Parallel 3D Temperature Image Reconstruction Using Multi-Color Magnetic Particle Imaging (MPI)

Klaus N. Quelhas,[1, 2, a)] Mark-Alexander Henn,[1, 3, b)] Ricardo Farias,[2, c)]
Weston L. Tew,[1, d)] and Solomon I. Woods[1, e)]

[1]*National Institute of Standards and Technology (NIST), 100 Bureau Drive, Gaithersburg, MD, USA.*
[2]*Systems Engineering and Computer Science Program (PESC), Federal University of Rio de Janeiro, Rio de Janeiro, RJ, Brazil.*
[3]*College of Computer, Mathematical & Natural Sciences, University of Maryland, 7998 Regents Dr, College Park, MD, USA.*

[a)]Corresponding author: klaus.quelhas@nist.gov
[b)]mark.henn@nist.gov
[c)]rfarias@cos.ufrj.br
[d)]weston.tew@nist.gov
[e)]solomon.woods@nist.gov

Abstract. Magnetic Particle Imaging (MPI) is a novel technique developed for remotely detecting magnetic nanoparticle (MNP) tracers, with great potential for biomedical imaging (as an alternative to traditional methods like MRI or CT), cell tracking, targeted drug delivery and magnetic hyperthermia. More recently, MPI has been studied as a potential method of non-contact temperature measurement. This work presents a simulation study of the *multi-color* MPI method tailored for 3D temperature imaging, discusses the feasibility of the method for 3D temperature measurements, and shows a parallel implementation of the *multi-color* T-MPI reconstruction algorithm in *graphics processing unit* (GPU). While the use of the parallel algorithm resulted in executions about 40x faster when compared to the serial implementation, the method exhibited serious limitations in accurately resolving particle temperatures between the calibration temperatures by interpolation.

INTRODUCTION

Since its inception [1], Magnetic Particle Imaging (MPI) has been mostly studied as a biomedical imaging method, such as magnetic resonance imaging (MRI), computed tomography (CT) or positron emission tomography (PET). In comparison to these traditional methods, MPI directly images tracer material without ionizing radiation, and achieves high contrast, fast acquisition times and high spatial resolutions [2]. MPI also finds use in other biomedical applications like targeted drug delivery, cancer detection and treatment with magnetic hyperthermia [3].

MPI works by measuring the nonlinear magnetization response of magnetic nanoparticle (MNP) tracers when exposed to a combination of static and oscillating magnetic fields generated by magnets and electromagnets, that allows scanning the desired field-of-view (FOV) at high frequencies - usually between 10 kHz and 25 kHz. The MNP magnetization of ideal particles follows the Langevin superparamagnetic function, which describes the balance between the magnetic Zeeman energy from the applied field that attempts to align the particles and the thermal energy kT that disturbs this alignment. Therefore, it did not take long until an application of MNPs as temperature sensors was proposed in the literature. In Weaver et al. [4], it was demonstrated for the first time how to measure temperature with MNPs and magnetic particle spectroscopy (MPS) techniques by taking the ratios of the harmonics found in the magnetization response spectrum. Later, it was shown how to obtain temperature maps using MPI measurement and reconstruction techniques, by employing the so-called *multi-color* method [5].

These and other examples [6, 7, 8, 9, 10, 11] were successful on demonstrating the potential of MPI temperature imaging (or T-MPI), but the reported uncertainties were still too high, mostly of the order of tenths of kelvin and above. In addition, the accuracy of MNP thermometry, especially in the imaging context, is still questionable, due to the lack of standards in the previous studies. In this context, NIST started the *Thermal MagIC* (Thermal Magnetic Imaging and Control) project, which aims to establish the technologies for remote 3D temperature measurement and control with microscale spatial resolution and an uncertainty of 25 mK or better, traceable to the ITS-90 [12]. This

will be achieved by means of a) new MNPs with improved signal response and thermal sensitivity [13, 14], and b) improved instrumentation for high-accuracy magnetometry and thermometry [15]. Recently, it has been demonstrated how to obtain 3D maps of the MNP harmonic signals, and how these harmonic maps relate to temperature [16].

Image reconstruction plays a critical role in MPI, since it allows the effective spatial visualization of the raw measurement data, and the adopted reconstruction method depends strongly on how the data are acquired. State-of-the-art reconstructions require numerically solving massive systems of equations, and the execution times are often a limitation factor for real-time applications (e.g., temperature measurement and control, the main goal of NIST's *Thermal MagIC* project), since the solution of such systems can take from minutes to hours. The optimization of the reconstruction methods is a major focus of research within the MPI community [17, 18], and recently it has been show that it is possible to accelerate the execution of 2D MPI reconstructions by implementing the algorithms in parallel in *graphics processing unit* (GPU) [19].

