Published in final edited form as:

J Magn Reson. 2013 December ; 237: 100–109. doi:10.1016/j.jmr.2013.09.005.

Accelerating multidimensional NMR and MRI experiments using iterated maps

Merideth A. Frey^{a,1}, Zachary M. Sethna^{b,1}, Gregory A. Manley^c, Suvrajit Sengupta^{c,2}, Kurt W. Zilm^c, J. Patrick Loria^d, and Sean E. Barrett^{a,*}

^aDepartment of Physics, Yale University, New Haven, CT 06511

^bDepartment of Physics, Princeton University, Princeton, NJ 08540

^cDepartment of Chemistry, Yale University, New Haven, CT 06511

^dDepartment of Chemistry and Department of Molecular Biophysics and Biochemistry, Yale University, New Haven CT 06511

Abstract

Techniques that accelerate data acquisition without sacrificing the advantages of fast Fourier transform (FFT) reconstruction could benefit a wide variety of magnetic resonance experiments. Here we discuss an approach for reconstructing multidimensional nuclear magnetic resonance (NMR) spectra and MR images from sparsely-sampled time domain data, by way of iterated maps. This method exploits the computational speed of the FFT algorithm and is done in a deterministic way, by reformulating any *a priori* knowledge or constraints into projections, and then iterating. In this paper we explain the motivation behind this approach, the formulation of the specific projections, the benefits of using a ‘QUasi-Even Sampling, plus jiTter’ (QUEST) sampling schedule, and various methods for handling noise. Applying the iterated maps method to real 2D NMR and 3D MRI of solids data, we show that it is flexible and robust enough to handle large data sets with significant noise and artifacts.

Keywords

sparse sampling; iterative maps; multi-dimensional nuclear magnetic resonance; magnetic resonance imaging

1. Introduction

For more than 40 years, pulsed Fourier transform spectroscopy [1] has been the dominant approach to nuclear magnetic resonance (NMR), and more recently, to magnetic resonance

© 2013 Elsevier Inc. All rights reserved.

*Corresponding author. sean.barrett@yale.edu (Sean E. Barrett).

¹Both authors contributed equally to this paper.

²Present Address: Department of Chemistry, University of California, Irvine, Irvine, CA 92697

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

imaging (MRI). The fast Fourier transform (FFT) algorithm has enabled rapid computation of the spectra obtained from MR experiments, even as they were extended to 2D, 3D, and higher dimensions [2]. Practically speaking, the bottleneck in exploiting the higher information content of multidimensional spectra is the time required to acquire the MR data. For example, in most 2D NMR experiments, only a single 1D ‘row’ of N uniformly-spaced discrete samples of the time-dependent, complex signal is measured per $\propto T_1$ (spin-lattice relaxation time). Filling the densely-sampled 2D Cartesian grid of $M \times N$ data points requires M additional experiments, for a total acquisition time $T_{\text{acq}} \propto M \times T_1$. One way to accelerate the 2D experiment is to measure only some of the M rows (i.e., to use sparse sampling). Unfortunately, this results in a low quality spectrum, unless non-Fourier reconstruction techniques such as l_1 minimization [3], entropy maximization (e.g., MaxEnt [4, 5], FM [6], or MINT [7]), or multi-dimensional decomposition (e.g., MDD [8]) are used [9, 10, 11, 12].

We became interested in this problem while developing a 3D MRI of solids technique, since the long T_1 of phosphorus-31 in bone meant that one dense data set required $T_{\text{acq}} > 45$ hours [13]. Taking the first step towards sparse MRI, or compressed-sensing (CS) MRI [14, 15], k space was sparsely-sampled in a pseudo-random way for $T_{\text{acq}} < 1.5$ hours. We then Fourier transformed this data, yielding an image of modest quality (Figure S3 in [13]) *but in much less time*. Typically, the next step to improve upon this image is to adopt non-Fourier reconstruction algorithms like l_1 minimization [16, 17, 18]. However, these methods were not readily applied to our 3D MRI of solids data, since the data sets were sizeable (e.g. $64 \times 64 \times 64$ complex points) and the spectra suffered from both large artifacts and relatively poor signal-to-noise ratios. We sought an easy-to-implement, familiar, and robust reconstruction method that exploited the computational speed of the FFT algorithm.

In this paper we demonstrate the use of iterated maps to accelerate MR data acquisition. In this approach, a speedily-acquired sparse data set is deterministically converted into a high quality approximation of the full, dense data set. We show several examples drawn from 2D NMR of liquids, 2D NMR of solids, and 3D MRI of solids. The approximate output spectra are in excellent agreement with the target, dense spectra over the full dynamic range, for positive and negative signals. This approach is built upon the FFT, so it has the computational speed to handle large multidimensional data sets. As we show below, the time-savings afforded by this method can be used to improve signal-to-noise and frequency resolution compared to spectra obtained from conventional Fourier reconstruction of dense data sets.

Our iterated maps approach was inspired by Veit Elser’s use of the iterated ‘Difference Map’ (DM) algorithm [19, 20] to determine the phase of complex signal when only the magnitude is measured, given a set of *a priori* constraints that the Fourier reconstruction of the object must satisfy. We realized that Elser’s approach could be adapted to the MR problem by filling in the ‘missing’ complex points in a sparsely-sampled MR data set to approximate the dense MR data set. Since both the magnitude and phase are measured in MR, the MR problem may be solved using fewer input data points than Elser’s phase-retrieval problem. While preparing this paper, we became aware of the pioneering work of Herzfeld and Matsuiki, who recently developed SIFT (Spectroscopy by Integration of Frequency and Time Domain Information) [21, 22, 23] to accelerate multidimensional NMR

experiments. Our work is related to theirs, and both approaches may be traced back along distinct lineages to a common ancestor (see Supporting Information). The conclusions drawn independently by both groups are in great agreement. In the discussion that follows we will point out similarities and differences between SIFT and our iterated maps approach.

2. Results

Our approach offers speed-up advantages when signals occupy a limited region of the frequency domain, as we show using 2D NMR of liquids (Figures 1–3), 2D NMR of solids (Figure 4), and 3D MRI of solids (Figure 5). For each demonstration of the method, we start by describing the data sets that have been densely-sampled in the nD time-domain as a complex vector $\vec{T}(t)$; the corresponding target ‘dense spectrum’ $\vec{T}(f) = Ph(FFT(\vec{T}(t)))$ is obtained following nD FFT and a known phase correction Ph , which are reversible operations: $IFFT(Ph^{-1}(\vec{T}(f))) = \vec{T}(t)$ (see Supporting Information). In practice, our algorithm acts upon $\vec{T}(t)$ and $\vec{T}(f)$ in their natural form as nD matrices.

