Published in final edited form as:

Magn Reson Med. 2014 June ; 71(6): 2139–2154. doi:10.1002/mrm.24883.

Design of k-Space Channel Combination Kernels and Integration with Parallel Imaging

Philip J. Beatty^{1,2,3,*}, Shaorong Chang⁴, James H. Holmes⁵, Kang Wang⁵, Anja C. S. Brau⁶, Scott B. Reeder^{7,8,9}, and Jean H. Brittain⁵

¹Global Applied Science Laboratory, GE Healthcare, Toronto, Canada

²Physical Sciences, Sunnybrook Research Institute, Toronto, Canada

³Department of Medical Biophysics, University of Toronto, Toronto, Canada

⁴MR Engineering, GE Healthcare, Waukesha, Wisconsin, USA

⁵Global Applied Science Laboratory, GE Healthcare, Madison, Wisconsin, USA

⁶Global Applied Science Laboratory, GE Healthcare, Menlo Park, California, USA

⁷Department of Radiology, University of Wisconsin, Madison, Wisconsin, USA

⁸Department of Medical Physics, University of Wisconsin, Madison, Wisconsin, USA

⁹Department of Biomedical Engineering and Medicine, University of Wisconsin, Madison, Wisconsin, USA

Abstract

Purpose—In this work, a new method is described for producing local k-space channel combination kernels using a small amount of low-resolution multichannel calibration data. Additionally, this work describes how these channel combination kernels can be combined with local k-space unaliasing kernels produced by the calibration phase of parallel imaging methods such as GRAPPA, PARS and ARC.

Methods—Experiments were conducted to evaluate both the image quality and computational efficiency of the proposed method compared to a channel-by-channel parallel imaging approach with image-space sum-of-squares channel combination.

Results—Results indicate comparable image quality overall, with some very minor differences seen in reduced field-of-view imaging. It was demonstrated that this method enables a speed up in computation time on the order of 3–16X for 32-channel data sets.

Conclusion—The proposed method enables high quality channel combination to occur earlier in the reconstruction pipeline, reducing computational and memory requirements for image reconstruction.

Keywords

channel combination; k-space kernel; parallel imaging; direct virtual coil; channel compression

Imaging using high channel-count coil arrays can benefit a broad range of clinical applications through improved signal-to-noise ratio (SNR), higher parallel imaging acceleration factors and increased anatomical coverage (1–3). However, maintaining clinically acceptable reconstruction times can be challenging when using high channel-count coil arrays. A variety of approaches have been explored to improve the efficiency of reconstruction algorithms for high channel-count imaging (4–14). Many of these techniques, including this work, have focused on combining or reducing the effective number of channels earlier in the reconstruction pipeline. In this way, the computational burden on subsequent steps in the reconstruction process is lessened.

Channel combination methods can be categorized based on the domain in which channel combination is performed: image-space, time-domain or k-space. Image-space channel combination methods are widely used due to their ability to optimize final image SNR and their ease of use (15). However, image-space solutions can be computationally and memory intensive for high channel-count coil arrays, since channel combination must take place relatively late in the reconstruction pipeline.

At the other end of the spectrum, time-domain channel combination (4) has received significant attention. Time-domain channel combination performs channel combination independently at each acquisition time point—at the beginning of the reconstruction pipeline. Time-domain channel combination is very effective in reducing the computation and memory requirements for the image reconstruction pipeline. The weaknesses of time-domain channel combination are (1) SNR can be degraded, sometimes significantly (16) and (2) phase cancellation artifacts are often unavoidable for some coil geometries (5). As a result, time-domain channel combination is not widely used to combine all receiver channels into a single channel data stream. Instead, the approach is typically used to reduce the channel-count significantly but without appreciably degrading image quality (6–12). Methods that take this approach are often referred to as either “array compression” or “coil selection” depending on whether or not they combine channel data streams during the channel reduction. Because these approaches do not output a single channel data set, they must be paired with another channel combination method later in the reconstruction pipeline. In this way, these approaches are complementary to other channel combination strategies, including image-space and k-space approaches (13,14).

In comparison to image-space channel combination and time-domain channel combination, k-space channel combination has received relatively less attention. k-Space channel combination can be performed by convolving the k-space data for each channel with its own k-space convolution kernel and summing across channels to create a single or reduced number of k-space data sets. Unlike time-domain channel combination, it is challenging to implement k-space channel combination in hardware as the convolution coefficients are dependent on the k-space sampling strategy. Unlike image-space channel combination, k-space channel combination is unable to optimize SNR independently at every pixel location.

However, k-space channel combination does have important advantages: it has more degrees of freedom compared to time-domain channel combination that allow it to achieve better SNR and it can be performed earlier in the reconstruction pipeline compared to image-space channel combination, potentially improving reconstruction efficiency.

Early work in k-space channel combination was carried out as part of parallel imaging reconstruction methods. Methods such as SMASH, AUTO-SMASH and VD-AUTO-SMASH combine the parallel imaging unaliasing operation and the channel combination operation into a single step, both in the generation and application of reconstruction coefficients (17–19). Two challenges for these methods are that the SNR of the final combined images can be sub-optimal and they are prone to phase cancellation artifacts (20). Both of these issues can be attributed to sub-optimal channel combination. The GRAPPA method improved on prior k-space parallel imaging methods by separating the unaliasing operation from the channel combination operation and performing channel combination in image-space (20). While this approach has been very successful, it does lead to an increased computational burden, since the unaliasing operation in GRAPPA scales as the square of the number of channels. There is also a large memory requirement as separate data sets for each channel are formed in image-space.

As the number of channels grows, a “channel-by-channel” approach becomes increasingly difficult to perform in clinically acceptable reconstruction times. The problem is exacerbated by the clinical motivations for moving to higher channel-counts: acquiring larger image matrices, and the use of more temporal phases for dynamic contrast enhanced imaging. For example, moving from a single phase eight-channel $256 \times 256 \times 64$ acquisition to a 20 phase 32-channel $320 \times 320 \times 128$ acquisition results in a 1000-fold increase in computation for the unaliasing operation of channel-by-channel parallel imaging with image-space channel combination. This increased computational load can be dealt with using both improvements in computing hardware and algorithmic improvements. Using more powerful reconstruction hardware without algorithmic changes has the advantage that this approach poses no risk to altering the image quality. In contrast, algorithmic changes can modify and potentially degrade image quality. However, hardware solutions can be more costly to roll out to a large installed base compared to software solutions. With such large potential increases in computation, it is likely that improvements in both computing hardware and algorithms will be needed. Because of this, it is important that improved algorithms work well with powerful hardware, which for the foreseeable future means that the algorithms must be amenable to parallel computing architectures.

This work is motivated by a desire to reduce the computational demands for high channel-count image reconstruction while retaining image quality and supporting sampling schemes acquiring irregularly spaced k-space points on a Cartesian grid. Irregularly spaced sampling schemes used in the experiments in this work consist of the inclusion of internal calibration lines in the reconstruction and the use of variable sample spacing (e.g., acquire two, skip one or skip one, acquire one, skip two to create sampling reductions, R , of 1.5 and 2.5 respectively). The purpose of this work is to describe and evaluate a new method for designing local k-space channel combination kernels and integrating them with parallel imaging. We refer to the approach as the direct virtual coil (DVC) method (21). This work

focuses on describing the technical methodology for DVC and relies on separate works to discuss specific clinical applications (22).

THEORY

Channel Combination Framework

The channel combination operation can be expressed in image-space as:

$$\hat{m}(x, y) = \sum_{j=1}^N c_j(x, y) \hat{m}_j(x, y), \quad [1]$$

where $\hat{m}_j(x, y)$ is the image reconstructed for channel j , $c_j(x, y)$ contains the channel combination coefficients for channel j and N denotes the number of channels. Whether channel combination actually occurs in image-space, the time-domain, k-space or another space, it can be understood as an equivalent operation in image-space, with image quality differences dictated by different formulations for the channel combination coefficients. Time-domain channel combination restricts the channel combination coefficients to be invariant over spatial position, whereas using k-space channel combination kernels restricts the channel combination coefficients to be Fourier Transforms of the k-space kernels:

$$c_j(x, y) = \sum_{q=1}^{N_w} w_q e^{i2\pi(k_{x,q}x + k_{y,q}y)}, \quad [2]$$

where $w_1 \dots w_{N_w}$ are the channel combination kernel coefficients for a kernel sampled at k-space locations $(k_{x,1}, k_{y,1}) \dots (k_{x,N_w}, k_{y,N_w})$ and N_w is the number of kernel coefficients.