This work presents a simulation study of the *multi-color* MPI method tailored for 3D temperature imaging, and discusses the feasibility of the method for 3D temperature measurements. The next sections describe the principles of MPI and *multi-color* T-MPI employed for generating the simulated data and the image reconstruction algorithms employed for obtaining the reconstructed 3D temperature maps. Additionally, based on the promising results from GPU-accelerated parallel MPI image reconstructions, this work presents a parallel implementation of the multi-color T-MPI method, showing that *speedups* of the order of 40x can be achieved.

PRINCIPLES OF MPI AND *MULTI-COLOR* T-MPI

MPI exploits the nonlinear magnetization response of MNPs. These particles are also often referred to as *superparamagnetic iron oxides*, which implies that they are typically composed by iron oxides (such as magnetite) and they exhibit superparamagnetic behavior. Superparamagnetism is observed in very small MNPs, usually with diameters smaller than 25 nm, in which the thermal energy kT is strong enough to surpass the magnetic anisotropy energy barrier and promote the reversal of their magnetic moment directions. This magnetization reversal is known as *Néel rotation* [2]. Since the orientations of the magnetic moments in a sample are constantly and randomly switching, the net magnetization of the sample is zero in equilibrium. Under a magnetic field, however, the MNPs tend to align in the same direction of the field, and the sample exhibits an observable magnetization. When the field is turned off, the magnetic moments scatter again due to the thermal agitation and the magnetization returns to zero. The magnetization of ideal superparamagnetic particles can be described by the Langevin function:

$$\mathbf{M}(\mathbf{H}) = \rho \cdot \mathscr{L}(\mathbf{H}) = \rho \cdot \left\{ \coth\left(\frac{\mu_0 m||\mathbf{H}||}{kT}\right) - \left(\frac{\mu_0 m||\mathbf{H}||}{kT}\right)^{-1} \right\} \cdot \frac{\mathbf{H}}{||\mathbf{H}||} \tag{1}$$

where ρ is the particle concentration (m^{-3}), μ_0 is the vacuum permeability ($\mathrm{H/m}$), m is the particle magnetic moment (Am^2), $\mathbf{H}$ is the applied magnetic field ($\mathrm{A/m}$), k is the Boltzmann constant ($\mathrm{J/K}$) and T is the temperature (K). The alignment of the moments with the applied field is not instantaneous, and the characteristic time of this alignment is known as *Néel relaxation time*. Larger MNPs in solutions can also exhibit superparamagnetic behavior, but in this case the particles physically rotate within the fluid to align with the applied field, being this form of magnetization reversal called *Brownian relaxation*. Both Néel and Brownian mechanisms take place simultaneously and contribute to the alignment of the MNPs with an applied field.

Static magnetic fields produced by a pair of magnets and/or electromagnets oriented in opposite directions result in a gradient field that vanishes at the center, being the point in which the fields cancel each other known as *field-free point* (FFP). In this configuration, all particles within this gradient field are saturated and will not respond to a superimposed oscillating field, except the ones that lie close to the FFP, and this allows selectively measuring the response of the MNPs over space by moving this FFP. Low and high frequency magnetic fields generated by electromagnets can steer the FFP in space, causing the particles to be rapidly magnetized and demagnetized along the FFP's trajectory. These fast changes in magnetization can be detected by inductive coils oriented with the x, y and z axes, and the measured signal $s(t)$ can be described by:

$$s(t) = m \cdot \mathbf{R} \cdot \int_{\mathbb{R}^3} \rho(x) \cdot \frac{d}{dt} \mathscr{L}\{\xi \cdot ||\mathbf{H}(x,t)||\} \cdot \frac{\mathbf{H}(x,t)}{||\mathbf{H}(x,t)||} dx \tag{2}$$

where $\mathbf{R}$ is the receive coil sensitivity ($\mathrm{T/A}$) and $\xi = \mu_0 m/kT$. When we express the field $\mathbf{H}$ as a function of the FFP trajectory $r(t)$:

$$\mathbf{H}(x,t) = \mathbf{G} \cdot (x - r(t)) \tag{3}$$

where $\mathbf{G}$ is the static gradient field (T/m), Equation 2 describes the signal generation as a convolution of the time derivative of the Langevin function over the spatial distribution of MNPs. In this model, the derivative of the Langevin function is the *point spread function* (PSF) that represents the interaction between the gradient field and the particles that surround the FFP [20]. In practice, the gradient field is such that the particles far from the FFP are indeed saturated, while in closer regions the field is not strong enough to completely saturate the MNPs, so they will show a weak response to the excitation field. Therefore, this signal contribution from surrounding particles causes the images obtained directly form the measured signals to be blurred, and image reconstructions techniques are employed to remove this blurring.

