Dynamic Range Requirements for MRI

RAYHAN BEHIN,[1] JONATHAN BISHOP,[1] R. MARK HENKELMAN[1,2]

[1] *Mouse Imaging Centre (MICe), The Hospital for Sick Children, 555 University Avenue, Toronto, Ontario M5G 1X8, Canada*
[2] *Department of Medical Biophysics, University of Toronto, Toronto, Canada*

ABSTRACT: In order to realize the full potential in an MR receiver, the digitizer must capture a signal magnitude range from the central *k*-space peak to the thermal noise floor of the system. This dynamic range can exceed the performance of standard 16-bit data converters. For example, a whole-body mouse scan in a 7 Tesla magnet requires 20 bits of dynamic range. A 3D high-resolution mouse scan at 75 μm isotropic voxel resolution using a 16-bit spectrometer shows an eightfold improvement in image SNR by gain stepping the receiver prior to digitization to cover the full magnitude dynamic range compared to a standard fixed gain approach. A method is presented to determine the dynamic range requirements of any experiment. Concepts Magn Reson Part B (Magn Reson Engineering) 26B: 28–35, 2005

KEY WORDS: MRI; high resolution; dynamic range; analog to digital; quantization noise; thermal noise; variable gain; signal-to-noise ratio (SNR); *k*-space sampling

INTRODUCTION

The frequency space (*k*-space) domain of magnetic resonance (MR) imaging is highly peaked at the center (low spatial frequencies) and falls off rapidly toward the periphery of *k*-space. Accurate digitization of this space requires representation at the central point and at the thermal noise level of the system. This range in signal intensity is typically referred to as the dynamic range (DR). High-resolution 3D imaging pushes the dynamic range requirements, ultimately to the thermal noise level of the receiver.

Magnetic resonance (MR) imaging systems typically use a quadrature pair of 16-bit analog-to-digital converters (ADC) operating at 1 MSps (megasamples per second). With low noise receivers and higher magnetic fields, these systems may be signal-to-noise ratio (SNR) limited by the bit resolution of the ADC.

This article illustrates through experiment the linear relationship between *k*-space log magnitude and log radius for a natural image. This relationship is used to determine the range in *k*-space signal magnitude for a given sample volume at any imaging resolution. Finally, a method is given to determine the dynamic range requirements for any MR setup.

THEORY

k-Space

MR imaging involves the capture of complex data in spatial frequency (*k*-space). This data is usually transformed by means of a Fourier transform to represent the magnitude image intensity as a function of spatial position—in other words, an image in real space. The units of measure in *k*-space is inverse distance (m^{-1}) and in the transformed real space image is distance (m). The extent of *k*-space coverage in each frequency axis is inversely proportional to the pixel size in real image space along the corresponding axis. The separation of *k*-space samples in each frequency axis is inversely

Received 23 November 2004; accepted 14 February 2005

Correspondence to: Rayhan Behin; E-mail: rbehin@sickkids.ca

Concepts in Magnetic Resonance Part B (Magnetic Resonance Engineering), Vol. 26B(1) 28–35 (2005)

Published online in Wiley InterScience (www.interscience.wiley.com). DOI 10.1002/cmr.b.20042

proportional to the field of view (FOV) along the corresponding spatial axis once transformed (*1*).

Power Spectra

It has been found that the ensemble of natural images have highly robust statistical features, in particular, roughly a $1/k^2$ power spectra (*2*, *3*). Fuderer (*4*) has shown the power spectrum of thin-slice MR images to decrease as $k_r^{-1.5}$, which has also been confirmed by Watts and Wang (*5*) with their phantom images. This phenomenon leads to an appreciation of a power law increase in dynamic range for increased resolution MR scans where *k*-space frequency samples extend farther.

Solutions for Increased Dynamic Range

There have been several methods described in the literature for capturing high dynamic range signals. One method, which is in commercial use with spectrometers designed by Philips, termed "profile-dependent amplification," involves adjusting the gain in the receiver prior to the ADC to absorb some of the dynamic range in the MR signal (*6*). This gain adjustment is usually implemented with a switchable attenuator. As the MR acquisition approaches the center of *k*-space, additional attenuation is switched into the signal path to reduce the peak signal, which the ADC must accurately sample. This method reduces the dynamic range requirements of the ADC by the amount of amplitude/phase calibrated gain control available. An analogy can be made to a gear box in a car which translates a limited usable range of engine speed to a much greater range in road speed.

