McMaster University

Advanced Optimization Laboratory

Title:

Optimal Fourier Transform Sampling Patterns

Author:

Christopher Kumar Anand and Anuroop Sharma

AdvOL-Report No. 2007/13

August, 2007
Hamilton, Ontario, Canada

Optimal Fourier Transform Sampling Patterns

Christopher Kumar Anand* and Anuroop Sharma

Department of Computing and Software
McMaster University, Hamilton, Canada

(30 July 2007)

A model for optimizing the sampling of the Fourier Transform of a function is presented, together with a sequential semi-definite trust region strategy, and an implementation using the open-source solver CSDP. Numerical performance correlates with increasing dimension and decreasing number of non-zero function values. Simulation of an application to multi-dimensional NMR (used to study to protein dynamics), predict a hundred-fold reduction in data acquisition time and increased signal to noise ratio when compared to the most common experimental method.

*Keywords:*NMR; MRI; Medical Imaging; Sparse k-Space Sampling; Nonuniform Fourier transform; SDP; Semi-Definite Optimization

1 Introduction

This paper presents a model and a sequential semi-definite solver for optimizing the discrete sampling of the Fourier Transform of a function with known support. This problem arises in many domains where the Fourier Transform or a closely-related transform, like the Radon Transform, is used, *e.g.* most types of medical imaging, in some types of non-destructive testing, and in multi-dimensional Nuclear Magnetic Resonance. In these domains, the domain of the sampled function is called either k-space or the time domain. The domain of the model function is either called the image domain or the frequency domain. Up to now, heuristics, best practices, and application of the Nyquist sampling criteria have been used to design such sampling patterns. For many applications, the *point spread function* (psf) gives a good multi-objective, which can be reduced to a single objective function in different ways for different applications. This is good for applications with small or no constraints on the model

*Corresponding author: Email: anandc@mcmaster.ca

function being estimated. In this case, any reduction in sampling from fully sampling the discretized function in k-space introduces reconstruction errors called aliasing artifacts. The goal is to minimize them.

The same approach has been applied to multi-dimensional NMR, with Coggins and Zhou (2007) having the best-optimized undersampling pattern based on a long evolution of ideas centred on radial sampling whose application to NMR can be traced back to Nagayama *et al.* (1978). The NMR problem is qualitatively different from 'imaging' problems, in that

(i) most of the expected model values are zero: only isolated peaks carry signal;
(ii) the dimensions do not correspond to physical space and are only constrained by the complexity of the pulse sequence. Kupce and Freeman (2006) have proposed using 10 dimensions, although not by imaging them in a single 10-dimensional array, since that would be too time consuming.

For mathematicians not familiar with these problems, Liang and Lauterbur (1999) provide a mathematical description of MRI, as do Hoch and Stern (1996) for NMR.

In all of these application areas, information about the support of the model function being estimated is usually available. For example, in imaging, the device has a known geometry which restricts the object being imaged to lie in a subset of the imaging area/volume. In NMR, the locations of all the peaks are either exactly or approximately known, or they are known to occupy some bands.

An application which would greatly benefit from optimization is the determination of protein dynamics from T_2 relaxation. In this experiment, the correspondence between peaks and bases in the protein's peptide chain is known, but functional information (including which parts of the protein are in contact with water) are not. Much can be inferred by the rate at which the signals corresponding to different bases decay. This is determined by making multiple measurements with variable delays inserted into the experiment to allow for decay following an exponential curve.

In addition to promising hundred-fold reductions in acquisition time for common 3D NMR, the remarkable thing about this method is that it gets more efficient as the dimension increases, whereas all conventional methods get more expensive as dimension increases.

New Model. A new way to optimize the efficiency of data collection for these inverse problems is to model the signal generation as a linear function of model variables, and minimize the amplification of noise caused by inverting it. For example, since the Fast Fourier Transform preserves the L^2 norm, identical, independent, normally distributed noise is transformed to independent, identical normally distributed noise. This is the best case if the model variables

come from a complete rectangular grid discretization. In general, inverting the linear signal generation transformation produces estimates for model variables with different variances, corresponding to condition numbers greater than one. In some cases, the transformation is not invertible at all, and some model variables cannot be estimated.

