Published in final edited form as:

IEEE Trans Med Imaging. 2009 July ; 28(7): 1093–1104. doi:10.1109/TMI.2008.2012030.

Improved Time Series Reconstruction for Dynamic Magnetic Resonance Imaging

Uygar Sümbül, Juan M. Santos, and John M. Pauly

Abstract

Time series of *in-vivo* magnetic resonance images exhibit high levels of temporal correlation. Higher temporal resolution reconstructions are obtained by acquiring data at a fraction of the Nyquist rate and resolving the resulting aliasing using the correlation information. The dynamic imaging experiment is modeled as a linear dynamical system. A Kalman filter based unaliasing reconstruction is described for accelerated dynamic magnetic resonance imaging (MRI). The algorithm handles arbitrary readout trajectories naturally. The reconstruction is causal and very fast, making it applicable to real-time imaging. *In-vivo* results are presented for cardiac MRI of healthy volunteers.

1 Introduction

Dynamic MRI is an important tool to monitor changes in tissue structure over time. It comprises a series of data acquisitions in the spatial frequency domain, known as k -space, from which a time series of images is formed. However, imaging speed often limits the ability to meet other imaging requirements of spatial resolution, temporal resolution, field of view (FOV) and signal-to-noise ratio (SNR). In *in-vivo* imaging, perhaps the most stringent requirements come from cardiac imaging applications. The complex motion of different parts of the heart such as the valves and the heart wall present visualization challenges. Also, free-breathing and untriggered scanning is desirable for both the physician and the patient, adding to the complexity of the problem. Since the speed of data acquisition is limited by both physical and physiological constraints, speed-up techniques in the reconstruction have become increasingly popular [1–3]. Classical sampling theory predicts that using a fraction of the full image data results in aliasing, severely degrading the image quality. Many speed-up techniques unalias the images by exploiting the redundancy in the temporal direction. Intuitively, two consecutive images from a time series should not differ by much if a satisfactory video is to be formed. This observation is at the heart of many estimation/prediction algorithms in general video processing.

MR images are formed from k -space data simply by two-dimensional Fourier inversion if the sampling pattern conforms to a Cartesian grid. There are, however, cases where non-Cartesian sampling is more advantageous due to the nature of data acquisition in MRI [4]. The gridding algorithm [5, 6] is an efficient way to obtain the values on the grid points.

Previous work in accelerated dynamic MRI include, among others, the UNFOLD method [7] which treats the problem from a k - t space packing perspective, the DIME method [8] which models the time variation of the image pixels using a parametric harmonic model, the PARADISE method [9, 10] which designs an optimal acquisition sequence using [11], Khalsa and Fessler's method [12] which treats the problem as a regularized reconstruction trading-off spatial and temporal resolutions, and the FOCUSS method [13] which presents a sparse reconstruction method, generalizing k - t BLAST/ k - t SENSE [1].

Among these methods, k - t BLAST/ k - t SENSE [1, 14] has received considerable attention. It tries to unalias the images of a time series by using individual pixel variances as prior information, employing the Wiener filter. The original method is practical only with Cartesian trajectories due to computational load. The authors proposed an iterative solution alleviating the problem [15], but Cartesian sampling remains much more favorable. Like the original k - t BLAST algorithm, the reconstruction is non-causal and uses the whole data set at once so that the per image reconstruction time depends on the duration of the experiment. k - t BLAST ignores any cross-correlations between pixels. In addition to the natural cross-correlations, aliased pixels are highly correlated in the Cartesian case. However, the results reported in [1, 14, 15] among others demonstrate that individual pixel variance information is enough to reasonably unalias the images in many applications. Using cross-correlation information would increase the computational complexity significantly. Another consideration is the impracticality of obtaining accurate joint estimates. The ratio of the number of observations to the number of variables is critical in arriving at accurate estimates [16].

Non-Cartesian trajectories do not create exact aliased replicas, and the aliased energy diffuses incoherently into many pixels. Such trajectories are thus more amenable to ignoring cross-correlations. Moreover, the flow and motion properties are generally superior [4]. The drawback, however, has been the increase in computational load. In this paper, we present a fast algorithm for arbitrary trajectories that shares the same statistical flavor as k - t BLAST. The algorithm is based on the well-known Kalman filter, which is a linear statistical filter like the Wiener filter used in k - t BLAST. Basically, the proposed algorithm can be interpreted as providing pixel-dependent causal reconstruction windows whose shapes depend on image statistics, instead of the symmetrical rectangular window of the sliding window reconstruction. Therefore, the algorithm should be appropriate whenever the sliding window reconstruction can be used. The proposed algorithm requires only two undersampled gridding and two Fourier transform operations per image, so it can be implemented efficiently. Another important aspect is causality, which makes real-time reconstruction possible. Real-time imaging allows for an interactive scanning experience, much like an ultrasound scan. Localization is easier and total scan time is reduced. Benefits of real-time cardiac MRI in a clinical setting are demonstrated in, for instance, [17]. Moreover, applications such as MR image guided therapy and catheter tracking will benefit from high-temporal-resolution reconstructions. When causality is not an issue, the estimate using all the available data can be obtained by combining the causal Kalman reconstruction with an anti-causal Kalman reconstruction running in the reverse direction.

In this study, we provide various theoretical comparisons between the proposed method and k - t BLAST due to the connection between the Kalman filter and the Wiener filter. For evaluation purposes, however, we compare our results to those of the sliding window reconstruction due to its real-time imaging ability. While the sliding window reconstruction is somewhat dated and many methods have since been proposed, it is used almost exclusively in clinical applications and its strengths and weaknesses are well known.

Cardiac imaging is used to motivate this paper. However, the algorithm does not use anything specific to cardiac imaging and hence should be directly applicable to any *in-vivo* time series MRI application. To better describe the basic algorithm and justify the various assumptions, we refrain from cardiac-specific extensions of the basic model. Similarly, quantitative studies on the results are out of the scope of this paper.

In the next section, we introduce the algorithm using a state-space formalism. Section 3 explains the experiments performed to test our method. In Section 4, we report the results of the *in-vivo* experiments performed. Section 5 concludes with a discussion of the method and

possible future work. In the Appendix, we provide a theoretical treatment of the steady state condition to aid the main text.

Throughout the paper, $\hat{a}$ represents an estimate of the generic parameter a , $\text{var}(a)$ denotes the covariance of a , and a^H denotes the Hermitian conjugate of a .