Siemens has incorporated a compressor/expander (compander) in their receiver to improve the dynamic range of their digitizer by 4 bits (*7*). Compression is a method used in nonuniform quantization where the analog signal passes through a compression (usually logarithm but in Siemens' case a third order root) amplifier and is then fed into a standard uniform quantizer (ADC). The digital signal is then expanded with higher bit resolution using a look up table to recover the original signal. This technique produces fine quantization steps for the frequently occurring low amplitudes in *k*-space and coarser steps for the less frequent large amplitudes around the center of *k*-space. The total distortion is reduced by decreasing the quantization noise (i.e., increasing the dynamic range) where the probability distribution function is large. The penalty in compressing an analog signal is spectral spreading requiring increased ADC sampling bandwidth.

Another method available involves sampling the MR signal at a higher rate than required to essentially reduce the ADC quantization noise over the bandwidth of interest (*8*). This is a common technique used to improve ADC dynamic range in modern digital receivers, especially when intermediate frequency (IF) sampling is used. Digital filters are used to low-pass filter the high-rate ADC samples, increasing the SNR over a smaller bandwidth where the signal of interest lies. The SNR improvement is directly related to the amount of noise rejected from the raw ADC samples by the digital filters. The noise processing gain (*PG*) in decibels (dB) can be calculated using:

$$PG = 10 \cdot \log_{10}\left[\frac{\left(\frac{F_S}{2}\right)}{BW}\right](\text{dB}) \qquad [1]$$

where F_s is the sampling rate of the ADC and BW is the bandwidth in Hz of the complex signal. Equation [1] yields the maximum possible noise-processing gain by bandwidth reduction using "brick wall" filters. The improvement in SNR can be stated in effective number of bits by dividing the processing gain by 6.02 dB/bit.

Finally, nonlinear gradient pulses have been proposed to compress the dynamic range requirements in MR imaging (*9*). The nonlinear gradient pulse converts the linear phase distribution of the subject into a nonlinear one, thus smoothing out the maximum peak of the *k*-space data. Dynamic range compression of an order of magnitude (3–4 bits) has been shown (*9*).

EXPERIMENT

An experiment was set up to capture the dynamic range in a whole-body mouse MR scan. Measurements were made using a Varian Unity*Inova* NMR spectrometer (Varian Inc., Palo Alto, CA) with a 7 Tesla magnet. A 3D scan was set up using a spin echo sequence, TE/TR 19.15/600 ms, FOV of 30 mm × 30 mm × 100 mm with isotropic voxel resolution of 75 μm. This produces 400 × 400 × 1330 complex *k*-space sample points in an imaging time of 26 hours. A fixed mouse perfused with gadolinium (Gd) was inserted into a Varian Millipede coil (30 mm diameter, 110 mm long) and imaged twice consecutively using a standard/low receiver gain setting to prevent saturating the ADCs and a much higher gain setting (42 dB higher) to allow proper sampling of the thermal noise floor while saturating the ADC for samples at and near the center of *k*-space. The in-phase and quadrature-phase DC levels were subtracted from both data sets.

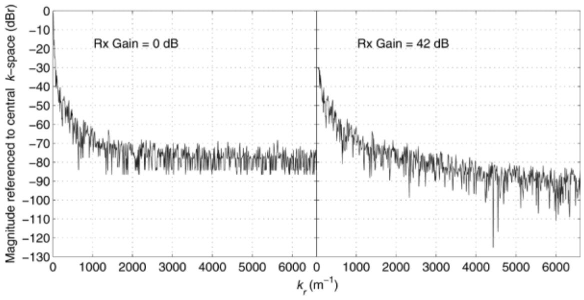

Figure 1 Magnitude of raw data samples from a line through central *k*-space. Plot on the left is from a standard/low-gain scan and on the right is from a high-gain scan. The low-gain scan shows accurate central *k*-space samples but is quickly floored because of the limited quantization range of the ADC. The high-gain scan shows clipping around the central *k*-space region but otherwise has accurate samples for the remainder.

RESULTS

Figure 1 shows the *k*-space magnitude plots of the two captures taken through the central *k*-space region with the high-gain capture shifted by its equivalent gain of 42 dB. The low-gain data set on the left shows accurate representation of the central points, but the data set is quickly floored because of the limited quantization range of the ADCs. In the second data set on the right, the first few points are clipped by the quadrature ADCs because of the 42 dB increase in gain, but now the higher spatial frequency (*k*-space) samples are represented accurately with no sign of a quantization floor in the samples. The high-gain data set also shows a lower thermal noise floor because of a system improvement in noise due to increased signal levels through the receiver path. Because of the rapid drop in *k*-space magnitude as mentioned, only a few points of the low-gain data set with the remaining points from the high-gain data set are needed to create a full *k*-space map. Overlapping data points with good SNR from the low-gain scan and unsaturated points from the high-gain scan are used to determine the amplitude and phase shift due to the gain switch. This single complex valued correction is applied to the low-gain data set and the central 0.0064% of total *k*-space samples from the low-gain data set is substituted into the high-gain data set to form a combined data set.