Anand *et al.* (2007) introduced a method of optimizing a *set* of steady-state MRI experiments with respect to expected noise. The same eigenvalue maximization approach using a semi-definite constraint applies to efficient k-space sampling, but the other (non-convex) constraints are quite different. The sampling problems in this paper are, in general, much larger than the problem considered in the previous paper, with no realistically-sized problem being solvable with a straight-forward model. All of the sampling problems we propose to solve carry a lot of additional structure. In this paper, the structure of the multi-dimensional NMR problem is used to decompose the natural problem into a series of sub-problems which are each optimizable by a trust-region method using a semi-definite/linear sub-problem solved by CSDP (Borchers 1999).

Organization. In section 2, the general model is presented in complex and real forms. It is translated into a problem with semidefinite constraints, linearized to form a trust region subproblem, and implemented using CSDP, in section 3. The next section, 4, details the application to multi-dimensional NMR, including a further decomposition into hyperplane subproblems. Section 5 follows with numerical results showing that realistic problem sizes can be solved by this method, comparing different strategies, and showing that the quality of solution increases with dimension.

2 Nonlinear Problem

Let $\{x_j\} \subset R^r$ be a set of discrete points of interest for a model function $f : R^r \to C$ (the real-valued case can be treated as a special case, or be solved by similar methods). The values of f are not directly measurable, but the values of

$$\tilde{f}(k_i) = \sum_{j=1}^{m} f(x_j) e^{\sqrt{-1}\langle k_i, x_j \rangle},$$

its Fourier Transform, are measurable. Let n be the number of such measurements, and m the number of model variables.

Notes: This assumes that the support is a discrete set of points. This may be strictly true, the support may be well-approximated by a discrete set of points,

or the contribution outside a discrete set may be removable by filtering the data. In cases such as projection imaging, the directly measured data is not the the Fourier Transform, but is equivalent to knowing the Fourier Transform on a restricted set.

Adding noise, we can write this as the affine transformation

$$\begin{pmatrix} \tilde{f}(k_1) \\ \vdots \\ \tilde{f}(k_n) \end{pmatrix} = S \begin{pmatrix} f(x_1) \\ \vdots \\ f(x_m) \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \vdots \\ \epsilon_n \end{pmatrix}$$

where S is the complex $n \times m$ matrix

$$S_{i,j} = e^{\sqrt{-1}\langle k_i, x_j \rangle}.$$

Without noise, the Moore-Penrose pseudo-inverse reconstructs the model f

$$\begin{pmatrix} f(x_1) \\ \vdots \\ f(x_m) \end{pmatrix} = (S^*S)^{-1}S^* \begin{pmatrix} \tilde{f}(k_1) \\ \vdots \\ \tilde{f}(k_n) \end{pmatrix} \tag{1}$$

exactly. But noise in the measurements is also transformed. The worst-case expected error (including correlated error involving multiple model variables) corresponds to the singular vector of $(S^*S)^{-1}S^*$ with the minimum singular value. Maximizing sampling efficiency therefore corresponds to maximizing the minimum eigenvalue of S^*S:

$$(S^*S)_{i,j} = \sum_{l=1}^{n} e^{\sqrt{-1}\langle k_l, x_j - x_i \rangle}$$

Note that all diagonal values are n, which is an upper bound on the minimum eigenvalue. If the minimum eigenvalue is equal to n, then S^*S must be diagonal. Very small problems could be solved in this form using a derivative-free optimizer. The minimum eigenvalue is a continuous function with discontinuities in its derivative at matrices with minimal eigenvalues with multiplicity greater than one, so smooth methods are not applicable. Non-smooth methods are generally slower, and were not tried.

The Real Problem. The Hermitian $m \times m$ matrix S^*S can be represented as a real symmetric $2m \times 2m$ matrix A with elements:

$$\begin{aligned} A_{2i-1,2j-1} &= \sum_{l=1}^{n} \cos\langle k_l, x_j - x_i \rangle \\ A_{2i,2j} &= \sum_{l=1}^{n} \cos\langle k_l, x_j - x_i \rangle \\ A_{2i,2j-1} &= \sum_{l=1}^{n} \sin\langle k_l, x_j - x_i \rangle \\ A_{2i-1,2j} &= -\sum_{l=1}^{n} \sin\langle k_l, x_j - x_i \rangle \end{aligned} \tag{2}$$

where $i = 1...m$ and $j = 1...m$.