As explained in the Supporting Information, even though the causal data set acquired using a ‘States’-like 2D NMR experiment (i.e., with $t_1 = 0, t_2 = 0$) can be converted into the conventional, purely-real (absorptive) 2D spectrum using ‘States’ processing [24], this cannot be easily reversed, which is a requirement of our iterated maps method. Instead, we followed the MRI model and used the acquired data to construct a $\vec{T}(t)$ that filled all four quadrants of the time-domain, with Hermitian symmetry about the origin, such that each component of the 2D matrix satisfies $T(t) = T^*(-t)$. A 2D complex Fourier transform (followed by a phase correction) yields a purely-real spectrum, $\vec{T}(f)$. The pseudo-echo transformation in 2D NMR [2] has a similar motivation.

To simulate the effect of skipping particular experiments, we construct an initial ‘sparsely-sampled’ data set $\vec{S}^0(t) = P_0 \vec{T}(t)$, where P_0 is a “projection” that sets to zero all skipped points, leaving the rest alone. In this paper, we follow Elser’s use of the term projection [20], as explained in Supporting Information. For example, Figure 1A shows $\vec{S}^0(t)$ if just 50.8% of the required t_1 rows are measured using States acquisition, which halves the normal T_{acq} . Unfortunately, the real part of the corresponding spectrum

$\tilde{S}_r^0(\vec{f}) = Re(Ph(FFT(\vec{S}^0(\vec{t}))))$ is a poor approximation to $\vec{T}_r(\vec{f})$, with pronounced sparse-sampling artifacts that smear along the f_2 columns, as shown in Figure 1B. To quantify the disagreement, Figure 1C (red circles) shows a point-by-point comparison of $\tilde{S}_r^0(\vec{f})$ to $\vec{T}_r(\vec{f})$ along six f_2 columns of Figure 1B, which shows that large signals in $\vec{T}_r(\vec{f})$ have only ~ 50.8% of their dense value in $\tilde{S}_r^0(\vec{f})$, while points that should be zero in the dense spectrum can have large positive or negative values in $\tilde{S}_r^0(\vec{f})$. To do better than this, we use iterated maps to deterministically convert $\vec{S}^0(f) \rightarrow \vec{F}^n(f) \approx \vec{T}(f)$.

Similar to SIFT, our iterated maps approach uses two projections which enforce *a priori* knowledge about the dense data in the time and frequency domains. The first projection (P_1) is a support constraint in the frequency domain. The simplest version (P_{1SIFT} as is used by SIFT) zeros all points outside of the support region and leaves the points inside the support intact. As we show in Figure S1, we further strengthen this constraint by ensuring proper

phasing of the time-domain data, so the spectrum $\vec{T}(f)$ in our stronger version of $\hat{P}_1$ has no imaginary part, and the real part has a known sign (+ in positive, – in negative, and $\pm$ in artifact support regions). The second projection ($\hat{P}_2$) resets all sampled time-domain points back to their measured values (i.e., the non-zeroed points in $S^0(t)$), as in SIFT. Both of these projections are altered slightly to allow for noise-handling with given error thresholds in one or both domains (see Supporting Information).

We combine these projections (and the identity operator $\mathbf{1}$) in a particular form of Elser's Difference Map algorithm [19, 20] which reproduces Fienup's hybrid input-output map [25]. Our map is given by $D = \hat{1} + \hat{P}_1(2\hat{P}_2 - \hat{1}) - \hat{P}_2$ (while SIFT uses $D_{SIFT} = \hat{P}_2\hat{P}_1$) and the iterative mapping becomes $S^{n+1} = DS^n$. The $\hat{P}_1$ and $\hat{P}_2$ projections are defined in the frequency domain and time domain, respectively, and we move from one domain to another using the FFT/IFFT (see Supporting Information). The final output of the algorithm after n iterations is given by $F^n = P_2 S^n$. For example, Figure 1D shows $F^{15}(t)$, and Figure 1E shows $\tilde{F}_r^{15}(\vec{f})$, which were calculated in six minutes on an iMac, using code running in IgorPro. Figure 1C,E,F show the excellent agreement between the dense $\vec{T}_r(f)$ and reconstructed $\tilde{F}_r^{15}(\vec{f})$ spectra. For example, Figure 1C (blue triangles) shows a point-by-point comparison of $\tilde{F}_r^{15}(\vec{f})$ to $\vec{T}_r(f)$ along six f_2 columns of Figure 1E. Clearly, all of these $\tilde{F}_r^{15}(\vec{f})$ points (not just the spectral peaks, and not just the positive signals) have reached $\approx 100\%$ of the target value in $\vec{T}_r(f)$. The Supporting Information provides additional details of the algorithm, projections, convergence, and noise-handling.

The successful demonstration in Figure 1 is for liquid-state 2D NMR of LuxU [26], where the dense spectrum $\vec{T}_r(f)$ had relatively little noise, and only a narrow water 'artifact'. To a first approximation, the $\hat{P}_1$ mask was constructed by placing a coarse bounding box on the 1D spectrum as a function of either f_1 or f_2 , and was then refined by placing negative signal masks at particular locations. A small noise-handling threshold ($\pm 0.3\%$ of max signal) was used for $\hat{P}_1$.