In the so-called *system function* method, the sample is excited by specific field sequences that move the FFP following a Lissajous 2D or 3D scanning trajectory over the field of view (FOV), and with enough measurements and a calibration matrix (containing the characteristic response of the MNPs to that particular MPI system), it is possible to resolve the unknown MNP concentration grid by solving a linear system of equations:

$$K \cdot c = s \tag{4}$$

where K is the (m x n) calibration matrix (or *convolution matrix*), c is the unknown particle concentration vector, and s is the vector containing the signal measured by the pickup coil(s). Each of the m rows of the calibration matrix contains the distribution of the PSF over the n *voxels* of the 3D grid following the FFP movement as it is driven over the FOV, so the PSF is scaled by the FFP speed. Since the FFP is moving over time, the position of the PSF changes accordingly, resulting in linearly independent rows in the convolution matrix. With enough measurements, the system is overdetermined and can be solved numerically, returning an estimate of the distributions of particles in space.

The calibration step that determines the convolution matrix consists in measuring the response of a small delta sample positioned in each of the n positions correspondent to the image *voxels*, being the measured signal stored as a column of K. This process can be tedious, and model-based methods can save time by computing the calibration matrix without the need of performing measurements. To do so, the elements that compose a MPI measurement are modelled and the calibration signals are simulated ans stored in the matrix [21, 22]. Regardless of being measurement-based or model-based, the *system function* method is the most widely adopted for imaging purposes, since it intrinsically removes the PSF blurring and other systematic effects by solving equation 4, resulting in better reconstructions, at the expense of numerically solving a large system of equations.

A similar approach can be used to estimate temperature with MPI techniques. In [23], it was demonstrated how the *multi-color* method can be employed to distinguish different particles within the same sample. The method relies on the fact that each kind of particle has a characteristic magnetization response, and therefore will produce unique calibration matrices. The same rationale could be applied for temperature measurement, since similar particles will have different responses at different temperatures. The approach proposed in [5] uses an extended system of equations:

$$\begin{vmatrix} K_{T_1} & K_{T_2} \end{vmatrix} \cdot \begin{vmatrix} c_{T_1} \\ c_{T_2} \end{vmatrix} = s \tag{5}$$

where K_{T_1} and K_{T_2} are the calibration matrices obtained at the calibration temperatures T_1 and T_2, and c_{T_1} and c_{T_2} are the unknown concentration vectors at T_1 and T_2. The final particle distribution is given by the sum of c_{T_1} and c_{T_2}, while the temperature would be given by linear interpolation of the concentrations in each channel. In principle, any number of calibration temperatures could be adopted, but the cost of it is solving increasingly massive systems of equations. These intensive calculations can be a serious limitation considering that a single calibration matrix can occupy several gigabytes of data in the computer memory, depending on the desired spatial resolution and number of measurements. In summary, each channel contains the number of particles at each temperature, and the resulting concentration is given by the sum of the concentrations in all channels, while the temperature would be given by:

$$T = \frac{c_{T_1} \cdot T_1 + c_{T_2} \cdot T_2 + c_{T_3} \cdot T_3}{c_{T_1} + c_{T_2} + c_{T_3}} \tag{6}$$

for three calibration temperatures, for example. Even though equation 1 tells one that the MNP magnetization is temperature dependent, the effect of temperature changes on the PSF is limited to slightly scaling and stretching it.

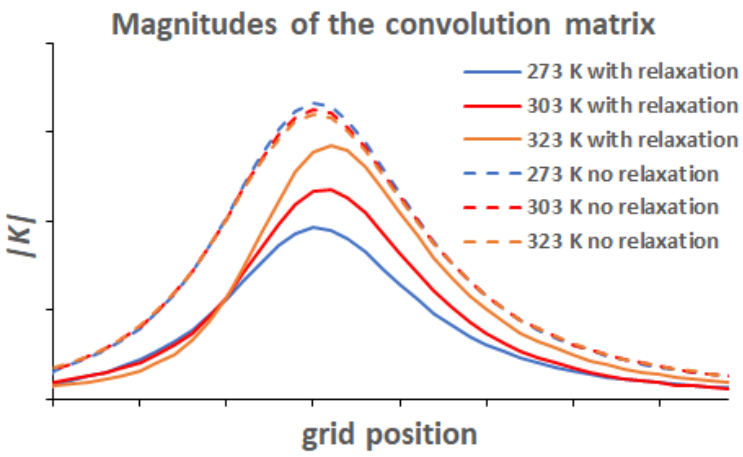

FIGURE 1. Simulated data plot showing the effects of temperature and relaxation over the convolution matrix that contains the PSF, at different temperatures. Solid lines show part of a single row of the convolution matrix with relaxation, while the dashed lines show the same region with no relaxation.