Figures 2 and 3 show the improvement in image SNR from the standard/low-gain 75 μm isotropic scan on the left compared with a 75 μm image composed from the combined data set on the right. Figure 3 shows a zoom-in on a kidney where one can see the significant benefits from the improvement in voxel SNR showing much greater detail in the kidney medulla and cortex. The voxel SNR, calculated as the mean of the whole-mouse body magnitude image (signal) divided by the mean of the artifact-free background magnitude image (noise) is 2.6 for the standard/low-gain scan image and 20.9 for the combined data set image which is an eightfold improvement.

This definition of voxel SNR for a magnitude image ensures an SNR of 1 is obtained when measuring noise in the image. Another common definition for voxel SNR is the mean of the magnitude image over the region of interest divided by the standard deviation of the magnitude image over the background noise. One can convert the voxel SNR values quoted above to a number according to this definition by multiplying with a factor of 1.913 (*10*). This factor arises from the ratio of the mean to the standard deviation of a Rayleigh distributed function defining the background noise of the magnitude image. The Rayleigh distribution arises from the magnitude operation on the complex Gaussian noise.

DISCUSSION

Two Data Sets

Figure 4 shows the mean magnitude *k*-space points for the low-and high-gain scans; the high-gain data set is level adjusted for comparison with the low-gain data set. The high-gain scan can be seen clipping around −30 dB below the central *k*-space point as also seen in Fig. 1. This can be easily confirmed because the central *k*-space point in the low-gain data set has a magnitude of 11,710 levels given from the quadrature 16-bit offset binary ADCs, which effectively puts it at $20 \cdot \log_{10}(11710/(\sqrt{2}*2^{16}/2)) = -12$ dB from full scale. Therefore, with a 42 dB gain increase in the second data set, we would expect clipping at −30 dBr (dB relative) from central *k*-space.

Characteristic Slope

The mean magnitude high-gain data set can be seen to extend far past the leveling off characteristic of the low-gain data set as we approach the 75 μm isotropic voxel resolution point. If we were to perform the same experiment with a much higher resolution, the high-gain data set would level off at the thermal noise floor of the system shown by the dashed line at −97.7 dB below the central *k*-space point . Note the characteristic slope of the log magnitude to log radius *k*-space response. This slope of −1.55, obtained in repeated mouse scans is similar to the results obtained by Fuderer (*4*) and Watts (*5*), namely the fall in *k*-space RMS signal as $k_r^{-1.5}$. This fall in energy can be attributed to the fact that there is more

