
The two papers are: JMR 83, 404 (1989) for pinging
JMR 84, 198 (1989) for COSY

1 Aug 02
Dear Eriks,
I had second thoughts about trying to tack several Hadamard matrices together. Here is a revision of that section for your comments (attachment).
Any pictures yet ?

2 Aug 02
Dear Eriks,
I hope you can make COSY work. Meanwhile here are some suggestions for possible 1D experiments:

1. Strychnine run at low definition (say) 32 channels only, spaced 0.125 ppm apart or 16 channels spaced 0.25 ppm apart. I would like to show this as an example of data compression. We could restrict the vertical scale to (say) 4 bits to reduce the information content even further. Just a "fingerprint" of the kind that might be stored in a library. Cover the same range as Figure 1.

2. Homonuclear decoupling in strychnine with blanking of the strongly irradiated region. Possibly decouple the multiplet at 2.3 to 2.4 ppm. May need to restrict the displayed frequency range so we can see changes in spin multiplets more clearly.

3. High-definition partial spectrum of strychnine to illustrate "zooming", possibly the region between 4.0 and 4.2 ppm. Long pulse sequence, good shimming;
512 by 512 matrix ?

4. A plot of the oscillating response from methyl iodide (MeJ_fid.tif) with the same response "rectified" by setting all negative excursions positive. To illustrate our processing scheme.

5. Double-quantum response from the AB spectrum at 3.6 ppm in the inosine spectrum (ht_inos.tif). High-definition spectrum (zoom) of the region between 3.5 and 3.7 ppm with very intense pulses in a few channels near the mid-point of the AB pattern (3.60 to 3.62 ppm), but normal pulse intensities elsewhere. Should see several DQ transitions (4 ?). DQ spectra are VERY sensitive to irradiation intensity (proportional to B one cubed).

6. Diagonal-suppressed COSY using spin pinging. Depends on how much trouble you are having with ordinary COSY.

7. FT spectrum of strychnine for comparison with Hadamard spectrum.

If you can get some of these spectra they would be great to illustrate many of the main advantages of Hadamard spectroscopy.

2 Aug 02
Dear Ray,
>I am struggling a bit with the COSY as the signals are very weak. It might
>be that the '+' element in the Hadamard matrix needs be replaced with '+ -'
>and the '-' element then would be '-+'. It's not easy.
>I am also concerned about the pulses. While the 'all +' pulse is easy to
>fix with the phase correction trick, the other combinations are more
>difficult to deal with. Hope to send you some pictures later today.

>Attached are a few spectra for comparison :
>a) ft1p - one point FT
>b) rectify - rectifying the real part
>c) magn - sum of magnitudes
>d) max - maximum value
>e) lastPT - last point

>I will make some nice pictures over the weekend.

3 Aug 02
Dear Eriks,

Thanks for the spectra. It is not clear to me why (b) (c) and (d) give absolute value line shapes while (a) and (e) do not. Maybe the selectivity is not high enough ?
Perhaps we should use the "synchronous detector" method, but why ?
If we want to finish by the end of next week, let's drop the 2D stuff for the time being; we can always write a second paper. But could you get some nice clean (1:2:1 filtered) spectra printed out (not from the screen) ?

5 Aug 02
Dear Eriks,
I have been thinking about the broad lineshapes in recent spectra. I think that the Hadamard decoding alone does not discriminate the phase of the signal. Thus signal from channels near to the "exact resonance" channel pick up the long dispersion tails (in addition to any absorption). This makes the result essentially an "absolute-value" lineshape. Multiplication by the frequency "F" and integration is the equivalent of a phase-sensitive detector, and the way you have been doing it, the phase is adjusted for pure absorption. Hence much better lineshapes.
The "rectifier" method sounded like a good idea, but I think we should abandon it now.
The trickier question then arises: why are absorption lines observed with "first point" and "last point" methods ? My guess is that they are ultra sensitive to the phase (or frequency) of the oscillation. We are sitting on the peak of a sine-wave and any change in phase or frequency allows us to slip down the side. Anyway the sensitivity is bad in these two modes.
I will rewrite that part of the text and send it to you.