In the time-domain, Figure 1A shows the 'QUasi-Evenly Spaced, plus jiTter' (QUEST) sampling schedule used to pick the particular t_1 rows used by $\hat{P}_0$ (see Supporting Information for more information). In its current form, the iterated maps approach can be used with any on-grid, non-uniform sampling (NUS) schedule (for a review, see [9]), and we studied many variations to assess their performance. However, in our experience, the iterated maps approach converges to essentially the same $\vec{F}^n(f) \approx \vec{T}(f)$ for all cases, *provided that enough samples are included in $S^0(t)$* . Lowering the number of samples (i.e. making the problem under-constrained) for each case revealed a common trend; the final $\vec{F}^n(f)$ started to deviate from $\vec{T}(f)$ whenever the gaps between the sampled points in $S^0(t)$ grew too large. Large gaps between samples were also observed to lower the quality of forward maximum entropy (FM) reconstruction [6]. We have a qualitative explanation of this effect, at least for our approach, based upon how it seems to work. Iterating the map causes information to diffuse from the sampled points in the time domain to particular unmeasured points, as determined by the point-spread function for a given $\hat{P}_1$ (as shown in Figure S1). In practice, if the gap between measured points in $S^0(t)$ is too large, it may never be filled in by this

process in the under-constrained case, and the final $\vec{F}^n(f)$ will differ from $\vec{T}(f)$. This picture suggests that QUEST is a good way to distribute a given number of sparse samples across the time domain, since the gaps between measurements are roughly uniform, and so it is used for Figures 1–4.

Figure 2 shows the method applied to another liquid-state 2D NMR data set. Using an Isoleucine, Leucine, Valine (ILV) ^{13}C -methyl labeled sample of imidazole glycerol phosphate synthase (IGPS), we were able to very accurately reconstruct the entire spectrum (Figure 2B–C) over a wide dynamic range, starting with just 58.6% of the time-domain data. In this case, no noise-handling is used for either $\hat{P}_1$ or $\hat{P}_2$. Note that the artifact domain (surrounded by blue in Figure 2A) is much larger than that in Figure 1B. In addition, this $\hat{P}_1$ mask has no negative support regions, and the positive support regions (surrounded by black in Figure 2A) are the portions of the dense spectrum with 0.7% of the maximum dense signal (MDS). Constructing such a tight mask would typically require the full dense spectrum, which is not available in all situations. However, in a serial experiment such as an NMR relaxation rate measurement, one dense data set (at maximum signal amplitude, to optimize the mask) could be followed by many sparsely-sampled ($x\%$) data sets acquired at various relaxation delays, requiring only $\sim x\%$ of the normal experimental time, as was recently demonstrated using SIFT [23].

While analyzing the IGPS data, we discovered that using the known frequency-space $\hat{P}_1$ mask along with the QUEST schedule enabled us to make very powerful predictions of just how sparsely we could sample the data. In Figure 3A we plot the Euclidean distance between the reconstructed and dense spectra ($l_2[\vec{F}^n(f) - \vec{T}(f)]$) for two different $\hat{P}_1$ masks, plotted versus the number of positive t_1 rows sampled (N_{t_1}). As expected (see Supporting Information), a looser mask (red curve) results in a larger difference between the reconstructed and dense spectra (larger l_2) than a tighter mask (blue curve), at each N_{t_1} .

To understand the striking peaks in Figure 3A at $N_{t_1} = 65$ and 85, we pick a $\hat{P}_1$ mask and study the aliasing that results from uniform undersampling (since this nicely approximates the artifacts from QUEST, see Supporting Information). Figure 3B shows the aliasing of the loose mask expected for $N_{t_1} = 85$. Zooming into the black dashed line region (in Figure 3C) shows that many strong peaks push into the support region (highlighted in yellow) as they are aliased from above (red) and below (blue). These aliasing artifacts will not be suppressed by our $\hat{P}_1$ projection when the loose mask is used, resulting in a worse reconstruction (larger l_2). However, when we do the same analysis with the tight mask (see Figure 3D), we see that most aliased peaks miss the support region so our $\hat{P}_1$ projection will efficiently suppress these artifacts, leading to a better reconstruction and the disappearance of the $N_{t_1} = 85$ peak for the tighter mask (red) data in Figure 3A. In this picture, if strong aliased peaks push into the support of the central mask (as N_{t_1} is lowered), then reconstruction quality suffers. The aliased masks are a convenient proxy for the aliased peaks, and they can be used to quickly estimate the minimum N_{t_1} for excellent reconstruction using QUEST. To do this, one can plot the overlap of the aliased masks with the central mask while lowering N_{t_1} ; all overlaps close to zero will yield excellent reconstructions.

Building upon this understanding, Figure 4 shows a very accurate reconstruction of a solid-state 2D NMR spectrum from a NCGB1 sample, despite starting with just 34.3% of the time-domain data and no noise-handling. This is a very large (4096×4096 complex points) 2D data set, with sparse spectral features (including artifacts of the magic angle spinning) that span a wide range of amplitudes. For the most-part, the $\hat{P}_1$ positive support region is a coarse set of blocks, as we used in Figure 1. In a crucial refinement, a ‘keyhole’ was cut into a signal-free region of the mask, to make room for an aliased rotational sideband (see Figure 4A). Adding the tiny keyhole drove the minimum N_{t1} from $1100 \rightarrow 702$ ($53.7\% \rightarrow 34.3\%$), while maintaining excellent quality of reconstruction over a wide dynamic range, as seen in Figure 4B–F.

In Figure 5, our reconstruction algorithm is applied to sparsely sampled 3D ^{31}P MRI of solids data (see [13] and Supporting Information for more information on how this data was acquired). In this experiment, the image is restricted to a portion of the full field of view (FOV) in order to avoid known artifacts (see Figure S2), which enables the construction of strong constraints in the $\hat{P}_1$ mask (where k and r in MRI map onto t and f in NMR). The dense data set for the image in Figure 5A and B took 35.2h to acquire. For Figure 5C, we sparsely sampled the time domain to reduce the required imaging time to 17% of the value for the dense data set (see Figure S3 and Supporting Information), which also lowered the reconstruction quality. To quantify the disagreement between the sparse $\tilde{S}_r^0(\vec{r})$ and dense $\tilde{T}_r(\vec{r})$ images, the signal amplitudes inside the positive support region are compared point-by-point in Figure 5C(inset). However, the agreement improves (Figure 5D) if we use our difference map reconstruction algorithm, with noise handling in both the $\hat{P}_1$ and $\hat{P}_2$ projections. Figure 5D(inset) shows the point-by-point comparison of signal amplitudes for our reconstruction $\tilde{F}_r^{20}(\vec{r})$ versus the dense $\tilde{T}_r(\vec{r})$ image, which looks good despite the ($\approx 10\%$) noise level of the dense spectrum. Considering the large noise and artifacts in this 3D MRI of solids data, this provides further evidence of our reconstruction algorithm’s robustness and usefulness for noisy experimental data.