2 Theory

2.1 Overview

In this section, we develop a practical accelerated reconstruction algorithm for arbitrary readout trajectories. The algorithm unaliases the images of a time series by constraining the undersampled data using temporal statistics as prior information. It is based on a linear state-space description of the dynamic imaging experiment, as detailed in Section 2.3. This model not only helps enhance our understanding of the dynamic imaging experiment, but also provides a natural framework for the Kalman filter. A straightforward implementation of the Kalman filter is highly impractical due to the large number of state variables. Every pixel in an image corresponds to a different state variable. Therefore, we introduce approximations to the Kalman iteration that simplify the calculations tremendously. These approximations and their consequences are discussed in Section 2.4. Section 2.5 extends the algorithm to multi-channel acquisitions. Section 2.6 describes how various initialization tasks can be performed.

2.2 The Kalman Filter

Consider the state-space description of a system with state s_t observed with measurements x_t at time t ,

$$\begin{aligned} s_{t+1} &= A_t s_t + B_t u_t \\ x_t &= H_t s_t + w_t. \end{aligned} \quad (1)$$

Here, s_t is the $m \times 1$ state vector. A_t is the state-transition matrix, describing the deterministic aspect of state transition. u_t is the $p \times 1$ system noise vector and $B_t u_t$ describes the stochastic disturbance to the state transition. The measurements are described by the $n \times 1$ observation vector x_t . H_t is the observation matrix relating the state to the measurements. w_t is the $n \times 1$ observation noise vector. A_t , B_t and H_t are known $m \times m$, $m \times p$ and $n \times m$ matrices, respectively with $1 \leq p, n \leq m$. The problem is to estimate the state sequence from the observations using the statistics governing the linear model of time evolution. When we adapt Eq. 1 to dynamic MRI in Section 2.3, s_t will denote the true image at time t and x_t will denote the measurements in k -space at time t .

The Kalman filtering equations provide a solution using $Q_t = \text{var}(u_t)$ and $R_t = \text{var}(w_t)$, $t = 1, 2, \dots$. The estimate of s_t obtained from the observations $x_0, x_1, \dots, x_t, \hat{s}_{t|t'}$, is calculated through updating $P_{t,t'}$, the covariance of the estimation error of the best linear estimator of s_t given all the observations up to $t' \leq t$. To simplify the notation, we will use $\hat{s}_t = \hat{s}_{t|t}$. The Kalman filter is known to give the least-squares optimal state sequence estimate when $\{u_t\}$ and $\{w_t\}$ are uncorrelated zero-mean white Gaussian noise sequences, and the initial state is uncorrelated to these noise sequences. It is widely used in real-time tracking applications. At each time step, all the relevant information is stored in the auxiliary variables and the previous estimate so that computations required at each step remain constant. For a general reference on Kalman filtering, see [18].

2.3 Dynamic Imaging Experiment as a Linear Dynamical System

MR image acquisition is inherently a linear process since the actual image is related to the observed data through the Fourier transform. Coupling this with a linear model of image

(state) evolution gives a linear state-space description of the dynamic MRI experiment. A model of an arbitrarily sampled dynamic MRI time series acquisition is

$$\begin{aligned} s_{t+1} &= s_t + u_t \\ x_t &= G_t F \Gamma^{-1} s_t + w_t, \end{aligned} \quad (2)$$

where s_t denotes the true noise-free image arranged into a one-dimensional vector, u_t denotes the difference between consecutive true images, w_t denotes the acquisition noise, x_t denotes the actual scanner data. For instance, in a Cartesian experiment with an acceleration factor of 4 and an image size of 100×100 , s_t denotes a vector of length $100 \times 100 = 10000$, and x_t and w_t denote vectors of length $100 \times 100/4 = 2500$ each. G_t denotes the inverse gridding matrix at time t , F denotes the Fourier transform matrix and Γ and Γ^{-1} denote the apodization and deapodization matrices. That is, G_t , F , and Γ are the matrix representations of the corresponding operators when the image is stored as a column vector. Note that Γ is a diagonal matrix satisfying $F G_t = \Gamma F$ due to the well-known convolution property of the Fourier transform when data points are the full Cartesian grid. Γ^{-1} exists whenever the diagonal entries $\Gamma(i, i) \neq 0$. Thus, in our case Γ^{-1} always exists since a finite extent conventional gridding kernel is used. The oversampling required in the gridding algorithm [19] is achieved by choosing an overdetermined F matrix.

The second part of the description in Eq. 2 reflects what is physically happening in data acquisition, whereas the first part is the model of the state evolution which describes the evolution of the object. Motion is considered as a random process and it is modeled through estimating the second moment of the *system noise*, the difference between consecutive images. Hence, linear estimation becomes possible. Higher moments of the temporal variations in pixels are ignored. This resembles truncation in Taylor expansion of functions. Arbitrary accuracy in the state evolution characterization is possible, at least in principle, as the time step decreases.

One might also consider using a more involved model for the state evolution (e.g., an autoregressive model). This would enhance the deterministic aspect of the model as opposed to the stochastic aspect that we are attempting to capture here. Such a model can be characterized by tracking derivatives of the state, in addition to the state itself. However, not only will it be more prone to noise enhancements, but cardiac diseases such as arrhythmia will most likely violate the model in a free-breathing and non-gated experiment, defying the purpose. The proposed model simplifies the computations. By not including any cardiac-specific terms, it is also directly applicable to other time-resolved studies.

The acquisition noise in MRI is thermal noise and can be modeled as a zero-mean white Gaussian process. It is uncorrelated to the images and their time evolution. Furthermore, due to its thermal nature, the variance of the acquisition noise is the same for all k -space samples. Then, $\text{var}(w_t) = \sigma^2 I$ for some scalar $\sigma > 0$. Finally, we can safely assume that the initial state is uncorrelated to the system and observation noise sequences due to asynchronous image acquisition, and the physical nature of thermal noise.

The whiteness of the system noise process $\{u_t\}$ is not guaranteed because of our simple model. In the Appendix, we suggest a more detailed model addressing this issue and point out the resulting penalties.

If we insert the specifics of our model into the general Kalman recursion [18], we get

$$\begin{aligned}
P_{0,0} &= \text{var}(s_0), \widehat{s}_0 = E(s_0) \\
P_{t,t-1} &= P_{t-1,t-1} + Q_{t-1} \\
K_t &= P_{t,t-1} H_t^H (H_t P_{t,t-1} H_t^H + \sigma^2 I)^{-1} \\
&= \sigma^{-2} P_{t,t} H_t^H \\
P_{t,t} &= (I - K_t H_t) P_{t,t-1} \\
\widehat{s}_t &= \widehat{s}_{t-1} + \sigma^{-2} P_{t,t} H_t^H (x_t - H_t \widehat{s}_{t-1}),
\end{aligned} \tag{3}$$

where $H_t = G_t F \Gamma^{-1}$. A few lines of algebra yields

$$P_{t,t} = P_{t,t-1} (I + \sigma^{-2} H_t^H H_t P_{t,t-1})^{-1}. \tag{4}$$

$P_{t,t}$ denotes the variance of the estimation error. Intuitively, it makes sense to update the state variables more aggressively when the corresponding error is large and more conservatively when the corresponding error is small. This is exactly how the Kalman filter utilizes $P_{t,t}$. As a result of the identity A matrix in Eq. 2, we have $\widehat{s}_{t-1} = \widehat{s}_{t-1}$. This means that we cannot further update our estimate without accessing the current observation. The contribution of the current observation is precisely the second term on the right-hand side of the last line in Eq. 3.