5 Aug 02
Dear Ray,
>Attached is a figure imbedded into the Word document. Please let me know
>whether you can read and see it.

>The fig is 3.6 Mb because I made it with good resolution. Attached is the
>compressed file. Let me know whether you can un-zip it.

>I have been thinking about lineshape too and have come to a similar
>conclusion. In FT the exponential decay directly transforms into Lorentzian
>lineshape with real and imaginary part. In our case there is no obvious
>connection to relaxation and one may wonder whether there should be
>lineshape at all. The result we see is probably a combination of relaxation
>and RF interference effects that are washed away using a phase sensitive
>detector.
>There is no particular magic about the last point, except that it is
>easier to phase (ideally should be in phase). Otherwise all the other
>points should be more less equivalent.

>Another try for Fig. 1. Can you print it?

>Attached is a low res figure just to show a potential Figure comparing

>strychnine HT spectra acquired with 8 Hz and 1 Hz resolution. Let me know
>what do you think.

5 Aug 02
Dear Eriks,

Here is a proposed paragraph about lineshapes; let me know if you agree.

After this decoding stage the signal in any given channel is monochromatic, with a uniform envelope except for a few early points where the spin system is still approaching its steady state (Figure 1). The Hadamard decoding process collects the total NMR signal in each channel but makes no distinction about the phase of this signal. Consequently nearby "off-resonance" channels pick up appreciable dispersion-mode signal components, with the result that the observed line shapes have long tails, as in an absolute-value display. The remedy is to employ the equivalent of phase-sensitive detection. Since each channel is associated with its characteristic excitation frequency F, this can now be used as a reference frequency to demodulate the signal in that channel. This involves multiplication of the NMR response by F, followed by integration -- equivalent to a single-point Fourier analysis. The phase of this reference signal is adjusted for pure absorption and the long dispersion tails consequently disappear.

Do you think we can finish this week ????????????????????????????????????
If not, will you be able to do some experiments in Palo Alto ?

We have to stop thinking in terms of Fourier transforms. The old-fashioned frequency sweep continuous-wave method gave the correct line shapes, provided that the sweep was slow enough. Multichannel irradiation monitors the same line shape. We observe the steady-state solutions to the Bloch equations rather than the transient solutions. Spin-lattice relaxation shows up as a saturation effect. Spin-spin relaxation determines the natural line width. Inhomogeneity determines the observed line width.
Your "time-domain signal" has nothing in common with a free induction decay. It is not really a transient signal but (almost) a steady-state response. In principle you could go on monitoring it forever . . . it does not decay.

I think I need more detail about the traces (a), (b) and (c) in Figure 1. I'm not sure what you mean by "Hadamard dimension" for example.
However it does seem to encapsulate our experiment and maybe we can go to press with just this (a "Communication") ??? But I would like to add solvent suppression and maybe some other applications (decoupling, double-quantum . . .)

5 Aug 02
Dear Eriks,

Yippee. Plots fine. What did you do ? I notice an antisymmetry about the centre of Figs 1(a), (b) and (c). Any comments ? The HT and FT spectra are beautiful.
(a) What exactly do you plot for the polychromatic pulse, and which one out of all 512 ?
I assume (b) is the NMR signal from a typical scan (which one ?).
I assume (c) is the NMR signal from a typical channel (Hadamard matrix column) ? Which channel ?

I think you are right -- there seems to be a slight frequency-dependent phase shift giving the negative glitches. Timing ?

>Just to let you know that the glitch turned out to be a folded peak. I am
>using presat to remove it.

You might even try 16 Hz resolution and/or a vertical dynamic range of one bit (zero or full intensity) to reduce the information content still further.
Maybe do it in two stages, 1Hz to 4 Hz to 16 Hz ???