In order to be able to reach a solution for the extended system of equations, the corresponding rows of the calibration matrices should be as close to orthogonal as possible, but the temperature dependence of the Langevin function is not enough for ensuring this orthogonality. On the other hand, since the alignment of the MNPs with the applied field is not instantaneous, the relaxation process governed by Brownian and Néel rotations is strongly temperature-dependent and results in a phase shift in the signal, represented by a shift of the PSF in space. Figure 1 shows how temperature has little effect on regimes without relaxation, while temperature-dependent relaxation effectively shifts the PSF, allowing for better convergence of the numerical solution. Relaxation, and thus its effect on the calibration matrix, depends strongly on the particle and on the media in which it is immersed.

Figure 2 illustrates a simulation example of a 3D reconstruction for 20 nm magnetite particles calibrated at three temperatures: the magnetization signals of a hypothetical phantom are measured by pickup coils in the x, y and z directions while being excited by a 3D Lissajous trajectory of the FFP, and these data are used to reconstruct the concentrations at the channels T_1, T_2 and T_3, which are combined to form the 3D temperature distribution. It is possible to notice that the concentration reconstruction is still not perfect, showing some remaining PSF blurring that could not be removed by the numerical deconvolution with a limited number of iterations (100, in this case). The accuracy

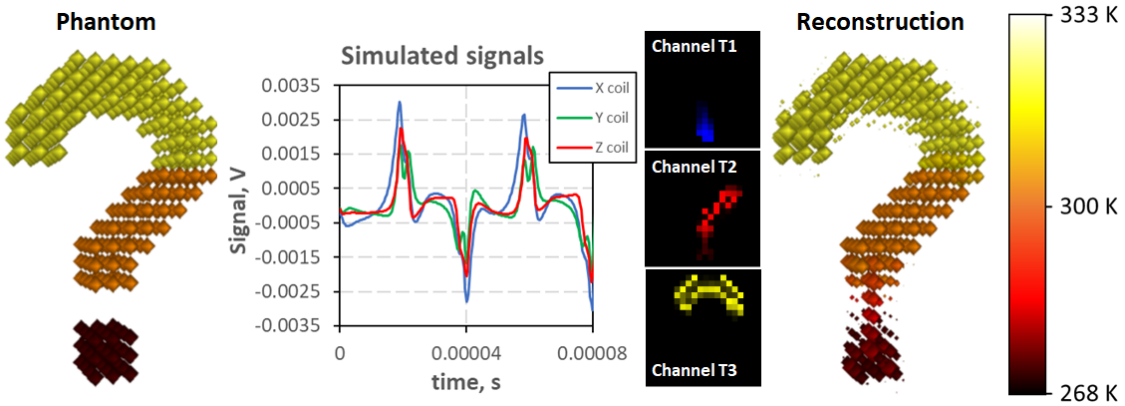

FIGURE 2. Illustration of a simulated *multi-color* T-MPI measurement. The magnetization signals generated by the hypothetical 3D phantom with *voxels* at three distinct temperatures are measured by pickup coils oriented in the x, y and z directions. These signals are used to reconstruct the concentrations at the calibration temperatures T_1, T_2 and T_3, and the concentrations of the three channels are used to compute the reconstructed concentration (represented by the size of the polygons) and temperatures (represented by the colors of the polygons). The image resolution is 20x20x20 *voxels*, and the total FOV volume is (10x10x10) mm^3.

of the reconstruction depends on a number of factors, including the gradient field strength, the scanning frequency, the particle sample being measured, the number of samples relative to the size of the unknown concentration vector, the desired image resolution, the reconstruction algorithm and the number of iterations. These parameters have to be carefully chosen to deliver the best results with the available resources, including the time to solve the system of equations.

RECONSTRUCTION ALGORITHMS

A number of algorithms have been tested for optimizing MPI image reconstructions. Among them, Kaczmarz's algorithm remains the most popular one, due to its simplicity and efficiency, which grants it great convergence rates. Another algorithm widely used is the conjugate gradients normal residuals (CGNR), which also has good convergence rate, but is slower to run than Kaczmarz. However, recently it has been demonstrated that the CGNR algorithm can be optimized to be run in parallel in *graphics processing units* (GPU), resulting in *speed-ups* of 10x for 2D MPI reconstructions, while the parallel version of Kaczmarz running in GPU is in fact slower than the serial version running in the CPU [19]. Therefore, this work studies the serial Kaczmarz's algorithm running in the CPU and the parallel CGNR algorithm running in the GPU.