The Kalman gain K_t plays an important role in the classical analysis of the filter. Since $\text{var}(w_t) = \sigma^2 I$ in our case, one can restate these results in terms of $P_{t,t}/\sigma^2$ for more insight. By the last line of Eq. 3, $y_t = H_t^H (x_t - H_t \widehat{s}_{t-1})$ is the raw update vector before applying the

statistical priors. Let $\nu(P_{t,t}/\sigma^2) = \frac{\|\sigma^{-2} P_{t,t} y_t\|_2}{\|y_t\|_2}$ denote the *amplification* for y_t and let us review three special cases that illustrate the role of Q :

- $\nu(P_{t,t}/\sigma^2)$ is very small: The reconstruction relies on the previous estimate with minimal contribution from the new measurement. The data might be changing very slowly so that the previous estimate is still accurate, or the current measurement may have too few samples. Another possibility is that the acquisition noise may be too large so that the current measurement is very unreliable.
- $\nu(P_{t,t}/\sigma^2)$ is very large: The reconstruction relies on the new measurement with minimal contribution from the previous estimate. The data may be changing very quickly so that the previous estimate has become almost irrelevant. Another possibility is that the current measurement may be accurate enough with relatively little noise that the estimate can be based almost entirely on the current measurement, making the previous estimate obsolete.
- The diagonal elements of $P_{t,t}$ are equal to a constant c : In our experiments, this corresponds to an unsuccessful Q estimate because the algorithm works by suppressing the changes in some pixels while accentuating others to resolve the aliasings in y_t created by incomplete measurements. The next section will make explicit use of this interpretation.

Equation 4 dominates the computational and memory requirements due to the matrix inversion. For a modest 100×100 image, one needs to invert a 10000×10000 matrix, clearly a very demanding task. Moreover, this has to be done for each image in a dynamic imaging experiment. Thus, we have to circumvent the matrix inversion step.

2.4 Diagonality Approximation

A key goal of this paper is real-time reconstruction of non-Cartesian sampling for dynamic imaging. We need to impose some structure on the matrices appearing in Eq. 4 to perform

the matrix inversion in a reasonable time for each frame of a time series. One such way that results in huge savings and offers insight into the mechanics is to ignore the off-diagonal entries of Q_t and $H_t^H H_t$. This simplification is motivated by the *in-vivo* MR experiment itself, as mentioned in the introduction.

Figure 1 shows a typical plot of the normalized absolute value of the autocorrelation estimate of the difference between two consecutive MR images as a function of pixel distance. The time between the images is 24 ms. For illustration purposes, we assumed spatial wide-sense stationarity. This enables the compact visualization in Fig. 1. The figure is obtained by first finding the power spectral density of the difference image and then inverse Fourier transforming the result due to the Wiener-Khinchine theorem [20]. As seen in Fig. 1, temporal subtraction already eliminates much of the cross-correlations. Hence, imposing diagonality on Q_t results in relatively little loss of information. Unlike *k-t* BLAST, diagonality is assumed on the covariance of the difference between consecutive images and not the covariance of the difference between an image and the “baseline estimate.” This is a much more robust approximation. Most of the cross-correlations in natural images are due to the low spatial-frequencies that result in many almost identical neighboring pixels. These low spatial frequency components change more slowly in time than the high spatial frequency components and hence subtraction leads to much more dominant diagonal entries. Even when components change quickly together in, for instance, perfusion studies, low spatial frequencies of adjacent frames tend to be more similar than the high spatial frequencies. When reasonable tracking is maintained, $P_{t,t-1} = \text{var}(s_t - \hat{s}_{t-1})$ and $P_{t,t} = \text{var}(s_t - \hat{s}_t)$ should also have dominant diagonals.

The matrix $H_t^H H_t$ in Eq. 4 contains the aliasing information. The off-diagonal terms of $H_t^H H_t$ are manifestations of the incurred aliasings. For instance, upon undersampling an image over a Cartesian grid it creates exact replicas, as required by the classical sampling theory. Therefore, while many elements of the matrix $H_t^H H_t$ are exactly zero, there are also off-diagonal stripes of elements that are as strong as the diagonal elements. If the sampling pattern does not conform to a grid, $H_t^H H_t$ does not create exact replicas and aliasing diffuses incoherently into many pixels. When spiral sampling is employed, many off-diagonal entries of $H_t^H H_t$ exhibit small non-zero values, rather than fewer non-zero entries with large values. This is the interpretation of the assertion that spirals have better point spread functions than Cartesian trajectories for dynamic imaging. The incoherency is conducive to ignoring the off-diagonal elements because aliasing adds seemingly noise-like components, instead of strong replicas [3,21]. The simple summation of many small incoherent complex terms is small due to cancellations even though their total energy is large. Figure 2 shows the absolute value of a row of $H_t^H H_t$ from the experiments reported in this paper, arranged in image format. Looking at a single row is sufficient due to shift-invariance. Perfect diagonality corresponds to a single non-zero value in the middle of the image. The large number of small non-zero off-diagonals suggests the incoherency of aliased energy. Considering the dramatic savings, the diagonality assumption looks quite acceptable. We investigate the effects of these assumptions at the end of this subsection. The difference between $H_t^H H_t$ and the point spread function is that $H_t^H H_t$ does not have a density compensation factor. Therefore, $H_t^H H_t$ also reflects the sampling density differences. Ignoring the off-diagonals then results in energy accumulation where *k*-space samples are closer. This becomes a problem especially if readout occurs during the slew-limited regime due to the huge differences in sampling density. This issue is addressed at the end of this section. We mention that *k-t* BLAST does not need any approximations in the corresponding computations.

