


Genealogy: a Method to Prevent Spikes and Ill-Conditioning in Fast Space-Domain Parallel MRI



M. Scheffe¹, G. P. Zientara¹

¹Radiology, Brigham & Women's Hospital, Boston, MA, United States

Abstract: Spiky noise and ill-conditioning are serious problems in space-domain MR reconstructions such as SENSE, which are based on fast (undersampled) multiple-coil data acquisitions. This paper demonstrates that these problems can arise when the aliasing process mixes higher-quality patient data (foreground = signal region) with low-quality background data (noise-only region). This problem can be solved simply by eliminating the background variables from the reconstruction equations at each pixel, resulting in a much more stable and higher-contrast output. The numerical condition number for this revised reconstruction technique can show an improvement by several orders of magnitude over naïve least-squares reconstruction.




Introduction: Fast parallel MRI algorithms such as SENSE [1] are the subject of intense current research activity, because of their potential for clinical applications such as near-real-time cardiac and interventional radiology. The processing for most fast space-domain reconstruction algorithms begins with undersampling the data in k-space, resulting in aliasing in the space domain. By collecting several independent aliased images at multiple receiver coils, it is possible for the least-squares fit to reconstruct an alias-free final image estimate. This reconstruction procedure includes an explicit description of the inhomogeneities in each coil image, by means of sensitivity functions, which can be assembled off-line. For illustrations in this paper, we will display examples from a cylindrical phantom, which includes both a dark central ring and a set of small, low-contrast discs at the center. The data was acquired on a GE 1.5T Signa MR scanner with a fast spin-echo sequence, using an experimental 4-coil array that wrapped tightly around the cylindrical phantom. The first reconstruction, shown on the top left, is based on fully-sampled data with no aliasing, i.e. a slow acquisition. Next, an example of fast aliased data from one of the 4 coils is shown at the right, with a decimation (acceleration) rate of 2 in each coordinate direction, i.e. horizontal and vertical. A second reconstruction was obtained from the 4 aliased coil images. Both reconstructions were based on the most straightforward, standard application of the least-squares fit in the space domain, including loading by the coil sensitivity functions, which were acquired separately.





Problem formulation: The image at the left shows some challenges that arise when the reconstruction process must deal with aliased inputs, rather than with slowly-acquired unaliased data. Note that a number of very large spikes have appeared in this fast reconstruction, despite removal of the highest ones by a simple thresholding procedure; these spikes did not occur in the reconstruction from fully-sampled, unaliased data. Also note that the low-contrast central discs are not visible in this aliased-data reconstruction. This spiky nature of the least-squares reconstruction from aliased data was recognized in [1], where it was addressed by a polynomial-fit smoothing algorithm. This approach could be thought of as a low-pass filter, which generally sacrifices spatial resolution. Hence it is worthwhile to search for alternative reconstruction methods that also mitigate the spikes, yet which retain as much of the clinically-significant high-frequency image features as possible.

Technical Approach and Methods: The figure to the right is intriguing because it suggests that the spikes are not really random noise, but rather that there is an underlying deterministic structure that causes them. More specifically, the quantity that is being displayed at each output pixel in this figure is basically the numerical condition number, taken from the system of 4 equations in 4 unknowns that describes the effect of aliasing. For each of the 4 coils, this system sums up the four "ancestor" or "pre-image" pixels in the underlying image that are combined into a single aliased output pixel. The



condition number is a measure of the numerical stability of this system of equations; unlike the determinant, it is independent of scale. The condition number becomes very large when the equations are essentially linearly dependent [2]. In order to see the overall trends without being sidetracked by a few isolated spikes, we have plotted the log of the condition number, rather than the condition number itself. What is striking here is that there is a central diamond-shaped region where the unaliasing equations are six orders of magnitude more stable than at the outlying locations. The diamond-shaped region has a straightforward explanation in terms of the origins (genealogy) of an aliased pixel. This genealogy consists of four ancestor pixels in the underlying unaliased image, as explained by the small colored figures on the right. The first figure has divided the unaliased data into two regions: foreground pixels which lie inside the array (patient signal), and background pixels located outside the array (noise only). After aliasing the first image, the second small colored image shows that each output pixel may have either 1, 2, 3 or 4 ancestors that came from the useful, high-information foreground region, rather than from the noise-only background region. It is rather surprising that although about 50% of the original unaliased image is foreground, only about 3 percent of the aliased output pixels have a full set of 4 high-information foreground ancestors. The region of pixels that possess the maximum of 4 high-quality ancestors is exactly the diamond-shaped region we saw earlier. Outside of this diamond, where there are at most 3 foreground ancestors for a given output, there is not sufficient information to solve 4 equations reliably for 4 unknowns, as reflected in the condition number. Thus the original least-squares fit software, which always solved 4 unaliasing equations in 4 unknowns, was replaced by a version that eliminated the background variables that contribute to each output pixel. For example, if there are only 2 foreground ancestors of a given pixel, then just 2 equations in 2 unknowns are solved.



Results: The simple expedient of limiting the reconstruction process to just the foreground variables results in this greatly-improved reconstruction. Unlike the earlier "naïve" reconstruction, no ad-hoc thresholding has been applied to clean up this image, and at least some of the central discs are now plainly discernible.

Discussion: An issue that could profit from some additional attention is the *boundaries* in the aliased foreground/background image, for example pixels that lie along the edge dividing the 3-ancestor region from the 2-ancestor region. Residual spikes often occur at such edges, because of straddling loss, interpolation loss and partial volume effects. But it is certainly possible to guard against unrealistically-large reconstructed values that might occur at such boundaries, because their location is known.

Conclusion: Outside of the already-well-behaved central diamond, this method reduced the ill-conditioning of the reconstruction equations generally by about six orders of magnitude, with the exception of a few outliers at the boundaries described above. Because of its simplicity, practicality and power, this process of eliminating the background variables from the unaliasing equations is useful for fast space-domain MR reconstruction.

Acknowledgements: Research supported by NIH T32-CA09536, P01-CA67165.

References:

- [1], Pruessmann KP et al., SENSE: Sensitivity encoding for fast MRI, MRM, 42:952-962 (1999)
- [2], Strang WG, Linear Algebra and Its Applications, 3rd edition, Wellesley-Cambridge Press, Wellesley, MA, 2003