

Real-Time Image Reconstruction for Pulse EPR Oxygen Imaging Using a GPU and Lookup Table Parameter Fitting

GAGE REDLER,¹ ZHIWEI QIAO,² BORIS EPEL,¹ HOWARD J. HALPERN¹

¹ Department of Radiation and Cellular Oncology, University of Chicago, Chicago, IL 60637

² School of Computer and Information Technology, Shanxi University, Taiyuan, Shanxi 030006, China

ABSTRACT: The importance of tissue oxygenation has led to a great interest in methods for imaging pO₂ *in vivo*. Electron paramagnetic resonance imaging (EPRI) provides noninvasive, near absolute 1 mm-resolved 3D images of pO₂ in the tissues and tumors of living animals. Current EPRI image reconstruction methods tend to be time consuming and preclude real-time visualization of information. Methods are presented to significantly accelerate the reconstruction process in order to enable real-time reconstruction of EPRI pO₂ images. These methods are image reconstruction using graphics processing unit (GPU)-based 3D filtered back-projection and lookup table parameter fitting. The combination of these methods leads to acceleration factors of over 650 compared to current methods and allows for real-time reconstruction of EPRI images of pO₂ *in vivo*. © 2015 Wiley Periodicals, Inc. Concepts Magn Reson Part B (Magn Reson Engineering) 45B: 46–57, 2015

KEY WORDS: electron paramagnetic resonance imaging; oxygen imaging; real-time image reconstruction; graphics processing unit; lookup table fitting

INTRODUCTION

Studies have shown that the oxygenation status of tumors has many important prognostic implications (1). Low oxygen concentration (pO₂), or hypoxia, has been found to result in increased resistance to radiation therapy for cancer cells (2,3) and in human patients there is strong correlation between electrode measurements of hypoxia and radiotherapy treatment failure (4). Hypoxic tumors have increased resistance to chemotherapy (5), faster tumor growth caused by abnormal proliferation (6), and higher potential for metastatic progression (7).

Due to these effects associated with hypoxia, there is an interest in methods for probing/imaging tissue pO₂. Advances in such methods are detailed in the literature (8–11). Electron paramagnetic resonance imaging (EPRI) has proven to be a robust modality for imaging tissue pO₂. EPRI provides noninvasive, highly resolved, both spatially (~1mm³ voxels) and in pO₂ (1–3 torr), near absolute quantitative 3D images of *in vivo* pO₂ (12–16). Low frequency (250MHz) EPRI excitation (comparable to those for 6T MRI) provide penetration and, therefore, imaging deep into tissue (>7cm). Standard EPRI image acquisition time is currently on the order of 10 minutes.

EPRI can be efficiently implemented as a tomographic imaging modality. EPR signals acquired under the conditions of applied magnetic field gradients can be viewed from a mathematical standpoint as Radon transformation projections acquired in a 3D spherical geometry. The inversion of the Radon transformation can be applied to recover a 3D image from these projections. Various reconstruction methods can be used. One of the common

Received 29 May 2014; revised 12 May 2015; accepted 12 May 2015

Correspondence to: Howard J. Halpern; E-mail: h-halpern@uchicago.edu and Gage Redler; E-mail: gage@uchicago.edu

Concepts in Magnetic Resonance Part B, Vol. 45B(1) 46–57 (2015)

Published online in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/cmrb.21281

© 2015 Wiley Periodicals, Inc.

analytic (those that provide explicit expressions in the form of analytic functions for the solution) reconstruction methods is filtered back-projection (FBP). For practical application, the analytic algorithm must be discretized to handle discrete data and the resulting sampled FBP is an approximation to the analytic solution.

This article describes an improvement to pulse EPRI methodology. A single 3D EPRI image provides the spin distribution in a sample. To acquire the pO_2 distribution, a series of 3D images are acquired with different pulse sequence parameters, which are chosen to monitor the evolution of spin magnetization as a function of time, the process of spin-lattice relaxation. The spin-lattice relaxation rate is then obtained in every voxel of a 3D image by fitting this evolution to a known function of time. The local spin-lattice relaxation rate is proportional to pO_2 and therefore provides a quantitative and near absolute measure of pO_2 at each voxel (17). 3D image reconstruction is a part of pO_2 image processing and one of the major contributors to the pO_2 image processing time.

Conventional, “multi-stage,” reconstruction of 3D tomographic EPRI images has been limited to particular projection schemes. This is a computationally efficient but highly restrictive method (18,19). In particular, this method does not allow correct image reconstruction from an incomplete or other arbitrary projection set. The most flexible FBP implementation for reconstruction of 3D images is a “single-stage” method, which involves direct back-projection of each projection through the entire 3D image. This method is computationally intensive and relatively slow on a conventional CPU. In this work, we present an implementation of single-stage reconstruction for EPRI using a graphics processing unit (GPU), which provides both versatility to reconstruct images from arbitrary projection distributions and speed. GPUs have been used in many imaging applications to improve reconstruction speed. In fact, the first non-graphics GPU application was to accelerate FBP for CT image recon-

struction (20). GPUs have also been used for many other image reconstruction applications (21).

An iterative fitting method can be used to characterize the exponential signal intensity decay rate measuring pO_2 at each voxel. However, this is a time consuming process. An alternative fitting method is described here, which implements a pre-defined lookup table (LUT) to match the observed signal decay to ideal signal decay curves. This LUT method may result in modestly reduced accuracy but is significantly faster. In certain circumstances, sacrificing image accuracy for convenience and speed is a favorable tradeoff. This is particularly true if offline image reconstruction will regain that accuracy. The method utilized in this work is similar to that used in magnetic resonance fingerprinting (22).

Conventional personal computer central processor unit (CPU)-based reconstruction and optimization-based parameter fitting are both time consuming processes and represent bottlenecks for the speed with which EPRI pO_2 images can be reconstructed. Using these methods, flexible real-time EPRI oxygen image reconstruction is extremely difficult. This work focuses on new methods and implementations of image reconstruction routines providing reconstruction speed enhancements, which enable a more flexible, real-time EPRI reconstruction.

Real-time image reconstruction has been implemented in several imaging modalities, such as MRI (23,24) and CT (25). *In vivo* imaging is often time sensitive, especially for human subjects. Real-time reconstruction, allowing for visualization of the buildup of an image as projection data are acquired, can help to reduce scanning time by resolving certain issues early. Such issues can include errors in setup of either the subject or imaging equipment, subject motion, imaging equipment failure, etc. Real-time reconstruction can allow an operator to potentially identify a particular region of interest (ROI) earlier during imaging. Furthermore, as the temporal resolution of EPRI continues to improve, real-time reconstruction will be necessary to eventually enable real-time visualization of *in vivo* pO_2 distributions, essentially analogous to fluoroscopic X-ray imaging.