Kaczmarz's algorithm

Kaczmarz's algorithm works by, for each iteration, projecting the current estimate over the corresponding line of the convolution matrix, taking the residual with relation to the measured signal, and then correcting the estimate by projecting the residual back. The $(i+1)$-th estimate is given by:

$$c^{(i+1)} = c^{(i)} + \frac{s_i - K_i \cdot c^{(i)}}{||K_i||^2} K_i^T \tag{7}$$

where K_i is the (i)-th row of the convolution matrix and s_i is the (i)-th measurement. Once the new estimate is computed, the algorithm moves to the next line. One sweep over all the lines of K is considered to be one Kaczmarz iteration.

Conjugate gradient normal residual - CGNR

The CGNR method is a adaptation of the conjugate gradients (CG) for non-symmetric matrices, which is the case in MPI reconstructions. In this scenario, the CGNR method is applied to solve the normal equation ($K^T K \cdot c = K^T \cdot s$) by iterating over a sequence of base vectors p^i that are mutually conjugate with respect to $(K^T K)$. The new estimate is given by

$$c^{(i+1)} = c^{(i)} + \frac{||K^T \cdot r^{(i)}||^2}{||K \cdot p^{(i)}||^2} p^{(i)} \tag{8}$$

where

$$p^{(i)} = K^T \cdot r^{(i)} + \frac{||K^T \cdot r^{(i)}||^2}{||K^T \cdot r^{(i-1)}||^2} p^{(i-1)} \tag{9}$$

and the residual $r^i = s - K \cdot c^i$. The method is guaranteed to converge after n iterations (where n is the number of columns of K), but usually it converges very quickly, reaching acceptable solutions in much fewer iterations. The dominance of matrix-vector multiplications is what grants CGNR its parallelization potential.

SIMULATIONS

Simulations were performed to investigate the efficiency of the *multi-color* method for temperature measurement with MPI, as well as to estimate the performance gains obtained by running the reconstructions in parallel in GPU. A hypothetical phantom, composed of 20x20x20 = 8,000 *voxels* and the shape of a 3D question mark ("?"), filled with 20 nm magnetite MNPs with 5 $\mathrm{mg(Fe)/ml}$ of concentration at different temperatures, was adopted for the simulations.

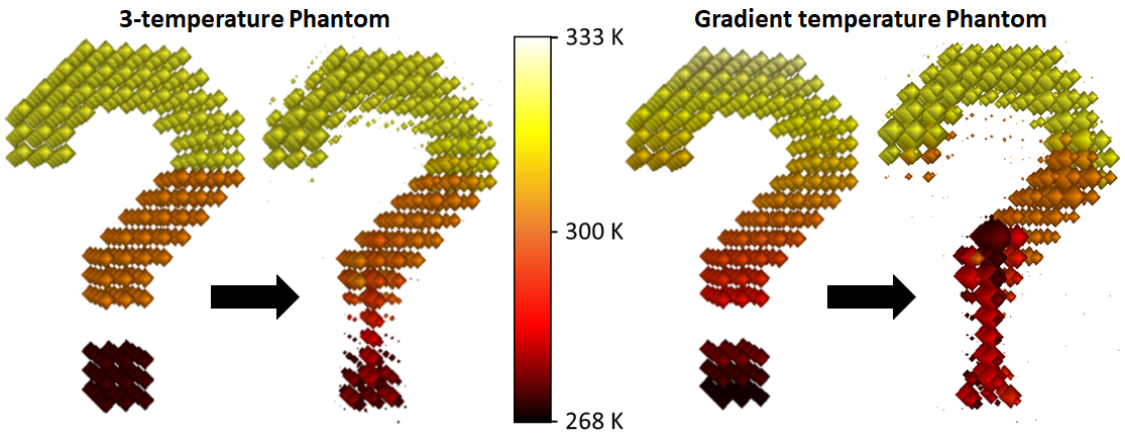

FIGURE 3. Reconstruction results for the phantom with the particles at three temperatures (273 K, 303 K and 323 K, left), and for the phantom with a vertical linear temperature gradient (right). In each case, the ground truth is shown before the arrow and the resulting image reconstruction is shown after the arrow.

The volume of the phantom was (10x10x10) mm^3, so the each cubic *voxel* had 0.5 mm of side length. The simulated scanner had characteristics similar to a typical MPI setup; its gradient field strength was 8 T/m, the FFP Lissajous trajectory was simulated with a base frequency of 250 Hz and with x, y and z multipliers equal to 101, 100 and 99, respectively. The calibration matrices were computed numerically from equation 2 at three different temperatures: 273 K, 303 K and 323 K. The adopted sampling frequency was 2.5 MHz, resulting in a total of 30,000 simulated measurements, corresponding to the x, y and z signals measured during a full scan.

For assessing the reconstruction accuracy, two different phantoms were employed in these simulations, one with the particles at exactly the three calibration temperatures (273 K at the bottom, 303 K in the central region and 323 K at the top), and another with with a linear temperature gradient. For both phantoms, the same calibration matrix, algorithm and number of iterations, were used. Figure 3 shows the two phantoms, as well as their correspondent reconstruction results after 100 iterations.