The real matrix A has the same eigenvalues as S^*S with each multiplicity doubled. So maximizing the minimum eigenvalue of A is equivalent to maximizing the minimum eigenvalue of S^*S.

2.1 *Nonlinear/Semidefinite Problem*

Maximizing the minimum eigenvalue of A can be formulated as a semidefinite programming problem: given parameters $x_i \in R^r$,

$$\min_{\{k_i\}} \quad -\lambda \tag{3}$$

$$\text{subject to} \quad A - \lambda I \succeq 0 \tag{4}$$

$$A \text{ satisfies (2)} \tag{5}$$

In different applications, the components of $k_i \in R^r$ must satisfy different constraints, given by the physical limits of the measurement hardware, basic physics, or practical limits on experiment time. If the constraints are simple bounds, which is the case for multi-dimensional NMR, the bounds can be added trivially, and will not be carried through the problem formulation.

3 Semidefinite/Linear Trust Region Subproblem

Trust Region methods (Conn *et al.* 2000) are commonly used for non-linear problems. Define a semidefinite/linear trust region subproblem by linearizing

the nonlinear constraints (5), using the first-order Taylor series for A at a previous guess $\tilde{k}$. Substituting it into the previous problem produces

$$\begin{aligned} \min_{k} \quad & -\lambda \\ \text{subject to} \quad & A|_{\tilde{k}} + \sum_{\substack{\alpha=1\ldots n \\ \beta=1\ldots r}} (k_{\alpha,\beta} - \tilde{k}_{\alpha,\beta}) \left.\frac{\partial A}{\partial k_{\alpha,\beta}}\right|_{\tilde{k}} - \lambda I \succeq 0, \end{aligned} \tag{6}$$

in which

$$\frac{\partial A_{2i-1,2j-1}}{\partial k_{\alpha,\beta}} = -(\sin\langle k_\alpha, x_j - x_i\rangle)(x_{j,\beta} - x_{i,\beta}) \tag{7}$$

$$\frac{\partial A_{2i,2j}}{\partial k_{\alpha,\beta}} = -(\sin\langle k_\alpha, x_j - x_i\rangle)(x_{j,\beta} - x_{i,\beta}) \tag{8}$$

$$\frac{\partial A_{2i,2j-1}}{\partial k_{\alpha,\beta}} = \quad (\cos\langle k_\alpha, x_j - x_i\rangle)(x_{j,\beta} - x_{i,\beta}) \tag{9}$$

$$\frac{\partial A_{2i-1,2j}}{\partial k_{\alpha,\beta}} = -(\cos\langle k_\alpha, x_j - x_i\rangle)(x_{j,\beta} - x_{i,\beta}). \tag{10}$$

In the trust region method, this problem is solved with additional constraints requiring the solution to stay within a region of trust where the linearization is a good estimate of the non-linear constraints. Commonly, the region is a sphere, which works well in general, but for the particular constraint (5), it is possible to tune the shape of the region to the problem. The set of k such that

$$|k_{\alpha,\beta} - \tilde{k}_{\alpha,\beta}| \leq \frac{\pi/2}{\max|x_{j,\beta} - x_{i,\beta}|} \tag{11}$$

has the property that the restrictions of the trigonometric components go through at most a quarter phase. This is a cheap way of scaling the trust-region to the curvature in different directions. Although it is not invariant under changes of co-ordinates, the co-ordinates have physical meaning for some problems, notably the NMR problem, and scale differently. Scaling this particular trust region builds in the relative sensitivity of A to changes in the components of k.

The trust region algorithm starting with an initial set of sample points $\tilde{k}$, at which $\tilde{\lambda}$ is the minimum eigenvalue of A, and trust region T defined by (11) is

(i) solve (6) plus bounds constraints and $k \in T$ for k with objective λ
(ii) calculate the minimum eigenvalue, $\lambda_{\min}$ of A at the point k

(iii) **if** $\lambda_{\min} < \tilde{\lambda}$
then shrink T by a factor of $\sqrt{2}$
elseif $\lambda_{\min} - \tilde{\lambda} > 3/4\,(\lambda - \tilde{\lambda})$
then grow T by a factor of $\sqrt{2}$
(iv) **if** $\lambda_{\min} > \tilde{\lambda}$ then replace $\tilde{k}$ by k, and $\tilde{\lambda}$ by $\lambda_{\min}$
(v) **while** $\left|k - \tilde{k}\right| > K$, repeat

where the constant K can be tuned to the problem.