This paper discusses the simultaneous implementation of GPU-based FBP reconstruction to more rapidly generate 3D images and the use of a LUT for rapid fitting of pO_2 related parameters. The combination of these two methods results in multiple orders of magnitude image reconstruction speed enhancement, which provides the capability to

Abbreviations

CUDA	compute unified device architecture
EPRI	Electron paramagnetic resonance imaging
ESE	Electron spin echo
FBP	filtered back-projection
GPU	graphics processing unit
IRESE	inversion recovery electron spin echo
LUT	lookup table
ROI	region of interest
SM	streaming multiprocessor
SNR	signal to noise ratio

reconstruct an updated image as each successive projection is acquired. This enables real-time reconstruction of EPRI pO₂ images.

METHODS

Pulsed EPRI Imaging Model

Pulse EPRI methods developed in our laboratory have been detailed elsewhere (26). A brief description follows. Electron spin echo (ESE) R_2 imaging was our first approach (27). More recently, inversion recovery electron spin echo (IRESE) R_1 imaging has been implemented, which reduced confounding spin probe concentration dependent self-broadening and therefore enabled essentially absolute pO₂ imaging (17). The pulse EPRI signal comes from an exogenous triarylmethyl (trityl) spin probe with an unpaired electron centrally located on the molecule that is remarkably difficult to reduce (16,28). Both R_1 and R_2 imaging methods provide oxygen information. The electron relaxation rates (R_1 and R_2) are proportional to pO₂ and therefore directly report local pO₂. Linear magnetic field gradients are applied to spatially encode an object. Gradients are applied along different directions in a 3D spherical geometry, producing radial k -space projections, which, when Fourier transformed, provide spatial projections. An inverse radon transform, such as FBP, from a set of spatial projections obtained from different angular views can reconstruct a 3D image of spin density. A series of 3D images is acquired at different times during magnetic relaxation to observe average relaxation rates from each image voxel. Fitting the signal relaxation rates at each voxel to known exponential functions governing the relaxation (refer to Eq. [6]), produces a 3D parametric image with pO₂ information from each voxel.

3D FBP Algorithm

Without derivation, the FBP algorithm (29) is described in Eq. [1–5]. Here, $f(x, y, z)$ is the 3D object being imaged; $p(t, \varphi, \theta)$ is a 1D spatial projection at angle (φ, θ) ; $g(t, \varphi, \theta)$ is the filtered 1D projection; and $h(t)$ is the unit impulse response for the parabolic filter. In Eq. [2], t is the location of a point (x, y, z) on a given projection at some angle (φ, θ) . In Eq. [4], the operator $F^{-1}\{-\}$ represents the inverse Fourier transform. In practice, these functions and variables are discretized, resulting in a sampled FBP algorithm that provides an approximation of the analytic solution.

$$f(x, y, z) = \int_0^{\frac{\pi}{2}} \int_0^{2\pi} g(t, \varphi, \theta) \sin \theta d\varphi d\theta \quad [1]$$

$$t = x \cos \varphi \sin \theta + y \sin \varphi \sin \theta + z \cos \theta \quad [2]$$

$$g(t, \varphi, \theta) = p(t, \varphi, \theta) * h(t) \quad [3]$$

$$h(t) = F^{-1}\{\omega^2\} = \int_{-\infty}^{+\infty} \omega^2 e^{i2\pi\omega t} d\omega \quad [4]$$

$$p(t, \varphi, \theta) = \iint \int_{-\infty}^{+\infty} f(x, y, z) \quad [5]$$

$$\delta(x \cos \varphi \sin \theta + y \sin \varphi \sin \theta + z \cos \theta - t) dx dy dz$$

GPU Implementation of FBP

The GPU is a processing unit specialized to rapidly execute many specific operations in parallel for accelerated creation of images intended for output to a display. The back-projection process in FBP is an analogous process and can take advantage of the excellent parallel computing abilities of the GPU using compute unified device architecture (CUDA) techniques. The GPU card used was a commercially available GeForce GTX 760 with compute capability 3.0, global memory 4GB, and 1,152 cores or streaming processors.

The GPU consists of a streaming multiprocessor (SM), which is comprised of many streaming processors (SPs). The GeForce GTX 760 has 6 SMs, each containing 192 SPs. To utilize the acceleration capability of the GPU for the time consuming back-projection process, a kernel function executing the necessary operations for back-projection must be written in CUDA C. The GPU runs the kernel function on a grid, which is a set of blocks, each consisting of a set of threads. A thread is a single execution of the kernel function with a specific index. Each thread uses its index to select the processed data and execution unit so that all threads may cooperatively process the entire data set in parallel.

An efficient parallelization strategy involves assigning a calculation for each specific voxel to an individual thread. In general, blocks and grids can be arranged into 1D, 2D, or 3D structures. For a reconstructed image matrix of size $[N, N, N]$, the thread block is chosen to be a 1D structure of size $[N]$ and the grid is chosen to be a 2D structure of size $[N, N]$. As such, the kernel function then performs the parallelized back-projection process so that the proper value or values from each projection point or points are back-projected onto each voxel. This is done for all voxels simultaneously. The results of the image reconstruction on the GPU are then retrieved.

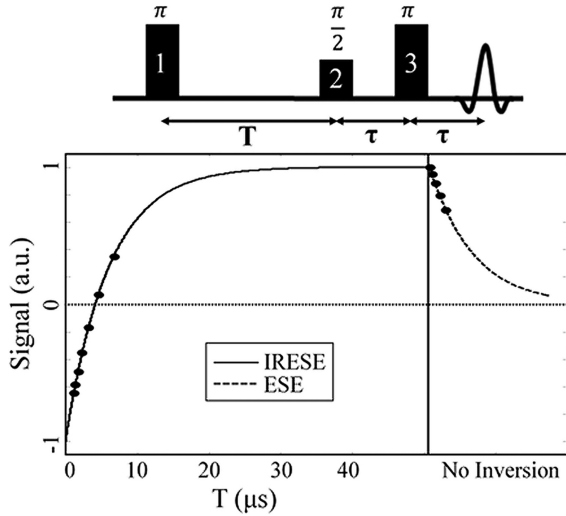


Figure 1 Diagram of pulse sequence used for EPRI is shown at top of figure. For ESE pulse sequences, only pulses 2 and 3 are used, whereas for IRESE pulse sequences all three pulses are used. At the end of the pulse sequence the echo can be read out after a delay (τ). The expected shape of the magnetic relaxation curve at the 12 different time points used for R_1/R_2 imaging is also plotted.