Figure 1 provides a high-level block diagram of the DVC method. The DVC method involves three main steps: (1) a calibration step that generates channel combination kernels; (2) a step that merges the channel combination kernels with unaliasing coefficients generated via a parallel imaging calibration procedure and (3) virtual coil data synthesis and image formation. The remainder of the theory section describes these operations in greater detail.

Operation 1: DVC Channel Combination Kernel Calibration

A diagram of the DVC calibration algorithm is illustrated in Figure 2. The algorithm first synthesizes a low-resolution complex-valued virtual coil image. It then computes the channel combination kernel coefficients through a least-squares fitting process between the virtual coil image and low-resolution images for each channel.

Low-Resolution Virtual Coil Image Synthesis—While the DVC concept itself is independent from the way the low-resolution virtual coil image is derived, a low-resolution complex-valued virtual coil image must be synthesized in order to calibrate the channel combination kernels. We have developed a method to do this where the virtual coil image is created from low-resolution calibration data (e.g., 32×32 matrix size) acquired from all of the receiver channels. These data are transformed into image-space to create low-resolution images for each channel. The magnitude of the virtual coil image is then computed by performing a sum-of-squares (SoS) combination of the low-resolution images across all

channels. As the SoS combination destroys all phase information, the phase is determined by a separate algorithm.

The virtual coil phase determination algorithm is designed to synthesize consistent phase information across the field-of-view (FOV) from the multi-channel array data. The algorithm deals with two challenges for phase determination: rapidly varying magnetization phase, found in out-of-phase images, and individual coil-elements that produce images dominated by noise over part of the FOV. When there is a phase discontinuity at a location in the individual channel images that have high local SNR, it suggests a phase discontinuity in the underlying magnetization phase. Consequently, the virtual coil phase should have a similar discontinuity to be consistent with this. When there is no phase discontinuity at a location in the individual channel images that have high local SNR, it suggests that the underlying magnetization has smoothly varying phase. At these locations, the virtual coil phase should not introduce a phase discontinuity that is inconsistent, or not seen, in the relevant channel images. We use the term “consistent phase” to refer to this desired relationship between the virtual coil phase and the phase in the individual channel images.

As illustrated in Figure 3, our approach consists of two basic steps. In the first step, sub-regions are defined across the FOV such that the collection of sub-regions covers the entire FOV and the sub-regions overlap (have regions of common voxels). The experiments in this work used rectangular regions that spanned one-fourth the FOV in each direction. Along each dimension, seven regions were evenly spaced, providing overlap of 1/8 the FOV and resulting in 49 regions for two-dimensional (2D) imaging and 343 regions for 3D imaging.

For each region, a single-vector (time-domain equivalent) channel combination is performed. This can be expressed mathematically as

$$\mathbf{r}_v = \mathbf{R}\mathbf{c} \quad [3]$$

where $\mathbf{R}$ contains the sub-region voxel values for all channels. Each column in $\mathbf{R}$ (size $n_v \times n_c$, where n_v is the number of voxels in the region and n_c is the number of channels) is multiplied by a channel combination coefficient in $\mathbf{c}$ (size $n_c \times 1$) and accumulated to create channel-combined voxel values, denoted by $\mathbf{r}_v$ (size $n_v \times 1$). For this work, $\mathbf{c}$ was computed as the eigenvector corresponding to the largest eigenvalue of $\mathbf{R}^H\mathbf{R}$, similar to (5). Other single-vector channel combination coefficients are possible (6) as long as they produce a complex-valued signal with no signal dropout within the subregion.

Using Walsh’s method (5), channel combination coefficients can be generated on a pixel-by-pixel, subregion-by-subregion or even entire FOV basis, enabling the generation of both image-space and time-domain combination coefficients. However, the phase of the channel combination coefficients in $\mathbf{c}$ is arbitrary and this can result in random and fast-changing phase between pixels or subregions. An example is presented in Figure 6 below. While this does not matter for image-space pixel-by-pixel channel combination, it does matter for k-space channel combination, because it is impossible to use a small k-space coil combination kernel to represent fast changing phase in image space. Therefore, a second step is needed to ensure consistent phase between the virtual coil and individual channels across subregions.

In the second step, the subregions are used to grow a region of consistent phase that covers the entire FOV. Although the phase within a subregion is internally consistent and a discontinuity in this phase indicates a corresponding discontinuity in the underlying magnetization phase, the phase between sub-regions is not necessarily consistent. When nonoverlapping subregions are stitched together without editing the phase of the subregions in any way, phase discontinuities can exist along the borders of the subregions, which is not desirable. When overlapping subregions are summed without editing the phase of the subregions, two overlapping subregions can have opposed-phase for the same voxel, which can also result in a poor phase estimate or cancellation effects. The DVC approach allows the phase of a subregion to be edited before combining it with the phase of another subregion. Specifically, an affine transformation is performed on the phase of a subregion that aligns the phase of two subregions in the overlapping voxels before the regions are combined to create a larger subregion. This is described in more detail in Appendix A. Once the phase is determined over the entire FOV, it is combined with the previously computed magnitude image and the complex-valued result is used to determine the channel combination kernels via least-squares fitting.

Least Squares Fitting—The low-resolution complex-valued virtual coil image $m_{LR,v}(x, y)$ (Fig. 2) can be modeled as a low-resolution magnetization image $m_{LR}(x, y)$ modulated by a virtual coil sensitivity function $s_v(x, y)$:

$$m_{LR,v}(x, y) = m_{LR}(x, y) s_v(x, y). \quad [4]$$

For each source channel, a least-squares fitting is performed between $m_{LR,v}(x, y)$ and a low-resolution channel image $m_{LR,j}(x, y)$ to determine the channel combination kernels:

$$\min \| m_{LR,v}^*(x, y) \sum_{q=1}^{N_w} w_q e^{i2\pi(k_{x,q}x + k_{y,q}y)} - \frac{m_{LR,j}^*(x, y)}{\sigma_j^2} \|, \quad [5]$$

where σ_j^2 is the noise variance for channel j and w_q are the channel combination kernel coefficients from Eq. 2.

Insight into the kernel calibration procedure can be aided by approximating $m_{LR,j}(x, y)$ as $m_{LR}(x, y) s_j(x, y)$ and using this to reformulate Eq. 5 as:

$$\min \| |m_{LR,v}(x, y)| \left[c_j(x, y) - \frac{1}{\sigma_j^2} \frac{|s_j(x, y)|}{|s_v(x, y)|} e^{-i2\pi(\angle s_j - \angle s_v)} \right] \|. \quad [6]$$

In Equation 6, $c_j(x, y)$ is the Fourier Transform of the channel combination kernel for channel j , as in Eq. 2. This reformulation makes it easier to see the appropriateness of the fit performed in Eq. 5. For example, when $s_v(x, y) \approx 1$, $c_j(x, y)$ will be fit to $s_j^*(x, y)/\sigma_j^2$. Equation 6 also highlights the importance of the virtual coil phase to the fit. Any spatial discontinuities in $\angle s_j - \angle s_v$ will be difficult to fit using a low-resolution channel combination kernel. In the worst case, using channel combination kernels resulting from

poor least squares fits can manifest in the final reconstructed image as regions of signal drop out in the image. This provides a gauge for the suitability of the virtual coil phase: the phase determination algorithm should generate virtual coil phase that minimizes spatial discontinuities in $\angle s_j - \angle s_v$ across all channels. Moreover, greater importance should be placed on minimizing discontinuities at locations with larger $|m_{LR,j}(x, y)|$, since these locations will have a greater impact on the least squares fit. This reiterates the goal of achieving virtual coil phase that is consistent with the phase in the individual channels and gives the extent to which we can apply significance to the virtual coil phase. For example, we would be overstepping to interpret the virtual coil phase as the absolute phase of the magnetization. Initially, this ambiguity can cause concern over the combination of DVC channel combination with phase sensitive applications, such as chemical shift encoded species separation or phase-based temperature mapping. However, these applications rely on the relative phase between echoes and time-points respectively, not the absolute phase of the magnetization. To combine DVC with these applications, it is sufficient to use the same channel combination coefficients across all echoes or time points.