3.1 *Implementation*

This algorithm was implemented in C, calling CSDP (Borchers 1999) both to solve the subproblem, and to find the true minimum eigenvalue at each guess k. CSDP's native definition format is dual to the formulation given, so the subproblem was first converted into dual form. Each of the linear inequalities (for the trust region and for bounds/nonnegativity constraints) results in a block in the block-diagonal matrix formulation used. Together this is an easy reformulation, and the CSDP interface matches the problem closely.

CSDP is written in portable C, and problems were solved on various single to 32-core machines. CSDP supports internal parallelization using OpenMP. This support was not used, because the NMR problems involve multiple such optimizations which can be solved in parallel using threads. Using standard unix pthreads requires a modification to the code to eliminate one static local variable used to support OpenMP.

4 Multi-Dimensional NMR

In NMR, a signal is created by a sequence of radio-frequency (rf) pulses that excite the spin states of common nuclei. The most abundant nuclei in organic molecules is hydrogen, so it has the strongest signal, and is the most common measurement target. Alone, each nuclei resonates at a different frequency. Bonds within a molecule modify this frequency and create coupling between neighbouring spins. This gives rise to a spectrum of frequencies, which serve as signatures for molecules and molecular substructures. The spectrum is calculated by Fourier transforming a time-series measured using an antenna.

The problem with this approach is that proteins (and other molecules) have so many atoms that the signatures overlap. To solve this problem, heteronuclear multi-dimensional experiments involve chains of specialized rf pulses capable of not only exciting spins but transferring their states from one nuclei to another. Once transferred, the spins evolve at different frequencies. By transferring spins to one or more additional nuclei and back, the resulting

signals are phase modulated, with the phase depending on the products of the frequency of the various host nuclear spins and time between transfers.

The data collected from such a scheme is a sampling of the Fourier transform of a function $R^r \rightarrow C$. (Note that multidimensional NMR experiments are usually designed so that the signals are real (or pure imaginary), or carry a mixture of signals from two different experiments in the real and imaginary parts. This reduces peak width in the spectrum, which is not an issue with the proposed method. For the purposes of exposition only the complex-valued case is considered, but the real-valued case requires only minor modifications.)

The structure of the experiment makes sampling in one direction (the directly-detected dimension—corresponding to hydrogen) cost-free, because they are sampled microseconds apart. The series of such samples is called a free-induction decay (FID). It also imposes non-negativity constraints on all the sampling and delay times, which, in this case, are the co-ordinates of the samples in k-space. In this paper, the signal locations in frequency space are assumed to be known and given by $x_i \in R^r$.

It is possible at this point to formulate an optimization problem with this information, but the resulting problems are too large to be solved practically.

4.1 *One densely sampled direction*

The dimension of the subproblems can be significantly reduced by using the known structure: samples are organized in dense lines parallel to the directly sampled direction. In practice, the sample spacing and number are known, so only the first sample point, $(0, k_{j,2}, k_{j,3}, ...)$ in each FID needs to be modelled, assuming without loss of generality that the first dimension is the directly detected dimension. The other points in the FID are $(1, k_{j,2}, k_{j,3}, ...)$, $(2, k_{j,2}, k_{j,3}, ...)$, ... in the appropriate units.

It follows that the matrix S^*S becomes

$$(S^*S)_{i,j} = \left(\sum_{p=0}^{n_p-1} e^{\sqrt{-1} \cdot n_p \cdot \delta \cdot (x_{j,0} - x_{i,0})} \right) \left(\sum_{l=1}^{n/n_p} e^{\sqrt{-1} \langle k_l, x_j - x_i \rangle} \right) \tag{12}$$

where n_p is the number of samples per FID, and δ is the inter-sample spacing. This reduces the problem in size to $n(r-1)$, and introduces a constant which can be computed once before the trust-region iteration.

This problem is still too large for biologically-interesting proteins.