For assessing the speed of the reconstruction algorithms, a sampling rate of 625 kHz was adopted, and the mean square errors (MSEs) and execution times were taken after the first iteration and after every 10 iterations up to the 100th. For simplicity and fair comparison purposes, the simulations were performed without the presence of noise. The simulations were performed on a computer equipped with an Intel(R) Xeon(R) W-2145 processor, 32 GB of RAM memory, a NVIDIA Quadro RTX 4000 GPU containing 2,304 CUDA cores and 8 GB of dedicated memory. The Kaczmarz's algorithm was serially executed in the CPU, while the CGNR was executed in parallel in the GPU using the CUDA API [24].

RESULTS AND DISCUSSION

Reconstruction accuracy

In Figure 3, the simulation performed with the 3-temperature phantom clearly resulted in a significantly better reconstruction, both in concentration (denoted by the size of the polygons that compose the 3D image) and temperature distribution (denoted by the color of the polygons, following the scale in the center of the figure). The reconstructed image from the 3-temperature phantom has three clear distinct regions, the top at a higher temperature, the middle at an intermediate temperature, and the bottom at a lower temperature, as expected. However, the concentration and temperature accuracy are poorer in the lower region of the image, with an estimated temperature higher than the true temperature of the phantom. Examining the PSF plot of figure 1, one notices that the PSFs at 273 K and 303 K are more similar than the PSF at 323 K. Also, the magnitude of the PSF at 273 K is smaller and this might contribute to poorer convergence at lower temperatures.

The simulation performed with the gradient temperature phantom resulted in a much less accurate reconstruction, both in concentration and temperature. The temperature at the bottom is higher than expected, and also there are concentration artifacts in the reconstructed images, denoted by the bigger polygons in the central portion of the resulting 3D image. These artifacts show that the algorithm did not converge properly for the gradient temperature

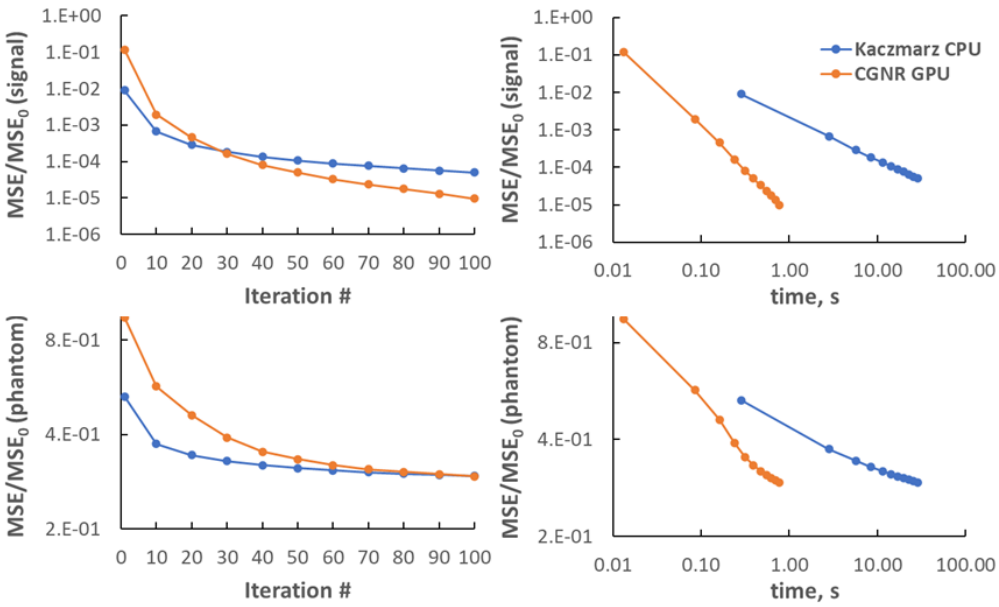

FIGURE 4. Convergence analysis of the algorithms studied in this work, serial Kaczmarz running in the CPU (blue) and parallel CGNR running in the GPU (orange). The top two plots show the convergence rates as functions of iteration number (left) and time (right) by means of the mean square error (MSE) values with relation to the measured signal. The bottom two plots show the MSEs as function of iteration number (left) and time (right) with relation to the known phantom concentration. The points in all plots correspond to the results of the first iteration, followed by the iterations in multiples of 10, until the 100th iteration.

phantom, and suggests limited interpolation capabilities of the *multi-color* T-MPI method. In other words, the method seems to not be capable of resolving intermediate temperatures between the calibration points.