The diagonal model is consistent in that the Kalman iterations preserve the diagonality of these matrices. It also simplifies Eq. 4 tremendously. In effect, each pixel is independently processed at the most time consuming steps. Moreover, only the diagonal elements of $H_t^H H_t$ are needed and they can be precomputed because $H_t^H H_t$ does not depend on the data. Since $H_t^H H_t$ is shift-invariant, computing just one of the diagonal elements suffices, greatly simplifying the precomputation. When a unitary Fourier transform is used, this single value equals the number of k -space points in an observation divided by the number of pixels in an image. Let $\times$ and $\div$ denote elementwise multiplication and division, respectively. When the diagonal elements of $P_{t,b}$, Q_t and $Z_t = H_t^H H_t$ are read into images the same size as the reconstructed image, the following algorithm implements the Kalman iteration efficiently:

$$\begin{aligned} P_{0,0} &\leftarrow \text{var}(s_0), \hat{s}_0 = E(s_0) \\ P_{t,t} &\leftarrow P_{t-1,t-1} + Q_{t-1} \\ P_{t,t} &\leftarrow P_{t,t} \div (\mathbf{1} + P_{t,t} \times Z_t / \sigma^2) \\ \hat{s}_t &\leftarrow \hat{s}_{t-1} + \sigma^{-2} P_{t,t} \times [H_t^H (x_t - H_t \hat{s}_{t-1})], \end{aligned} \quad (5)$$

where $\mathbf{1}$ is an $n \times n$ matrix, all of whose entries are equal to 1. Note that gridding is indeed performed in state update, whereas covariance update uses a precomputed Z_t . Thus, the dominant operations in the iteration become the gridding and inverse gridding reconstructions.

In many dynamic imaging applications, the time span is short enough that Q_t remains approximately constant during the experiment. By replacing Q_t with a constant Q , the above-given algorithm will reach a practical steady state exponentially fast (See Appendix and Chapter 6 of [18]). Exact steady state is not reached because the observation matrix is time-dependent. In the Appendix, we suggest a concatenated state-space description satisfying $H_t = H$ and reaching exact steady state. In our experiments, we observed that a steady state is practically reached after only a few frames. Therefore, we can analyze the effects of the various assumptions under the steady state condition. The steady state estimation error covariance $\hat{P}$ that our algorithm will converge to is exactly that of a system with the diagonal Q and the diagonal $\sigma^{-2} H_t^H H_t$. Let the computed estimation error covariance converge to $\hat{P}$ instead of the actual P . Then, Eq. 3 suggests that the error introduced by $P - \hat{P}$ at each step and the contribution of each observation will decay in an exponential envelope and the error will appear as a combination of temporal blurring and unresolved aliasing depending on the acquisition noise estimate and the fidelity of the diagonality assumption. This is the window interpretation suggested in the Introduction. Instead of the crude rectangular window of the sliding window reconstruction, different sets of temporal coefficients are used for different pixels of a frame.

For a divergence analysis in steady state, we first note that the errors in both Q and $\sigma^{-2} H^H H$ can be reflected to Q only, via Eq. 4. One can then assume that $\sigma^{-2} H^H H$ is calculated correctly and the only source of error is the estimation of Q , $\hat{Q}$. When the computed error covariance converges to $\hat{P}$, the apparent $\hat{Q}$ is found as $\hat{Q} = [I - \hat{P}(\sigma^{-2} H^H H)]^{-1} \hat{P} - \hat{P}$, after a few lines of algebra. Ref. [22, 23] analyze the convergence behavior of the Kalman filter under incorrect noise covariances. In particular, if there exists a vector in the left null space of $\hat{Q}$ that is not in the left null space of Q , then the reconstruction diverges [22, Theorem 4.2]. In our experiments, we observed divergence only when the sampling density variation is large. This corresponds to a highly erroneous $H_t^H H_t$ after diagonality is imposed, which in turn corresponds to a highly erroneous apparent $\hat{Q}$.

Many non-Cartesian trajectories exhibit large variations in the sampling density due to the slew-limited regime. This is viewed as a benefit in many applications because data

acquisition becomes overdetermined around the origin. The Kalman filter will also use all the available samples to come up with the least-squares optimal solution. Obtaining more samples cannot hurt. On the other hand, ignoring the off-diagonals results in an accumulation where k -space samples are closer. Therefore, the Kalman filter will diverge, nullifying the development. The exact solution to this problem is to resample uniformly along the trajectory. However, if the reconstruction budget is tight (i.e, real-time), one can ignore samples at a rate proportional to the sampling density to obtain an approximately uniform density. This is a rather benign operation because SNR is already very high around k -space origin, where samples are discarded. One might even consider the reduced reconstruction time as a benefit. We point out that both of these solutions can handle arbitrary readout trajectories. The behavior of the Kalman filter and its divergence have been studied extensively over the years. The reader may refer to [22–24] and the references therein.

2.5 Multi-Coil Case

We are now ready to extend the basic algorithm to the multi-coil case. Let $C_1, \dots, C_c$ denote the sensitivity matrices of individual coils obtained by reading the corresponding sensitivity maps into diagonal matrices. The state-space representation of Eq. 2 then becomes

$$\begin{aligned} s_{t+1} &= s_t + u_t \\ x_t &= \underbrace{\begin{bmatrix} G_t F \Gamma^{-1} C_1 \\ \vdots \\ G_t F \Gamma^{-1} C_c \end{bmatrix}}_{H_t} s_t + w_t. \end{aligned} \quad (6)$$

By relabeling H_t as the observation matrix in Eq. 6, the algorithm given in Section 2.4 becomes directly applicable. Sensitivity maps can be used in elementwise multiplication, doing away with the redundancy of the formal representation. In effect, single-coil reconstruction is performed once for each coil without any extra cost.

Let $\bar{Z}_t = H_t^H R_t^{-1} H_t$ for the multi-coil case, where $R_t = \text{var}(w_t)$ is a diagonal matrix whose diagonals are given by $\rho = \text{diag}(R_t) = [\sigma_1^2 \dots \sigma_1^2 \dots \sigma_c^2 \dots \sigma_c^2]$. That is, we allow individual coils to have different noise variances. Coil noise cross-correlations are ignored, which are usually small [4]. We obtain, after a few lines of algebra,

$$\bar{Z}_t = \sum_{i=1}^c \sigma_i^{-2} C_i^H Z_i C_i, \quad (7)$$

where Z_i is as defined in the single-coil case. Therefore, $\bar{Z}_t$ can be computed at virtually no additional cost, compared to the single-coil case.