4.2 *Hyperplane Subproblems*

Since the directly acquired dimension is densely sampled, it is possible to assume the Fourier transform has been applied in this dimension and consider each frequency in this dimension separately. This reduces the dimension of the problem by $(r-1)/r \cdot 1/n_p$, on the sampling (k) side, and on the peak (x) side by the ratio of peaks in that plane to overall peaks. Such problems can be solved. This does not take into account that a single set of FIDs is acquired, so the the k values are not independent in separate subproblems. Another problem is that we have to 'round' peaks which occur at arbitrary points to discrete frequency planes. Given a method of rounding peaks to planes, this can be used to define a block-diagonal A:

$$(S^*S)_{i,j} = \begin{cases} \sum_{l=0}^{n} e^{\sqrt{-1}\langle k_l, x_j - x_i\rangle} & \text{if } x_j, x_i \text{ belongs to plane } l \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

where $x_j, k_i \in \mathbb{R}^{r-1}$.

The resulting problem is still impracticably slow in CSDP. Perhaps a solver with more complete support for sparsity would do better, but it turns out that there is little incentive to try, as seen by the results of a simpler approach.

The simpler approach is simply to partition the peaks x_i into disjoint planes, and solve the problem associated with (12) for each subset in the partition. In the current implementation, a hierarchical clustering indexed by the frequency in the directly-acquired direction is performed. Hierarchical clustering works well for the benchmark problem, but it does not guarantee overlapping peaks will be be in a common plane, so an algorithm tailored to this problem be better, or even required for larger proteins.

Once optimal sampling co-ordinates are found for each plane, they can be unioned, eliminating nearly-identical points, if desired. Two approaches have been tried: (1) apportioning samples to planes in advance, greedily generating a random starting point and optimizing, (2) incrementally adding random points and re-optimizing the full set until a target efficiency is reached. The second method could require arbitrarily many samples. It is also slower.

For the purposes of benchmarking, *random greedy* search was performed by randomly selecting 64 initial points, computing the minimum eigenvalues and taking the best set. To these points were added 8 random samples, one point at a time, and the best augmented set was retained. This was optionally followed by the trust region algorithm.

Greedy trust is a variation on this, in which after each additional candidate is added, the trust region algorithm is applied, and the best set after optimization is retained for the next iteration. This has the advantage that the efficiency after optimization can be used as termination criterion, but involves much

more computation.

4.3 *Sentinel Points*

Using prior information about the model in optimizing the solution of an inverse problem introduces the risk that the prior information may become invalid. The more efficient the optimized solution, the greater the danger. One way of detecting model failures in this problem is to add sentinel points: x_is not corresponding to defined peaks. If the model is valid, the estimated signals at these points will be pure noise. Depending on the possible sources of failure, the co-ordinates of the sentinel points can be chosen to detect such problems. To test for machine failures, arbitrary points can be used. In the numerical examples in this paper, two points were added to each of the clustered planes. The projections onto the hyperplane orthogonal to the directly acquired direction were the same for all planes, and the points were chosen to be well within the convex hull of the projected points, but well separated from actual peak values.

5 Numerical Results

To test these methods, all algorithms were implemented in C, calling CSDP 2.3. Resonance frequencies from an assignment table for the protein RIa (119-244) were used. Not all frequencies for all residues were given, and these residues were ignored, leaving 113 points. Some frequencies could not be determined because of overlap, these can usually be resolved in higher dimensions. If the peaks are not of interest, one or more peaks can be added to estimate the overlapped peak so that their signals do not contaminate other peak measurements, but this is not required for benchmarking the optimization algorithm.

Effect of Dimension. In the first test, the seventeen bases with six resonance frequencies (corresponding to hydrogen, nitrogen, and four carbons) were projected onto lower-dimensional subspaces by dropping co-ordinates. Figure 1 shows the minimum eigenvalue obtained by randomly selecting 34 points and then applying the trust region algorithm. Four different starting sets were taken for each dimension. The figure shows a near linear increase in the efficiency of the optimized sampling patterns as the dimension increases, while the efficiency of random samples stays near 20%. As one would expect for a problem containing periodic functions, this problem has many local minima, indicated by the different efficiency values found by the trust region method. This shows that solution quality is remarkably consistent, and becomes more consistent with increasing dimension, even without multiple random starting