Another factor that can affect the quality of the least squares fit is the location of the k-space sample locations in Eq. 2. When these locations are separated by k , the channel combination coefficients in image-space, $c_j(x, y)$, repeat every $1/k$. Because the channel combination kernels contain only low-frequency components, there cannot be any sharp edges or discontinuities, even across repetition boundaries. Without oversampling, the repetition distance of $c_j(x, y)$ is equal to the FOV. For channels that are very sensitive at one edge of the FOV and not sensitive at the opposite edge of the FOV, a perfect fit in Eq. 6 would require a discontinuity in $c_j(x, y)$ at the repetition boundary, which is not possible. This leads to less optimal channel combination coefficients in this case. The issue can be solved by moving the kernel samples closer together in k-space (oversampling), which creates a transition band between FOV edges and allows a good fit from Eq. 6 without forcing a discontinuity in $c_j(x, y)$. Because of other design constraints explained below, the implementation presented in this work uses an oversampled kernel in the readout direction but not the phase-encode directions.

Operation 2: Merging Channel Combination Kernels with PI Unaliasing Kernels

Merging channel combination coefficients with parallel imaging coefficients enables both parallel imaging and channel combination operations to be performed together, greatly reducing the required computation for data synthesis. Figure 4 demonstrates this with a small 2-channel example. In the example, the parallel imaging data synthesis requires eight complex multiplication and addition operations (cmad operations), the channel combination synthesis requires six cmad operations, while the combined operation with the merged coefficients only requires four cmad operations. In general, the net reduction in cmad operations depends on the number of channels, k-space sampling pattern and kernel sizes, and can be difficult to predict.

The kernel merge operation can be implemented by grouping all possible paths from each accelerated multichannel sample locations to each virtual coil sample location and accumulating the coefficients as shown in Figure 4. When the channel combination kernel

and PI unaliasing kernel sample spacing is not aligned, the number of coincident paths is greatly reduced, leading to an increase in the number of sample locations in the merged kernel. For this reason, the implementation used in this work aligns the sample spacing for all kernels.

Operation 3: Virtual Coil Data Synthesis and Image Formation

After creating the merged kernels, virtual coil data synthesis is performed in a manner similar to coil-by-coil parallel imaging data synthesis, except that there is only one output coil.

We perform the data synthesis operation in hybrid (x, k_y, k_z) -space (23), which gives flexibility in choosing the kernel sample spacing along k_x . Regardless of the width and sample spacing of the kernel along k_x , when the kernel is transformed to hybrid (x, k_y, k_z) -space, it can be applied as a unit-width kernel at each x -location. For this reason, using an oversampled kernel along k_x is feasible without increasing the data synthesis computation. Since this approach cannot be used along the accelerated phase-encode directions, our implementation uses a kernel that is oversampled by $1.5\times$ in the readout direction but not oversampled in the phase encode directions.

As illustrated in Figure 1, the image formation process for DVC also includes a shading correction step. To understand the rationale for including a shading correction step, it helps to think of the channel combination coefficients, $c_j(x, y)$ as trying to satisfy three objectives: (1) to align the phase of the individual channel images $m_j(x, y)$ so that they add constructively, (2) to provide relative weighting of $m_j(x, y)$ that maximizes SNR, and (3) to provide an absolute weighting that gives a desired sensitivity profile over the entire FOV. The first two objectives must be met as the multichannel data are combined; in the case of DVC, this means that they must be met by the channel combination kernels. However, the third objective is different. If it is not met by the channel combination kernels, we have another opportunity to achieve it by applying a shading correction on the channel-combined data set in image space. We describe the shading correction algorithm that we use in Appendix B; this algorithm is designed to impose a desired shading profile on the virtual coil image.

METHODS

The DVC method was evaluated using simulated data, acquired phantom data and *in vivo* acquired data. Numerically simulated data (2D) were generated in MATLAB (Mathworks, Natick, MA) for a test object that contained water and fat components and a region of signal void within the object. The simulation modeled a 1.5T imaging experiment and included field inhomogeneity in the range of ± 500 Hz. Data was acquired at echo times of 2.2 and 3.7 ms to achieve a relative phase between fat and water of π and $5/3\pi$. Coil sensitivities for a 32-element coil were synthesized with each element having significant sensitivity over a limited portion of the overall imaging volume. Acquired data sets included: (1) a 3D data set of a resolution phantom acquired using an eight-channel body array (GE Healthcare, Waukesha, WI) (axial GRE, TE/TR 1.9/3.9 ms, $320 \times 224 \times 48$ acquisition matrix, 2D accelerated $R = 2.6$, partial k-space in slice-direction as shown in Fig. 7d); (2) 2D dual-echo torso data sets acquired using a prototype 32-channel torso array (24) (coronal GRE, TR 160

ms, 256×256 acquisition matrix, 10 slices, 1D accelerated $R = 1.8, 2.6, 3.4$); (3) a 3D contrast-enhanced brain data set acquired using an 8HR Brain array (GE Healthcare, Waukesha, WI) (sagittal GRE, TE/TR 1.7/4.8 ms, $292 \times 384 \times 160$ acquisition matrix, 2D accelerated $R = 3.6$, partial k-space in readout direction); (4) a 3D precontrast reduced rFOV torso data set acquired using a 32-channel torso array (Neocoil, Pewaukee, WI) (sagittal GRE, TE/TR 1.2/3.4 ms, $196 \times 136 \times 180$ acquisition matrix, 2D accelerated $R = 3.7$, partial k-space in the readout direction).

A version of the DVC algorithm was implemented in C++ and integrated with ARC, a data-driven parallel imaging implementation (GE Healthcare, Waukesha, WI) (25). This implementation used a k-space oversampling factor of 1.5 in the readout direction for both the parallel imaging and channel combination kernels. The channel combination kernels were 5×5 and $5 \times 5 \times 5$ for all 2D and 3D data sets respectively. Unaliasing kernel sizes for the acquired data sets were as follows: (1) $3 \times 7 \times 7$, (2) $3 \times 3, 3 \times 5, 3 \times 7$, (3) $3 \times 7 \times 7$, and (4) $3 \times 3 \times 3$. All kernel sizes are relative to $1/(1.5 \text{ FOV}_x)$ in the readout direction and $1/\text{FOV}_{y,z}$ in the phase encode directions. In the phase encode directions, kernel sizes were chosen to include at least four source points along each direction for eight channel data sets and at least two source points along each direction for 32 channel data sets. The simulated data was reconstructed using the DVC method, with SNR-optimal channel combination as proposed by Roemer et al. (15) and with the adaptive reconstruction method proposed by Walsh et al. (5). For the adaptive reconstruction, channel combination coefficients were generated from the same 32×32 central k-space samples used for DVC calibration and the phase was aligned with the phase of the first channel calibration image. The channel combination coefficients at each voxel location were determined separately using a 15×15 smoothing block. For the acquired data sets, the data was reconstructed using both the DVC method and a channel-by-channel/ SoS approach. The SoS approach was selected as the standard for comparison because phase cancellation artifacts do not occur and, outside of low SNR regions, SNR is close to optimal without the need to estimate the coil sensitivity functions (26). In addition, SoS channel combination has well understood behavior in rFOV imaging. Walsh's channel combination method was also used with the 2D dual-echo torso data set to evaluate the relative phase between echoes computed by DVC.