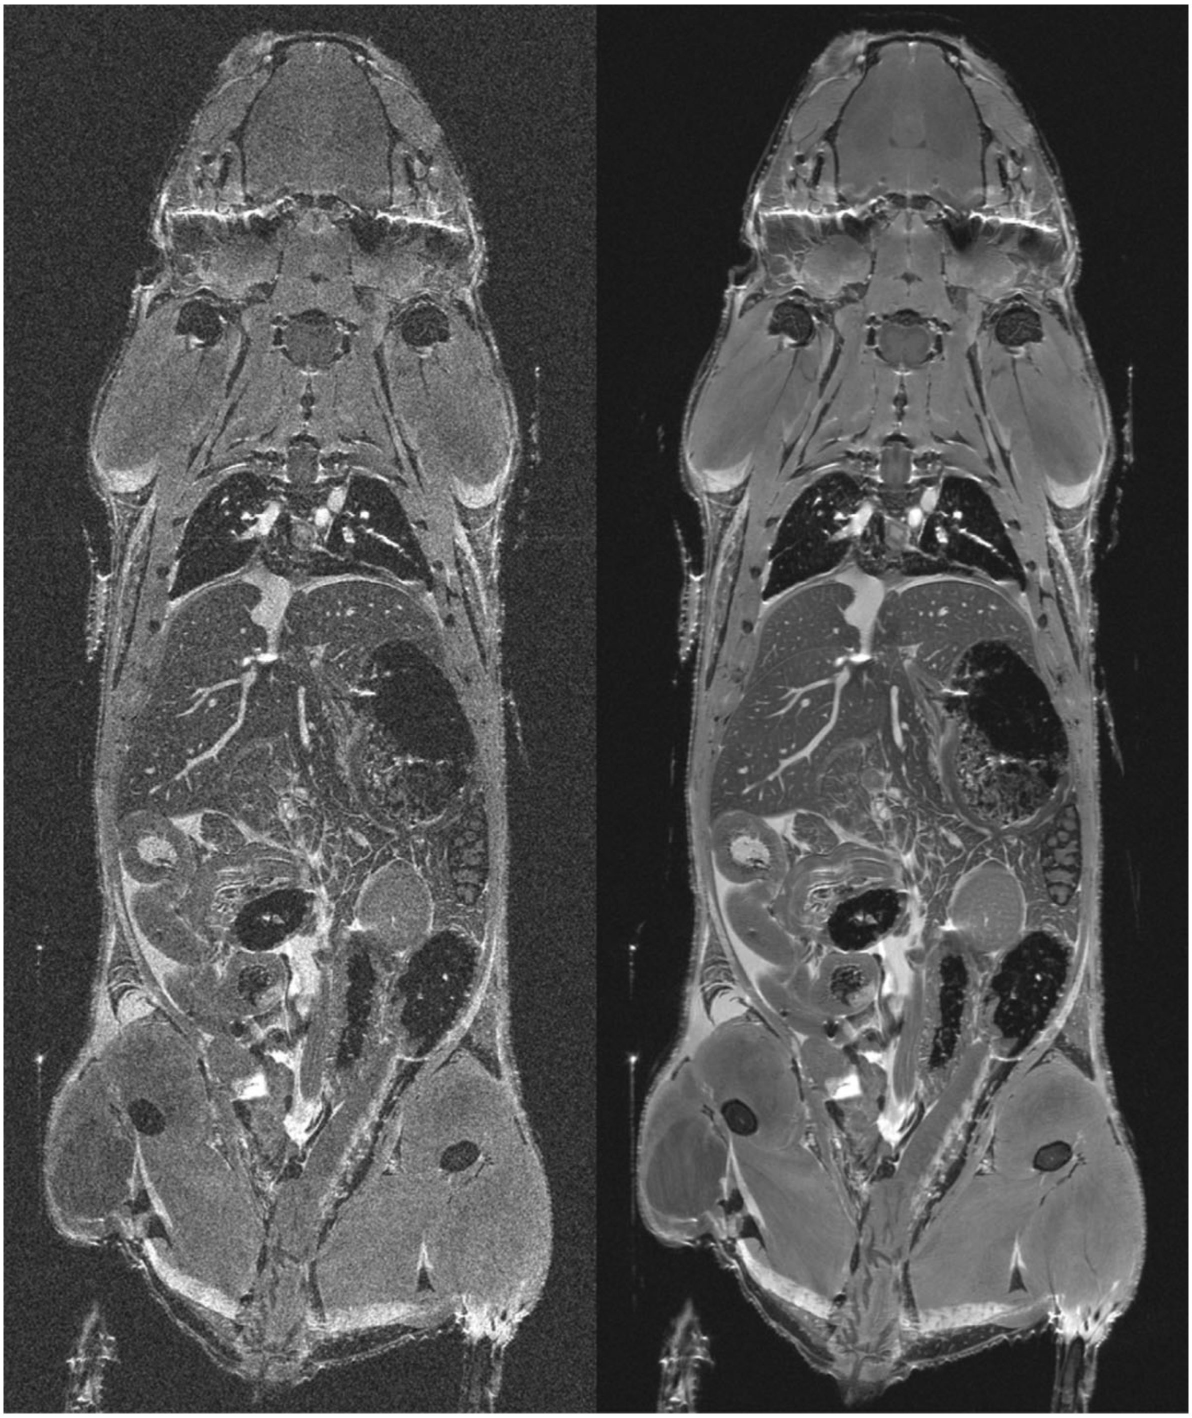

Figure 2 A 75 μm isotropic whole-body mouse scan with a standard/low-gain setting on the left and a combined data set using high-gain samples with amplitude/phase adjusted low-gain central *k*-space samples on the right. The voxel SNR on the left is 2.6 and on the right is 20.9.

signal energy in the large-scale features of the image versus the higher frequency edges. The composite nature of the whole-body mouse can therefore be attributed to the slope of -1.55, and one can generalize that any whole-body mammal would have a similar response, assuming the distribution of large-scale features to the smallest is similar, the absolute whole-body size not being a factor.

