

Accelerated Parallel and Distributed Iterative Coordinate Descent (ICD) for X-ray CT

Madison G. McGaffin Jeffrey A. Fessler

Dept. of Electrical Engineering and Computer Science, University of Michigan

Abstract—This paper describes a new distributed algorithm for accelerating model-based image reconstruction in X-ray CT using iterated coordinate descent (ICD). The key novel component is a majorizer whose Hessian involves a block-diagonal matrix with triangular blocks (BDTriB). The resulting majorize-minimize algorithm combines aspects of ICD and the distributed block-separable surrogates (DBSS) algorithm for CT reconstruction [1]. Unlike traditional ICD, the proposed algorithm is also amenable to acceleration using Nesterov’s momentum [2] and the optimized gradient method (OGM) [3]. A simple preliminary experiment indicates potential for significant acceleration over traditional ICD and promising performance for distributed computing.

I. INTRODUCTION

Model-based image reconstruction (MBIR) in X-ray CT may improve image quality over direct reconstruction methods like filtered backprojection, but long reconstruction times impede widespread clinical use. Accelerating model-based reconstructions involves improving the mathematical structure of the numerical optimization algorithms [3], [4] and exploiting both modern hardware [5], [6] and distributed computing [1], [7]. This paper describes a parallelizable version of ICD.

Consider the following penalized weighted least-squares (PWLS) image reconstruction problem [8]:

$$\hat{\mathbf{x}} = \underset{\mathbf{x} \geq 0}{\operatorname{argmin}} \Psi(\mathbf{x}), \quad \Psi(\mathbf{x}) = \mathbf{L}(\mathbf{A}\mathbf{x}) + \mathbf{R}(\mathbf{C}\mathbf{x}), \quad (1)$$

with CT system matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$, finite differences matrix $\mathbf{C} \in \mathbb{R}^{K \times N}$, and data-fit and regularizer terms \mathbf{L} and \mathbf{R} :

$$\mathbf{L}(\mathbf{p}) = \sum_{i=1}^M \frac{w_i}{2} (p_i - y_i)^2, \quad \mathbf{R}(\mathbf{d}) = \sum_{k=1}^K \beta_k \psi(d_k), \quad (2)$$

with nonnegative statistical weights $\{w_i\}$ and regularization parameters $\{\beta_k\}$, where $\{y_i\}$ denotes the measured sinogram (log) data. We assume the convex potential function ψ is smooth with bounded curvature.

Suppose that we have B compute nodes that communicate via some interconnect; *e.g.*, multiple GPUs or processors on a single computer or multiple computers connected by a network. Data communication over this interconnect typically is slower than communication and computation within each node. Accordingly, we solve (1) by finding a majorizer consisting of a sum of B components that we minimize in parallel and then communicate the results to update the image \mathbf{x} .

Supported in part by NIH grant U01 EB018753 and Intel equipment donations. Email: {mcgaffin | fessler}@umich.edu .

II. METHODS

The proposed method combines the distributed block-separable surrogate (DBSS) [1] in a majorize-minimize (MM) framework [9] using a form of ICD for the inner minimizations and momentum for acceleration [2], [3].

A. Distributed block-separable surrogates

We start with a majorizer based on the distributed block-separable surrogate (DBSS) [1]. We partition the image \mathbf{x} into B blocks, $\{\mathbf{x}_b\}$, *i.e.*, $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_B)$, typically by axial slabs. The corresponding components of the CT system matrix \mathbf{A} and the finite differencing matrix \mathbf{C} are $\{\mathbf{A}_b\}$ and $\{\mathbf{C}_b\}$, respectively, where, *e.g.*, $\mathbf{A} = [\mathbf{A}_1, \dots, \mathbf{A}_B]$ and \mathbf{A}_b is the submatrix of \mathbf{A} having columns correspond to the pixels in \mathbf{x}_b . During each outer iteration, each computational node updates one of these sub-images, after which the sub-images are communicated between blocks.