DVC and SoS magnitude images were compared by taking difference images and by computing the relative root mean square (RRMS) difference (23). A "masked RRMS" was also computed by limiting the RRMS calculation to consider only voxels with a magnitude of at least 10% of the maximum magnitude in the SoS reconstruction. The masked RRMS was computed to minimize the effect on the RRMS of coherent addition in noise-dominated voxels in SoS reconstructions. For the synthesized data set, the DVC kernels were transformed into image-space and compared with the image-space channel combination coefficients generated by the Roemer and Walsh methods.

Image quality was also assessed by a board certified radiologist. For this study, fourteen anonymized rFOV data sets acquired using a 32-channel torso array (Neo-coil, Pewaukee, WI) were reconstructed using both SoS and DVC reconstructions. The radiologist was blinded to the reconstruction method used and compared the two reconstructions according to the following scale: 0: no difference or trivial differences outside of the diagnostically

relevant region, 1: very minor differences, diagnostically irrelevant, 2: moderate differences, 3: substantial differences in image quality.

To measure computational performance, the reconstructions were run on an m3.2xlarge Amazon Web Services virtual machine (Amazon.com, Seattle, WA) using both one and eight computational threads. Execution times for the calibration procedures, data synthesis and reconstruction operations post data synthesis (e.g., Fourier Transformation, Homodyne processing for partial k-space acquisitions, channel combination for SoS reconstructions) were recorded separately. To provide a machine and implementation independent measure of performance, the number of complex multiply and add (cmad) operations performed during the application of the data synthesis kernels in hybrid (x, k_y, k_z) -space was computed for each data set. For the channel-by-channel reconstruction, this included the number of operations required to fill in the missing data due to parallel imaging acceleration. For the DVC reconstruction, this included the number of operations required to fill in the virtual coil data set.

RESULTS

Numerical simulation results are shown in Figures 5 and 6. The image reconstruction results for echo 1 shown in Figure 5, row 1, and associated difference images shown in Figure 6c illustrate that DVC was able to avoid SNR degradation and phase cancellation artifacts even while calibrating to a data set with sharply varying magnetization phase and highly localized coil sensitivities. Figure 5 illustrates that the phase of the virtual coil sensitivity is not uniform, but DVC still retains the relative phase between echoes, making it possible to combine DVC with phase sensitive imaging techniques.

Figure 6 looks in more detail at the channel combination coefficients produced by the DVC method. The cross sectional plots illustrate that DVC does not generate coefficients that match the optimal coefficients with high fidelity in the phase encode direction, however, channel combination is very resilient to less-than-optimal combination coefficients and DVC takes advantage of this to still achieve near-optimal reconstructions, as illustrated in the difference images in Figure 6a, which are brightened 10× compared to row 1 of Figure 5. Comparison of the phase of the channel combination coefficient images generated using the adaptive reconstruction method proposed by Walsh et al. and the DVC method show a significant difference. Since adaptive reconstruction performs channel combination in image space, there is no reason to try and limit rapid phase variations in the channel combination coefficient images. Limiting rapid phase variation is an added restriction for DVC in order to enable small k-space channel combination kernels. The smoothly varying phase for the DVC channel combination coefficient images shown in Figure 6 is a result of generating virtual coil phase that is consistent with the phase in the individual channels.

Reconstruction results for acquired data sets are shown in Figures 7–10. An image quality comparison for the resolution phantom reconstruction is shown in Figure 7. For this data set the RRMS difference was 0.068 and the masked RRMS difference was 0.012. Figure 8 displays one slice from the dual-echo 2D torso acquisitions. The average RRMS difference across all slices, echoes and acquisitions was 0.062 and the masked RRMS difference was

0.024. Figure 8 also shows that the relative phase between echoes for the DVC reconstructions is identical to the relative phase produced using Walsh's adaptive combine method. The brain data set, shown in Figure 9, is also nearly identical with the largest difference at the edge of the phase encode direction (A/P), where there is a slight wrap of the back of the skull and nose in the tight FOV. The RRMS and masked RRMS differences were 0.088 and 0.027, respectively. The largest difference in image quality is apparent in Figure 10, which shows the result from the rFOV torso acquisition. For this data set the RRMS and masked RRMS were much closer at 0.19 and 0.14, respectively. The differences are limited to the regions of signal wrap, but the behavior in these regions is different between the two reconstruction techniques. Signal wrap in a single channel of an rFOV acquisition causes sudden changes in signal intensity that is different between channels and is hard for low-resolution channel combination kernels to mimic. This leads to a significantly different weighting of the channels in the overlapping regions. While both reconstructions exhibit aliasing artifacts, the DVC reconstruction suppresses the aliasing arm slightly better within the body.

Except for the rFOV case, the majority of the image quality differences can be attributed to noise-dominated voxels. In these voxels, SoS channel combination aligns the phase before summing across channels and this leads to a lighter background signal compared to channel combination methods such as DVC that do not sum the noise across channels coherently. This view is corroborated by the significant drop in RRMS difference when the low-signal voxels are masked out.

The blinded radiologist review study found no clinically relevant image quality differences between the SoS and DVC reconstructed images. Specifically, all 14 cases were graded as category 0: no difference or trivial differences outside of the diagnostically relevant region.

Tables 1 and 2 summarize the results of the computational performance experiments. Table 1 reports reconstruction times for the first echo or temporal phase, which includes the calibration time necessary to generate the parallel imaging unaliasing coefficients plus, for DVC, the time necessary to generate the channel combination kernels and merge them with the unaliasing coefficients. Table 1 also reports reconstruction times for additional echoes or temporal phases, where the calibration time is not included. These results would also be applicable to cases where the calibration steps are carried out beforehand. Because applications with multiple echoes/temporal phases tend to be the most computationally demanding, we focus on the computational performance for additional echoes or temporal phases. While the overall computational time depends on many factors including the underlying computer architecture and software implementation as well as the k-space sampling pattern, kernel sizes and number of channels, some trends do emerge from the data. The first is that DVC tends to have a much larger impact on performance at higher channel counts. For the eight channel data sets, DVC resulted in an execution time reduction of 2.2–2.9 $\times$ while for the 32 channel data sets, the execution time reduction ranged from 8.6 to 16.4 $\times$ for additional echoes/ phases.

When moving from a single thread to eight threads to take advantage of parallel computing, reconstruction times were reduced between 1.8 and 3.4 $\times$ for the channel-by-channel/SoS

reconstructions; for the DVC reconstructions, times were reduced by a similar range (1.7–3.3×). These results indicate that DVC can benefit from parallel computing to a similar extent as channel-by-channel reconstructions.

The results reported in Table 2 indicate that the number of cmad operations is roughly indicative of the relative execution time for the data synthesis operations, but tends to overestimate the speedup factor of DVC data synthesis for the implementation used in this work. This indicates that factors other than the amount of computation affect the execution time of the data synthesis operation. Considering the results in Table 2 relative to Table 1, it is clear that the net effect of DVC on performance depends on the relative computational time of the data synthesis step and the subsequent steps in the reconstruction pipeline. For example, in the rFOV torso case, DVC reduces the number of cmad operations in the data synthesis step by 4.9× which leads to an execution time reduction of 3.4× (eight threads). However, DVC also reduces the subsequent coil-by-coil data processing (a homodyne partial k-space reconstruction in this case) which leads to a net execution time reduction from 47.7 to 2.9 s (16.4×) for additional echoes/ phases with eight threads. This computation time reduction enables much faster feedback to the scan operator and helps the reconstruction system keep up with data acquisition, which is especially important during multiecho/multiphase acquisitions.

Finally, an interesting trend can be seen by examining the computational results for the 2D dual echo data sets: DVC exhibits larger computational reduction factors when higher acceleration factors are used. Higher acceleration factors resulted in more missing data for each channel. This requires the channel-by-channel approaches to synthesize more data, adding to the reconstruction computation. For DVC, more missing data means less channel data is used to synthesize the virtual-coil data, reducing the computation. Using eight threads, the SoS reconstruction was able to reconstruct additional images in 214–229 ms, depending on acceleration factor, while DVC was able to reconstruct additional images in 22–23 ms and was much less dependent on acceleration factor. For real time imaging, this is a significant difference; while the SoS reconstruction can achieve four frames/second, DVC enables over 40 frames/second.