In the previous results, we aimed for nearly ideal output spectra where $\tilde{F}^n(\vec{f}) \approx \tilde{T}(\vec{f})$. In some cases, a reduced quality output spectrum $\tilde{F}^n(\vec{f})$ may be sufficient for the user’s needs, which means that even less data (and a smaller T_{acq}) is needed. Of course, in order to determine how much data is required for the task at hand, one first needs to understand how the output spectrum is affected by using less data at the input. As we pushed down to even lower sampling percentages for 2D NMR, we found that QUEST helped us to identify the particular regions of the output spectrum that degrade first, due to aliasing artifacts (see Figures 6–9 and Supporting Information). Using the LuxU data as an example, the same panels A–F found in Figure 1 are shown again for different row numbers: $N_{t1} = 65$ (Figure 6), $N_{t1} = 50$ (Figure 7), $N_{t1} = 35$ (Figure 8), $N_{t1} = 20$ (Figure 9). In addition, panel G shows how the output vector becomes more ‘parallel’ and less ‘perpendicular’ to the target vector when the algorithm works as desired (see Supporting Information). Lastly, panel H shows a plot in the style of Figure 3C, showing the central LuxU mask (dark yellow), the $-1 \times BW_1$ mask (dark red) poking in from above, and the $+1 \times BW_1$ mask (dark blue) poking in from below, as explained in the Supporting Information.

Comparing Figure 6–9, we can summarize how the algorithm seems to work, and what happens as the input data is reduced. The red points in panel *C* show that most large features start at only $\approx x\%$ of their dense value, where $x\%$ is the sparse sampling percentage; at the same time, pixels that should be zero in the dense signal are non-zero, due to the artifacts of sparse sampling. As the algorithm iterates, the artifacts are driven towards zero, while the true signals push up toward their dense values. The total area under the signal is a conserved quantity, since we always include the data point at $t = 0$. Panel *H* shows that artifacts (blue features) that survive to the end of the algorithm are located in the portions of the frequency domain where the aliased and central masks overlap. The consequence of artifact survival is a poorer output quality, since it is also correlated with additional undershoot (red features) of the true signal amplitude in panel *H*. For clarity, panel *H* only shows the location of the $M = 1$ aliases at specific N_{t1} values, but the formulas in the Supporting Information indicate when $M = 2$ and $M = 3$ aliases will start to matter. For example, at $N_{t1} = (65/2) \approx 33$ the $M = 2$ aliases will just start to poke into the central mask (as shown in Figure 6*H* for $M = 1$). On the other hand, at $N_{t1} = (50/2) = 25$ the $M = 2$ aliases will just touch in the central mask (as shown in Figure 7*H* for $M = 1$). Finally, at $N_{t1} = (65/3) \approx 22$ the $M = 3$ aliases will just start to poke into the central mask (as shown in Figure 6*H* for $M = 1$). These ‘special’ numbers are consistent with the trends in Figure 6–9, with reasonably high quality fits from $N_{t1} = 65$ down to $N_{t1} = 35$, and a noticeably lower quality fit for $N_{t1} = 20$.

3. Discussion

SIFT can handle even ‘unphased’ input data, which is a clear advantage whenever the experimental parameters are difficult to change. On the other hand, whenever input data can be ‘phased’ (see Supporting Information), this additional information about the signal can be used to strengthen the projections, enabling our iterated maps approach to push down to even sparser sampling. A few 1D experiments should be sufficient to quickly set up the proper conditions to run an nD experiment that can be accelerated using iterated maps.

For the 2D NMR experiments described here, we used the stronger P_1 projection that makes use of the known sign of the real part of the signal (+ amplitude in positive, – amplitude in negative support regions). In cases where the signal sign is not known in advance, our so-called artifact mask should be used instead (with $\pm$ real signal amplitude allowed in the artifact support regions). This will also work, but since less *a priori* information is fed into the algorithm, a larger number of N_{t1} rows will typically be required to reach the same high quality output (e.g., see Figures S4–S5). The precise increase in the required number of rows will depend upon the details of the signal, the mask, the sampling schedule, and the noise-handling method. In general, the user should try to take advantage of all known characteristics of the dense data set.

In Figures 1, 2, 4, and 5, we aimed for nearly ideal output spectra $\vec{F}^n(f) \approx \vec{T}(f)$, since results of that quality can be used for any application. In the case of 2D NMR, we found that QUEST helped us to determine the minimum number of samples consistent with that goal. At first glance, our sparse sampling percentages may not seem that low, but in fact they appear to be quite close to the minimum necessary for a constrained linear system. For example, the corresponding sparse sampling percentages used in Figures 1,2,4 are: (50.8%,

58.6%,34.3%), which are similar to the largest percentage of positive or negative support pixels along the f_2 columns for each mask: (55.1%,44.5%,32.6%). Since our current 2D NMR experiments use sparse sampling in the single indirect dimension (by ‘skipping’ some t_1 values), the required sparse sampling percentages should drop quickly as this method is applied in 3D NMR, 4D NMR, etc.

In addition to the Difference Map $\hat{D}$ shown here, we tried several other iterated maps, including D_{flip} which flips $\hat{P}_1 \leftrightarrow \hat{P}_2$, $D_{AltProjs} = \hat{P}_1 \hat{P}_2$, and Elser’s ‘Divide and Concur’ [27] which accommodates more than two projections. All of the maps worked, and they were relatively easy to implement. The output spectra $\vec{F}^n(f)$ have better signal-to-noise ratios than $\vec{T}(f)$, since fewer noisy samples are used at the input. The ability to sample at very long times, without requiring the acquisition of all intermediate times, can be used to achieve higher spectral resolution in less acquisition time. The time savings offered by iterated maps could be leveraged to allow practical acquisition of higher dimensional (5D and 6D) NMR experiments, which are currently not feasible in all but the most ideal circumstances (for a current review, see [28]). Moreover, iterated maps offer another approach to further accelerate [29] ultrafast 2D NMR [30].