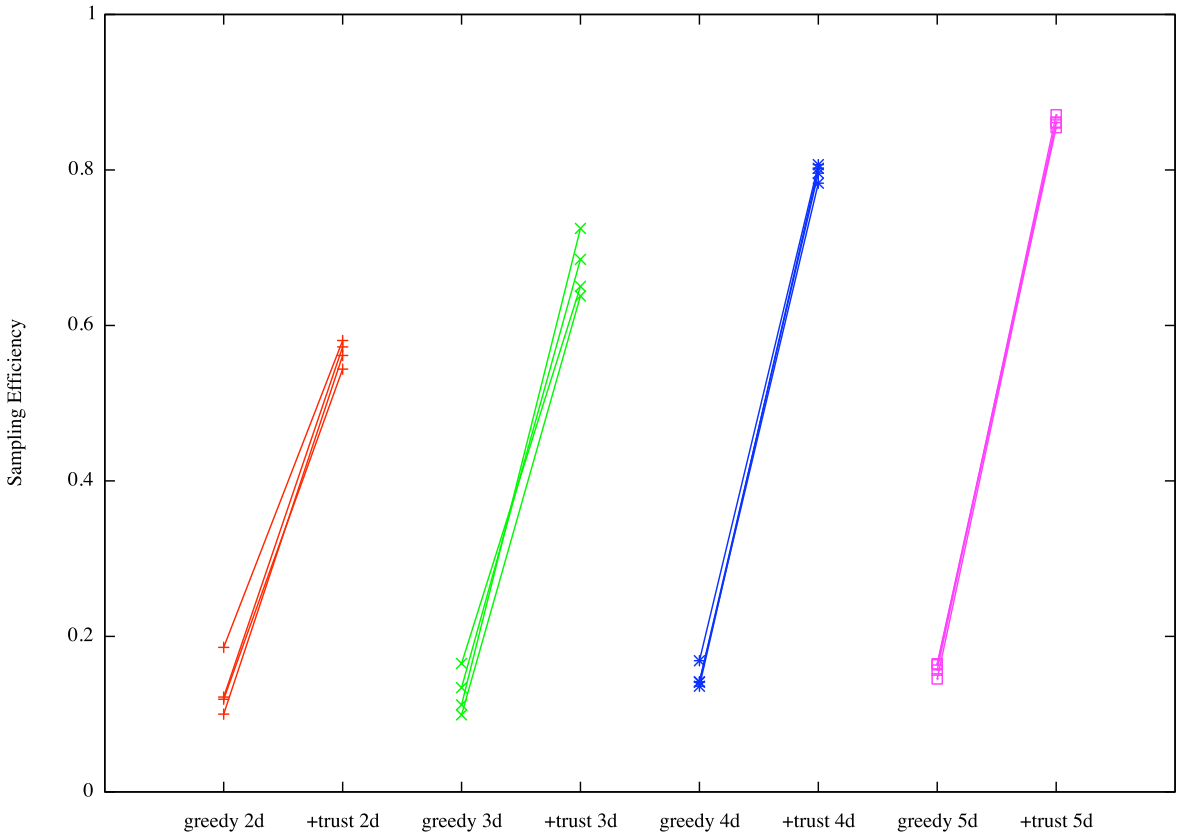

Figure 1. Efficiency of optimized sampling grows almost linearly in this range of dimensions.

points, and suggests that higher-dimensional NMR be used when optimizing sampling patterns.

Single plane. The full set of 113 points with measured H, N and C resonances was used to test the clustered hyperplane approach. To compare the relative efficiency and variability of the three methods, a clustered plane with 7 peaks (and the 2 sentinel peaks) was used in 8 runs of greedy random optimization, 3 runs of greedy trust optimization, and 16 runs of greedy random + single trust region step optimization, to show the variability of each method. Figure 2 shows that without continuous optimization 80% is the limit of the expected efficiency. Both single- and multiple-step trust region methods reach 80% efficiency with a third as many sample points. They also show less variability, with the variability decreasing as more points are added.

