Scalable Regularized Tomography without Repeated Projections

Jonas August Takeo Kanade
The Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213
{jonas,tk}@cs.cmu.edu

Abstract

X-ray computerized tomography (CT) and related imaging modalities (e.g., PET) are notorious for their excessive computational demands, especially when noise-resistant probabilistic methods such as regularized tomography are used. The basic idea of regularizated tomography is to compute a smooth image whose simulated projections (line integrals) approximate the observed, noisy X-ray projections. The computational expense in previous methods stems from explicitly applying a large sparse projection matrix to enforce these smoothness and data fidelity constraints during each of many iterations of the algorithm. Here we review our recent work in regularized tomography in which the smoothness constraint is analytically transformed from the image to the projection domain, before any computations begin. As a result, iterations take place entirely in the projection domain, avoiding the repeated sparse matrix-vector products. A more surprising benefit is the decoupling of a large system of regularization equations into many small systems of simpler independent equations, whose solution requires an "embarassingly parallel" computation. In this paper, we demonstrate that this method provides linear speedup of regularized tomography for up to 20 compute nodes (Pentium 4, 1.5 GHz) on a 100 Mb/s network using a Matlab MPI implementation.

1 Introduction

1.1 The Computational Burden of Regularized Tomography

X-ray computerized tomography (CT) and other related imaging modalities such as positron emission tomography (PET) are notorious for their excessive computational demands. The earliest algorithms in tomography, such as filtered backprojection (FBP), are now trivial in two-dimensions and scalable in three-dimensions due to their "pleasantly" or "embarassingly" parallel nature [18]. Unfortunately, FBP can produce streaking artifacts in high-noise situations because it does not model the statistics of the noise. The more noise-resistant probabilistic (e.g., regularized) tomography methods are still computationally prohibitive, especially in three-dimensions and at high resolutions. In this paper we consider a new class of algorithms in which probabilistic tomography becomes "pleasantly parallel" as well.

To understand how this works, recall that tomographic images are produced by converting observed projections (data) into an image. For example, in X-ray CT imaging, X-ray beams are directed at a patient and become attenuated by various amounts by the different materials within the body. On the opposite side of the patient the attenuated beams are detected as an array of measurements called a *projection*. (Such projections are produced at many different angles around the patient.) Not only are these measurements noisy, the relative noise level depends on the amount of attenuation: projections through dense materials such as bone and especially metal have lower signal-to-noise ratio than projections through only tissue or water (Fig. 1). Coping with the large and spatially-varying fluctuations in the number of detected photons often requires a probabilistic smoothing technique (also known as regularization) to improve the image. The basic idea of regularization is to infer a smooth image whose simulated projections approximate the observed (but noisy) projections. The difficulty is that the standard regularization specifies the two basic constraints in different *domains*: (1) smoothness is an image property specified in the *image domain*, while (2) closeness to the data is determined in the *projection domain*. Enforcing these constraints computationally requires a very large sparse matrix-vector multiplication (projection or backprojection) to convert between the two domains during each iteration of the algorithm (§2.2). This multiplication places great demands on system memory and memory bandwidth, and requires significant network usage when parallelized. Projection and backprojection are thus considered to be the largest expense in iterative methods for tomography [13].

The novelty in our proposed formulation of regularized tomography is in the *analytical* conversion of the smoothness constraint itself from the image to the projection do-

Proceedings of the 18th International Parallel and Distributed Processing Symposium (IPDPS'04)

0-7695-2132-0/04/$17.00 (C) 2004 IEEE

main, before any computations begin. Regularized CT thus becomes a two-stage process of (nonhomogeneous) smoothing of the projections followed by filtered backprojection. As a by-product, the repeated forward and backprojections common in iterative image reconstruction are eliminated. Despite the computational simplification, we demonstrate that this method can be used to reduce metal artifacts in X-ray CT images (Fig. 4). But a more fascinating benefit of the conversion of the smoothness constraint to the projection domain is the *decoupling* of a large system of regularization equations into many small systems of simpler regularization equations, each for a separate projection. Thus the difficult computation of regularized tomography becomes pleasantly parallel, so that latency tolerant and ideally scalable parallel computations are possible, as we report in §4.2 in two-dimensions.