We now provide the multi-coil extension of the algorithm given in Section 2.4:

$$\begin{aligned} P_{0,0} &\leftarrow \text{var}(s_0), \widehat{s}_0 = E(s_0) \\ P_{t,t} &\leftarrow P_{t-1,t-1} + Q_{t-1} \\ P_{t,t} &\leftarrow P_{t,t} \div (\mathbf{1} + P_{t,t} \times \bar{Z}_t) \\ \widehat{s}_t &\leftarrow \widehat{s}_{t-1} + P_{t,t} \times [H_t^H ((x_t - H_t \widehat{s}_{t-1}) \div \rho)], \end{aligned} \quad (8)$$

The multi-coil extension described above combines all the coil data at a single time point to reconstruct the image at that time point. The method requires explicit sensitivity information. The advantage is that the localized sensitivities of the coils enhance the

diagonality assumption on $\bar{Z}_t = H_t^H R_t^{-1} H_t$. Sensitivity estimates can be obtained from an initialization scan or from the actual data and we have found that the algorithm is robust against variations in these estimates. The effects of using incorrect observation matrices may be analyzed under steady state as described in the previous subsection. This *combine-then-reconstruct* approach possesses some of the features of SENSE-like [25] parallel image reconstruction methods. It is also possible to reconstruct each coil data independently according to the single-coil algorithm, and combine the individual coil images by a sum-of-squares reconstruction. The obvious advantage of this method is that the need for sensitivity estimates is eliminated. This *reconstruct-then-combine* approach possesses some of the features of GRAPPA-like [26] parallel image reconstruction methods. Both of these methods are viable for parallel computation (i.e., each node reconstructing a single coil image), the *reconstruct-then-combine* approach being more straightforward. In this paper, we implement the *combine-then-reconstruct* algorithm in Eq. 8.

Equation 8 suggests that the algorithm requires essentially only two undersampled gridding and two $2 - D$ Fourier operations per image per coil, twice that of the sliding window reconstruction. Latency of the proposed algorithm is very low and at most equal to the reconstruction time, unlike the sliding window algorithm which waits for neighboring future data. Equation 8 provides a pseudocode using only standard routines and the four operations of arithmetic. Therefore, complexity and machine-specific timing calculations can be performed via Eq. 8.

2.6 Obtaining Signal Estimates

The algorithm described in the previous subsections requires the knowledge of the expected starting state $E(s_0)$, the covariance of $\{u_t\}$, $\text{var}(u_t) = Q_t$, the acquisition noise covariance $\text{var}(w_t)$, and the initial error covariance $P_{0,0}$. These estimates may be obtained through separate initialization scans or the data itself depending on the application. With the common ergodicity assumption, we use the available temporal data to come up with the estimates, as widely employed in the medical imaging literature.

It is important to maintain correct normalizations of the operators throughout initialization and reconstruction.

Initial State Estimation—The expected starting frame $E(s_0)$ can be obtained by conventionally reconstructing the first few interleaves through gridding. One can also obtain coil sensitivity estimates as a by-product.

System Noise Covariance Estimation—In an auto-calibrating mode, Q_t can be estimated by computing the sample covariance of $\{u_t\}$ within a sliding window over a fully sampled centric disc. In this study, we are interested in scans that are around 10 seconds long and we assume that within that time interval Q_t does not change significantly. For the offline reconstructions reported in this paper, we acquired separate initialization data and obtained a single estimate $\hat{Q}_t = \hat{Q}$.

Hansen *et al* [28] suggested that only the low spatial harmonics are adequate in characterizing individual pixel variances, as employed in the k - t BLAST method [1]. We also observed a similar behavior in our data sets. We mention, though, that our low resolution scans cover a centric disc in k -space and capture a greater amount of meaningful temporal information. Therefore, one low spatial resolution (and high temporal resolution) initialization scan suffices. This suggests robustness of the filter to small variations in the Q estimate since contribution of higher spatial harmonics have little energy (See Fig. 4(a)).

In this paper, we also pursue obtaining the signal estimates from all (or a greater portion) of k -space although the contribution of multiple initialization scans will be vital only when cross-correlations are not ignored. Obviously, a straightforward strategy is not adequate to obtain such initialization data. Otherwise, one would not need speed-up techniques in the first place. Being *natural images*, *in-vivo* MR images are asymptotically decorrelated by the Fourier transform, except for the conjugate symmetry redundancy [29, 30]. That is, k -space data exhibits rapidly decaying correlations as a function of k -space distance. To a very good degree of approximation, k -space data is thus uncorrelated. Consider then, dividing k -space into r chunks as in Fig. 3. These chunks can have arbitrary shapes so long as conjugate points are included in the same chunk. Let $S_{t,i}$ denote the images obtained from only the i^{th} chunk. Then,

$$\begin{aligned} Q &= \text{var}(u_t) = \text{var}(S_t - S_{t-1}) \\ &= \sum_{i=1}^r \text{var}(S_{t,i} - S_{t-1,i}), \end{aligned} \quad (9)$$

where we used uncorrelatedness of the k -space chunks. k -space is divided into non-overlapping chunks instead of interleaved acquisitions in order not to compromise the uncorrelatedness. Thus, although separate time series of these chunks cannot be combined to reconstruct images due to asynchronous scans, they can be used to come up with first- and second-order signal statistics. The repetition times of these readout trajectories are designed to be the exact same as that of the actual scan. Each initialization scan is reconstructed independently and the calculated diagonal entries of the individual Q matrices are simply summed up to give $\hat{Q}_{raw}$. We note that the efficiency of this method decreases as the acceleration factor increases because the number of initialization scans needed to cover k -space equals roughly the acceleration factor. This also means that the first initialization scan, which indeed covers a disc, goes out to roughly (acceleration factor) $^{-1/2}$ of the maximum k -space radius. In that case, obtaining a few initialization scans to cover part of k -space might make more sense and the widely used method of obtaining a single initialization scan becomes a special case.

Figure 4(a) shows typical individual Q estimates of five non-overlapping spiral rings that cover k -space, obtained from a GRE experiment. Notice that the effect of the acquisition noise becomes more apparent as the signal level drops towards the edge of k -space. The low resolution data seems successful in obtaining the estimate and may be used by itself. Yet, there are still some contributions coming from the outer spiral rings that have different intensity distributions. Figure 4(b) shows the low resolution Q estimate from another GRE experiment. Note that between the two experiments, the imaging plane and the temporal characteristics change. Lastly, as a comparison, Fig. 4(c) shows the low resolution Q estimate from a steady state free precession (SSFP) experiment. As expected, the temporal behavior is completely different from the GRE experiments. We infer that the signal from the blood pool varies rapidly in time in the GRE experiment. In contrast, the edges of the myocardium are the brightest parts of the temporal variation map in the SSFP experiment.

The Q estimate obtained from scanner data $\hat{Q}_{raw}$ is corrupted by the acquisition noise, as seen in Fig. 4. Assuming that a background pixel with vanishing variance exists in the noise-free case, subtracting the offset above zero eliminates the noise contribution. In practice, due mainly to disturbing swirl artifacts of spirals, subtracting the offset above zero may not be enough. A practical solution is to subtract twice the offset and set all values smaller than a small positive value to that small positive value to preserve positive-definiteness of Q .

Acquisition Noise Covariance Estimation—There are a few convenient ways to estimate the acquisition noise covariance. For our purposes a rough estimate suffices

because we treat the acquisition noise covariance as a parameter trading-off image noise for faster tracking of moving parts and usually end up using larger values.