DISCUSSION

This work describes a new method for designing local k-space channel combination kernels and integrating them with parallel imaging. It has been demonstrated that the proposed DVC method achieves very similar image quality to image-space based channel combination across a set of challenging reconstruction situations. DVC enables channel combination to be performed earlier in the reconstruction pipeline and significantly reduces memory and computational requirements for image reconstruction. For multiecho, multiphase data sets, available computer memory can restrict clinicians from prescribing the desired resolution or coverage; DVC significantly reduces this problem, which is a huge benefit of DVC for 3D data sets. When k-space (or time-domain) channel combination such as DVC is used, storing all acquired multichannel data of all k_y – k_z encodings in memory is not necessary; instead, it is possible to store only multichannel data of a single k_y – k_z encoding in memory, process with DVC and save the processed data to a single channel 3D memory block. In this way,

the “acquisition memory” of a scan prescription can be reduced. Additionally, the computational benefits of DVC increase when imaging with higher channel-count coil arrays and higher acceleration factors.

While the image quality differences between DVC and SoS reconstructions were minimal, two minor differences were identified: (1) DVC tended to have a darker background signal, especially with higher channel-count coils and (2) differences were observed in the regions of wrapped signal in rFOV imaging. SoS reconstructions are known to increase the signal in background pixels by constructively summing noise in these pixels (27); DVC does not share this disadvantage. It is challenging to determine which channel combination approach has better image quality in the rFOV case, since signal wrap contaminates both reconstructions. However, in this limited number of data sets, signal wrap was more suppressed in the DVC reconstructions. It is unclear what impact this will have, if any, on the interpretation of clinical imaging studies.

In the proposed method, separate calibrations are performed to generate k-space unaliasing and channel combination coefficients. These coefficients are then merged and transformed into (x, k_y, k_z) -space before being applied. Although it is possible to calibrate in a single step (28,29), separate unaliasing and channel combination calibrations enable better optimization of channel combination when using local k -space kernels. However, synthesis using a single set of merged unaliasing/channel combination coefficients is significantly more computationally efficient than performing separate unaliasing and channel combination operations. Intermediate manipulation of calibration results [e.g., merging kernels and transforming kernels from k -space to hybrid (x, k_y, k_z) -space] enables more flexibility in choosing both calibration and data synthesis strategies and the proposed approach takes advantage of this to generate high quality reconstruction coefficients and apply them efficiently.

Wang et al. has proposed a similar separate calibration with joint synthesis approach where the parallel imaging operation is moved into image-space and merged with image-space channel combination coefficients (30). This is an effective solution when uniform k -space undersampling is used (31). However, when nonuniform Cartesian k -space sampling patterns are employed, such as in Figure 7d, it is challenging to move the parallel imaging operation into image-space. To support nonuniform Cartesian k -space sampling patterns, the channel combination operation must be moved to hybrid (x, k_y, k_z) -space or k -space, motivating our desire to produce high quality channel combination kernel coefficients.

Although outside the scope of this work, it should be noted that DVC channel combination kernels have a small region of support in k -space, making them useful for applications where channel combination is desired for only a portion of k -space. Prospective motion correction with multichannel coils is one potential application (32). DVC can reconstruct complex-valued images, potentially making it useful for phase-sensitive imaging such as phase velocity imaging and chemical shift encoded imaging where the preservation of relative phase between acquisitions is important.

Although we have achieved promising results with the proposed method, a potential limitation of the approach is sensitivity to the phase determination algorithm during the DVC calibration. Failure of the phase determination algorithm could lead to poor quality channel combination kernels that may cause phase cancellation artifacts. Much of the effort in developing the DVC method has focused on the phase determination algorithm and our results indicate that the algorithm presented here is quite resilient. This is evidenced by the challenging cases demonstrated in this article including high channel-count out-of-phase imaging, low SNR imaging and reduced FOV imaging.

The phase determination algorithm shares significant similarities to the channel combination method proposed by Walsh et al. (5), where reconstruction vectors are computed and applied in a block-by-block manner. However, the phase determination algorithm proposed here uses overlapping regions (blocks) and the affine phase correction technique described in Appendix A to solve the phase consistency problem across regions.

The challenging phase determination aspect of DVC calibration could be avoided by using calibration data from a calibration scan that is designed to produce constant phase within the imaged magnetization, leaving all of the phase within $m_j(x, y)$ attributable to $s_j(x, y)$.

Another strategy for avoiding phase determination is to employ a separate calibration scan and repeat the scan, acquiring once with the multichannel receive array and once with a receive coil that has fairly homogeneous sensitivity over the imaging volume. The phase, or the phase and the magnitude, for the virtual coil could then be extracted from the homogeneous coil, since the goal of the virtual coil is to approximate data acquired from such a homogeneous coil. However, both of these approaches would require additional calibration scans.

There are also a number of parameters that need to be chosen as part of the DVC reconstruction: the readout direction oversampling factor, the size of the channel combination kernels and the subregions used in the phase determination algorithm.

Although there are many reasonable choices for these parameters that will achieve good image quality, a detailed study of the effect of these parameters has not been performed and it is not known whether these parameters should be varied as a function of other imaging parameters.

This work describes the design of local k-space channel combination kernels and their integration with Cartesian parallel imaging. Additionally, the proposed method can be combined with an earlier step of time-domain channel compression/selection. Further study is needed to determine the most efficient strategy considering multiple channel reduction techniques. The k-space channel combination kernels would also be useful for non-Cartesian imaging, where they could be combined with gridding kernels instead of unaliasing coefficients; efficient ways of merging these kernels are still being explored (33).

As mentioned in the introduction, an alternative approach to reducing reconstruction times is through advanced hardware, such as GPUs (34–38), rather than through more efficient and accelerated reconstruction algorithms, such as DVC. Further, GPU implementations of DVC are also possible, although the amount of additional speedup is unknown until such

implementations have been performed. Table 3 summarizes reconstruction times and speed up factors for several algorithms published over the last few years that have been implemented on GPU hardware. Because different algorithms and imaging parameters are used and because the published acceleration factors range from $4.2\times$ to $640\times$, comparison with DVC alone is difficult. However, these examples are helpful in putting the DVC method's results into context. If hardware solutions alone could achieve the necessary accelerations, they would eliminate any risk associated with potential artifacts that could be introduced by new algorithms. It will be important for future work to compare the relative benefits and risks of these two approaches as well as to explore the implementation of efficient algorithms, such as DVC, with high performance hardware, such as GPUs.

In summary, this work describes a new method for designing local k-space channel combination kernels and integrating them with Cartesian parallel imaging in a way that supports the reconstruction of data acquired with irregular undersampling of k-space points on a Cartesian grid. This method achieves good image quality while significantly reducing the required computation and memory usage.

Acknowledgments

Grant sponsor: National Institutes of Health; Grant numbers: R01 DK083380, R01 DK088925, RC1 EB010384 and R01 DK096169.