1.2 Applications of Fast Regularized Tomography

A wide variety of applications in physics, biology, and medicine can benefit from such a scalable probabilistic algorithm for tomographic imagery.

1.2.1 Higher Resolution Tomography

Current tomographic problems typically have up to on the order of a million unknowns in the volumetric image. There is always a demand for higher resolution, and thus more unknowns. For example, a tenfold increase in resolution in each dimension (1000-by-1000-by-1000) would give rise to a problem in a billion unknowns. One can already see nearly such resolutions emerging in electron microscopy [18] for studying biological specimens as well as in CT.

1.2.2 A Hospital's Workstations as a Supercomputer for Tomography

An algorithm for tomographic imaging that is "embarassingly parallel" would be beneficial for hospitals, as it would make high performance computing readily accessible: the increasingly powerful desktop computers throughout the hospital could be used as a cluster. This is due to the tolerance of embarassingly parallel applications to network latency and variations in node processor speed.

Previous methods for parallelizing probabilistic tomography cannot easily exploit available workstations [5, 13] as they are based on iterated projection and backprojection. Recall that the projection matrix is a discretization of a linear operator taking the line-integral of the image over the X-ray path; this $O(n^2)$-by-$O(n^2)$ sparse matrix has $O(n)$ nonzero entries per row for an n-by-n 2-dimensional image. This growth of the number of nonzero entries with resolution means that tomographic projection has much greater