For data acquisition we used the RTHawk real-time system [31, 32], whose flexibility makes it very convenient to obtain a direct estimate of the acquisition noise covariance. After acquiring data, we set the flip angle to 0° very briefly, without even stopping the scan. Those samples will be purely noise and an estimate is obtained by computing the sample covariance. Perhaps the easiest way to obtain a rough estimate of the acquisition noise covariance is to compute the sample covariance of the outermost k -space sample of the actual scan. Although this method neglects the signal in that sample, the resulting estimate is adequate for our purposes.

Initial Error Covariance Estimation—Solving the $P_{t,t}$ recursion (line 3, Eq. 8) assuming that a steady state is achieved is a very effective way to initialize $P_{0,0}$. It merely requires solving a quadratic equation when diagonality is enforced. This way, the Kalman filter provides good reconstructions immediately. In the Appendix, we address the steady state condition.

3 Experiments

Free-breathing, untriggered dynamic cardiac imaging experiments were performed on a 1.5T GE Signa system, using the GE 8-element cardiac array that wraps around the torso. Informed consent was obtained from three healthy volunteers before *in-vivo* experiments. A fat-suppressed GRE excitation with 30° flip angle was used. Slice thickness was 4.7 mm. Initialization data was acquired with the same scanner configuration. Individual scan duration is chosen as approximately 10 seconds to include common bulk motions such as breathing. A full experiment takes about one minute with the current setup. Reconstruction was done offline. We do not report auto-calibrating experiments in this paper.

We designed 8-, 12- and 16-interleaf spiral trajectories. Each interleaf corresponds to a different time point so that a new image is obtained using a single interleaf, corresponding to $8\times$, $12\times$, and $16\times$ accelerated reconstructions. We conducted the twelve-interleaf experiment on two volunteers. One of these experiments was reconstructed to yield $6\times$ acceleration by combining every two interleaves, in addition to the regular $12\times$ acceleration. The Q matrix for the $6\times$ reconstruction was obtained by averaging every two images in the initialization data.

We used an oblique slice going through the four-chamber view of the heart so that cardiac valves move in and out of the imaging slice. The rapid motion of the valves helped us compare reconstructions. The oblique slice goes through all of the upper torso. The FOV is thus set as 42 cm. The in-plane resolution is 2 mm. Remaining imaging parameters are given in Table 1.

As described at the end of Section 2.4, we discarded some of the acquired samples to arrive at approximately uniform sampling densities. The discarded samples come from the already oversampled, high SNR, slew-limited central part. This suboptimal method is used for a fast reconstruction. The corresponding SNR loss can be eliminated by first resampling uniformly along the trajectory (See Section 2.4). The percentage of discarded samples depends on how quickly gradient-limited regime is achieved (Table 1). No data was discarded and a symmetric temporal window is used in the sliding window reconstruction. Therefore, comparison is slightly biased in favor of the sliding window reconstruction. Lastly, the sliding window reconstruction also accessed the available sensitivity and coil noise estimates to achieve SNR optimality.

We employed three initialization scans going out to approximately 40% of the maximum k -space extent when combined. The lower two channels at the back of the patient did not contribute any signal in some cases due to the distance to the excited plane. In those cases, we turned them off for faster reconstruction, although they cannot be harmful in the *combine-then-reconstruct* method when coil sensitivity estimates are reasonably accurate.

4 Results

Videos pertaining to the results reported in this section can be downloaded at <http://www-mrsrl.stanford.edu/publicfiles/kalman/>.

Figure 5 shows three consecutive images in time from the eight-interleaf experiment. The top row shows the images obtained by the sliding window reconstruction and the bottom row shows the corresponding Kalman reconstruction. The heart is in diastole and the tricuspid valve is open. The arrows point out the tip of the valve, where fastest motion occurs. This is blurred out in the sliding window reconstruction, whereas it is visible in the Kalman reconstruction.

Figure 6 shows consecutive images from the twelve-interleaf experiment in the same format. The heart is again in diastole. We observe that while the sharpness of the images of the two rows are comparable in general, the valve leaflets are quite blurry in the top row, but well depicted in the bottom row. This is expected since the speed of the leaflets can be quite high. Compare the motion due to breathing with the swings of the valve leaflets in diastole due to blood flow.

Figure 7 shows images from the same twelve-interleaf setup with a different volunteer. The heart is in systole, just before the valve opens up. The left image is obtained by the $12\times$ sliding window reconstruction, the middle image is obtained by the corresponding $6\times$ Kalman reconstruction and the right image is obtained by the $12\times$ Kalman reconstruction. The $6\times$ Kalman reconstruction is obtained by combining every two interleaves as a single observation. That is, it is a combination of the two approaches and the figure confirms that. The depiction of the valve leaflets is sharpest in the right image and most blurred in the left image.

Figure 8 shows consecutive images from the sixteen-interleaf experiment. The subject is the same as that of the first twelve-interleaf experiment. The experiments were performed during the same session with the same imaging slice. The heart is in diastole. Not only does the sliding window reconstruction result in more blur around the valve leaflets, but the position of the valve is different between the two reconstructions. This is due to the extremely long temporal window (around 0.3 seconds) of the sliding window reconstruction.

The video from the $12\times$ experiment looks better than that of the $16\times$ experiment in terms of temporal response. While this is an isolated observation, it is reasonable that an optimum speed-up factor exists for a particular reconstruction. As the amount of data in each observation decreases, the resulting image depends more and more on previous observations. In addition, the quality of the diagonality assumption on $H_t^H H_t$ degrades. This also suggests that enhancements such as better motion maps or non-diagonal algorithms can make valuable contributions. In general, the competitive advantage of the Kalman method over the sliding window reconstruction increases as the speed-up factor increases, as expected.

Figure 9 shows the temporal variation of a pixel close to the tip of a valve leaflet in diastole and the temporal variation of the vertical line containing that pixel for the $8\times$ and $12\times$

experiments. There is a time shift between the Kalman reconstruction and the sliding window reconstruction due to the symmetric temporal window of the sliding window reconstruction. This lag is approximately adjusted by introducing a shift of four time points in the 8 $\times$ experiment and six time points in the 12 $\times$ experiment. The range of the pixel amplitudes is [1, 256]. Both the coinciding positions of the peaks and the large differences in peak amplitudes suggest that the differences cannot be due to noise. The peak-to-peak variation is larger and the swinging motion of the leaflets are better depicted in the Kalman reconstruction. Overall temporal blurring of the sliding window reconstruction is perhaps better depicted in the line profiles.