In each outer iteration n , we form a block separable surrogate $\Phi^{(n)}$ as follows:

$$\Phi^{(n)}(\mathbf{x}) \triangleq \frac{1}{2} \left\| \mathbf{x} - \mathbf{x}^{(n)} \right\|_M^2 + \left(\mathbf{x} - \mathbf{x}^{(n)} \right)' \nabla \Psi \left(\mathbf{x}^{(n)} \right) + \Psi \left(\mathbf{x}^{(n)} \right), \quad (3)$$

where the iteration-invariant, block-diagonal, $N \times N$ matrix \mathbf{M} majorizes the Hessian of the cost function, *i.e.*, $\mathbf{M} \succeq \nabla^2 \Psi$. This is a tangent majorizer [9], *i.e.*, it satisfies

$$\begin{aligned} \Phi^{(n)} \left(\mathbf{x}^{(n)} \right) &= \Psi \left(\mathbf{x}^{(n)} \right), \\ \Phi^{(n)}(\mathbf{x}) &\geq \Psi(\mathbf{x}) \quad \forall \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (4)$$

Therefore, any $\mathbf{x}^{(n+1)}$ that descends the surrogate $\Phi^{(n)}$ will also descend the original cost function Ψ . For the MBIR problem (1), this property ensures convergence to $\hat{\mathbf{x}}$ [9].

The DBSS [1] has the block-diagonal Hessian $\mathbf{M}_{\text{DBSS}} \triangleq \operatorname{diag}\{\mathbf{M}_{\text{DBSS},b}\}$, where

$$\begin{aligned} \mathbf{M}_{\text{DBSS},b} &\triangleq \mathbf{A}'_b \mathbf{\Lambda}_b \mathbf{W} \mathbf{A}_b + \mathbf{C}'_b \mathbf{K}_b \mathbf{B} \mathbf{C}_b, \\ \mathbf{W} &= \operatorname{diag}\{w_i\}, \quad \mathbf{B} = \operatorname{diag}\left\{ \beta_k \cdot \max_d \psi''(d) \right\}, \end{aligned} \quad (5)$$

where $\mathbf{\Lambda}_b$ and \mathbf{K}_b are determined by the partition of \mathbf{x} :

$$\mathbf{\Lambda}_b \triangleq \operatorname{diag}\left\{ \begin{bmatrix} [\mathbf{A}\mathbf{1}]_i \\ [\mathbf{A}_b\mathbf{1}]_i \end{bmatrix} \right\}, \quad \mathbf{K}_b \triangleq \operatorname{diag}\left\{ \begin{bmatrix} \|\mathbf{C}\mathbf{1}\|_k \\ \|\mathbf{C}_b\mathbf{1}\|_k \end{bmatrix} \right\}, \quad (6)$$

where $|\mathcal{C}|$ denotes the element-wise absolute value of \mathcal{C} . This majorizer is block-separable, *i.e.*, it decomposes into B independent functions over different groups of pixels $\{\mathbf{x}_b\}$:

$$\Phi_{\text{DBSS}}^{(n)}(\mathbf{x}) = \sum_{b=1}^B \Phi_{\text{DBSS},b}^{(n)}(\mathbf{x}_b), \quad (7)$$

The DBSS algorithm descends each block surrogate $\Phi_{\text{DBSS},b}^{(n)}$ using ordered subsets with momentum acceleration (OS-MOM) [1], [3].

B. Distributed iterated coordinate descent

The DBSS algorithm does not *exactly* minimize each block-separable surrogate $\Phi_{\text{DBSS},b}^{(n)}$ because that would require too many inner iterations of OS-MOM to be practical. Although merely descending each surrogate is sufficient to ensure convergence, mere descent precludes using momentum-based techniques to accelerate the outer iterations.

To design a majorizer for which ICD can perform *exact* minimization, we propose to further majorize the block-separable matrix M_{DBSS} with a block-diagonal matrix having (lower) triangular blocks (BDTriB): $\mathbf{D} + \mathbf{T}$. We invert each of the triangular blocks of the BDTriB matrix *exactly* simply using back-substitution, *i.e.*, one sweep of ICD. By a proof similar to the one in [3], this design allows us to accelerate the outer iterations using momentum.