R_1/R_2 Hybrid Imaging

Equation [6] below shows the function describing the EPR signal (S) as a function of inversion delay time (T) and echo delay time (τ), as well as spin probe concentration (A), R_1 relaxation, and R_2 relaxation. While R_1 imaging is the preferred EPRI technique, both R_1 and R_2 imaging have advantages and disadvantages.

For R_2 imaging, signal is not a function of inversion delay time because no inversion is used (equivalent to $T = \infty$), which results in an exponential decay as seen in the right side of the curve in Fig. 1. The disadvantage of R_2 imaging is that accuracy of reported pO_2 is decreased by self-broadening effects for higher concentrations of spin probe (17). The advantage of R_2 imaging is that it provides an absolute measure of the voxel amplitude, A , which is a measure of spin probe concentration.

For R_1 imaging, signal is not a function of echo delay (held constant). This results in exponential recovery as seen in the left side of the curve in Fig. 1. The advantage of using R_1 imaging is that it ameliorates the detrimental self-broadening effects present in R_2 imaging, thereby providing an absolute measure of pO_2 . The disadvantage is that A measured using R_1 imaging is not a true measure of A because it has a confounding dependence on the R_2 relaxation.

It is useful to have an accurate measurement of both spin probe concentration and pO_2 . This is why a hybrid R_1/R_2 imaging methodology is introduced (Fig. 1), where both inversion recovery electron spin echo (IRESE) and ESE pulse sequences are used to measure both R_1 and R_2 , and subsequently use Eq. [6] to simultaneously fit A , R_1 , and R_2 .

$$S(T, \tau) = A(e^{-2\tau R_2})(1 - 2e^{-TR_1}) \quad [6]$$

Standard R_2 imaging uses five values for τ (logarithmically spaced from $0.730 \mu s$ to $3.000 \mu s$) to obtain a fit. For R_2 imaging, the ESE pulse sequence consists of a $\pi/2$ pulse, followed by a delay (τ), then a π pulse and the same delay (τ) before the echo is read out. This is the pulse sequence shown in Fig. 1 if only pulses 2 and 3 are used.

Standard R_1 imaging uses eight values for T (the first seven are logarithmically spaced from $0.430 \mu s$ to $6.000 \mu s$; the final point has no inversion, i.e., $T = 0.000 \mu s$) and a fixed τ value ($0.730 \mu s$) to obtain a fit. For R_1 imaging the IRESE pulse sequence consists of an inversion π pulse, followed by a delay (T) for longitudinal magnetization recovery, and an ESE pulse sequence to read out the echo. To simulate infinite T , no inversion pulse is used (pulse 1, π pulse) for the final point.

Since no inversion is used for R_2 points, the final R_1 point and the initial R_2 point are equivalent. Therefore, for R_1/R_2 hybrid imaging, 12 points (as opposed to the combined 13 R_1 and R_2 points) are used in total for the fit (Fig. 1). T values used for the first eight points are those used for standard R_1 imaging and, for the final four points, no inversion is used (i.e., $T = 0.000 \mu s$). The τ values for the final five points are those used for standard R_2 imaging and, for the first seven points, $\tau = 0.730 \mu s$. Acquiring more points (12 as opposed to 8 or 5) results in an R_1/R_2 image with lower signal to noise ratio (SNR) compared to an R_1 image of the same time duration (SNR for R_1/R_2 image is 80% that of an R_1 image). With this tradeoff in SNR, R_1/R_2 imaging provides images with accurate measures of both pO_2 and spin probe concentration.

Lookup Table Fitting Algorithm

Currently, voxel-by-voxel curve fitting is performed using built in MatlabTM (Mathworks, MA) functions. For R_1/R_2 imaging, the difference between acquired data points and a curve calculated from the fitting function given in Eq. [6], with guesses for the fit parameters (A , R_1 , and R_2) is iteratively

minimized using the “fminsearch” function in Matlab™. Unconstrained nonlinear optimization is used to iteratively update the fit parameters to best fit the data (30). This method can be time consuming (see Table 4). This fitting method will be referred to as the FMS method.

An alternative fitting method uses a LUT to match acquired data to curves generated using different fit parameter values. The curve generated from known parameters that provides the closest match gives an approximation of the corresponding parameters for the acquired data. This can also be referred to as using a dictionary. Previous studies have used such methods for parameter mapping in MRI applications. Complete dictionaries (number of dictionary elements equal to data dimensionality) can be generated using independent component analysis or principal component analysis (31). Overcomplete dictionaries have been used as well (32,33). These methods however are relatively complicated and dictionary generation and fitting can be time consuming (on the order of tens of minutes).

In this work, a less complicated and faster LUT fitting method is used, similar to other previously used methods (22,34). The algorithm consists of two main parts: LUT generation and data fitting. To generate the necessary LUTs, parameters from the data to be fit are determined: number of time points, τ and T values used for each time point, and number of time points corresponding to either R_1 or R_2 portions of the data.

A linear exponential decay fit is used for both R_1 and R_2 determination (Eq. [7] and [8], respectively). For R_1/R_2 imaging, R_1 and R_2 LUTs are generated and used separately. The range and resolution of R_2 values (\propto pO₂) suitable for the data are determined for R_2 LUT generation. Here, the R_2 values used were linearly spaced from ΔR_2 to $1.61 \mu\text{s}^{-1}$, with ΔR_2 equal to 0.001, 0.01, or $0.1 \mu\text{s}^{-1}$. This range corresponds to a pO₂ range of 0–150 torr, which includes values observed in previous EPRI experiments. Unit vectors are necessary for the LUT entries because degree of matching between data and LUT entries is determined based on dot product amplitude.

Unit vector R_2 LUT entries, \hat{L}_{R_2} , are generated using Eq. [7], (derived from Eq. [6] with $A = 1$ and $T = \infty$). The R_1 LUT is produced in the same way. The range of R_1 values (\propto pO₂) used is the same as that used for the R_2 values. Equation [8] shows the exponential used to generate the vectors that are normalized to be unit vectors and used as the entries in the R_1 LUT, \hat{L}_{R_1} .

$$\hat{L}_{R_2} = \frac{e^{-2\bar{\tau}R_2}}{|e^{-2\bar{\tau}R_2}|} \quad [7]$$

$$\hat{L}_{R_1} = \frac{e^{-\bar{T}R_1}}{|e^{-\bar{T}R_1}|} \quad [8]$$

Following LUT generation, R_2 is determined first. The R_2 portion of the data is selected (final 5 data points) and its dot products with all R_2 LUT entries are computed. The entry resulting in the maximum dot product is found and used to determine the corresponding R_2 value.