6 Aug 02
Dear Eriks,
Here is my first attempt to write figure captions (attached). Comments ?

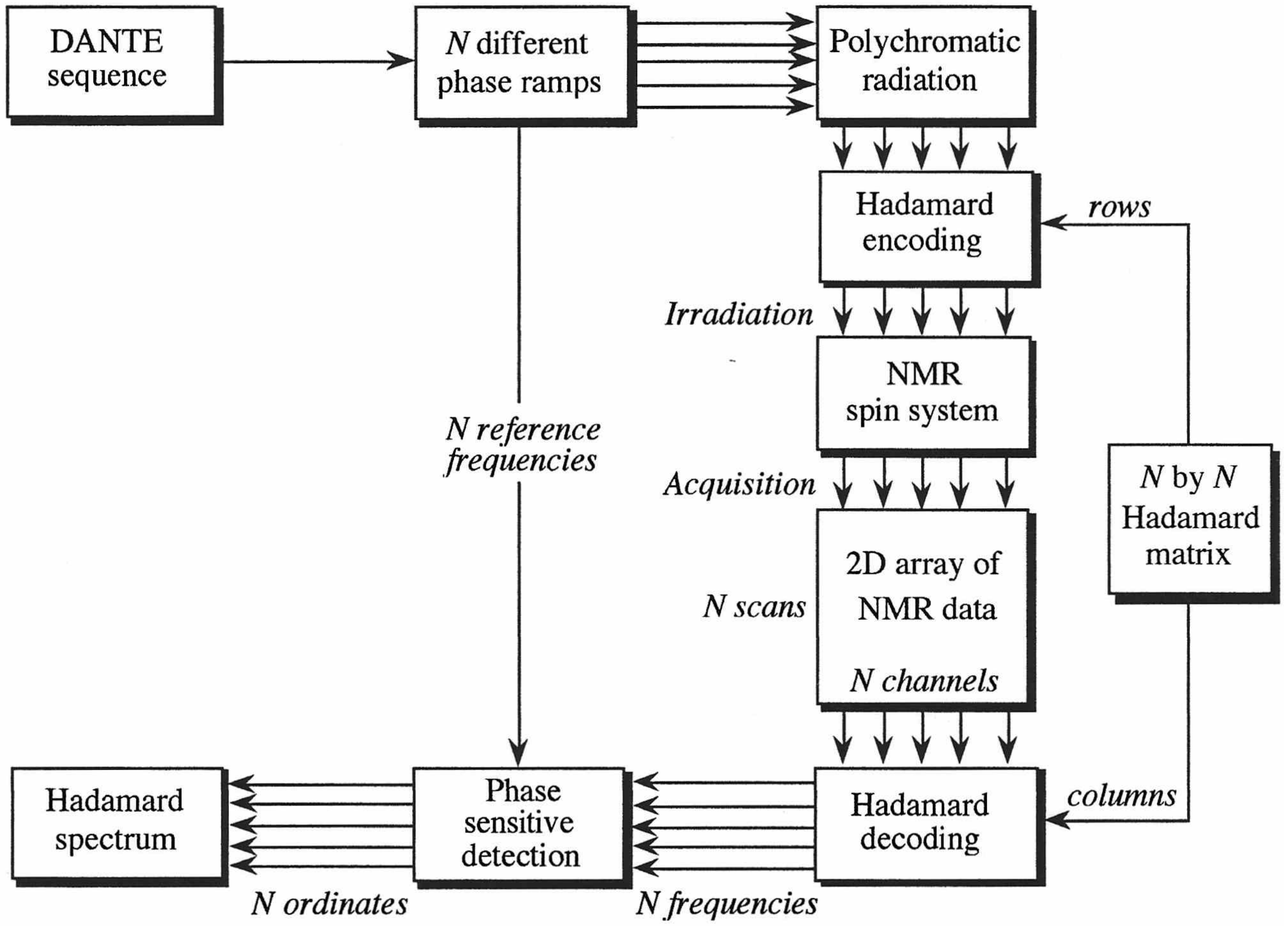

7 Aug 02
>Dear Ray,
>The Hadamard trace is Fig 1c is from the central line of the triplet at
>7.57 ppm in the aromatics region of 2-Ethyl-1-indanone spectrum.
>I am looking at different flip-angles today.
>My PC refuses to recognize MacDraw. Can you save the diagram as a
>PostScript file? Just trying to establish a common language for diagrams.
>Will send you the final strychnine spectra later today.