For derivation (but not implementation), define the $N \times N$ block diagonal matrix $\mathbf{T} = \text{diag}\{\mathbf{T}_b\}$ where each block \mathbf{T}_b is lower-triangular and defined as follows:

$$[\mathbf{T}_b]_{ij} = \begin{cases} [\mathbf{A}'_b \mathbf{W} \mathbf{A}_b + \mathbf{C}'_b \mathbf{B} \mathbf{C}_b]_{ij}, & i \geq j \\ 0, & \text{else.} \end{cases} \quad (8)$$

This definition retains the block structure of M_{DBSS} . We choose the diagonal matrix \mathbf{D} such that for all $\mathbf{z} \in \mathbb{R}^N$, the following majorization condition holds:

$$\mathbf{z}'(\mathbf{D} + \mathbf{T})\mathbf{z} = \mathbf{z}'\left(\mathbf{D} + \frac{1}{2}(\mathbf{T} + \mathbf{T}')\right)\mathbf{z} \geq \mathbf{z}'M_{\text{DBSS}}\mathbf{z}. \quad (9)$$

In particular, in each block we design diagonal \mathbf{D}_b such that:

$$\mathbf{z}'_b\left(\mathbf{D}_b + \frac{1}{2}(\mathbf{T}_b + \mathbf{T}'_b)\right)\mathbf{z}_b \geq \mathbf{z}'_b M_{\text{DBSS},b} \mathbf{z}_b. \quad (10)$$

Expanding the definition of \mathbf{T} (8), we design \mathbf{D}_b such that:

$$\begin{aligned} \mathbf{D}_b &\succeq M_{\text{DBSS},b} - \frac{1}{2}(\mathbf{T}_b + \mathbf{T}'_b) \\ &= \mathbf{A}'_b \mathbf{\Lambda}_b \mathbf{W} \mathbf{A}_b + \mathbf{C}'_b \mathbf{K}_b \mathbf{B} \mathbf{C}_b \\ &\quad - \frac{1}{2}(\mathbf{A}'_b \mathbf{W} \mathbf{A}_b + \mathbf{C}'_b \mathbf{B} \mathbf{C}_b + \mathbf{G}_b) \\ &= \mathbf{A}'_b \mathbf{W} \left(\mathbf{\Lambda}_b - \frac{1}{2}\mathbf{I}\right) \mathbf{A}_b + \mathbf{C}'_b \mathbf{B} \left(\mathbf{K}_b - \frac{1}{2}\mathbf{I}\right) \mathbf{C}_b - \frac{1}{2}\mathbf{G}_b, \end{aligned} \quad (11)$$

where \mathbf{G}_b contains the diagonal of $\mathbf{A}'_b \mathbf{W} \mathbf{A}_b + \mathbf{C}'_b \mathbf{B} \mathbf{C}_b$. The entries of $\mathbf{\Lambda}_b$ and \mathbf{K}_b (6) are greater than or equal to unity if

nonzero. We majorize the nondiagonal terms in (12) with the following ‘‘SQS-like’’ majorizer [10]:

$$\begin{aligned} D_{\text{SQS},b} &\triangleq \text{diag} \left\{ \left(\mathbf{A}'_b \mathbf{W} \left(\mathbf{\Lambda}_b - \frac{1}{2}\mathbf{I} \right) \mathbf{A}_b \right. \right. \\ &\quad \left. \left. + |\mathbf{C}_b|' \mathbf{B} \left(\mathbf{K}_b - \frac{1}{2}\mathbf{I} \right) |\mathbf{C}_b| \right) \mathbf{1} \right\}. \end{aligned} \quad (13)$$

Thus, our final diagonal component for the b th block is

$$\mathbf{D}_b \triangleq D_{\text{SQS},b} - \frac{1}{2}\mathbf{G}_b, \quad (14)$$

which one can verify is nonnegative. Computing this majorizer requires no more time than computing the diagonal majorizer for the SQS-MOM inner step of the DBSS algorithm. In other applications, *e.g.*, phase-contrast CT, the gram matrix $\mathbf{A}'\mathbf{\Lambda}_b\mathbf{A}$ may have negative entries. In these cases, the SQS majorizer (13) may be very loose or difficult to compute, and one can use another technique to find a diagonal majorizer, *e.g.*, the memory-efficient algorithm in [11].