A is then determined. Again, only the R_2 portion of the data is used. From Eq. [6] it can be seen that, for ideal data, A is equal to the R_2 data points, $S_{R_2,i}$, divided by the corresponding factor accounting for exponential decay due to R_2 relaxation. Ideally, this quotient would be constant for all data points. However, due to noise, this is not necessarily true and A is approximately determined for non-ideal data by finding this quotient for each of the 5 data points and taking the average, as in Eq. [9], where N is the number of R_2 data points (e.g., 5), $S_{R_2,i}$ is the i th entry in the R_2 data vector, τ_i is the τ delay associated with this i th R_2 data point, and the denominator represents the corresponding factor accounting for exponential signal decay due to R_2 relaxation.

$$A = \frac{1}{N} \sum_{i=1}^N \frac{S_{R_2,i}}{e^{-2\tau_i R_2}} \quad [9]$$

The final step is to determine R_1 . Similar to the R_2 fitting process, only the R_1 portion of the data is considered. Equation [10] shows how the R_1 data is conditioned to the functional form in Eq. [8], where \bar{S}_{R_1} is the original R_1 portion of the data and \bar{S}'_{R_1} is the conditioned data. The dot products of \bar{S}'_{R_1} with all R_1 LUT entries are then computed. The entry resulting in the maximum dot product determines the corresponding R_1 value. It should be noted that the coefficient with a value of 2 in Eq. [6] is for an ideal inversion pulse that rotates the magnetization by exactly π . However, in practice there are spatial inhomogeneities in the applied pulse, resulting in imperfect magnetization inversion, causing this coefficient to deviate from the ideal value of 2 based on spatial location. However, by conditioning the affected R_1 data as in Eq. [10], the true coefficient is retained as a scaling of the magnitude of the data vector \bar{S}'_{R_1} corresponding to a specific spatial location. This additional scalar multiplicative factor will not alter which LUT entry corresponds to the maximum dot product with \bar{S}'_{R_1} as this is a relative comparison.




Figure 2 (a) 2D slice through the 3D mathematical phantom used to generate simulated projections. The spherical regions are labeled. (b) Photograph of the 1.0 mM trityl bottle phantom used for imaging.

$$\bar{S}'_{R_1} = 1 - \left(\frac{\bar{S}_{R_1}}{Ae^{-2\tau R_2}} \right) \quad [10]$$

Simulated Model

Simulated experiments with and without noise are used to compare the accuracy of different fitting methods. A mathematical phantom consisting of six spheres: five non-overlapping spheres of equal size (0.5 cm radius) embedded into a larger sphere (2.5 cm radius). Fig. 2(a) shows a slice through the 3D mathematical phantom.

The model used allows for arbitrary choices of A , R_1 , and R_2 . The parameters used for simulated experiments are shown in Table 1. The R_1 values used correspond to an approximate pO_2 range of 2–25 torr. To avoid unwanted inaccuracies from reconstruction artifacts, a large number of projections were used for simulated experiments (6,368; equal solid angle distribution). Projections were calculated using the Radon transform. Reconstructed images were $64 \times 64 \times 64$ voxels, corresponding to a $10 \times 10 \times 10 \text{ cm}^3$ field of view.

Noisy data were simulated by adding random Gaussian noise to projections. Noise levels in exper-

imental data were used to determine the magnitude of the simulated noise. Noise levels corresponding to 60, 10, 5, 2, and 1 minute images were used. Projection SNR corresponding to these imaging times for 6,368 projection images with 12 time points was 29.17, 21.39, 18.38, 14.40, and 11.39 dB respectively. SNR in dB was calculated using Eq. [11], where S is signal and σ_N is the standard deviation of the noise.

$$\text{SNR} = 20 \log_{10} \left(\frac{\max(S)}{\sigma_N} \right) \quad [11]$$

Phantom Experiments

A bottle phantom was imaged to test the new methods on experimental data. This phantom was a flat-bottomed borosilicate glass cylinder of 9.5mm inner diameter filled with a 1mM solution of trityl as shown in Fig. 2(b). The sample was deoxygenated using a multiple-cycle freeze-pump-thaw technique and flame sealed to ensure internal pO_2 of ~ 0 torr.

Images were standard R_1/R_2 , 10 minute images. Images were reconstructed from 208 projections with equal solid angle spacing onto $64 \times 64 \times 64$, $100 \times 100 \times 100$, or $128 \times 128 \times 128$ voxel image matrices with a physical field of view of $(3\sqrt{2} \text{ cm})^3$. The standard time points for R_1/R_2 imaging discussed in the R_1/R_2 Hybrid Imaging subsection of the Methods section were used (12 time points).

Table 1 Parameters Used for Generation of Simulated Data From the Digital Phantom

Phantom Region	A (a.u.)	R_1 (μs^{-1})	R_2 (μs^{-1})
1	0.06	0.33	0.67
2	0.02	0.20	0.29
3	0.04	0.22	0.33
4	0.08	0.25	0.40
5	0.10	0.29	0.50
6	0.15	0.40	1.00

RESULTS AND DISCUSSION

GPU Accelerated 3D Reconstruction

The speeds of three different 3D image reconstruction algorithms were compared to evaluate the advantage of GPU-based reconstruction. These three

Table 2 Comparison of Time Required for the Back-Projection Process for 3D Image Reconstruction Using Three Different Algorithms

Image Matrix Size (voxels)	GPU Single-Stage (GS)			CPU Multi-Stage (CM)	CPU Single-Stage (CS)
	Compute Time (s)	Transfer Time (s)	Total Reconstruction Time (s)	Reconstruction Time (s)	Reconstruction Time (s)
64 ³	0.011	0.049	0.060	0.203	13.0
100 ³	0.016	0.229	0.245	0.624	44.0
128 ³	0.022	0.378	0.400	1.108	83.0

The process on the GPU can be split into 2 portions: compute time and transfer time.

algorithms were: GPU-based single-stage 3D FBP (GS), CPU-based single-stage 3D FBP (CS), and CPU-based multi-stage 3D FBP (CM). A bottle phantom was imaged and reconstructed onto 64³, 100³, and 128³ image matrices, as described previously.