The *multi-color* method is successful on detecting different particles because the solution of the extended linear system of equations allows determining how many particles of a specific type are present within the measured sample, without requiring any kind of interpolation. Having different PSFs is key for *multi-color* T-MPI, but is also its main weakness, since it only allows determining the number of particles at the calibration temperatures only, without being able to resolve it for intermediate temperatures. Improved accuracy would require significantly more than three different temperature calibration matrices, more closely spaced in temperature, which is not computationally feasible in memory requirements (the calibration matrices for this experiment occupied about 1.3 GB of data for 625 kHz of sampling rate and about 5.4 GB for 2.5 MHz) and in execution times. Additionally, the PSFs in the calibration matrices should be nearly orthogonal to allow the algorithm to converge, and this would require particle samples extremely sensitive to temperature changes to allow a reasonable temperature resolution. Basically the temperature resolution of the method is tied to the ability of obtaining distinct PSFs in the smallest temperature interval possible. Even though it is not likely that the *multi-color* T-MPI method will be used for general temperature measurement, it might find good use in specific applications, in which temperatures at very well defined values must be detected.

Convergence and *speed-up*

Figure 4 shows the convergence rates of Kaczmarz's algorithm and CGNR as a function the number of iterations (left plots) and time (right plots). The metrics adopted were the mean square error (MSE) values with relation to the measured signal (top plots) and with relation to the known phantom concentration (bottom plots). The convergence rate of Kaczmarz's algorithm is better in the first iterations, but it slows down after about 20 iterations. CGNR takes more iterations than Kaczmarz's, but evolves in a more consistent pace, and reaches lower MSE values with relation to the measured signal by the 30th iteration. Both algorithms seem to converge to the same MSE with relation to the previously know concentration (the ground truth), with Kaczmarz converging faster than CGNR. It is interesting to note that the algorithms converge differently, depending on the referential (the measured signal or the ground truth), which demonstrates that both algorithms converge to the same solution following different paths.

In time domain, however, one sees a more dramatic difference between the two algorithms, due to the speed boost per iteration associated with the parallel computation. By the time the Kaczmarz's algorithm finished its very first

iteration, the parallel CGNR algorithm already finished nearly 40 iterations, reaching MSE values about two orders of magnitude lower than the one resulting from a single Kaczmarz iteration. This is true for both the MSEs with relation to the measured signal and the ground truth, and demonstrates how the parallelization of the CGNR algorithm in the GPU helped obtain a much more efficient reconstruction, with an average *speed-up* of 36.6x in comparison to the serial Kaczmarz's algorithm.

CONCLUSION

This work served two purposes: to test the feasibility of the *multi-color* method for temperature measurements, and to measure the potential *speed-up* of the parallel CGNR algorithm for such a massive numerical problem. While the *multi-color* T-MPI method showed serious limitations in resolving intermediate temperatures (i.e., temperature values between the calibration temperatures) by interpolation, it worked well in the case of particles at well defined temperatures, and for which there were calibration matrices available. This shows that the method is only successful at detecting the amount of particles at each of the calibration temperatures instead of determining the temperature of the particles, which would be expected from a MNP-based thermometer. For measurements at interpolated temperatures, it is more likely that the harmonic ratio [16] approach will be more successful, and this will be the focus of future investigations.

Regarding the parallel CGNR, this work successfully demonstrated that the paralellization of the algorithm resulted in higher *speed-ups* when the size of the problem increased from a simple 2D concentration model [19] to a 3-temperature 3D reconstruction. The algorithm will be further investigated to deliver better performance.

Acknowledgments

The authors acknowledge the funding from NIST's Innovations in Measurement Science (IMS) grant.