C. Minimizing the new surrogate

The new surrogate for the b th block is

$$\Phi_b^{(n)}(\mathbf{x}_b) = \frac{1}{2} \left\| \mathbf{x}_b - \mathbf{x}_b^{(n)} \right\|_{\mathbf{D}_b + \mathbf{T}_b}^2 + \mathbf{x}'_b \nabla_{\mathbf{x}_b} \Psi(\mathbf{x}^{(n)}). \quad (15)$$

The matrix $\mathbf{D}_b + \mathbf{T}_b$ couples the entries of \mathbf{x}_b together, but because it is lower triangular, we minimize $\Phi_b^{(n)}$ *exactly* using back substitution with a nonnegativity constraint. That is, we loop though each pixel x_j of \mathbf{x}_b in a predetermined order and for each pixel solve the 1D minimization problem

$$\begin{aligned} x_j^{(n+1)} &= \underset{x_j \geq 0}{\text{argmin}} \frac{\omega_j}{2} (x_j - x_j^{(n)})^2 + x_j g_j, \\ &= \max\left(x_j^{(n)} - \frac{1}{\omega_j} g_j, 0\right), \end{aligned} \quad (16)$$

where

$$\begin{aligned} \omega_j &\triangleq [\mathbf{D}]_{jj} + \mathbf{a}'_j \mathbf{W} \mathbf{a}_j + \mathbf{c}'_j \mathbf{B} \mathbf{c}_j, \\ g_j &\triangleq \left[\nabla \mathbf{R}(\mathbf{x}^{(n)}) \right]_j + \mathbf{a}'_j \mathbf{W} (\mathbf{r}^{(n)} + \mathbf{r}_b) \\ &\quad + \mathbf{c}'_j \mathbf{B} \mathbf{C}_b (\mathbf{x}_b^{(n+)} - \mathbf{x}_b^{(n)}), \\ \mathbf{r}^{(n)} &\triangleq \mathbf{A} \mathbf{x}^{(n)} - \mathbf{y}, \quad \mathbf{r}_b \triangleq \mathbf{A}_b (\mathbf{x}_b^{(n+)} - \mathbf{x}_b^{(n)}), \end{aligned} \quad (17)$$

and \mathbf{a}_j and \mathbf{c}_j denote the j th columns of \mathbf{A} and \mathbf{C} , respectively. The vector $\mathbf{x}_b^{(n+)}$ contains the state of \mathbf{x}_b after the all the pixels before the j th pixel have been updated. Updating the j th pixel involves:

- computing the column vectors \mathbf{a}_j and \mathbf{c}_j ;
- computing ω_j and g_j , which involves finite differences (for the second term of g_j) and inner products;
- solving the one-pixel update problem (16);
- and finally updating the residual buffer

$$\mathbf{r}_b \leftarrow \mathbf{r}_b + \mathbf{a}_j (x_j^{(n+1)} - x_j^{(n)}). \quad (18)$$

These are the same steps as ICD [12] applied to (7) with the minor addition of the ‘‘relaxation’’ $[\mathbf{D}]_{jj}$ in ω_j .