>Here is a hypothetical Hadamard TOCSY experiment.
>In a classical selective TOCSY experiment we prepare the initial condition
>by inverting one of the spins in alternating scans and then get the final
>spectrum by taking the sum of the two experiments. For a two spin system
>this corresponds to a 2x2 Hadamard matrix:
>+ +
>+ -
>The mixing period converts it into
>+ +
>- +
>After the Hadamard transform we get the "spin density" matrix
>0 2
>2 0
>If the spins were not coupled we would have a diagonal spectrum:
>2 0
>0 2
>This translated directly for more complicate spin systems (I suppose).
>Note that the phase cycling is essentially built into the Hadamard
>experiment. We just need to store all the experiments separately instead of
>adding the alternating scans.
>We could try this on strychnine with a 16x16 Hadamard matrix, which would

>be a great simplification as compared to the regular 2D experiment.
>What do you think ?

(Already we are spending more time on the applications of Hadamard spectroscopy to two-dimensional experiments, where the aim is to speed up data gathering)

>Attached is Fig. 3. Let me know if it requires any changes.
>I have received your first block diagram. It looks OK. Perhaps we need to
>somehow illustrate the point that the same Hadamard matrix is used for both
>en-coding and de-coding.

8 Aug 02
Dear Eriks,
Thanks for the two diagrams. What I would like to show for Figure 3 is the high-resolution spectrum (a) and a modified version of the low-resolution spectrum (c). Delete (b). The modification would convert the vertical scale to 0 or 1. rather like taking the scissors to the present spectrum at about the 30% level. OK, it doesn't preserve the integrals but it should show clearly which ppm ranges carry signals. Could you do this ?
I had hoped you would receive the second version of the block diagram where I have indeed included a Hadamard matrix (and made some other changes).
For the moment I am confused about your statement that in Figure 2, trace (b) is the Fourier transform of the response from (a).
I may need to make further changes to the second block diagram with respect to the NMR data array. I think I made an error there.
Anything on the water suppression result ?

8 Aug 02
>Dear Ray,
>Attached is the corrected version of Fig. 2. The trace (b) was indeed the
>Fourier transform of the response from (a). The rest is unchanged.

DO YOU REALLY PERFORM A FOURIER TRANSFORM ?

>The answer is no.
>It was Fourier transform in the previous version of Fig. 2 (I used a wrong
>command and saved the Fourier transformed trace instead of the time-domain
>signal). In the new Fig. 2 this is corrected and the trace now shows the
>time-domain signal, as it was supposed to.

GOOD. I THOUGHT I HAD LOST THE PLOT . . .

>I will change the Fig. 3 as you suggested. Do you also want to include the
>FT spectrum or would it be just two Hadamard spectra - high resolution and
>low resolution, uniform intensity?

JUST TWO HADAMARD SPECTRA. WHO NEEDS FT ?

8 Aug 02
Dear Eriks,
TOCSY looks very promising (I just knew you couldn't resist trying it). It certainly would strengthen our arguments about the usefulness of low-definition Hadamard spectra but could we save the detailed 2D stuff until the second paper?

8 Aug 02
Eriks: Modified block diagram embedded in a Word 6 document.

Here is the tricky section in a new draft (attached). Let me know of any "howlers".