Full experiment. Applying the greedy random + trust algorithm to all planes (with 2 sentinel peaks), using a single set of samples with triple the cardinality of the plane, resulted in efficiencies shown in figure 3. Efficiency decreases as the cardinality of the plane increases. This computation took three hours on a 2.6GHz, 8 Dual Core Opteron server in a shared environment. To optimize NMR experiments which can run days on expensive spectrometers, this is more

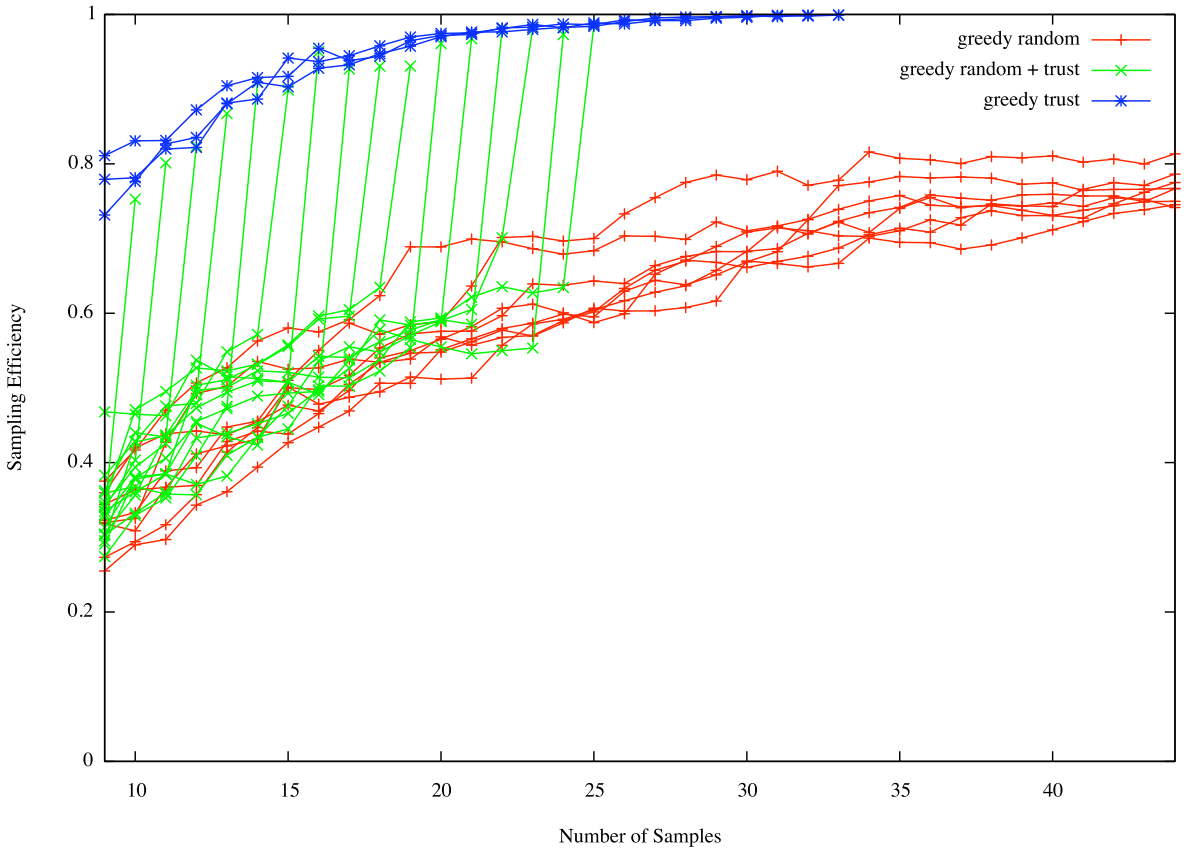

Figure 2. Trust region method does much better than random sampling, whether applied once in a final step, or after each additional random point.

plane	number of peaks	number of samples	efficiency
1	7+2	26	0.99
2	38+2	119	0.72
3	26+2	83	0.82
4	15+2	50	0.87
5	22+2	71	0.80
6	2+2	11	1.00

Figure 3. Efficiencies found by greedy random + trust method for all planes clustered from 3d RIa (119-244) peaks.

than justified.

When 512-point FIDS with μs sampling are collected for the union of these 360 k-space points, the sampling efficiency as measured by the minimal eigenvalue of A will be 88.2%. This number is higher than the efficiency for the 38-peak plane because (1) the clustered peaks are only partially overlapping, and (2) more than half the total points optimized for other planes also contribute.