It can be seen in Tables 2 and 3 that the GS method enables faster image reconstruction compared to the other two methods. For a 128³ reconstructed object, the CS method required 83 seconds to reconstruct one 3D image. The CM method required 1.108 seconds for reconstruction. However, the GS method was the fastest, requiring only 0.4 seconds for reconstruction, including both computing time and time to transfer data back to the CPU. If only computing time is considered, the GS reconstruction method requires 0.022 seconds, which is 3,772 times faster than the CS method and 50.4 times faster than the CM method.

For a more comprehensive comparison, the three average reconstruction speeds for different reconstructed image sizes were compared. Table 3 shows that, on average, the GPU-based reconstruction provides acceleration of the 3D reconstruction by a factor of 2.88 compared to the CM method and 201 compared to the CS method.

While the CM method is significantly faster than the CS method and nearly as fast as the GS method, it is restrictively specialized only for the case of equal linear angle spaced projections. Other projection distributions can be reconstructed using the CM method by interpolating onto equal linear angle spaced projections. This works well for distributions such as equal solid angle spaced projections, but for arbitrary projection distributions this may create appreciable artifacts in the reconstructed image. To accurately reconstruct images from arbitrary projection distributions, single-stage reconstruction is necessary. However, the CS method is relatively slow. The GS method presented in this paper is preferable to both methods. It offers accurate reconstruction from arbitrary projection sets using single-stage reconstruction and takes advantage of the highly parallelizable nature of the FBP algorithm to provide extremely fast image reconstruction using the GPU.

Lookup Table Accelerated Fitting

A comparison of the accuracy of fits for R_1 , R_2 , and spin probe concentration using different fitting methods is shown in Fig. 3. The error for the tested

Table 3 Acceleration Factors for GPU Based Single-Stage Reconstruction Versus Previously Used CPU Based Reconstruction Methods

Image Matrix Size (voxels)	GS vs. CM	GS vs. CM	GS vs. CS	GS vs. CS
	<i>Not considering transfer time</i>	<i>Considering transfer time</i>	<i>Not considering transfer time</i>	<i>Considering transfer time</i>
64 ³	18.4	3.40	1,181	216
100 ³	39.0	2.55	2,750	180
128 ³	50.4	2.70	3,772	207
<i>Average Result</i>	35.9	2.88	2,568	201




Figure 3 Plots of mean error in the fitted parameters versus SNR of simulated projection data, using different fitting methods. (a) Mean R_1 error, (b) mean R_2 error, and (c) mean spin probe concentration error. Note curves for LUT 0.01 and 0.001 are nearly equivalent.

fitting methods (LUT method with LUT parameter resolution of 0.1, 0.01, and 0.001 μs^{-1} , as well as the FMS method) asymptotically decreases towards the error level of the noiseless data as SNR increases. Figure 3 shows that the fit accuracy using the LUT method depends on parameter resolution for the LUT. Over the range of SNR values used here, LUTs with R_1 and R_2 resolution of 0.01 and 0.001 μs^{-1} result in essentially equivalent accuracy independent of SNR (see red and black curves, respectively in Fig. 3, which essentially lie on top of one another), while resolution of 0.1 μs^{-1} gives

lower accuracy. This implies that, for this application, 0.01 μs^{-1} resolution may be sufficient.

Figure 4 provides a visual comparison of images reconstructed from the mathematical phantom described in the Methods section, with noise equivalent to a 10 minute image (SNR = 21.39 dB), using different fitting methods. The ideal data shown are the mathematical phantom discretized onto the same $64 \times 64 \times 64$ matrix as the reconstructed images. It can be seen that, when a LUT with sufficient parameters resolution is used, the LUT method provides similar results to the FMS

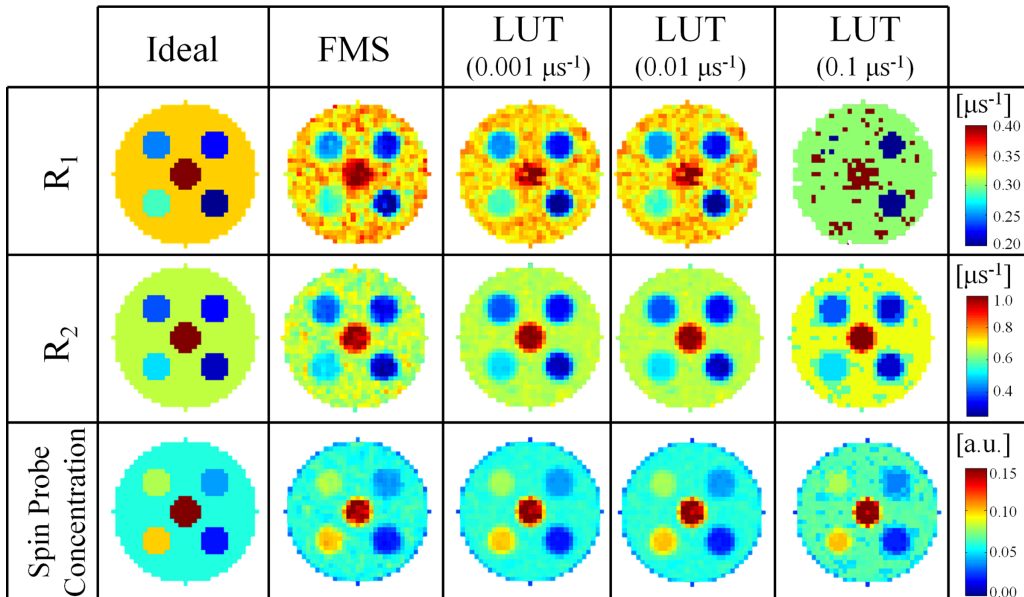


Figure 4 Visual comparison of results using different fitting methods. Simulated images using the mathematical phantom are reconstructed using the GS method and either the FMS method or the LUT method with different resolution LUTs. Noise equivalent to a 10 minutes image (i.e., SNR = 21.39dB) is added and 12 time points are used for the fit. The ideal images are the discretized mathematical phantom.

Table 4 Comparison of Fit Times for the Different Fitting Methods

	FMS	LUT ($0.001 \mu\text{s}^{-1}$)	LUT ($0.01 \mu\text{s}^{-1}$)
Fit Time (s)	282	0.264	0.035
Acceleration Factor	–	1,068	8,057

Acceleration factor is compared to standard FMS method.

method. However, LUTs with decreased parameter resolution eventually fail to appropriately fit the data, which is most apparent in the R_1 image fit using the LUT method with $0.1 \mu\text{s}^{-1}$ resolution.