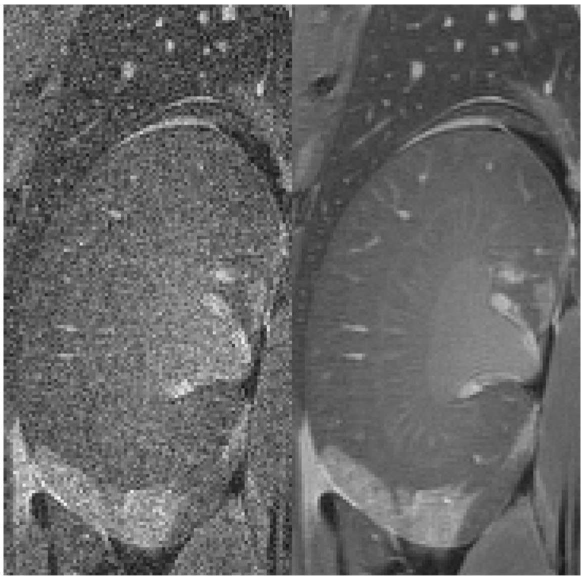

Figure 3 Zoom-in on a kidney of the same 75 μm isotropic scan with a standard/low-gain setting on the left and a combined data set using high-gain samples with amplitude/phase-adjusted low-gain central *k*-space samples on the right. The voxel SNR on the left is 2.6 and on the right is 20.9.

Maybe a more intuitive measure to the design engineer for this slope would be a 30.9 dB drop in *k*-space power magnitude as the resolution is increased 10-fold. This relationship can be used to determine the change in dynamic range of the magnitude signal as the voxel resolution changes for the whole-body mouse.

Noise Factors

The mean magnitude low-gain data set points level off at the noise floor shown by the dotted line in Fig. 4. This noise floor is primarily due to the quantization noise from the ADCs, as we will discover. The measured thermal noise figure (NF) of the receiver actually increases from 2.5 dB to approximately 13.2 dB as we switch from the high-gain setting to the low-gain setting. Because we measured the thermal noise floor with the high-gain setting at −97.7 dBr (confirmed later), the thermal noise at the low-gain setting should increase to −87.0 dBr (−97.7 + 13.2 − 2.5) due to the NF increase. The noise capture at the low-gain setting gave us a total noise level of −79.5 dBr as seen in Fig. 4. Therefore, the quantization noise floor, obtained as the linear power subtraction of the thermal noise floor from the total measured noise,

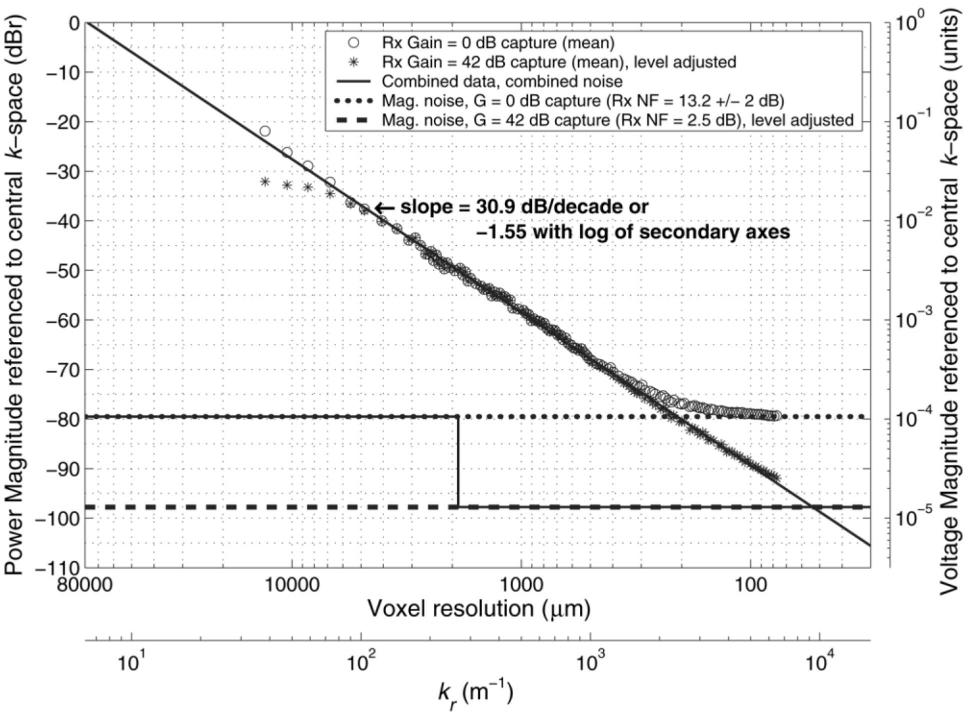

Figure 4 Log magnitude versus log *k*-space radius for a whole-body mouse scan. The low-gain and high-gain mean magnitudes are plotted with a best-fit line associated with the combined data set. The power spectra falls off as $k_r^{-1.55}$. The low-and high-gain scan mean magnitude noise floors are shown overlaid by the combined data set noise floor.