5 Conclusions

Although non-Cartesian readout trajectories provide unique advantages in dynamic MRI such as better flow and motion properties, and efficient use of the gradients, the computational overhead has limited their use. We have presented a solution for non-Cartesian dynamic MRI by introducing a state-space formalism and applying the well-known Kalman filter. To arrive at a practical algorithm, we proposed various approximations, most notably diagonality of the large covariance matrices. The resulting algorithm is fast and causal, thereby a good candidate for real-time dynamic MRI. The algorithm is not iterative and the computationally dominant components are two undersampled gridding and two 2-D Fourier transform operations per image.

The algorithm requires the first and second order statistics of the time series. There are several ways of obtaining adequate statistical estimates. We proposed a method using multiple initialization scans. This obtains more accurate estimates, and allows for visualizing the contribution of each initialization scan. It is also possible to achieve an auto-calibration ability by fully sampling a small centric disc for each frame. The overhead of the fully sampled centric disc is minimal because of the slew-limited regime. This extension was out of the scope of our paper.

The results reported in this study demonstrate that our algorithm results in sharper reconstructions both in the temporal dimension and in the spatial dimensions when compared to the sliding window reconstruction, which is a basic and fast dynamic MRI acceleration algorithm. During our *in-vivo* experiments, we chose the imaging plane to include a fast moving cardiac valve so that the merit could be better assessed. While the valve almost disappeared in some frames in the sliding window reconstruction, it was nicely depicted in the Kalman reconstruction. We presented reconstructions with speed-up factors of 6, 8, 12, and 16. Upon qualitatively examining the resulting videos, we found that the highest quality reconstructions were obtained with the 8 $\times$ and 12 $\times$ reconstructions, although many factors come into play and different experiments may arrive at different results. Our main motivation was to test the algorithm at very high acceleration factors. Yet, those very high acceleration factors might be more useful in, for instance, SSFP experiments or 3-D scans.

As future work, a quantitative phantom study examining the effects of various parameters as well as comparing the Kalman method to other relevant methods should be helpful. The state update equation of the state-space formalism can be tailored to cardiac imaging to ensure the whiteness of $\{u_i\}$. Alternatively, a colored noise model can be employed. On a separate note, the Kalman algorithm also preserves the structure of block diagonal matrices. Therefore, such matrices trivially generalize our algorithm. Lastly, we note that the auto-calibrating mode makes the method instantly applicable to many time-resolved MRI applications.

Acknowledgments

This work was supported by NIH grants R01 HL067161, R01 HL074332, and R21 EB007715.

Appendix

Steady State

Let the integer a denote the acceleration factor for a particular experiment. Assume that H_t denotes the observation matrices applied in succession so that, when combined, the samples due to $H_t = H_{(mod\ a)}$ uniquely determine the whole rectangular k -space of the $2D$ discrete

Fourier transform. Then, consider tracking $\bar{s}_t = [s_{at+r}^H s_{at+r-1}^H \dots s_{a(t-1)+r+1}^H]^H$ instead of s_t , $0 < r < a$. The matrices in such a state-space description are time-independent and the noise sequences are still zero-mean Gaussian white, uncorrelated from each other. The system is also observable and controllable. Therefore, it will converge in ℓ_2 -norm to a steady state exponentially fast [18]. Moreover, the reconstruction should be identical to that of Eq. 2 (or Eq. 6) for s_{at+r} because the Kalman filter finds the least-squares optimal solution based on all the available past and present data, and $\{u_t\}$ is white. Hence, the error covariance matrix should also be identical for s_{at+r} . By varying r over $\{1, 2, \dots, a\}$, we find that $P_{t,t}$ of the original model converges to a periodic sequence of period a . Since we are merely rotating a spiral trajectory for each observation, the diagonals of $P_{t,b}$ only change by very small amounts within a period in *in-vivo* imaging, practically converging to a steady state exponentially fast.

In our experiments, we observed that the steady state is reached within the first few frames even though a circular region in k -space is covered. Noise-free zero measurements correspond to the basic imaging assumption of ignoring the data outside of k -space coverage.

Colored System Noise

The Kalman filter provides least-squares optimal reconstructions when the system noise sequence $\{u_t\}$ is white, besides other conditions. In this paper, we did not attempt to ensure whiteness due to the computational simplicity afforded by Eq. 2. Let the system noise $\{u_t\}$

be a colored sequence modeled by an autoregressive process $u_{t+1} = \sum_{i=0}^{t_0} V_i u_{t-i} + \beta_t$, where t_0 is an integer and $\{\beta_t\}$ is an uncorrelated zero-mean white Gaussian process. Then, a more elaborate model that whitens the system noise can be obtained by tracking

$[s_t^H u_t^H \dots u_{t-t_0}^H]^H$ instead of s_t .

References

1. Tsao J, Boesiger P, Pruessmann KP. k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Magn. Res. Med.* 2003; vol. 50(no. 5):1031–1042.
2. Mistretta CA, Wieben O, Velikina J, Block W, Perry J, Wu Y, Johnson K, Wu Y. Highly constrained backprojection for time-resolved MRI. *Magn. Res. Med.* 2006; vol. 55(no. 1):30–40.
3. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Res. Med.* 2007; vol. 58(no. 6):1182–1195.
4. Bernstein, MA.; King, KF.; Zhou, XJ. *Handbook of MRI Pulse Sequences*. San Diego: Elsevier; 2004.
5. Schomberg H, Timmer J. The gridding method for image-reconstruction by Fourier transformation. *IEEE Trans. Med. Imag.* 1995; vol. 14(no. 3):596–607.