TABLE I
 MOMENTUM-ACCELERATED DISTRIBUTED ICD ALGORITHMS

1)	Distribute $\{\mathbf{x}_b^{(0)}\}$ to the B computational nodes, initialize $\mathbf{z}^{(0)} = \mathbf{x}^{(0)}$. Compute $\mathbf{r}^{(0)} = \mathbf{A}\mathbf{x}^{(0)} - \mathbf{y}$ and $\{\mathbf{D}_b\}$ (14). Set $t^{(0)} = 1$.
2)	Loop outer iteration $n = 1, \dots, N_{\text{iter}}$: <ul style="list-style-type: none"> a) $t^{(n+1)} = \frac{1}{2} \left(1 + \sqrt{1 + 4(t^{(n)})^2} \right)$ b) In parallel for $b = 1, \dots, B$: <ul style="list-style-type: none"> i) Minimize block surrogate (15) using ICD to compute $\mathbf{x}_b^{(n+1)}$. ii) Momentum update: <ul style="list-style-type: none"> No momentum : $\mathbf{z}_b^{(n+1)} = \mathbf{x}_b^{(n+1)}$, FGM : $\mathbf{z}_b^{(n+1)} = \mathbf{x}_b^{(n+1)} + \frac{t^{(n)} - 1}{t^{(n+1)}} (\mathbf{x}_b^{(n+1)} - \mathbf{x}_b^{(n)})$ OGM : $\mathbf{z}_b^{(n+1)} = \mathbf{x}_b^{(n+1)} + \frac{t^{(n)} - 1}{t^{(n+1)}} (\mathbf{x}_b^{(n+1)} - \mathbf{x}_b^{(n)}) + \frac{t^{(n)}}{t^{(n+1)}} (\mathbf{x}_b^{(n+1)} - \mathbf{z}_b^{(n)})$ c) Broadcast node residuals $\{\mathbf{r}_b\}$ to compute $\mathbf{r}^{(n+1)}$ and edge slices of $\mathbf{x}_b^{(n+1)}$.
3)	Output: $\mathbf{z}^{(N_{\text{iter}})}$.

Between outer iterations, we synchronize the regularizer gradient $\nabla R(\mathbf{x}^{(n)})$ and residual $\mathbf{r}^{(n)} = \mathbf{A}\mathbf{x}^{(n)} - \mathbf{y}$ between computational nodes. The latter needs no additional projections or backprojections because

$$\mathbf{r}^{(n+1)} - \mathbf{r}^{(n)} = \sum_{b=1}^B \mathbf{r}_b, \quad (19)$$

so each node needs only to communicate residual updates $\{\mathbf{r}_b\}$ and edge voxels to compute $\nabla_{\mathbf{x}_b} R(\mathbf{x}^{(n)})$.

D. Momentum-based acceleration

In each outer iteration, the proposed algorithm forms and exactly minimizes a block-separable quadratic surrogate for the original cost function. This places the proposed algorithm in the same category as iterative shrinkage and thresholding (ISTA) and momentum-accelerated SQS algorithms. Thus, the distributed ICD algorithm can be improved with momentum-based acceleration *e.g.*, Nesterov's fast gradient methods (FGM) [2] or the optimized gradient method (OGM) [13]. Table I summarizes these algorithms.

Traditional ICD methods are not in general amenable to momentum-based acceleration [14]. For example, simply running ICD on the block-separable surrogate $\Phi_{\text{DBSS}}^{(n)}$ in (7) is equivalent to the proposed algorithm with each $\mathbf{D}_b = \mathbf{0}$. Although this basic approach descends the surrogate $\Phi^{(n)}$ and converges, it is incompatible with momentum-based acceleration; applying momentum-based acceleration to this basic combination of DBSS with ICD may cause divergence. The additional under-relaxation provided by \mathbf{D}_b (14) allows us to use ICD for exact minimization of a quadratic majorizer, making the algorithm compatible with momentum-based acceleration. This property is the key contribution of this work.

III. PRELIMINARY EXPERIMENT

To illustrate the concept, we simplified the CT reconstruction problem (1) to a one-dimensional 512-“pixel” problem.

We formed the system matrix \mathbf{A} using a Toeplitz matrix with point spread function $\frac{1}{r^{1/2}}$ so the Gram matrix $\mathbf{A}'\mathbf{A}$ has response proportional to $\frac{1}{r}$ [15]. The data-fit weights were uniform $w_i = 1$, and regularizer weights β_k were also uniform. We used a quadratic regularizer potential function $\psi(d) = \frac{1}{2}d^2$. This is an extreme simplification of the CT reconstruction problem for a preliminary investigation.

A. Single-node relaxation and acceleration

We compared 6 methods for $B = 1$: conventional ICD and ICD with the proposed relaxation (RICD), with and without FGM [2] ((R)ICD+FGM), or OGM [13] ((R)ICD+OGM). In this “one-node” setting, $\mathbf{A} = \mathbf{W} = \mathbf{I}$ in (6) and (1).