References

1. Hardy CJ, Giaquinto RO, Piel JE, et al. 128-channel body MRI with a flexible high-density receiver-coil array. *J Magn Reson Imaging*. 2008; 28:1219–1225. [PubMed: 18972330]
2. Schmitt M, Potthast A, Sosnovik DE, Polimeni JR, Wiggins GC, Triantafyllou C, Wald LL. A 128-channel receive-only cardiac coil for highly accelerated cardiac MRI at 3 Tesla. *Magn Reson Med*. 2008; 59:1431–1439.
3. Wiggins GC, Polimeni JR, Potthast A, Schmitt M, Alagappan V, Wald LL. 96-Channel receive-only head coil for 3 Tesla: design optimization and evaluation. *Magn Reson Med*. 2009; 62:754–762. [PubMed: 19623621]
4. Debbins JP, Felmlee JP, Riederer SJ. Phase alignment of multiple surface coil data for reduced bandwidth and reconstruction requirements in volumetric MRI applications. *Magn Reson Med*. 1997; 38:1003–1011. [PubMed: 9402202]
5. Walsh DO, Gmitro AF, Marcellin MW. Adaptive reconstruction of phased array MR imagery. *Magn Reson Med*. 2000; 43:682–690. [PubMed: 10800033]
6. Buehrer M, Pruessmann KP, Boesiger P, Kozerke S. Array compression for MRI with large coil arrays. *Magn Reson Med*. 2007; 57:1131–1139. [PubMed: 17534913]
7. Marinelli, L.; Hardy, CJ. Coil Selection optimization using mean-field annealing and its application to 128-channel imaging. *Proceedings of the 15th Annual Meeting of ISMRM*; Berlin, Germany. 2007; p. Abstract 750
8. Doneva M, Börner P. Automatic coil selection for channel reduction in SENSE-based parallel imaging. *Magn Reson Mater Phys*. 2008; 21:187–196.
9. Huang F, Vijayakumar S, Li Y, Hertel S, Duensing GR. A software channel compression technique for faster reconstruction with many channels. *Magn Reson Imag*. 2008; 26:133–141.
10. King SB, Varosi SM, Duensing GR. Optimum SNR data compression in hardware using an Eigencoil array. *Magn Reson Med*. 2010; 63:1346–1356. [PubMed: 20432305]
11. Feng S, Zhu Y, Ji J. Efficient large-array k-domain parallel MRI using channel-by-channel array reduction. *Magn Reson Imag*. 2011; 29:209–215.

12. Zhang T, Pauly JM, Vasanawala SS, Lustig M. Coil compression for accelerated imaging with Cartesian sampling. *Magn Reson Med*. 2013; 69:571–582. [PubMed: 22488589]
13. Stemmer, A.; Jellus, V.; Kannengiesser, S.; Kiefer, B. Channel compression for BLADE. Proceedings of the 16th Annual Meeting of ISMRM; Toronto, Canada. 2008; p. Abstract 1274
14. Huang, F.; Lin, W.; Duensing, GR.; Reykowski, A. Flexible Virtual Coils (FVC) for faster channel-by-channel partially parallel imaging. Proceedings of the 19th Annual Meeting of ISMRM; Montreal, Canada. 2011; p. Abstract 2877
15. Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. *Magn Reson Med*. 1990; 16:192–225. [PubMed: 2266841]
16. Sodickson DK, Griswold MA, Jakob PM, Edelman RR, Manning WJ. Signal-to-noise ratio and signal-to-noise efficiency in SMASH imaging. *Magn Reson Med*. 1999; 41:1009–1022. [PubMed: 10332885]
17. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med*. 1997; 38:591–603. [PubMed: 9324327]
18. Jakob PM, Griswold MA, Edelman RR, Sodickson DK. AUTO-SMASH: a self-calibrating technique for SMASH imaging. *MAGMA*. 1998; 7:42–54. [PubMed: 9877459]
19. Heidemann RM, Griswold MA, Haase A, Jakob PM. VD-AUTO-SMASH imaging. *Magn Reson Med*. 2001; 45:1066–1074. [PubMed: 11378885]
20. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisition (GRAPPA). *Magn Reson Med*. 2002; 47:1202–1210. [PubMed: 12111967]
21. Beatty, PJ.; Sun, W.; Brau, ACS. Direct Virtual Coil (DVC) reconstruction for data-driven parallel imaging. Proceedings of the 16th Annual Meeting of ISMRM; Toronto, Canada. 2008; p. Abstract 8
22. Wang K, Beatty PJ, Nagle SK, Reeder SB, Holmes JH, Rahimi MS, Bell LC, Korosec FR, Brittain JH. Application of direct virtual coil to dynamic contrast-enhanced MRI and MR angiography with data-driven parallel imaging. *Magn Reson Med*. 2014; 71:783–789.
23. Brau ACS, Beatty PJ, Skare S, Bammer R. Comparison of reconstruction accuracy and efficiency among autocalibrating data-driven parallel imaging methods. *Magn Reson Med*. 2008; 59:382–395. [PubMed: 18228603]
24. Zhu Y, Hardy CJ, Sodickson DK, Giaquinto RO, Dumoulin CL, Kenwood G, Niendorf T, Lejay H, McKenzie CA, Ohliger MA, Rofsky NM. Highly parallel volumetric imaging with a 32-element RF coil array. *Magn Reson Med*. 2004; 52:869–877. [PubMed: 15389961]
25. Beatty, PJ.; Brau, AC.; Chang, S.; Joshi, SM.; Michelich, CR.; Bayram, E.; Nelson, TE.; Herfkens, RJ.; Brittain, JH. A method for autocalibrating 2-D accelerated volumetric parallel imaging with clinically practical reconstruction times. Proceedings of the 15th Annual Meeting of ISMRM; Berlin, Germany. 2007; p. Abstract 1749
26. Larsson EG, Erdogmus D, Yan R, Principe JC, Fitzsimmons JR. SNR-optimality of sum-of-squares reconstruction for phased-array magnetic resonance imaging. *J Magn Reson*. 2003; 163:121–123. [PubMed: 12852915]
27. Constantinides CD, Atalar E, McVeigh ER. Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magn Reson Med*. 1997; 38:852–857. [PubMed: 9358462]
28. Beatty, PJ.; Brau, ACS. Uniform virtual coil reconstruction for autocalibrating parallel imaging. Proceedings of the 15th Annual Meeting of ISMRM; Berlin, Germany. 2007; p. Abstract 1751
29. Chen, W.; Hu, P.; Meyer, CH. Rapid partially parallel reconstruction using a single synthetic target coil. Proceedings of the 16th Annual Meeting of ISMRM; Toronto, Canada. 2008; p. Abstract 1296
30. Wang, J.; Zhang, B.; Zhong, K.; Zhuo, Y. Image domain based fast GRAPPA reconstruction and relative SNR degradation factor. Proceedings of the 13th Annual Meeting of ISMRM; Miami, Florida, USA. 2005; p. Abstract 2428
31. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*. 1999; 42:952–962. [PubMed: 10542355]

32. Banerjee S, Beatty PJ, Zhang JZ, Shankaranarayanan A. Parallel and partial Fourier imaging with prospective motion correction. *Magn Reson Med*. 2013; 69:421–433. [PubMed: 22488750]
33. Beatty, PJ.; Takahashi, A.; Johnson, KM.; Brittain, JH. k-Space channel combination for non-cartesian acquisitions using direct virtual coil (DVC) calibration. Proceedings of the 19th Annual Meeting of ISMRM; Montreal, Canada. 2011; p. Abstract 2858
34. Hansen MS, Atkinson D, Sorensen TS. Cartesian SENSE and k - t SENSE reconstruction using commodity graphics hardware. *Magn Reson Med*. 2008; 59:463–468. [PubMed: 18306398]
35. Wu, XL.; Gai, J.; Lam, F.; Fu, M.; Haldar, JP.; Zhuo, Y.; Liang, ZP.; Hwu, WM.; Sutton, BP. IMPATIENT MRI: Illinois massively parallel acceleration toolkit for image reconstruction with enhanced throughput in MRI. Proceedings of the 19th Annual Meeting of ISMRM; Montreal, Canada. 2011; p. Abstract 4396
36. Smith DS, Gore JC, Yankeelov TE, Welch EB. Real-time compressive sensing MRI reconstruction using GPU computing and split bregman methods. *Int J Biomed Imaging*. 2012; 2012:864827. [PubMed: 22481908]
37. Vasanawala, SS.; Murphy, MJ.; Alley, MT.; Lai, P.; Keutzer, K.; Pauly, JM.; Lustig, M. Practical parallel imaging compressed sensing MRI: summary of two years of experience in accelerating body MRI of pediatric patients. 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro; Chicago, Illinois, USA. 2011; p. 1039
38. Nam S, Akcakaya M, Basha T, Stehning C, Manning WJ, Tarokh V, Nezafat R. Compressed sensing reconstruction for whole-heart imaging with 3D radial trajectories: a graphics Processing Unit Implementation. *Magn Reson Med*. 2013; 69:91–102. [PubMed: 22392604]
39. Kellman P, McVeigh ER. Image reconstruction in SNR units: a general method for SNR measurement. *Magn Reson Med*. 2005; 54:1439–1447. [PubMed: 16261576]
40. McKenzie CA, Yeh EN, Ohliger MA, Price MD, Sodickson DK. Self-calibrating parallel imaging with automatic coil sensitivity extraction. *Magn Reson Med*. 2002; 47:529–538. [PubMed: 11870840]