6. Jackson JI, Meyer CH, Nishimura DG, Macovski A. Selection of a convolution function for Fourier inversion using gridding. *IEEE Trans. Med. Imag.* 1991; vol. 10(no. 3):473–478.
7. Madore B, Glover GH, Pelc NJ. Unaliasing by Fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI. *Magn. Res. Med.* 1999; vol. 42(no. 5): 813–828.
8. Liang ZP, Jiang H, Hess CP, Lauterbur PC. Dynamic imaging by model estimation. *Int. J. Imag. Syst. Technol.* 1997; vol. 8(no. 6):551–557.
9. Zhao, Q.; Aggarwal, N.; Bresler, Y. Dynamic imaging of time-varying objects. *Proc. ISMRM & ESMRMB Joint Annual Meeting; International Society for Magnetic Resonance in Medicine; Glasgow.* 2001. p. 1776
10. Sharif, B.; Bresler, Y. Optimal multi-channel time-sequential acquisition in dynamic MRI with parallel coils. *3rd IEEE International Symposium on Biomedical Imaging: From Nano to Macro; Institute of Electrical and Electronics Engineers; Arlington.* 2006. p. 45–48.
11. Willis NP, Bresler Y. Lattice-theoretic analysis of time-sequential sampling of spatiotemporal signals. *IEEE Trans. Info. Theory.* 1997; vol. 43(no. 1):190–207.
12. Khalsa, KA.; Fessler, JA. *4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro.* Washington: Institute of Electrical and Electronics Engineers, Metro; 2007. Resolution properties in regularized dynamic MRI reconstruction, in; p. 456–459.
13. Jung H, Ye JC, Kim EY. Improved k-t BLAST and k-t SENSE using FOCUSS. *Phys. Med. Biol.* 2007; vol. 52(no. 11):3201–3226. [PubMed: 17505098]
14. Tsao J, Kozerke S, Boesiger P, Pruessmann KP. Optimizing spatiotemporal sampling for k-t BLAST and k-t SENSE: Applications to high-resolution real-time cardiac steady-state free precession. *Magn. Res. Med.* 2005; vol. 53(no. 6):1372–1382.
15. Hansen MS, Baltes C, Tsao J, Kozerke S, Pruessmann KP, Eggers H. k-t BLAST reconstruction from non-Cartesian k-t space sampling. *Magn. Res. Med.* 2006; vol. 55(no. 1):85–91.
16. Schäfer J, Strimmer K. A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Statistical Applications in Genetics and Molecular Biology.* 2005; vol. 4(no.1):32. art.
17. Yang PC, Santos JM, Nguyen PK, Scott GC, Engvall J, McConnell MV, Wright GA, Nishimura DG, Pauly JM, Hu BS. Dynamic real-time architecture in magnetic resonance coronary Angiography - A prospective clinical trial. *Journal of Cardiovascular Magnetic Resonance.* 2004; vol. 6(no. 4):885–894. [PubMed: 15646892]
18. Chui, CK.; Chen, G. *Kalman Filtering with Real-Time Applications.* Berlin: Springer-Verlag; 1987.
19. Beatty PJ, Nishimura DG, Pauly JM. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Trans. Med. Imag.* 2005; vol. 24(no. 6):799–808.
20. Proakis, JG.; Salehi, M. *Digital Communications.* Boston: Prentice-Hall; 2001.
21. Donoho DL, Elad M, Temlyakov VN. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Trans. Info. Theory.* 2006; vol. 52(no. 1):6–18.
22. Sangsuk-Iam S, Bullock TE. Analysis of discrete-time Kalman filtering under incorrect noise covariances. *IEEE Trans. Automatic Control.* 1990; vol. 35(no. 12):1304–1309.
23. Sangsuk-Iam, S. Divergence of the discrete-time Kalman filter under incorrect noise co-variances for linear periodic systems; *Proceedings of the American Control Conference; Baltimore, Maryland;* 1994. p. 1190–1194.
24. Fitzgerald RJ. Divergence of Kalman filter. *IEEE Trans. Automatic Control.* 1971; vol. AC-16(no. 6):736–747.
25. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. *Magn. Res. Med.* 1999; vol. 42(no. 5):952–962.
26. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Res. Med.* 2002; vol. 47(no. 6):1202–1210.
27. Vanvaals JJ, Brummer ME, Dixon WT, Tuithof HH, Engels H, Nelson RC, Gerety BM, Chezmar JL, Denboer JA. Keyhole method for accelerating imaging of contrast agent uptake. *J. Magn. Reson. Imaging.* 1993; vol. 3(no. 6):671–675.

28. Hansen MS, Kozerke S, Pruessmann KP, Boesiger P, Pedersen EM, Tsao J. On the influence of training data quality in k-t BLAST reconstruction. *Magn. Res. Med.* 2004; vol. 52(no. 5):1175–1183.
29. Fuderer M. The information-content of MR images. *IEEE Trans. Med. Imag.* 1988; vol. 7(no. 4): 368–380.
30. Simoncelli EP, Olshausen BA. Natural image statistics and neural representation. *Ann. Rev. Neurosci.* 2001; vol. 24:1193–1216. [PubMed: 11520932]
31. Santos, JM.; Wright, GA.; Pauly, JM. Flexible Real-Time Magnetic Resonance Imaging Framework. 26th Annual Int. Conference IEEE EMBS; Institute of Electrical and Electronics Engineers; San Francisco. 2004. p. 1048
32. Santos JM, Cunningham CH, Lustig M, Hargreaves BA, Hu BS, Nishimura DG, Pauly JM. Single breath-hold whole-heart MRA using variable-density spirals at 3T. *Magn. Res. Med.* 2006; vol. 55(no. 2):371–379.

Figure 1. Normalized absolute autocorrelation function estimate of the difference between two cardiac images that are 24 ms apart - Obtained from the power spectral density by using the Wiener-Khinchine theorem [20] assuming wide-sense stationarity

(a)

(b)

(c)

(d)

Figure 2.

Absolute value of one row of $H_t^H H_t$ arranged in image format: (a) 4× undersampling: Peak value is 0.197. (b) 8× undersampling: Peak value is 0.082. (c),(d) With 3× saturated brightness levels emphasizing the side lobes of (a) and (b), respectively. – Note that in case of Cartesian Nyquist sampling, there would be a single non-zero value, equal to 1.

Figure 3.
Dividing k -space into r chunks to be sampled separately by spiral rings

Figure 4.

Only the diagonals of Q are estimated and shown here in image format. (a) Individual Q estimates of five non-overlapping spiral rings from a GRE experiment. As we go out in k -space (left to right), the intensity decreases and the intensity distribution changes. The numbers at the top left corners indicate the scaling factors used for better visualization. (b) Low resolution Q estimate from a different GRE experiment (c) Low resolution Q estimate from an SSFP experiment.

Figure 5. Three consecutive time points from the 8-interleaf experiment - top: sliding window reconstruction, bottom: Kalman reconstruction

Figure 6.
Three consecutive time points from the 12-interleaf experiment - top: sliding window reconstruction, bottom: Kalman reconstruction

Figure 7.

Three different reconstructions of a single time point - left: 12× sliding window reconstruction, middle: 6× Kalman reconstruction, right: 12× Kalman reconstruction

Figure 8.
Three consecutive time points from the 16-interleaf experiment - top: sliding window reconstruction, bottom: Kalman reconstruction

Figure 9.
Time-shift adjusted pixel(a,b) and line profiles(c,d) for the 8×(a,c) and 12×(b,d) experiments
– 8×: pixel (150,102), vertical line 102, 12×: pixel (149,99), vertical line 99

Table 1

Imaging parameters – TE: echo time, TR: repetition time, fps: frames per second for accelerated reconstructions

	TE/TR (ms)	Readout duration (ms)	frame rate (fps)	% data discarded (%)
6-int.	3.8/20.6	11.8×2=23.6	24.3	1.8
8-int.	3.8/23.9	15	41.8	3.3
12-int.	3.8/20.6	11.8	48.5	1.8
16-int.	3.8/17.7	8.9	56.5	1.9