REFERENCES

1. B. Gleich and J. Weizenecker, "Tomographic imaging using the nonlinear response of magnetic particles," Nature **435**, 1214–1217 (2005).
2. T. Knopp and T. M. Buzug, *Magnetic Particle Imaging* (Springer, 2012).
3. T. Knopp, N. Gdaniec, and M. Moddel, "Magnetic particle imaging: from prooof of principle to preclinical applications," Physics in Medicine and Biology **62**, R124–R178 (2017).
4. J. B. Weaver, A. M. Rauwerdink, and E. W. Hansen, "Magnetic nanoparticle temperature estimation," Medical Physics **36**, 1822–1829 (2009).
5. C. Stehning, B. Gleich, and J. Rahmer, "Simultaneous magnetic particle imaging (mpi) and temperature mapping using multi-color mpi," International Journal on Magnetic Particle Imaging **2**, 6pp (2016).
6. A. M. Rauwerdink, E. W. Hansen, and J. B. Weaver, "Nanoparticle temperature estimation in combined ac and dc magnetic fields," Phys Med Biol. **54**, 6pp (2009).
7. J. Zhong, W. Liu, Z. Du, P. C. Morais, Q. Xiang, and Q. Xie, "A noninvasive, remote and precise method for temperature and concentration estimation using magnetic nanoparticles," Nanotechnology **29**, 9pp (2012).
8. I. M. Perreard, D. B. Reeves, X. Zhang, E. Kuehlert, E. R. Forauer, and J. B. Weaver, "Temperature of the magnetic nanoparticle microenvironment: estimation from relaxation times," Phys. Med. Biol **59**, 1109–1119 (2014).
9. M. Zhou, J. Zhong, W. Liu, Z. Du, Z. Huang, M. Yang, and P. C. Morais, "Study of magnetic nanoparticle spectrum for magnetic nanothermometry," IEEE TRANSACTIONS ON MAGNETICS **51**, 6pp (2015).
10. J. Zhong, M. Schilling, and F. Ludwig, "Magnetic nanoparticle temperature imaging with a scanning magnetic particle spectrometer," Measurement Science and Technology **29**, 9pp (2018).
11. J. Salamon, J. Dieckhoff, M. G. Kaul, C. Jung, G. Adam, M. Möddel, T. Knopp, S. Draack, F. Ludwig, and H. Ittrich, "Visualization of spatial and temporal temperature distributions with magnetic particle imaging for liver tumor ablation therapy," Scientific Reports **10**, 11pp (2020).
12. NIST, "Thermal MagIC: An SI-Traceable Method for 3D Thermal Magnetic Imaging and Control (online)," https://www.nist.gov/programs-projects/thermal-magic-si-traceable-method-3d-thermal-magnetic-imaging-and-control (accessed 10/11/2023).

13. A. J. Biacchi1, T. Q. Bui, C. L. Dennis, S. I. Woods, and A. R. H. Walker, "Design and engineering colloidal magneticparticles for nanoscale thermometry," International Journal on Magnetic Particle Imaging **6** (2020), https://doi.org/10.18416/IJMPI.2020.2009068.
14. F. M. Abel, E. L. Correa, A. J. Biacchi, T. Q. Bui, S. I. Woods, A. R. H. Walker, and C. L. Dennis, "Thermosensitivity through Exchange Coupling in Ferrimagnetic/Antiferromagnetic Nano-Objects for Magnetic-Based Thermometry," ACS Applied Materials Interfaces **15** (2023), https://doi.org/10.1021/acsami.2c19673.
15. T. Q. Bui, W. L. Tew, and S. I. Woods, "Ac magnetometry with active stabilization and harmonic suppression for magnetic nanoparticle spectroscopy and thermometry," Journal of Applied Physics **128** (2020), https://doi.org/10.1063/5.0031451.
16. T. Q. Bui, M.-A. Henn, W. L. Tew, M. A. Catterton, and S. I. Woods, "Harmonic dependence of thermal magnetic particle imaging," Scientific Reports **13** (2023), https://doi.org/10.1038/s41598-023-42620-1.
17. J. Weizenecker, B. Gleich, J. Rahmer, H. Dahnke, and J. Borget, "Three-dimensional real-time in vivo magnetic particle imaging," Physics in Medicine and Biology **54**, L1–L10 (2009).
18. P. Vogel, S. Herz, T. Kampf, M. A. Rückert, T. A. Bley, and V. C. Behr, "Low latency real-time reconstruction for MPI systems," International Journal on Magnetic Particle Imaging **3**, 8 (2017).
19. K. N. Quelhas, M.-A. Henn, R. C. Farias, W. L. Tew, and S. I. Woods, "Parallel MPI image reconstructions in GPU using CUDA," International Journal on Magnetic Particle Imaging **9** (2023), https://doi.org/10.18416/ijmpi.2023.2303043.
20. P. W. Goodwill and S. M. Conolly, "The x-space formulation of the magnetic particle imaging process: 1-d signal, resolution, bandwidth, snr, sar, and magnetostimulation," IEEE TRANSACTIONS ON MEDICAL IMAGING **29**, 1851–1859 (2010).
21. T. Knopp, S. Biederer, and T. F. Sattel, "2D Model-Based Reconstruction for Magnetic Particle Imaging," Medical Physics **37**, 485–491 (2010).
22. K. N. Quelhas, M.-A. Henn, T. Q. Bui, H. R. Wagese, W. L. Tew, and S. I. Woods, "Flexible Software for Rigorous Simulations of Magnetic Particle Imaging Systems," International Journal on Magnetic Particle Imaging **8** (2022), https://doi.org/10.18416/IJMPI.2022.2203081.
23. J. Rahmer, A. Halkola, B. Gleich, I. Schmale, and J. Borgert, "First experimental evidence of the feasibility of multi-color magnetic particle imaging," Physics in Medicine and Biology **60**, 1775–1791 (2015).
24. Disclaimer: Certain commercial materials are identified in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology or any of the institutions associated with the authors, nor is it intended to imply that the materials are necessarily the best available for the purpose.