8 Aug 02
Dear Ray,
>Attached is the first attempt to produce the Hadamard TOCSY - "quick and
>dirty". It used 14 frequencies that were generated automatically from the

>low resolution strychnine spectrum. The 16 scan experiment (16x16 Hadamard)
>matrix was really quick and gives all the connections. Looks promising.

>... with 1 scan per increment it takes 40 seconds to complete !!!!

8 Aug 02
Dear Ray,
>Attached please find the modified figures. Pleas let me know if they need
>any further polishing.

9 Aug 02
Dear Eriks,

Great diagrams. Just one minor thing. In the data reduction trace (Figure 3b) would it be possible to go the whole way to a one-bit intensity scale ? Anything below 20% intensity put to zero. Should give a completely flat baseline and steep-sided responses. Might merge the two multiplets at 4.1 ppm but who cares ? I know how hard it is to throw away information but it is in a good cause.

9 Aug 02
FROM TOSHI, FOR INFORMATION (*Forwarded from Dr Toshiaki Nishida, Stockholm*)
>Dear Ray;
>We wrote two papers;
>a) T. Nishida, G. Widmalm and P. Sandor, Hadamard Long-Range Proton-Carbon
>Coupling Constant Measurements with Band-Selective Proton Decoupling, Magn
>Reson Chem 1995, 33, 596 - 599.
>b) T. Nishida, G. Widmalm and P. Sandor, Hadamard Long-Range Proton-Carbon
>Coupling Constant Measurements with Pulsed Field Gradients, Magn Reson Chem
>1996, 34, 377 - 382.
>Rather recently I became aware of the following paper;
>Jan Schraml, Herman van Halbeck, Andre De Bruyn, Roland Contreras, Marleen
>Maras, and Piet Herdewijn, Hadamard 1D 1H TOCSY and its Application to
>Oligosaccharides, Magn Reson Chem 1997, 35, 883 - 888.
>They referenced "Hadamard" to your paper (with V. Blechta) in Chem Phys
>Lett 1993, 215, 341 and your Book, "Spin Choreography", pp. 177 -183.
>Krish Krishnamurthy, Hadamard Excitation Sculpting, J Magn Reson 2001, 153,
>144 - 150. Maybe this is the latest application?
>I use Hadamard for band selective gHMBC (a single band selective work was
>originally shown by Jean_Marc Nuzillard -- He had a poster at "Spin
>Choreography Symposium" and then published later in J Magn Reson 1999, 139,
>454 - 459). The paper by Krish above shows Hadamard band selective gHSQC.
>We feel Hadamard band selective gHMBC is more practically useful.

(This brought us up-to-date with the literature on related Hadamard experiments)

9 Aug 02
Dear Ray,
>Attached are Figures 3 and 5. In Fig. 5 trace (a) is FT spectrum and trace
>(b) is Hadamard spectrum. Let me know if any changes required.

9 Aug 02
Dear Eriks,

Excellent Figures ! Is the strychnine "going off" ? Seems to have a weak impurity subspectrum.

Forgot to mention the magic word ! What is model name of your Varian spectrometer ?

Attached is the most recent draft. Was it an Inova spectrometer ?

Have a relaxed weekend ! Attached is a proposed cover letter for JMR.

Professor Stanley J. Opella,
Editor, Journal of Magnetic Resonance,
Editorial Office,
525 B Street,
San Diego,
California 92101-4495,
USA.

12 August 2002

Dear Stan,

We would like to submit the attached manuscript "Frequency-Domain Hadamard Spectroscopy" by Kupce and Freeman for publication as a regular article in the *Journal of Magnetic Resonance*. This describes a new way to excite high-resolution NMR spectra and many of the accepted ideas of Fourier transform NMR do not apply. So I hope you can find referees that will appreciate this.

Hard copies and a diskette are enclosed.

Now that I am formally retired, the College address (above) is more appropriate than the University Chemistry Department.

Best wishes,

Ray Freeman

****

12 Aug 02
Dear Eriks,

I think I shall submit the paper tomorrow (Tuesday).