The iterative maps approach has many characteristics that should be familiar to magnetic resonance practitioners, such as its use of on-grid sampling, the FFT/IFFT, and final outputs which look the same as dense data sets in both the time and frequency domains. It complements existing methods to reconstruct spectra from sparsely-sampled data, and should find further applications in NMR and MRI of solids. More generally, any data acquisition and image modalities which make use of two reciprocal spaces related by a transformation can use this technique to harness *a priori* knowledge to fill in under-sampled data. The speed, simplicity, and robustness to error makes this an ideal technique for fast analysis of noisy experimental data.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank J. Sethna for suggesting that V. Elser’s work could be adapted to our MR problem. We thank D. Ulrich for acquiring the LuxU data set. We thank R. Blum, T. Constable, R. de Graaf, M. Devoret, J. Duncan, S. Eltrington, M. Lustig, D. Rothman, J. Rovny, and P. Zhan for helpful discussions. This work was supported in part by the NSF (DMR-0653377, DMR-1310274) along with seed grants from YINQE, the Yale CRISP (an NSF MRSEC, DMR-1119826), and the Yale MRRC, and NIH (1P30NS052519-01A1) Yale Core Center for Quantitative Neuroscience with Magnetic Resonance, a NINDS-funded P30 Core Center (M.F., Z.S., S.B.), along with NSF (CHE-1012573) (S.S., K.Z.) and NSF (MCB-1121372) (G.M., P.L.). M.F. is an NSF fellow and this material is based upon work supported by the NSF GRF under Grant No. (DGE-0644492).

References

1. Ernst RR, Anderson WA. Application of Fourier transform spectroscopy to magnetic resonance. *Rev Sci Instrum.* 1966; 37:93–102.
2. Ernst, RR.; Bodenhausen, G.; Wokaun, A. *Principles of Nuclear Magnetic Resonance in One and Two Dimensions.* Oxford University Press; Oxford: 1987.

3. Stern AS, Donoho DL, Hoch JC. NMR data processing using iterative thresholding and minimum $l(1)$ -norm reconstruction. *J Magn Reson.* 2007; 188:295–300. [PubMed: 17723313]
4. Hoch, JC.; Stern, AS. *NMR Data Processing.* Wiley-Liss; New York: 1996.
5. Mobli M, Hoch JC. Maximum entropy spectral reconstruction of non-uniformly sampled data. *Concept Magn Reson Part A.* 2008; 32A:436178.
6. Hyberts SG, Takeuchi K, Wagner G. Poisson-Gap Sampling and Forward Maximum Entropy Reconstruction for Enhancing the Resolution and Sensitivity of Protein NMR Data. *J Am Chem Soc.* 2010; 132:2145–2147. [PubMed: 20121194]
7. Paramasivam SS, Suiter CL, Hou GG, Sun SS, Palmer MM, Hoch JC, Rovnyak DD, Polenova TT. Enhanced sensitivity by nonuniform sampling enables multidimensional MAS NMR spectroscopy of protein assemblies. *J Phys Chem B.* 2012; 116:7416–7427. [PubMed: 22667827]
8. Orekhov VY, Jaravine VA. Analysis of non-uniformly sampled spectra with multi-dimensional decomposition. *Progr NMR Spectr.* 2011; 59:271292.
9. Mobli M, Maciejewski MW, Schuyler AD, Stern AS, Hoch JC. Sparse sampling methods in multidimensional NMR. *Phys Chem Chem Phys.* 2012; 14:10835–10843. [PubMed: 22481242]
10. Stanek J, Augustyniak R, Kozminski W. Suppression of sampling artefacts in high-resolution four-dimensional NMR spectra using signal separation algorithm. *J Magn Reson.* 2012; 214:91–102. [PubMed: 22070970]
11. Kazimierczuk K, Stanek J, Zawadzka-Kazimierczuk A, Kozminski W. Random sampling in multidimensional NMR spectroscopy. *Progr NMR Spectr.* 2010; 57:420434.
12. Kazimierczuk K, Orekhov VY. A comparison of convex and non-convex compressed sensing applied to multidimensional NMR. *J Magn Reson.* 2012; 223:1–10. [PubMed: 22960668]
13. Frey MA, Michaud M, VanHouten JN, Insogna KL, Madri JA, Barrett SE. Phosphorus-31 MRI of hard and soft solids using quadratic echo line-narrowing. *Proc Natl Acad Sci USA.* 2012; 109:5190–5195. [PubMed: 22431609]
14. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med.* 2007; 58:1182–1195. [PubMed: 17969013]
15. Candes E, Wakin MB. An introduction to compressive sampling. *IEEE Signal Proc Mag.* 2008; 25:21–30.
16. Candes E, Romberg J, Tao T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theory.* 2006; 52:489–509.
17. Candes EJ, Wakin MB, Boyd SP. Enhancing sparsity by reweighted l_1 minimization. *J Fourier Anal Appl.* 2008; 14:877–905.
18. Hu S, Lustig M, Balakrishnan A, Larson PE, Bok R, Kurhanewicz J, Nelson SJ, Goga A, Pauly JM, Vigneron DB. 3D compressed sensing for highly accelerated hyperpolarized $(13)\text{C}$ MRSI with in vivo applications to transgenic mouse models of cancer. *Magn Reson Med.* 2010; 63:312–321. [PubMed: 20017160]
19. Elser V. Phase retrieval by iterated projections. *J Opt Soc Am A.* 2003; 20:40–55.
20. Elser V, Thibault P, Rankenburg I. Search with iterated maps. *Proc Natl Acad Sci USA.* 2007; 104:418–423. [PubMed: 17202267]
21. Matsuki Y, Eddy MT, Herzfeld J. Spectroscopy by integration of frequency and time domain information (SIFT) for fast acquisition of high resolution dark spectra. *J Am Chem Soc.* 2009; 131:4648–4656. [PubMed: 19284727]
22. Matsuki Y, Eddy MT, Griffin RG, Herzfeld J. Rapid three-dimensional MAS NMR spectroscopy at critical sensitivity. *Angew Chem Int Ed Engl.* 2010; 49:9512–9518.
23. Matsuki Y, Konuma T, Fujiwara T, Sugase K. Boosting protein dynamics studies using quantitative nonuniform sampling NMR spectroscopy. *J Phys Chem B.* 2011; 115:13740–13745. [PubMed: 21992609]
24. States DJ, Haberkorn RA, Ruben DJ. A Two-Dimensional Nuclear Overhauser Experiment with Pure Absorption Phase in Four Quadrants. *J Magn Reson.* 1982; 48:286–292.
25. Fienup JR. Phase retrieval algorithms: a comparison. *Appl Opt.* 1982; 21:2758–2769. [PubMed: 20396114]