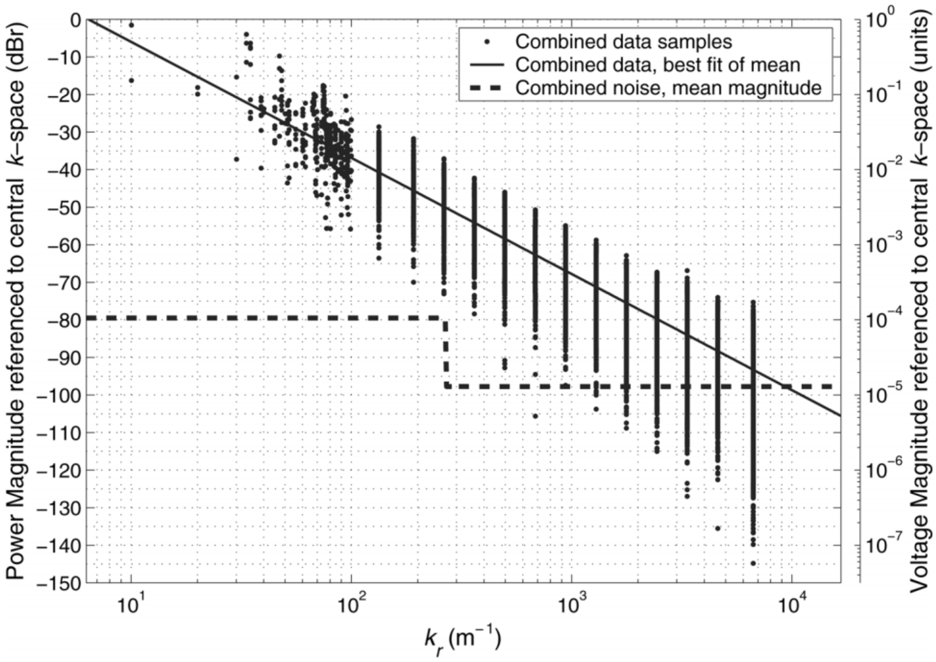

Figure 5 Magnitude k-space samples up to k_r of 100, after which the magnitude samples are shown at select radii for the combined data set.

is −80.4 dBr. Because the magnitude level for the central k-space point is 12 dB below quadrature ADC full scale as was calculated above, the absolute quantization noise floor should be 92.4 dB below quadrature ADC full scale. The data sheet for the ADCs used by the Varian Unity*Inova* NMR spectrometer (Analogic ADC4320) shows a typical SNR of around 90 dB. Used in quadrature, an SNR of 93 dB should be obtainable, which is approximately what we have shown. It is important to check the SNR specified for the ADC in question at the particular input frequency and signal level and not to assume the standard sine wave SNR for an n-bit ADC given by Eq. [2].

$$SNR = 6.02 \cdot n + 1.8 \text{ (dB)} \qquad [2]$$

To calculate the noise composition in the high-gain data set, we use the full-scale magnitude of the quadrature ADC combination at −30 dBr combined with the 93 dB of quadrature ADC SNR to obtain a quantization noise level of −123 dBr. This quantization noise level is far below the measured noise level of −97.7 dBr, contributing almost nothing to the overall noise and thus confirming that our high-gain data set is in fact receiver thermal noise limited.

Gain Step and *k*-space Magnitude Spread

Finding the best gain transition radius in forming the combined data set involves choosing the shortest k-space radius while ensuring no clipped data is included from the high-gain data set. Choosing the shortest radius ensures minimum impact of the higher noise level in the low-gain data set to the overall combined data set. Figure 5 shows all the k-space samples up to 100 k_r and after that the spread of k-space samples at discrete k_r steps to the maximum radius at 6,667 k_r. You will notice why the transition radius was chosen at 267 k_r or 4% of the full radius. We know from our discussion that the clipping level in the high-gain data set is around −30 dBr and therefore we would not want to be near that level when including the high-gain data set. From Fig. 5, the spread of magnitude k-space samples at 267 k_r shows a maximum magnitude of around −37 dBr, which is safely below the full-scale magnitude of the quadrature ADCs.

Another observation that can be collectively made from Figs. 4 and 5 is that although there are many points in the higher k-space radii that are below the mean noise floor, the mean of all the samples is still above the noise floor and is following the characteristic slope, therefore contributing more signal energy than noise with an improvement in resolution. There may even be justification for capturing k-space samples past the thermal noise floor of the receiver if the benefit of increased resolution is important despite the degradation in SNR due to the volumetric increase in noise.