APPENDIX A: ALIGNING THE PHASE OF TWO SUB-REGIONS USING AN AFFINE PHASE TRANSFORMATION

Consider two overlapping regions in image-space: $r_1(x, y)$ and $r_2(x, y)$. This analysis describes how to find an affine transformation of the phase of $r_2(x, y)$ that will match it closely to the phase of $r_1(x, y)$. A new quantity $h(x, y)$ can be computed with a phase that is the difference in phase between $r_1(x, y)$ and $r_2(x, y)$:

$$h(x, y) = r_1^*(x, y) \frac{r_2(x, y)}{|r_2(x, y)|} \quad [A1]$$

For this method, it is assumed that the phase of $h(x, y)$ is fairly close to affine. Given that the overlapping region is sufficiently small, this is a good assumption. $h(x, y)$ is then modeled as having affine phase:

$$h(x, y) \approx |h(x, y)| e^{i2\pi(\Delta + \Delta k_x x + \Delta k_y y)}. \quad [A2]$$

Taking the Fourier Transform of $h(x, y)$ gives

$$H(k_x, k_y) \approx e^{i2\pi\Delta} FT\{|h(x, y)|\} (k_x, k_y) * \delta(k_x - \Delta k_x, k_y - \Delta k_y). \quad [A3]$$

Note that $FT\{|h(x, y)|\}(0, 0)$ is positive and real valued (i.e. zero phase). $FT\{|h(x, y)|\}(0, 0)$ also contains the maximum value of $FT\{|h(x, y)|\}(k_x, k_y)$, since all values in the Fourier transform integral are summed coherently when $k_x = k_x = 0$. Thus, k_x and k_y can be determined from the maximum value of $H(k_x, k_y)$: k_x is determined from the phase of the maximum value and (k_x, k_y) is determined by the maximum value location in k-space. The affine transformation can then be applied to $r_2(x, y)$:

$$r_2'(x, y) = r_2(x, y) e^{-i2\pi(\Delta + \Delta k_x x + \Delta k_y y)}. \quad [\text{A4}]$$

For the implementation used in this work, inverse gridding is used to efficiently find the location of the maximum value of $H(k_x, k_y)$. After weighting $h(x, y)$ by the gridding kernel apodization correction function, a DFT is computed and the location of the maximum of the result is found. Convolution interpolation is then used in the area around the maximum to further refine the maximum value location.

APPENDIX B: CORRECTING THE SHADING ENVELOPE

In this analysis, channel noise is assumed to be independent (i.e., the noise covariance matrix has non-zero values only along the diagonal). For cases where the noise correlation between coils must be taken into account, one can apply a prewhitening transform on the acquired channels to create a set of noise independent synthetic channels before computing the channel combination coefficients (39).

For a set of image-space channel combination coefficients, $c_j(x, y)$, the shading envelope $p(x, y)$ can be determined by computing:

$$p(x, y) = \left| \sum_{j=1}^N c_j(x, y) s_j(x, y) \right|. \quad [\text{B1}]$$

A set of channel combination coefficients with the ideal *relative* phase and magnitude but unrestricted shading envelope can be modeled as

$$c_j(x, y) = \lambda(x, y) \frac{s_j^*(x, y)}{\sigma_j^2}, \quad [\text{B2}]$$

where $\lambda(x, y)$ accounts for the freedom in the shading envelope. Substituting Eq. B2 into Eq. B1 gives the following expression for the shading envelope:

$$p(x, y) = \lambda(x, y) \sum_{j=1}^N \frac{|s_j(x, y)|^2}{\sigma_j^2} \quad [\text{B3}]$$

For SoS channel combination, the shading envelope can be expressed as

$$p_{\text{SOS}}(x, y) = \sqrt{\sum_{j=1}^N \frac{|s_j(x, y)|^2}{\sigma_j^2}}; \quad [\text{B4}]$$

this is an extension of Eq. B3 in Ref. 40 when the noise variance is not identical across channels.

In the context of this work, correcting the shading envelope refers to finding a correction function $\gamma(x, y)$ that can be multiplied with a channel-combined image to give a shading envelope similar to the shading envelope of a SoS combination. Correcting the shading envelope in this way allows the channel-combined image to replace a SoS-combined image in a reconstruction pipeline. A similar analysis can be performed to match other shading envelopes.

Mathematically the correction function can be expressed as:

$$\gamma(x, y) = \frac{p_{\text{SOS}}(x, y)}{p(x, y)}. \quad [\text{B5}]$$

Substituting Eqs. B3 and B4 into Eq. B5, we can express the correction function as

$$\gamma(x, y) = 1 / \sqrt{\sum_{j=1}^N \sigma_j^2 |c_j(x, y)|^2}, \quad [\text{B6}]$$

which can be calculated from the channel combination coefficients and used to correct the shading envelope in image-space.

DVC Method Overview

1) DVC channel combination kernel calibration

2) Merging channel combination kernels with PI 'unaliasing kernels'

3) Virtual coil data synthesis & image formation

FIG. 1.

DVC channel combination is comprised of three operations: (1) a calibration operation that generates channel combination kernels, (2) a merge operation, used when DVC is combined with parallel imaging, that merges channel combination kernels with parallel imaging kernels and (3) a data synthesis and image formation operation that generates a virtual coil image. When parallel imaging is not used, the merge operation can be skipped and the channel combination kernels can be used directly to synthesize virtual coil data in operation three.

FIG. 2.

Illustration of the DVC calibration operation. Using low-resolution multichannel calibration data, a low-resolution complex-valued virtual coil image $m_{LR,v}(x,y)$ is formed and used in a least-squares fitting process to generate channel combination kernels. The magnitude and phase of $m_{LR,v}(x,y)$ are determined separately.

FIG. 3.

The phase determination algorithm used in the DVC calibration operation consists of two steps. Step 1: a single-vector channel combination (Eq. 3) is used to combine channel data for subregions of the image. These subregions are designed to overlap. Step 2: the overlapping subregions of channel-combined data are stitched together. A region of consistent phase is grown by joining subregions after applying affine transformations to the phase of the subregions (Appendix A).

Merge Operation

(a) Parallel Imaging Coil-by-Coil Data Synthesis

(b) Channel Combination Data Synthesis

(c) Data Synthesis with Merged Coefficients

FIG. 4.

An illustration of how the sequential application of parallel imaging data synthesis in (a) and channel combination data synthesis in (b) can be replaced, with equivalent results, by a single operation shown in (c). The merge operation is the operation that computes the merged kernel coefficients $w_1 \dots w_4$ from the parallel imaging coefficients $u_1 \dots u_8$ and the channel combination coefficients $v_1 \dots v_6$. In each case, data synthesis is performed as a weighted sum of data from channels 1 and 2 using the kernel coefficients shown for weights.

FIG. 5.

Simulation results. Row 1 contains magnitude image results for echo 1 (2.2 ms) for each channel combination method (Roemer, Walsh, DVC). Rows 2 and 3 show that the phase images for the echoes differ significantly depending on the reconstruction approach. Only the Roemer reconstruction, which uses the ground truth absolute coil sensitivities reconstructs the absolute phase of the underlying magnetization (which is a function of the fat and water chemical species and the field inhomogeneity). Row 4 plots the relative phase between echoes to demonstrate that even though the Walsh and DVC methods do not determine the absolute phase on an echo-by-echo basis, they can be used to determine the relative phase between echoes, enabling their use with chemical shift encoding.