The 3D experiment used to determine the peak positions used 64×64 sam-

ples in the C and N dimensions, for a total of $2097152 = 64 \times 64 \times 512$ samples. Using a Fast Fourier Transform, equal measurement noise (call it σ^2) translates to equal noise in the reconstructed spectrum. The extracted 113 peak values have expected noise variance σ^2. If peak areas are used, this is equivalent to taking a weighted average, which would reduce the noise variance by the area under the peak. This depends on the resolution and width of the individual peaks, but for the purposes of approximation, can generously be estimated to be 10. Call the expected noise variance $\sigma^2/10$.

For the optimized sampling 184320 samples are linearly combined to produce 125 estimates (including 12 sentinels). This is again equivalent to a weighted average, but the efficiency is bounded below by the minimum eigenvalue, and is 0.88. It follows that the expected variance is $\sigma^2/1298$—an expected 100-fold reduction in the variance.

The acquisition time is $64 \times 64/360 = 11.4$ times shorter with optimized sampling. Since the 8 or 16 averages usually used with this sequence are not necessary (given the reduction in noise variance for the optimized reconstruction), an overall hundred-fold reduction in acquisition time is possible if the averaging is eliminated.

In comparison, the most efficient sampling patterns designed for blind spectroscopy (where the peak locations are not known) offer a 10-fold reduction in acquisition time relative to dense sampling, but with a greater increase in apparent noise variance (since the number of reconstructed frequencies is the same, and aliasing artifacts will appear as systematic errors).

5.1 *Simulation*

As a check on correctness of the reconstruction, and a verification of the reported noise estimates. Data corresponding to 113 peaks with additional noise was added, and the peak values were estimated by solving (1) using the conjugate gradient method. Note that the effect of maxim.

The k-space picture presented in this paper is not a complete model. The most important feature excluded is relaxation during the sampling time. Relaxation made no difference to the simulated results.

6 Conclusion

This paper proposes a general model for optimizing the sampling of the Fourier Transform of a primary model function. This situation occurs in numerous applications. Repeated multi-dimensional NMR is the extreme example, in terms of sparsity of the model function, range in number of dimensions, benefit in reducing sampling. Structure in this problem (the densely sampled directly

acquired dimension) can be used to reduce the overall size of the problem and decompose it into smaller subproblems. These subproblems can then be solved using a novel general trust region method with a linear/semidefinite subproblem. Numerical tests using actual protein peak data show that the optimization method works on real-world problem sizes, and both analysis and simulation show that 100-fold reductions in acquisition time for $3D$ NMR can be expected. Higher-dimensional experiments require too much acquisition time to be used for such measurements of protein dynamics, so this method opens up the possibility of using such measurements for the first time. Since higher dimensions can be optimized even more effectively than $3D$, even fewer sample points are needed to get the same signal to noise ratio. This suggests a change in the operation of such experiments, but this will require that complete multi-dimensional peak information be pre-determined, which is not current practice.

Building on these results, the model and decomposition methods should be extended to solve other problems in imaging and multi-dimensional NMR. In particular, this method should be modified to accelerate the determination of full higher-dimensional peak locations.

Acknowledgements

The authors acknowledge Alex Bain, Giuseppe Melacini and Rahul Das who explained the details of the problem, and provided data for benchmarking. NSERC, CFI and OIT provided research support.

REFERENCES

Anand C.K., Sotirov R., Terlaky T. and Zheng Z. Magnetic resonance tissue quantification using optimal bssfp pulse-sequence design. *Optimization and Engineering*, 2007. **8**(2), 215–238.

Borchers B. Csdp, a c library for semidefinite programming. *Optimization Methods and Software*, 1999. **11**, 613–623.

Coggins B.E. and Zhou P. Sampling of the nmr time domain along concentric rings. *Journal of Magnetic Resonance*, 2007. **184**, 207–221.

Conn A.R., Gould N.I.M. and Toint P.L. *Trust Region Methods* (MPS–SIAM Series on Optimization), 2000.

Hoch J.C. and Stern A.S. *NMR Data Processing* (Wiley-Liss), 1996.

Kupce E. and Freeman R. Hyperdimensional NMR spectroscopy. *J Am Chem Soc*, 2006. **128**(18), 6020–6021.

Liang Z.P. and Lauterbur P.C. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective* (Wiley-IEEE Press), 1999.

Nagayama K., Bachmann P., Wüthrich K. and Ernst R. The use of cross-sections and of projections in two-dimensional nmr spectroscopy. *J. Magn. Reson.*, 1978. **31**, 133–148.