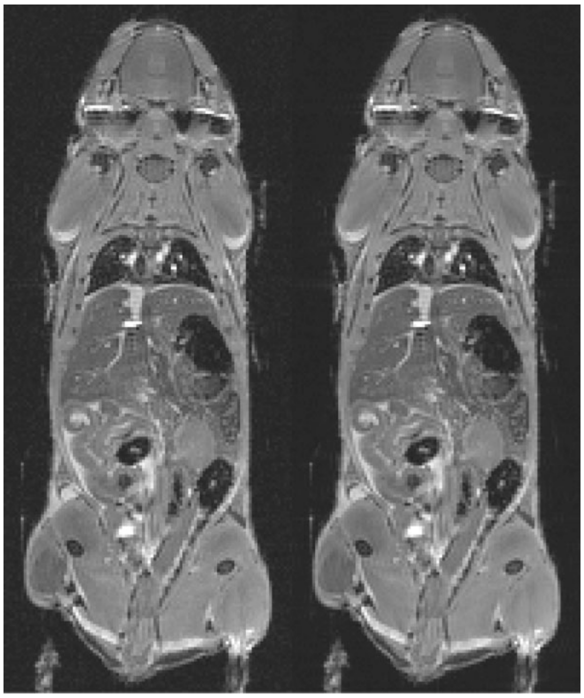

Figure 6 A 300 μm isotropic whole-body mouse scan with a standard/low-gain setting on the left and a combined data set using high-gain samples with amplitude/phase-adjusted low-gain central *k*-space samples on the right. The voxel SNR on the left is 18.8 and on the right is 102.4. Visually, the image quality between the two is nearly identical.

Relaxed Dynamic Range Case

Figure 6 shows the same whole-body mouse slice, this time with 300 μm voxel resolution. The image on the left is that composed with only the standard/low-gain scan, and the image on the right is composed from the combined data set. Although the voxel SNR of 102.4 is much better with the combined data set on the right, there is almost no visual difference compared to the standard/low-gain data set on the left having a voxel SNR of 18.8 At this *k*-space cut-off radius (Fig. 4) generating the 300 μm image one can observe the low-gain data set and the combined data set samples following the same characteristic slope and above their respective noise levels, though the low-gain data set samples have a much lower SNR per sample due to the proximity of the quantization noise level. This is a reason why systems with limited dynamic range do not have an observable problem with lower resolution images. Also, MR systems using rapid imaging (e.g., fMRI), SENSE, selective excitation, a lower magnetic field, or much smaller sample sizes have a lower peak signal and thus may not suffer from a limited dynamic range.

Optimum Dynamic Range

To determine the dynamic range requirements to digitize the full range of magnitude signal from a particular MR setup, the maximum signal and the noise floor of the receiver need to be known. For this experiment, the maximum signal at the receiver front end (i.e., preamplifier input after the quadrature combination of the orthogonal coils) is −25.7 dBm (dB relative to 1 milliwatt). This maximum signal was measured by comparing ADC readings of the central *k*-space point with that from a calibrated signal source. The noise figure of the receiver at 50 Ω without the coil was measured at 2.5 dB in the high-gain setting using standard techniques. Another experiment showed a 1.1 dB degradation in noise figure with a mouse-loaded coil connected to the preamplifier as compared with a 50 Ω load. This increase in noise figure can be attributed to a slight 50 Ω mismatch of the loaded coil. Equation [3] shows the standard thermal noise (N_{out}) equation referenced to ambient temperature as a function of noise figure (NF) and signal bandwidth (BW).

$$N_{out} = -174 + 10 \cdot \log_{10} BW + NF \text{ (dBm)} \quad [3]$$

The dual-gain experiment was performed with a sampling rate of 100 kSps, thus giving us a signal bandwidth near 50 kHz. Using an adjusted noise figure of 3.6 dB gives us a thermal noise floor of −123.4 dBm. Therefore, the thermal noise floor is 97.7 dB below the maximum signal. This is exactly what we see in Fig. 4 as the dashed line representing the high-gain noise capture, thus confirming this method of quantifying the thermal noise floor.

A final step in this analysis is to determine the placement of the quantization noise floor relative to the thermal noise floor. For minimal impact to the overall noise floor due to quantization noise and to maintain a reasonable design goal, a common practice is to place the quantization noise floor 10 dB below the thermal noise floor. This placement raises the overall noise floor 0.4 dB (using linear power addition) above the thermal noise floor of the receiver, which should be acceptable.