Fig. 2(a) plots the cost function per iteration for all 6 algorithms. Clearly, applying momentum to basic ICD without additional relaxation is infeasible.

The additional relaxation from the diagonal matrices $\{\mathbf{D}_b\}$ (14) slows the convergence of RICD compared to ICD. However, because every loop through all the pixels in RICD corresponds to minimizing a quadratic surrogate, we can apply momentum. With momentum acceleration, RICD converges faster than regular ICD. Furthermore RICD+FGM is provably convergent.

B. Distributed accelerated ICD

We also implemented the proposed distributed ICD algorithms on a simulated network of 2, 4 and 8 nodes. As the number of nodes B in the network increases, the block-separable majorizer becomes looser; Fig. 1(d) shows this effect in the increasing entries of the diagonal majorizer \mathbf{D} .

Fig. 2(b) shows the value of the cost function *vs.* iteration for normal ICD and RICD+OGM for 1, 2, 4, and 8 nodes. The ICD algorithms use the block-separable surrogate [1], and RICD-OGM uses the proposed BDTriB majorizer. As B increases, the larger majorizer values $\{\mathbf{D}_b\}$ slow convergence on a per-iteration basis. However the increase of the majorizer values as the number of nodes doubles is less than a factor of 2, so there is opportunity to accelerate the algorithm with distributed computing provided the communication overhead is not too high.

We simulated the time behavior of the distributed algorithms by

$$\Delta t_{\text{iter}} = \begin{cases} 1, & B = 1 \\ \alpha_{\text{overhead}} B + \frac{1}{B}, & B > 1. \end{cases} \quad (20)$$

Each node beyond the first ($B > 1$) adds some overhead due to communication and synchronization, α_{overhead} , but reduces how long it takes to compute the parallelizable workload, $\frac{1}{B}$. In this experiment, we assumed $\alpha_{\text{overhead}} = 0.05$. This is a pessimistic estimate; in other experiments with non-ICD algorithms, we found $\alpha_{\text{overhead}} \approx 0.01$ for multiple GPUs connected to the same computer or $\alpha_{\text{overhead}} \approx 0.03$ for computers connected by Ethernet.

Fig. 2(c) plots the value of the cost function *vs.* estimate time for the distributed (R)ICD algorithms. The distributed ICD algorithms reach peak performance at only two nodes, and is still slower than RICD on one node. The RICD algorithm

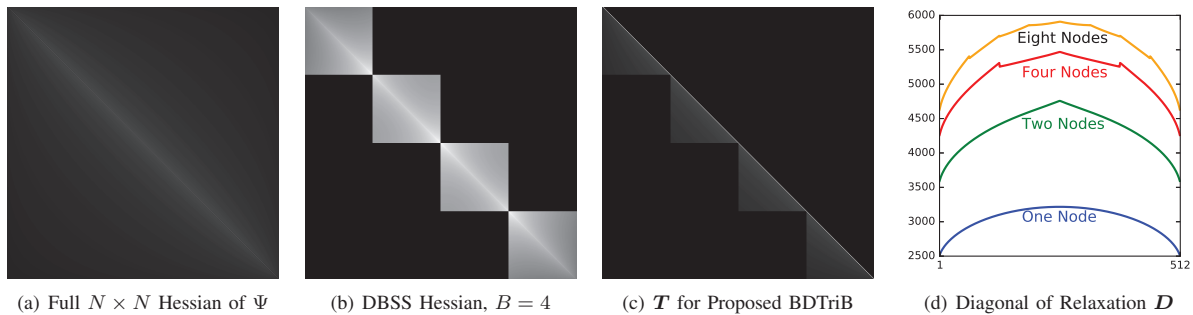


Fig. 1. Full Hessian of Ψ , Hessian of block-separable majorizer for $B = 4$, and T and D for BDTriB majorizer for the simplified reconstruction problem in Section III.