For simulated 10 minute images with 12 time points, average R_1 error using the LUT method with $0.01 \mu\text{s}^{-1}$ resolution is 10.5% compared to 3.62% for the FMS method, whereas the R_2 error for the two methods is nearly equivalent (5.4% for $0.01 \mu\text{s}^{-1}$ resolution LUT method; 5.3% for FMS method). The spin probe concentration error is 12.82% for the $0.01 \mu\text{s}^{-1}$ resolution LUT method and 9.04% for the FMS method. Figure 3 shows that, as SNR increases and fit errors for the various methods approach that of the noiseless cases, the accuracy for both the $0.01 \mu\text{s}^{-1}$ resolution LUT method and the FMS method approach similar values (1.39%, 1.61%, 5.8% vs. 0.63%, 1.13%, 5.14% for the R_1 , R_2 , and spin probe concentration error using the $0.01 \mu\text{s}^{-1}$ resolution LUT method vs. using the FMS method respectively).

While accuracy may decrease when using the LUT fitting method, fitting speed increases significantly. Table 4 shows that, when using $0.01 \mu\text{s}^{-1}$ resolution LUT method there is nearly a 4 order of magnitude acceleration compared to the FMS method. If the application required a LUT with higher parameter resolution, the acceleration factor is just over 3 orders of magnitude using the $0.001 \mu\text{s}^{-1}$ resolution LUT method compared to the FMS method. Tradeoffs in accuracy may be acceptable for real-time imaging applications where only a rough idea of the image is required. Following an experiment, higher-quality images with increased accuracy can be produced using other, more robust, yet time intensive methods.

In the future, LUTs incorporating multiple parameters into a single LUT, rather than using the multiple LUT algorithm described in the methods, may provide improved accuracy. For the current application however, such multi-parameter LUTs became unacceptably large.

Real-Time $p\text{O}_2$ Image Reconstruction

To achieve real-time reconstruction, a new, updated image must be reconstructed as each new projection is acquired. Current standard R_1/R_2 EPRI oxygen images acquire full sets of 208 projections for 12 different time points in 10 minutes. Each projection is acquired at all of the time points before the next projection is acquired. Therefore, for a standard 10 minute, 208 projection image, each projection is acquired in 2.88 seconds. Accordingly, for real-time image reconstruction, reconstruction speeds of at least 2.88 seconds/image are necessary. Using standard reconstruction methods this is very difficult. Accelerations of both 3D image reconstruction and fitting of voxel-by-voxel signal relaxation are necessary. Table 5 summarizes the total $p\text{O}_2$ image reconstruction times when using different combinations of the 3D FBP implementations and fitting methods described above.

Using the CM method for 3D reconstruction and FMS method for fitting, requires 0.203 seconds per 3D reconstruction at each time point for standard $64 \times 64 \times 64$ R_1/R_2 EPRI $p\text{O}_2$ images, plus 282 seconds for fitting. In this case, 3D reconstruction is relatively fast while fitting limits the overall reconstruction speed. Using the LUT fitting method and CM reconstruction method, $p\text{O}_2$ image reconstruction time is faster than projection acquisition speed. However, multi-stage reconstruction (e.g., the CM method) cannot handle arbitrary projection sets and therefore is not a viable candidate for real-time

Table 5 Summary of Total $p\text{O}_2$ Image Reconstruction Times When Using Different Combinations of 3D FBP Implementations and Fitting Methods

Reconstruction Method (3D FBP + Fitting)	Minimum Reconstruction Time (s)	Maximum Reconstruction Time (s)
CM + FMS	–	284
CM + LUT 0.001	–	2.70
CM + LUT 0.01	–	2.47
CS + FMS	283	503
CS + LUT 0.001	1.33	221
CS + LUT 0.01	1.10	221
GS + LUT 0.001	0.85	0.98
GS + LUT 0.01	0.62	0.76

Minimum reconstruction time is for reconstruction from a single projection. Maximum reconstruction time is for reconstruction from the full 208 projections. The CM method requires reconstruction from a full projection set and therefore only one reconstruction time is reported. Numbers shown next to LUT denote parameter resolution used in μs^{-1} .

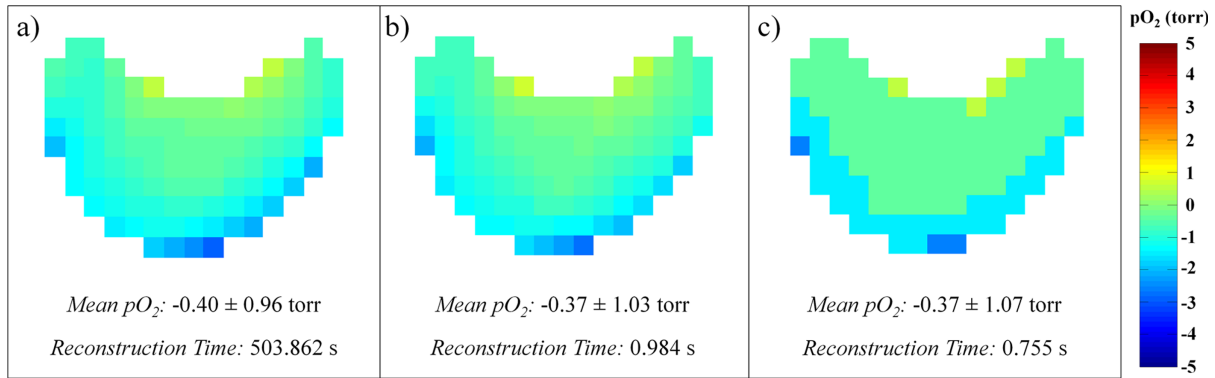


Figure 5 Axial slice through EPRI pO_2 image of bottle phantom showing meniscus in trityl solution. (a) Image reconstructed using previous methods: CS+FMS. (b) Image reconstructed using accelerated methods: GS + LUT $0.001 \mu s^{-1}$ resolution. (c) Image reconstructed using accelerated methods: GS + LUT $0.01 \mu s^{-1}$ resolution.

reconstruction from incomplete sets of projections during acquisition.

Single-stage reconstruction, however, can be used for any arbitrary set of projections and therefore can be used to reconstruct images from incomplete projection sets, but the current implementation on the CPU is restrictively slow. 3D image reconstruction speed is proportional to the number of projections in the image. Total pO_2 image reconstruction time for single-stage methods can be calculated as the time to back-project each projection multiplied by both the total number of projections used and the number of time points used (e.g., 12 for standard R_1/R_2 imaging) plus the time to fit the spectral dimension. When using a GPU, time to transfer 3D images between GPU and CPU must be added as well. When using the CS method and FMS method, both 3D reconstruction and fitting can be time consuming and real-time pO_2 image reconstruction is not possible. If the CS method and LUT method are used, for smaller numbers of projections, reconstruction is fast enough to keep up with data acquisition; however as imaging progresses, real-time pO_2 image reconstruction cannot be achieved.