(On 19 August 2002, our paper was received by the Journal of Magnetic Resonance, and was published in the May issue of 2003 -- J. Magn. Reson. ***162****, 158-165 (2003) with the two-dimensional idea mentioned in the final paragraph of that paper. There were delays in finding suitable referees. The published version is attached here. Subsequently we wrote three more papers and a review on multidimensional Hadamard spectroscopy)*

ACADEMIC
PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Magnetic Resonance 162 (2003) 158–165

JMR
Journal of
Magnetic Resonance

www.elsevier.com/locate/jmr

Frequency-domain Hadamard spectroscopy

Ēriks Kupče[a] and Ray Freeman[b,*]

[a] Varian Inc., Eynsham, Oxford, UK
[b] Jesus College, Cambridge University, Jesus Lane, Cambridge CB5 8BL, UK

Received 19 August 2002

Abstract

A new technique is proposed for multichannel excitation and detection of NMR signals in the frequency domain, an alternative to the widely used pulse-excited Fourier transform method. An extensive array of *N* radiofrequency irradiation channels covers the spectrum of interest. A selective radiofrequency pulse sequence is applied to each channel, generating a steady-state NMR response acquired one-point-at-a-time in the intervals between pulses. The excitation pattern is repeated *N* times, phase-encoded according to a Hadamard matrix, and the corresponding *N* composite responses are decoded by reference to the same matrix. This multiplex technique offers the same sensitivity advantage as conventional Fourier transform spectroscopy. The irradiation pattern may be tailored to concentrate on interesting spectral regions, to facilitate homonuclear double resonance, or to avoid exciting strong solvent peaks. As no free induction decay is involved, the new method avoids problems of pulse breakthrough or lineshape distortion by premature termination of the time-domain signal.

Keywords: Frequency-domain spectroscopy; Multichannel excitation; Hadamard spectroscopy; Multiplex advantage; Solvent peak suppression

1. Introduction

The frequency domain is the natural mode for observing NMR spectra, and the spectroscopist always visualizes high-resolution spectra in this dimension. Time-domain acquisition followed by Fourier transformation is essentially a roundabout process for achieving the same end. For many years the continuous-wave slow-passage method was universally accepted, until it was seen to be inefficient from the point of view of sensitivity. To counter this deficiency, Anderson [1] set out to explore the possibility of a multiplex spectrometer which employed a regular array of "channels"—each defined by a stationary excitation frequency and a matching synchronous detector. In principle, this provides a multiplex advantage [2] that improves the signal-to-noise ratio by a factor of $\sqrt{N}$, where N is the number of independent channels. Now, four decades later, Anderson's mechanical devices for generating a regular "comb" of closely spaced excitation frequencies and for demodulating the resultant signals may seem a little cumbersome, but the idea nevertheless served as the spur to the invention of the pulse-excited Fourier transform method [3] which revolutionized the practice of high-resolution NMR. The Anderson "Prayer Wheel" was quietly abandoned and consigned to the Smithsonian Institution in Washington, DC.

The extraordinary success of the Fourier transform technique has left very few viable alternatives for recording NMR spectra of high information content. The inherently poor sensitivity of magnetic resonance has been largely redressed by the orders of magnitude improvement afforded by pulse excitation and Fourier transformation. The few remaining shortcomings are generally accepted as inevitable; the goose is still laying the golden eggs. The main practical disadvantage is the serious dynamic range problem associated with aqueous solutions, a widespread limitation that has inspired the invention of an entire armoury of techniques for solvent peak suppression. Furthermore, wide-band excitation with a hard radiofrequency pulse imposes certain restrictions on the way NMR spectra are recorded by the conventional Fourier transform method. It is difficult to

---

[*] Corresponding author. Fax: +44-1223-336362.
E-mail address: rf110@cam.ac.uk (R. Freeman).

1090-7807/03/$ - see front matter © 2003 Elsevier Science (USA). All rights reserved.
doi:10.1016/S1090-7807(02)00194-5