26. Ulrich DL, Kojetin D, Bassler B, Cavanagh J, Loria JP. Solution Structure and Dynamics of LuxU from *Vibrio harveyi*, a Phosphotransferase Protein Involved in Bacterial Quorum Sensing. *J Mol Biol.* 2005; 347:297–307. [PubMed: 15740742]
27. Gravel S, Elser V. Divide and conquer: a general approach to constraint satisfaction. *Phys Rev E.* 2008; 78:036706.
28. Kazimierczuk K, Stanek J, Zawadzka-Kazimierczuk A, Kozminski W. High-Dimensional NMR Spectra for Structural Studies of Biomolecules. *ChemPhysChem.* 2013;10.1002/cphc.201300277
29. Shrot Y, Frydman L. Compressed sensing and the reconstruction of ultrafast 2D NMR data: Principles and biomolecular applications. *J Magn Reson.* 2011; 209:352–358. [PubMed: 21316276]
30. Frydman L, Scherf T, Lupulescu A. The acquisition of multidimensional NMR spectra within a single scan. *Proc Natl Acad Sci USA.* 2002; 99:15858–15862. [PubMed: 12461169]
31. Lipchock JM, Loria JP. Nanometer Propagation of Millisecond Motions in V-Type Allostery. *Structure.* 2010; 18:1596–1607. [PubMed: 21134639]
32. Korzhnev DM, Klotzner K, Kay LE. Multiple-Quantum Relaxation Dispersion NMR Spectroscopy Probing Millisecond Time-Scale Dynamics in Proteins: Theory and Application. *J Am Chem Soc.* 2004; 126:7320–7329. [PubMed: 15186169]
33. Morcombe CR, Gaponenko V, Byrd RA, Zilm KW. Diluting abundant spins by isotope edited radio frequency field assisted diffusion. *J Am Chem Soc.* 2004; 126:7196–7197. [PubMed: 15186155]
34. Franks WT, Zhou DH, Wylie BJ, Money BG, Graesser DT, Frericks HL, Sahota G, Rienstra CM. Magic-angle spinning solid-state NMR spectroscopy of the beta1 immunoglobulin binding domain of protein G (GB1): ^{15}N and ^{13}C chemical shift assignments and conformational analysis. *J Am Chem Soc.* 2005; 127:12291–12305. [PubMed: 16131207]
35. Schmidt HL, Sperling LJ, Gao YG, Wylie BJ, Boettcher JM, Wilson SR, Rienstra CM. Crystal polymorphism of protein GB1 examined by solid-state NMR spectroscopy and X-ray diffraction. *J Phys Chem B.* 2007; 111:14362–14369. [PubMed: 18052145]

Appendix A. Materials

The spectral acquisition parameters used throughout this paper were determined by the goals of each experiment, and they were not adjusted to optimize the performance of the algorithm.

The ^{15}N HSQC of LuxU in Fig. 1 was collected at 14.1 T and 20°C with a ^1H spectral width of 12000 Hz and a ^{15}N spectral width of 2500 Hz [26]. For Figs. 2–3, the ^2H , ^{13}C -methyl ILV labeled IGPS (*t. maritima*) NMR sample was prepared as previously described [31]. The zero time point of a $^{13}\text{CH}_3$ multiple quantum CPMG relaxation dispersion experiment was collected at 14.1 T and 30°C with 120 t_1 increments using a previously published sequence [32]. The ^1H carrier was centered at 4.70 ppm with a spectral width of 8500 Hz, while the ^{13}C carrier was centered at 19.5 ppm with a spectral width of 3200 Hz.

The ^{13}C - ^{13}C correlation spectrum for a microcrystalline sample of uniformly $^{13}\text{C}/^{15}\text{N}$ enriched GB1 shown in Fig. 4 was taken using a RAD mixing sequence [33] on a Varian VNMRs spectrometer operating at a ^1H frequency of 798.89 MHz, a MAS rate of 18.2 kHz, and 4 scans per t_1 point. A spectral width of 100 kHz was used in both the direct and indirect dimensions. The sample was expressed and purified according to a previously reported protocol [34], using the pET-11a plasmid incorporating the gene for the T2Q mutant of GB1 (kindly provided by Dr. Angela M. Gronenborn, University of Pittsburgh) into BL21 *E. coli* cells. For solid-state NMR, microcrystalline GB1 was precipitated from a solution

containing 25 mg/mL GB1 in 50 mM Na_2HPO_4 at pH 5.5, and 50% (v/v) Methyl-2,4-pentanediol (MPD), 25% (v/v) Isopropyl alcohol (IPA) as reported by Rienstra and co-workers [35].

The sample used in Fig. 5 is a marrow-filled section of a pork rib bone, which was cut from a whole fresh pork rib. The rib bone section was placed in a cryotube vial, which was then filled to volume with phosphate buffered saline (PBS) solution to keep the sample hydrated. The ^{31}P MRI data was acquired at the Yale University Medical School's Magnetic Resonance Research Center, using the Bruker Avance 4.0 Tesla / 31 cm animal system, running ParaVision 3.0.1. See reference [13] for more information on how this data was acquired.

Appendix B. Supporting Information

In addition to the figures and details referred to in the main text, the Supporting Information includes several figures referred to in that section. Figure S6 shows the mapping of 2D NMR 'States'-like data into the four quadrants of $\vec{T}(t)$. Figure S7 is a schematic depiction of the behavior of the difference map near a gap between the subspaces defined by the projections. Figure S8 shows how different types of noise-handling affect a slice of the output, and the metrics we can use to follow the iterations of the algorithm.

We also include seven Quicktime movies, which show the algorithm iterating from the beginning to the end (Movie S1), how the final output compares to the dense target (Movie S2), and a slice-by-slice comparison of the output before and after the difference map to the dense target spectrum (Movie S3), all for the case of the 2D NMR of liquids LuxU data (from Figure 1).

Movie S4 is a slice-by-slice comparison of the output before and after the difference map to the dense target spectrum, for the case of the 2D NMR of liquids IGPS data (from Figure 2). Movie S5 is a slice-by-slice comparison of the output before and after the difference map to the dense target spectrum, for the case of the 2D NMR of solids $^{13}\text{C}/^{15}\text{N}$ enriched GB1 data (from Figure 4).