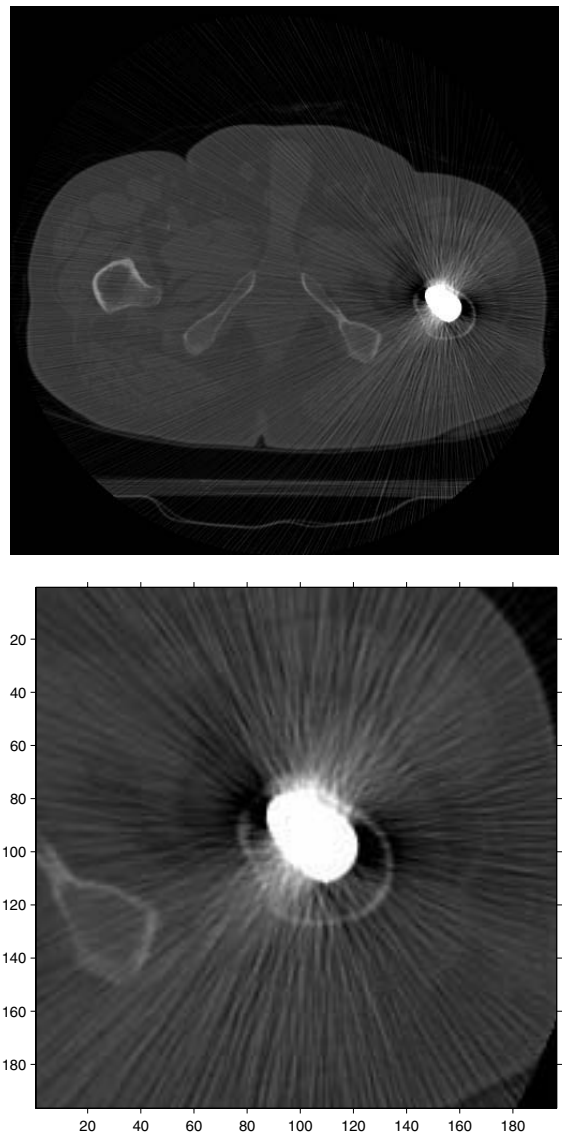

Figure 1. Streaking artifacts in this X-ray CT (top; bottom: detail around hip to right) of a human hip are due to the presence of a metal hip implant (bright oval). Such streaks interfere with post-operative examinations following total hip replacement surgery as well as with the automatic extraction from the CT image, or segmentation, of the patient's other hip (top left). CT metal artifacts are also caused by pins that secure bones together, by bullets and shrapnel in combat casualties, and by luggage screened at airports (an application we hope to develop using our method). Similar artifacts occur in PET and SPECT scans as well because the common cause is the random, particulate nature of the Poisson-distributed projection data.

0-7695-2132-0/04/$17.00 (C) 2004 IEEE

network bandwidth needs than the fixed local-neighborhood communication requirements for partial differential operators. Some methods based on projection matrices even require a global summation of a volumetric dataset across the network, where each processor sends and/or receives the entire dataset. Not only is this communication a serial bottleneck, it can become the dominant component of execution time. In addition, when exploiting the sparsity of projection data or when using processors of different speeds, load balancing can be a problem. Also, even after significant attempts at exploiting symmetries and sparsity have been made [13], these matrices can require a lot of memory (e.g., 220 MB for the projection matrix for a *two*-dimensional 256-by-256 image; more realistic image sizes and three-dimensional images can be much worse). Our embarassingly parallel technique requires far less communication (distributing data and gathering final results) and scales almost perfectly at higher resolutions (see §4.2). In addition, our memory requirements are trivial (kilobytes) in two-dimensions and modest (megabytes) in three-dimensions.

1.2.3 Reduction of Radiation Exposure

A tantalizing application for faster probabilistic methods in tomography is the reduction of the amount of X-ray radiation or radio-pharmaceutical to which the patient would have to be exposed. The reason this is possible stems from the faithfulness of the probabilistic model to the physics of the imaging, and thus enabling a faithful reconstruction from noisier data, which in turn allows for lower X-ray radiation. Outside of medicine, applications in biology would benefit from less exposure to specimens for examination by microCT (the application of CT to small objects) or electron microscopy.

1.2.4 Other Applications

Another application of faster probabilistic tomography is in fully 3-dimensional tomography (as opposed to a stack of 2-dimensional tomographs, or *slices*). Three-dimensional tomography would allow the exploitation of smoothness between the slices and noise reductions in PET images. A related benefit would be the application of the C-arm (a rotating flouroscope) as an inexpensive CT device. On-line electron microscopy, where the resulting tomograph is formed during the scan, would also benefit from faster probabilistic tomography. Somewhat surprisingly, probabilistic tomography can be used to increase the resolution of astronomical images [3]. Here one rotates a binocular telescope (which has higher resolution in one of the two dimensions), and then reconstructs a single high-resolution image.

1.3 Overview of Paper

The remainder of this paper is organized as follows. In §2.1 we provide background on CT. In §2.2 we present the theory of regularized CT, ending with the equations that one typically must solve in standard approaches. In §3, we then review the mathematical techniques needed to carry the smoothness constraint from the image domain to the projection domain and thus decouple the regularized CT equations. Our parallel algorithm is described in §4.1 and its scalability is demonstrated in §4.2.

2 Review of Regularized Tomography

2.1 Background in CT

In the absence of noise, the basic problem of computerized tomography (CT) is to determine an unknown image $f = f(x, y)$ from its (forward) projections, or *Radon transform* $\mathcal{R}f$, where

$$(\mathcal{R}f)(t,\theta) := \int\int f(x,y)\delta(t - x\cos\theta - y\sin\theta)dxdy,$$