FIG. 6.

Channel combination kernel investigation using simulated data. **(a)** Channel combination coefficients for channel 27 ($ccc(27)$). Channel 27 has a large sensitivity variation across the phase encode direction and the DVC channel combination coefficients cannot match the optimal channel combination coefficients, as illustrated in the cross section plot in **(e)**. **(b)** Channel combination coefficients for channel 15 ($ccc(15)$). Channel 15 has a large sensitivity variation across the readout direction. By using an oversampled kernel in the readout direction, DVC is better able to handle this situation, as illustrated by the cross section plot in **(f)**. The ultimate result of this is shown in **(c)**, which shows magnitude difference images relative to the Roemer reconstruction for echo 1, windowed $10\times$ brighter compared to Row 1, Figure 5. The deviation of the DVC channel combination coefficients from the optimal coefficients in the phase encode direction results in a slight enhancement in the difference images at the edges of the FOV along the phase encode direction. **(d)** Plots of the phase of the channel combination coefficients for channel 15, illustrating the phase discontinuities in the channel combination coefficients determined using the method

proposed by Walsh et al. (5), making this approach unsuitable for generating small k-space channel combination kernels.

FIG. 7.

Imaging results from the resolution phantom with an eight-channel body coil. Both SoS (column 1) and DVC (column 2) images were retrospectively reconstructed using the same data set. Difference images (column 3) were computed as $|\text{SoS}| - |\text{DVC}|$ and windowed 10 $\times$ brighter. The 3D data set was acquired in the axial plane (**a**) and reformatted in the coronal (**b**) and sagittal (**c**) planes from the DICOM images. The images are nearly identical, except for the lower background intensity seen in the DVC images. (**d**) Irregularly spaced k-space sampling pattern used for this experiment.

FIG. 8.

Imaging results from a dual-echo 32-channel torso data set. **(a)** Magnitude images for SoS (column 1) and DVC (column 2). Difference images (column 3) were computed as $|\text{SoS}| - |\text{DVC}|$ and windowed $10\times$ brighter. **(b)** Comparison of phase images computing using both the method proposed by Walsh et al. (5) and DVC. While the phase for individual echoes differs greatly, the relative phase between echoes is almost identical between the channel combination methods. This is confirmed by checking the difference image between the relative phases computing using both the Walsh and DVC methods.

FIG. 9. (a) Brain results from eight-channel head coil. Sagittal difference image (b) illustrates slight difference in signal intensity in the A/P direction due to the tight FOV (back of head and nose). These differences are slight and hard to see without the difference image ($|\text{SoS}| - |\text{DVC}|$, windowed 10 $\times$ brighter).

FIG. 10.

Imaging results for the rFOV data set acquired with a 32-ch body coil. Significant differences can be seen in the axial and coronal reformats (**b** and **c**) where an arm aliases into the torso. Although the aliasing artifacts are significant for both SoS and DVC reconstructions, the DVC reconstruction appears to suppress the aliasing arm slightly better within the body.

Table 1
Results of Computation Run Time Experiments on m3.2xlarge Amazon Web Services Virtual Machine

	Resolution phantom	Dual-echo torso				Brain	rFOV torso
		R = 1.8	R = 2.6	R = 3.4			
Acquisition Matrix	320 × 224 × 48	256 × 256	256 × 256	256 × 256	292 × 384 × 160	196 × 136 × 180	
Reconstruction Matrix	512 × 512 × 48	512 × 512	512 × 512	512 × 512	512 × 512 × 160	512 × 512 × 360	
Channels	8	32	32	32	8	32	
Serial computation (one thread)							
Reconstruction time for first echo or temporal phase							
SoS	13.6 s	540 ms	585 ms	627 ms	66.2	123.9 s	
DVC	9.2 s	201 ms	190 ms	192 ms	31.8	11.9 s	
Speed up factor	1.5×	2.7×	3.1×	3.3×	2.1×	10.4×	
Reconstruction time for additional echoes or temporal phases							
SoS	9.1 s	426 ms	474 ms	509 ms	53.4 s	121.1 s	
DVC	4.0 s	50 ms	43 ms	39 ms	18.3 s	7.7 s	
Speed up factor	2.3×	8.6×	11.0×	12.9×	2.9×	15.7×	
Parallel computation (eight threads)							
Reconstruction time for first echo or temporal phase							
SoS	4.7 s	290 ms	292 ms	301 ms	22.0 s	48.6 s	
DVC	3.6 s	118 ms	112 ms	111 ms	12.3	4.5 s	
Speed up factor	1.3×	2.5×	2.6×	2.7×	1.8×	10.8×	
Reconstruction time for additional echoes or temporal phases							
SoS	2.7 s	214 ms	220 ms	229 ms	16.4 s	47.7 s	
DVC	1.2 s	23 ms	22 ms	22 ms	6.4 s	2.9 s	
Speed up factor	2.2×	9.1×	10.1×	10.6×	2.6×	16.4×	

Since calibration operations can be carried out once and applied to multiple echoes or temporal phases, the reconstruction times for additional echoes and/or temporal phases can be reduced. Reconstruction times for the 2D data sets are for a single echo, single slice. As expected, the computational benefit of DVC increases with the number of channels. The computational speed ups achieved with DVC, compared to SoS, hold fairly constant moving from 1 to 8 threads, demonstrating the ability of DVC to take advantage of parallel computing

Table 2
Computation Run Time Results for the Data Synthesis Operation and Comparison to the Number of Complex Multiply and-Add (cmad) Operations for the Data Synthesis Step

	Resolution phantom	Dual-echo torso				Brain	rFOV torso
		R = 1.8	R = 2.6	R = 3.4			
Acquisition Matrix	320 × 224 × 48	256 × 256	256 × 256	256 × 256	292 × 384 × 160	196 × 136 × 180	
Channels	8	32	32	32	8	32	
Data synthesis time (eight threads)							
SoS	1.3 s	60 ms	71 ms	81 ms	8.8 s	6.5 s	
DVC	1.1 s	20 ms	18 ms	18 ms	5.7 s	1.9 s	
Speed up factor	1.2×	3.0×	3.9×	4.5×	1.5×	3.4×	
Number of complex multiply and-add (cmad) operations during data synthesis							
SoS	2.03 mil	60.3 K	81.3 K	92.8 K	11.0 mil	7.45 mil	
DVC	1.07 mil	7.57 K	6.51 K	5.96 K	4.30 mil	1.51 mil	
Reduction factor	1.89×	7.97×	12.48×	15.56×	2.56×	4.94×	

Computation times for the 2D data sets are for a single echo, single slice. Operations after data synthesis, such as transformation to image space, account for a larger portion of the reconstruction times for SoS reconstructions, especially for the rFOV torso data set. This explains how the overall computation time reduction for DVC can be significantly greater than the reduction in the data synthesis operation. The reduction in cmad operations overestimates the speed up factor for the implementation used in this work, indicating that the implementation run time is not completely tied to the amount of computation.

Table 3
Summary of Computation Times and Speed Up Factors for a Selection of Published Algorithms Implemented on GPU Hardware

Algorithm	Reconstruction matrix	Channels	CPU time	GPU time	Speed up factor
2D SENSE (34)	256 × 256	16	622 ms	26 ms	26.1×
	512 × 512	32	5.65 s	155 ms	36.4×
Non-cartesian SENSE (35)	256 × 256	8	9.4 h	53.1 s	640×
CS MRI (36)	256 × 256	1	0.28 s	66 ms	4×
	512 × 512	1	1.1 s	120 ms	9.2×
LI-Spirit POCS (37)	4096 × 4096	1	68 s	8.3 s	8.19×
	172 × 230 × 188	32	–	169.9 s	–
	192 × 320 × 66	8	–	24.4 s	–
3D FFT (38)	392 × 392 × 392	1	5.0s	0.26 s	18.7×

Speed up factors vary from 4.2 to 640×. Comparison of computation times between problems of similar matrix size and channel count give context to the results presented in Tables 2 and 3.