The methods presented in this paper (GPU-based single-stage 3D reconstruction and LUT fitting), when used in concert, provide the versatility of single-stage reconstruction and the speed necessary for real-time pO_2 image reconstruction. Using the GS method and $0.001 \mu s^{-1}$ resolution LUT method results in reconstruction speed acceleration by up to a factor of 511 compared to current methods. Using the GS method and $0.01 \mu s^{-1}$ resolution LUT method results in reconstruction speed acceleration by up to over a factor of 650 compared to current methods.

Figure 5 shows that, if a LUT with appropriate resolution is used, the accelerated methods for EPRI pO_2 image reconstruction presented provide

essentially the same result as previously used methods. Only very subtle differences can be visually discerned between the image reconstructed using the CS + FMS methods and that reconstructed using the GS+LUT with $0.001 \mu s^{-1}$ resolution (average voxel absolute pO_2 difference was 0.08 torr). Using a $0.01 \mu s^{-1}$ resolution LUT still provides a reasonable image, however the effect of decreased LUT resolution can clearly be observed as a decrease in pO_2 resolution of the image (average voxel absolute pO_2 difference was 0.28 torr). With these accelerated reconstruction speeds, the new methods enable real-time reconstruction of pO_2 images for imaging times as short as 3.4 minutes when using $0.001 \mu s^{-1}$ resolution or 2.6 minutes when using $0.01 \mu s^{-1}$ resolution LUTs. As imaging times decrease, reconstruction speed must increase accordingly to enable real-time reconstruction. More powerful, higher-end GPU cards exist that could increase reconstruction speeds. However, GPU-based single-stage 3D reconstruction and LUT fitting together provide EPRI pO_2 image reconstruction accelerations sufficient to enable real-time reconstruction for current imaging applications.

CONCLUSIONS

Methods providing EPRI pO_2 image reconstruction acceleration are presented. The acceleration is sufficient to enable real-time pO_2 image reconstruction. It should be emphasized that any sacrifice in reconstruction quality that is associated with the techniques discussed in this work can be recovered post acquisition with more standard reconstruction techniques, for example, avoiding the loss of pO_2 resolution with the LUT technique.

Real-time reconstruction is advantageous for different reasons. For example, for *in vivo* imaging applications, particularly in the patient setting, procedures are time sensitive. Real-time image reconstruction can help operators recognize and resolve issues with an EPRI scan early and limit wasted time if the operator must wait for imaging to be completed to discover these issues. For example, if an image becomes corrupted during imaging, (due to improper setup of the subject and/or imaging equipment, subject motion, imaging equipment failure etc.) real-time reconstruction allows detection of such issues prior to scan completion. Real-time reconstruction can also identify a ROI earlier during imaging.

Projection acquisition schemes have been developed that provide approximately uniform sampling throughout the imaging process. Such acquisition schemes enable reconstruction of useful intermediate images from incomplete projection data, which closely approximate the final image reconstructed from complete projection data (35,36). For EPRI data acquired using these robust techniques, the real-time reconstruction presented in this work allows for the visualization of the buildup of an oxygen image as data are acquired.

Studies suggest that both chronic and transient states of hypoxia may exist in tumors and that they may have differing physiological and clinical implications (11). This dynamic nature of the *in vivo* oxygen environment has led to increased interest in time resolved measurements and visualization of pO_2 in order to help clarify the differences between chronic and transient hypoxia. To this end, efforts are being made to increase EPRI temporal resolution (37–41).

As EPRI image acquisition times approach real-time acquisition, the real-time image reconstruction discussed herein will eventually enable real-time visualization of temporal pO_2 fluctuations. This could allow for real-time differentiation of chronically hypoxic regions and transiently hypoxic regions. One potential future approach is a sliding window or view-sharing technique, which could provide quasi-real-time pO_2 imaging based on rapid reconstruction of images from temporally overlapping projection subsets throughout acquisition. Such methods have been used for real-time CT (25) and MRI (34,42) applications. In addition to visualizing transient hypoxia, future real-time EPRI pO_2 imaging could enable image guided treatment adaptation based on spatiotemporally heterogeneous tumor pO_2 , EPRI-guided biopsies, and real-time subject positioning and motion monitoring.

The methods presented in this work provide EPRI pO_2 image reconstruction acceleration factors of over 650. This enables real-time reconstruction of noninvasively acquired EPRI images of *in vivo* oxygen concentrations for the first time.

ACKNOWLEDGMENTS

This work was supported by grants from the NIH (P41 EB002034 and R01 CA98575). The authors are very grateful for the insight provided by Dr. Xiaochuan Pan and his research group in helping to develop and refine this work.

REFERENCES

1. Overgaard J. 2007. Hypoxic radiosensitization: adored and ignored. *J Clin Oncol* 25:4066–4074.
2. Schwarz G. 1909. Über desensibilisierung gegen Rontgen- und radiumstrahlen. *Munchner Medizinische Wochenschrift* 56:1217–1218.
3. Hall EJ. 2000. *Radiobiology for the Radiologist*, 5th Ed. Philadelphia: Lippincott Williams & Wilkins.
4. Hockel M, Schlenger K, Aral B, Mitze M, Schaffer U, Vaupel P. 1996. Association between tumor hypoxia and malignant progression in advanced cancer of the uterine cervix. *Cancer Res* 56:4509–4515.
5. Shannon AM, Bouchier-Hayes DJ, Condron CM, Toomey D. 2003. Tumour hypoxia, chemotherapeutic resistance and hypoxia-related therapies. *Cancer Treat Rev* 29:297–307.
6. Carmeliet P, Dor Y, Herbert JM, Fukumura D, Brusselmans K, Dewerchin M, et al. 1998. Role of HIF-1 α in hypoxia-mediated apoptosis, cell proliferation and tumour angiogenesis. *Nature* 394:485–490.
7. Rofstad EK. 2000. Microenvironment-induced cancer metastasis. *Int J Radiat Biol* 76:589–605.
8. Dewhirst MW, Klitzman B, Braun RD, Brizel DM, Haroon ZA, Secomb TW. 2000. Review of methods used to study oxygen transport at the microcirculatory level. *Int J Cancer* 90:237–255.
9. Zhao DW, Jiang L, Mason RP. 2004. Measuring changes in tumor oxygenation. *Methods Enzymol* 386:378–418.
10. Tatum JL. 2006. Hypoxia: importance in tumor biology, noninvasive measurement by imaging, and value of its measurement in the management of cancer therapy. *Int J Rad Biol* 82:699–757.
11. Bayer C, Vaupel P. 2012. Acute versus chronic hypoxia in tumors: controversial data concerning time frames and biological consequences. *Strahlenther Onkol* 188:616–627.