For our particular system, the best noise figure possible at a high-gain setting is 2.4 dB, which in a 50 kHz bandwidth on a perfectly tuned loaded coil gives us a noise floor of −124.6 dBm. Therefore, the quantization noise floor should be set at −134.6 dBm. The maximum signal obtainable with our 7 Tesla magnet and a millipede coil on the largest fixed Gd infused mouse is −20 dBm. Therefore, the dynamic range required to properly digitize any scan resolution pos-

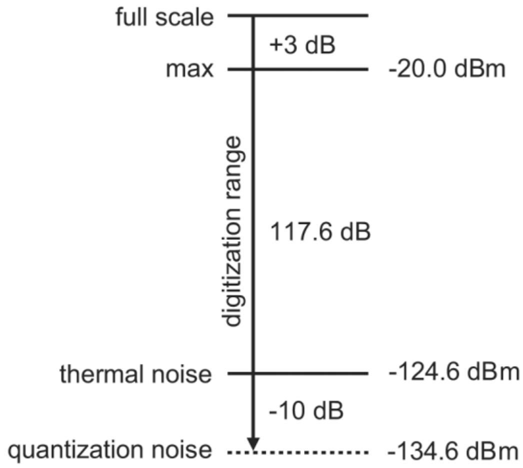

Figure 7 The dynamic range requirements for digitization of a whole-body mouse in a millipede coil at 7 Tesla field strength over a signal bandwidth of 50 kHz with a receiver noise figure of 2.4 dB.

sible using 50 kHz of signal bandwidth is 114.6 dB. For a practical design, when setting the central *k*-space point below the ADC full scale with the variable gain stage of the receiver, margin must be put in to allow for the step resolution plus any variation in gain from step to step. A practical realization is to use 2 dB gain steps and a 1 dB margin for gain variation. This increases the dynamic range requirement to 117.6 dB, summarized in Fig. 7. Therefore, using Eq. [2], we reach an overall requirement for 20 bits of resolution. Of course, the 93 dB SNR achievable with the Varian Unity*Inova* NMR spectrometer is not quite adequate and an easy solution would be to apply gain switching with a calibrated amplitude and phase adjustment for the critical high-signal magnitude central *k*-space region. Higher magnetic fields and larger samples will increase the dynamic range requirements even further.

CONCLUSION

Ideally, an MR receiver should be able to digitize the full range of signals from the maximum at the central *k*-space point to the thermal noise level of the receiver. The standard 16-bit conversion may not be sufficient and its limitation becomes apparent in high-resolution scans. For a whole-body mouse scan at 7 Tesla field strength, a dynamic range of 20 bits was determined for maximal SNR. Solutions for improving the dynamic range have been mentioned, including gain stepping, oversampling, and nonlinear gradient pulses. A simple method has been presented to determine the dynamic range requirements of any experiment.

ACKNOWLEDGMENTS

The Mouse Imaging Centre acknowledges funding from the Canada Foundation for Innovation and the Ontario Innovation Trust for providing facilities along with The Hospital for Sick Children. Operating funds from the Burroughs Wellcome Fund, the Canadian Institutes of Health Research, the National Cancer Institute of Canada—Terry Fox Program Projects, the National Institutes of Health, and the Ontario Research and Development Challenge Fund are gratefully acknowledged. R. Mark Henkelman holds a Canada Research Chair in Imaging.

REFERENCES

1. Paschal CB, Morris HD. 2004. *k*-Space in the clinic. J Magn Reson Imaging 19:145–159.
2. Field DJ. 1987. Relations between the statistics of natural images and the response properties of cortical cells. J Opt Soc Am A 4:2379–2394.
3. Ruderman DL, Bialek W. 1994. Statistics of natural images: scaling in the woods. Phys Rev Lett 73:814–817.
4. Fuderer M. 1988. The information content of MR images. IEEE Trans Med Imaging 7:368–380.
5. Watts R, Wang Y. 2002. k-Space interpretation of the Rose model: noise limitation on the detectable resolution in MRI. Magn Reson Med 48:550–554.
6. Elliott MA, Insko EK, Greenman RL, Leigh JS. 1998. Improved resolution and signal-to-noise ratio in MRI via enhanced signal digitization. J Magn Reson 130: 300–304.
7. Bollenbeck J, Vester M, Oppelt R, Kroeckel H, Schnell W. A high performance multi-channel RF receiver for magnet resonance imaging systems. In: Proceedings of the 13th Annual Meeting of ISMRM, Miami Beach; 2005, p 860.
8. Delsuc MA, Lallemand JY. 1986. Improvement of dynamic range in NMR by oversampling. J Magn Reson 69:504–507.
9. Wedeen VJ, Chao YS, Ackerman JL. 1988. Dynamic range compression in MRI by means of a nonlinear gradient pulse. Magn Reson Med 6:287–295.
10. Henkelman RM. 1985. Measurement of signal intensities in the presence of noise in MR images. Med Phys 12:232–233.