Movies S6 and S7 show the algorithm iterating from beginning to end for the case of the 2D NMR of liquids LuxU data, using two different sampling schedules. Movie S6 uses $N_{t_1} = 65$ and the QUEST schedule (just as in Figure 6), while Movie S7 uses a larger $N_{t_1} = 71$ and a very different sampling schedule that skips contiguous rows in the intermediate t_1 range.

Highlights

An iterated maps algorithm is applied to sparsely-sampled time domain data.

Used to reconstruct spectra from noisy 2D NMR and 3D MRI of solids data.

High quality results achieved with sparse sampling approaching theoretical minimum.

We use the QUEST sampling schedule and discuss its benefits for 2D NMR data.

FFT-based method is computationally fast, simple to implement, and robust.

Figure 1.

Reconstructing liquid-state 2D NMR data of the ^{15}N -LuxU sample. (A) Magnitude plot of 50.8% sparsely-sampled (t_2, t_1) -domain NMR data, $|S^0(t)|$, (256 rows $\times$ 4096 columns). Pink shows the t_1 rows set to zero by P_0 . (B) Real part of phase-corrected FFT of the sparsely-sampled data shown in A, $\tilde{S}_r^0(\vec{f})$, which shows the resulting aliasing artifacts along f_2 columns. Red pixels are 6% of the maximum dense signal (MDS) in $T_r(f)$ (not shown). The P_1 mask is shown, where blue surrounds the artifact region, black surrounds the positive support regions, and green surrounds the negative support regions. (C) Plot of pixel values along the six columns marked by red arrows in B ($\tilde{S}_r^0(\vec{f}) = \tilde{F}_r^0(\vec{f})$) and the blue arrows in E ($\tilde{F}_r^{15}(\vec{f})$) versus corresponding pixels from $T_r(\vec{f})$. The long dashed line $y = (0.508)x$ shows the poor quality of the fit before reconstruction (red open circles are $\tilde{F}_r^0(\vec{f})$ vs. $T_r(\vec{f})$). The short dashed line $y = x$ shows the excellent agreement after reconstruction (blue open triangles are $\tilde{F}_r^{15}(\vec{f})$ vs. $T_r(\vec{f})$), over the full range of positive and negative pixel values. (D) The resulting time-domain data after 15 iterations of our difference map (DM) algorithm, $F^{15}(t)$, using a small value of P_1 noise-handling ($\pm 0.3\%$ of the MDS). (E) Real part of phase corrected FFT, $\tilde{F}_r^{15}(\vec{f})$, of the DM reconstructed data shown in D. Red pixels are 6% of the MDS. (F) A portion of a contour plot showing the close match between the real parts of the dense spectrum (dashed) and the reconstructed spectrum (solid), using the same color scale and contour values (6% $\rightarrow$ 72% of the MDS, in 6% steps) for both $T_r(f)$ and $\tilde{F}_r^{15}(\vec{f})$.

Figure 2.

Reconstructing liquid-state 2D NMR data of the IGPS sample. (A) Real part of phase-corrected FFT $\tilde{S}_r^0(\vec{f})$ of 58.6% sparsely-sampled (t_2, t_1) -domain data, which shows the resulting aliasing artifacts along the f_2 columns (256 rows $\times$ 4096 columns). Red pixels are 6% of the MDS. The P_1 mask is shown, where blue surrounds the artifact region, and black surrounds the positive support regions (0.7% of the MDS). (B) The reconstructed (f_2, f_1) -domain data after 15 iterations of our difference map algorithm, $\tilde{F}_r^{15}(\vec{f})$, without using any noise-handling. (B, Inset) Plot of pixel values along the three columns marked by blue arrows in B versus corresponding pixels from full dense spectrum. The short dashed line $y = x$ shows the excellent agreement after reconstruction (blue open triangles are $\tilde{F}_r^{15}(\vec{f})$ vs. $\tilde{T}_r(\vec{f})$). (C) A contour plot of a region comparing the real parts of the full dense spectrum $\tilde{T}_r(\vec{f})$ (black) to the reconstructed spectrum $\tilde{F}_r^{15}(\vec{f})$ (colored), using identical contour values (3%, 8%, 17%, 25%, 34%, 67% of the MDS) for both spectra.

Figure 3.

Comparing performance of the difference map algorithm using two different $\hat{P}_1$ masks for the IGPS data (from Figure 2). (A) The l_2 difference between the reconstructed and dense spectra (inside the positive support) plotted versus the number of positive t_1 rows sampled (N_{t_1}), where $N_{t_1} = 128$, for both looser ($> 0.14\%$ of MDS, red triangles) and tighter ($> 0.70\%$ of MDS, blue triangles) $\hat{P}_1$ masks. The tighter mask (blue) was used in Figure 2. Both traces have a peak in l_2 at $N_{t_1} = 65$, but only the looser mask (red) has a peak at $N_{t_1} = 85$. (B) The aliasing of the looser mask due to uniform sampling at $N_{t_1} = 85$. (C) Zooming into the region outlined by the black dashed line in B. Multiple, strong ($> 20\%$ of MDS), aliased peaks (in red and blue) are inside the positive support of the looser $\hat{P}_1$ (highlighted in yellow), and they result in a worse reconstruction. (D) The same region shown in C, but now using the tighter $\hat{P}_1$, which makes most aliased peaks fall outside the tighter support region, improving the reconstruction quality and explaining the disappearance of the peak in l_2 at $N_{t_1} = 85$ for this mask in A (blue).

Figure 4.