Fig. 2. Cost function curves for the experiment in Section III.

can exploit further parallelism, and convges most quickly on four nodes; the overhead from the eight-node configuration slows overall convergence. Although this preliminary experiment suggests only modest parallelization is useful, our selection for α_{overhead} is pessimistic and the time-characteristics of a real distributed system are difficult to estimate *a priori*.

IV. SUMMARY

We explored accelerating ICD using momentum and distributed computing. With additional relaxation based on matrix majorization, we can combine ICD [3] with the momentum methods that have been effective in accelerating ordered-subsets methods. A proof-of-concept experiment suggests that this approach can accelerate ICD. The proposed method also shows promising accelerations via distributed computing, by combining momentum-based acceleration with block-separable surrogates [1]. The voxels on the boundary between blocks might converge slower than others; this effect could be mitigated by dithering the block boundaries [1].

The sequential nature of ICD algorithms seems not well matched to computing hardware that becomes increasingly parallel. Nevertheless, the general techniques for accelerating and distributing ICD described here are also relevant to accelerating dual coordinate ascent algorithms for CT, *e.g.*, [5], that are amenable to implementation on modern hardware.

REFERENCES

- [1] D. Kim and J. A. Fessler. Distributed block-separable ordered subsets for helical X-ray CT image reconstruction. In *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.*, pp. 138–41, 2015.
- [2] Y. Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Math. Dokl.*, 27(2):372–76, 1983.
- [3] D. Kim, S. Ramani, and J. A. Fessler. Combining ordered subsets and momentum for accelerated X-ray CT image reconstruction. *IEEE Trans. Med. Imag.*, 34(1):167–78, January 2015.
- [4] H. Nien and J. A. Fessler. Fast X-ray CT image reconstruction using a linearized augmented Lagrangian method with ordered subsets. *IEEE Trans. Med. Imag.*, 34(2):388–99, February 2015.
- [5] M. McGaffin and J. A. Fessler. Alternating dual updates algorithm for X-ray CT reconstruction on the GPU. *IEEE Trans. Computational Imaging*, 1(3):186–99, September 2015.
- [6] R. Sampson, M. G. McGaffin, T. F. Wenisch, and J. A. Fessler. Investigating multi-threaded SIMD for helical CT reconstruction on a CPU. In *Proc. 4th Intl. Mtg. on image formation in X-ray CT*, 2016.
- [7] J. M. Rosen, J. Wu, T. F. Wenisch, and J. A. Fessler. Iterative helical CT reconstruction in the cloud for ten dollars in five minutes. In *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.*, pp. 241–4, 2013.
- [8] J-B. Thibault, K. Sauer, C. Bouman, and J. Hsieh. A three-dimensional statistical approach to improved image quality for multi-slice helical CT. *Med. Phys.*, 34(11):4526–44, November 2007.
- [9] M. W. Jacobson and J. A. Fessler. An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms. *IEEE Trans. Im. Proc.*, 16(10):2411–22, October 2007.
- [10] H. Erdoğan and J. A. Fessler. Ordered subsets algorithms for transmission tomography. *Phys. Med. Biol.*, 44(11):2835–51, November 1999.
- [11] M. G. McGaffin and J. A. Fessler. Algorithmic design of majorizers for large-scale inverse problems, 2015. arxiv 1508.02958.
- [12] Z. Yu, J-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh. Fast model-based X-ray CT reconstruction using spatially non-homogeneous ICD optimization. *IEEE Trans. Im. Proc.*, 20(1):161–75, January 2011.
- [13] D. Kim and J. A. Fessler. Optimized first-order methods for smooth convex minimization. *Mathematical Programming*, 2016. To appear.
- [14] A. Chambolle and T. Pock. A remark on accelerated block coordinate descent for computing the proximity operators of a sum of convex functions. *SMAI J. of Computational Mathematics*, 1:29–54, 2015.
- [15] N. H. Clinthorne, T. S. Pan, P. C. Chiao, W. L. Rogers, and J. A. Stamos. Preconditioning methods for improved convergence rates in iterative reconstructions. *IEEE Trans. Med. Imag.*, 12(1):78–83, March 1993.