12. Elas M, Ahn KH, Parasca A, Barth ED, Lee D, Haney C, et al. 2006. Electron paramagnetic resonance oxygen images correlate spatially and quantitatively with oxylyte oxygen measurements. *Clin Cancer Res* 12:4209–4217.
13. Epel B, Sundramoorthy SV, Mailer C, Halpern HJ. 2008. A versatile high speed 250-MHz pulse imager for biomedical applications. *Concept Magn Reson B Magn Reson Eng* 33B:163–176.
14. Epel B, Sundramoorthy SV, Barth ED, Mailer C, Halpern HJ. 2011. Comparison of 250 MHz electron spin echo and continuous wave oxygen EPR imaging methods for in vivo applications. *Med Phys* 38:2045–2052.
15. Epel B, Haney CR, Hleihel D, Wardrip C, Barth ED, Halpern HJ. 2010. Electron paramagnetic resonance oxygen imaging of a rabbit tumor using localized spin probe delivery. *Med Phys* 37:2553–2559.
16. Elas M, Williams BB, Parasca A, Mailer C, Pelizzari CA, Lewis MA, et al. 2003. Quantitative tumor oxymetric images from 4D electron paramagnetic resonance imaging (EPRI): methodology and comparison with blood oxygen level-dependent (BOLD) MRI. *Magn Reson Med* 49:682–691.
17. Epel B, Bowman MK, Mailer C, Halpern HJ. 2014. Absolute oxygen R1e imaging in vivo with pulse electron paramagnetic resonance. *Magn Reson Med* 72:362–368.
18. Lai CM, Lauterbur PC. 1980. A gradient control device for complete 3-dimensional nuclear magnetic resonance zeugmatographic imaging. *J Phys E Sci Instrum* 13:747–750.
19. Hyslop WB, Woods RK, Lauterbur PC. 1995. Four-dimensional spectral-spatial imaging using projection reconstruction. *IEEE Trans Med Imaging* 14:374–383.
20. Cabral B, Cam N, Foran J. 1994. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. *Proc Symp Volume Visualization* 91–98.
21. Pratz G, Xing L. 2011. GPU computing in medical physics: a review. *Med Phys* 38:2685–2697.
22. Ma D, Gulani V, Seiberlich N, Liu K, Sunshine JL, Duerk JL, et al. 2013. Magnetic resonance fingerprinting. *Nature* 495:187–192.
23. Eggers H, Proksa R. 1999. Multiprocessor system for real-time convolution interpolation reconstruction. *Proc Int Soc Magn Reson Med* 95.
24. Sorensen TS, Atkinson D, Schaeffter T, Hansen MS. 2009. Real-time reconstruction of sensitivity encoded radial magnetic resonance imaging using a graphics processing unit. *IEEE Trans Med Imaging* 28:1974–1985.
25. Xu F, Mueller K. 2007. Real-time 3D computed tomographic reconstruction using commodity graphics hardware. *Phys Med Biol* 52:3405–3419.
26. Seifi P. Scalability of time-domain electron paramagnetic resonance in vivo imaging. Ph.D. Thesis, University of Chicago, December 2010.
27. Mailer C, Sundramoorthy SV, Pelizzari CA, Halpern HJ. 2006. Spin echo spectroscopic electron paramagnetic resonance imaging. *Magn Reson Med* 55:904–912.
28. Ardenkjaer-Larsen JH, Laursen I, Leunbach I, Ehnholm G, Wistrand LG, Petersson JS, et al. 1998. EPR and DNP properties of certain novel single electron contrast agents intended for oximetric imaging. *J Magn Reson* 133:1–12.
29. Deans SR. 1983. *The Radon Transformation and Some of Its Applications*. New York, NY: John Wiley & Sons.
30. Lagarias JC, Reeds JA, Wright MH, Wright PE. 1998. Convergence properties of the Nelder-mead simplex method in low dimensions. *Siam J Optimiz* 9: 112–147.
31. Doneva M, Senegas J, Bornert P, Eggers H, Mertins A. 2009. Accelerated MR parameter mapping using compressed sensing with model-based sparsifying transform. *Proc Int Soc Magn Reson* 2812.
32. Aharon M, Elad M, Bruckstein A. 2006. K-SVD: an algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans Signal Process* 54: 4311–4322.
33. Doneva M, Bornert P, Eggers H, Stehning C, Senegas J, Mertins A. 2010. Compressed sensing reconstruction for magnetic resonance parameter mapping. *Magn Reson Med* 64:1114–1120.
34. Ehses P, Seiberlich N, Ma D, Breuer FA, Jakob PM, Griswold MA, et al. 2013. IR TrueFISP with a golden-ratio-based radial readout: fast quantification of t_1 , t_2 , and proton density. *Magn Reson Med* 69:71–81.
35. Chan RW, Ramsay EA, Cunningham CH, Plewes DB. 2009. Temporal stability of adaptive 3D radial MRI using multidimensional golden means. *Magn Reson Med* 61:354–363.
36. Redler G, Epel B, Halpern HJ. Maximally spaced projection sequencing in electron paramagnetic resonance imaging. *Concepts Magn Reson Part B*. DOI: 10.1002/cmr.b.21280.
37. Redler G, Epel B, Halpern HJ. 2014. EPR image based oxygen movies for transient hypoxia. *Adv Exp Med Biol* 812:127–133.
38. Redler G, Epel B, Halpern HJ. 2014. Principal component analysis enhances snr for dynamic electron paramagnetic resonance oxygen imaging of cycling hypoxia in vivo. *Magn Reson Med* 71:440–450.
39. Yasui H, Matsumoto S, Devasahayam N, Munasinghe JP, Choudhuri R, Saito K, et al. 2010. Low-field magnetic resonance imaging to visualize chronic and cycling hypoxia in tumor-bearing mice. *Cancer Res* 70:6427–6436.
40. Matsumoto S, Yasui H, Mitchell JB, Krishna MC. 2010. Imaging cycling tumor hypoxia. *Cancer Res* 70:10019–10023.
41. Redler G. Enhance dynamic electron paramagnetic resonance imaging of in vivo physiology. Ph.D. Thesis, University of Chicago, June 2014.
42. Uecker M, Zhang S, Frahm J. 2010. Nonlinear inverse reconstruction for real-time MRI of the human heart using undersampled radial FLASH. *Magn Reson Med* 63:1456–1462.