Reconstructing the ^{13}C - ^{13}C 2D MAS correlation spectrum for the $^{13}\text{C}/^{15}\text{N}$ enriched GB1

solid state sample. (A) Real part of phase-corrected FFT $\tilde{S}_r^0(\vec{f})$ of 34.3% sparsely-sampled (t_2, t_1) -domain data, which shows the resulting aliasing artifacts along the f_2 columns (4096 rows $\times$ 4096 columns). Red pixels are 2.7% of the MDS. The P_1 mask is shown, with artifact regions surrounded by blue, negative support regions by green, and positive support regions surrounded by black. As described in the main text, a tiny ‘keyhole’ was cut into the mask over the region $((17.7\text{kHz} \leq f_2 \leq 13.4\text{kHz}) \cap (9.5\text{kHz} \leq f_1 \leq -1.1\text{kHz}))$, which significantly reduced the requisite sampling percentage (53.7% $\rightarrow$ 34.3%). (B) The reconstructed (f_2, f_1) -domain data after 15 iterations of our difference map algorithm,

$\tilde{F}_r^{15}(\vec{f})$, without using any noise-handling. (C – F) Zoom-in contour plots of various regions comparing the full dense spectrum (dashed) to the reconstructed spectrum (solid), using identical color scales and contour values (1%, 2%, 3%, 4%, 8% of the MDS) for both spectra. C shows a high-signal region near the main diagonal; D shows a low-signal region far from the main diagonal; E shows another low-signal region far from the main diagonal; and F shows a low-signal region where there is a rotary resonance (MAS rate ≈ 18.2 kHz). (F, Inset) Plot of pixel values along the two columns marked by blue arrows in B versus corresponding pixels from the full dense spectrum. The short dashed line $y = x$ shows the

high-quality fit over a wide dynamic range after reconstruction (blue open triangles are

$$\tilde{F}_r^{15}(\vec{f}) \text{ vs. } \tilde{T}_r(\vec{f})).$$

Figure 5.

Reconstruction of sparsely-sampled solid-state MRI data. (A) Isosurface rendering of a portion of the 3D ($64 \times 64 \times 64$) $T_r(\vec{r})$ showing ^{31}P density in pork rib in PBS solution. The isosurface value is 33% of the MDS and shows the thick cortical bone ring. The spatial resolution is $(1.19\text{mm})^3$ and T_{acq} was 35.2h. (B) A 2D slice of $T_r(\vec{r})$ with thickness of 0.595mm. The support region for this 2D slice (where signal is expected to be positive) is outlined in blue. (C) The same 2D slice from $\tilde{S}_r^0(\vec{r})$ which would require only 17% of the imaging time used for the dense image. (Inset) Plot of $\tilde{S}_r^0(\vec{r})$ vs. $\tilde{T}_r(\vec{r})$ for voxels within the support region. The thick, dashed line $y = x$ shows the poor-quality fit prior to DM reconstruction. By definition, $\tilde{S}_r^0(\vec{r}) = \tilde{F}_r^0(\vec{r})$. (D) The same 2D slice of $\tilde{F}_r^{20}(\vec{r})$ after DM reconstruction with P_1 and P_2 error handling. (Inset) Plot of $\tilde{F}_r^{20}(\vec{r})$ vs. $\tilde{T}_r(\vec{r})$ for voxels within the support region. The thick, dashed line $y = x$ shows the high-quality fit after reconstruction, where most of the points are within the 10% noise level for the measured data.

Figure 6.

The first of four figures illustrating the specific ways in which the LuxU reconstruction quality drops, as fewer time points are sampled (i.e., as N_{t1} is systematically lowered over Figures 6–9). (A–F) Identical to the corresponding panels in Figure 1, using the lowest $N_{t1} =$

65 for a very high quality output $\tilde{\mathbf{F}}_r^{15}(\vec{f}) \approx \tilde{\mathbf{T}}_r(\vec{f})$. (G) As explained in Supporting Information, we calculate how ‘parallel’ and ‘perpendicular’ $\mathbf{F}^n(\vec{f})$ is to the target vector, $\mathbf{T}(\vec{f})$. The resulting parametric plot of $F_{\parallel \mathbf{T}}(n)$ vs. $F_{\perp \mathbf{T}}(n)$, from the $n = 0$ point (thin black arrow), to $n = 15$ (thick black arrow), shows the approach to the target (red arrow). (H). A plot like Figure 3C, showing the central LuxU mask (dark yellow), the $-1 \times BW_1$ mask (dark red) poking in from above, and the $+1 \times BW_1$ mask (dark blue) poking in from below, as explained in the Supporting Information. The black contours are at $n_{1\%} = 1\% \times MDS$ of $\mathbf{T}_r(\vec{f})$. The few blue pixels are where $\tilde{\mathbf{F}}_r^{15}(\vec{f}) - \tilde{\mathbf{T}}_r(\vec{f}) \geq n_{1\%}$, and the few red pixels are where $\tilde{\mathbf{T}}_r(\vec{f}) - \tilde{\mathbf{F}}_r^{15}(\vec{f}) \geq n_{1\%}$. The blue and red pixel density will grow as N_{t1} is lowered in Figures 7,8,9.

Figure 7.

Same as Figure 6, if we drop to $N_{t_1} = 50$ (note the increase in pink row density in A). The long dashed line in (C) has slope 50/128. Still a high quality output. (H) shows more blue pixels near the overlap of masks (top and bottom of yellow box), the residual traces of sparse sampling artifacts. The red pixels are concentrated in the black contours.

Figure 8.

Same as Figure 6, if we drop to $N_{t1} = 35$. The long dashed line in (C) has slope 35/128. This is approximately the lower end of the N_{t1} range for high quality outputs. In (H), we see that the dark red and dark blue masks just touch, and so blue pixels are scattered across the central mask. Red pixels fill more of the black contours. This $\tilde{F}_r^{15}(\vec{f})$ requires only 27% of the normal acquisition time of $T_r(\vec{f})$.

Figure 9.

Same as Figure 6, if we drop to $N_{t1} = 20$, which is a lower quality output. The long dashed line in (C) has slope 20/128. Compared to Figure 6, most blue points in (C) fail to reach the short dashed line (with a slope of 1). Note also that this is the first case where many blue points in (C) that should be zero to match the dense image have instead grown larger than the red point values they had at the start of the algorithm (look at points at zero on the horizontal axis, in Figs. 6C–9C). The ‘perpendicular’ component in G grows monotonically for $n = 0$ to $n = 15$. In D , dark stripes are noticeable, along with residual aliasing within the mask in E . For this $N_{t1} = 20$, H shows that the dark red and dark blue masks overlap by a lot, and the $\pm 2 \times BW1$ masks (not shown) overlap by a little, and the $\pm 3 \times BW1$ masks (not shown) are starting to poke into the central mask. As a result, even more blue pixels are scattered across the central mask. Red pixels fill even more of the black contours. Still, C shows that the final output is better than $\tilde{S}_r^0(\vec{f})$, and E shows that some $T_r(\vec{f})$ features are recognizable (F), in just 16% of the normal acquisition time.