where (x, y) are planar coordinates, t is the location along each projection, $\theta \in [0, \pi)$ is the orientation of the projection, and $\delta(\cdot)$ is the Dirac delta-function. (Here we focus on the two dimensional problem with standard parallel-beam geometry, but the ideas readily extend to three dimensions (see §3.1) and other scanning geometries.) Unfortunately, since measurements are never perfect, what we actually observe are the noisy projection data $g = g(t, \theta)$. (To emphasize the essentials of the tomography problem, we view the unknown $f = f(x, y)$ and the observation $g = g(t, \theta)$ as functions, although the implementation is discrete. See §3.2.) We make the standard independence and locality assumptions that the likelihood $P(g|f)$, or conditional distribution of g given f, equals $\prod p(g(t,\theta)|(\mathcal{R}f)(t,\theta))$, where the product is over all (t, θ). Since the likelihood depends on the particular imaging modality, we illustrate using X-ray CT for concreteness, although our technique applies more broadly (e.g., to PET). Here, therefore, $b\exp(-g(t,\theta))$ is the observed X-ray photon count, Poisson-distributed with mean $b\exp(-(\mathcal{R}f)(t,\theta))$, where b is the mean photon count before any attenuation and $f(x, y)$ is the attenuation coefficient at (x, y). Our task is to infer the image f given the noisy projections g.

2.2 Regularized CT

Because this inverse problem is ill-posed [16], one typically imposes extra constraints on f. In penalized maximum likelihood [8], or *regularization*, inferring f amounts

to finding that f which minimizes $-\ln P(g|f) + \rho(f)$, where $\rho(f)$ characterizes the extra constraint on f. Here we impose the standard smoothness constraint that uses the image gradient $\nabla f = (\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y})$ in the quadratic penalty $\rho(f) := \beta||\nabla f||^2$, where $\beta > 0$, $||f||^2 := \langle f, f\rangle$, and $\langle f_1, f_2\rangle := \int f_1 f_2$ is an inner product. Following [17] and to simplify the presentation, we approximate $-\ln P(g|f)$ with the quadratic form $||g - \mathcal{R}f||^2_{\mathcal{W}}$, where $||g||^2_{\mathcal{W}} = \langle g, \mathcal{W}g\rangle$ is a weighted norm with (diagonal) *weight operator* $\mathcal{W}$ satisfying $(\mathcal{W}g)(t,\theta) := w(t,\theta)g(t,\theta)$. The weight $w(t,\theta) := b\exp(-g(t,\theta))$ is small for those rays passing through dense materials such as bone or metal, and larger otherwise. This formulation of tomography requires that we solve the following "hard" optimization problem:

Problem 1 (Regularized CT). *Given the projection data g, find the image f that minimizes*

$$||g - \mathcal{R}f||^2_{\mathcal{W}} + \beta||\nabla f||^2.$$

To proceed, let $\triangle$ denote $\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$, the Laplacian in the plane. Using integration by parts and zero boundary conditions, we recall that $||\nabla f||^2 = \langle f, -\triangle f\rangle$. By taking the functional derivative with respect to f of the functional in Problem 1, we obtain the Euler-Lagrange equation

$$\mathcal{R}^*\mathcal{W}\mathcal{R}f - \beta\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)f = \mathcal{R}^*\mathcal{W}g, \qquad (1)$$

where f is unknown and A^* denotes the adjoint of linear operator A ($\mathcal{R}^*$ is also known as the backprojection operator). By examining (1), we see that Problem 1 is hard in two related ways.

First, the problem constraints occur in two different domains. Fidelity to the data ($||g - \mathcal{R}f||^2_{\mathcal{W}}$) is enforced in the projection domain $\{(t,\theta)\}$, while smoothness ($||\nabla f||^2$) is imposed in the image domain $\{(x,y)\}$. Thus we see in (1) the operators $\mathcal{R}$ and $\mathcal{R}^*$ for shuffling back and forth between these domains; iterative solution techniques typically compute these forward and backprojections explicitly and at great expense (§1).

Second, observe that (1) is a *coupled* equation in the two-variable function $f = f(x,y)$, i.e., in the large set of variables $\{f(x,y),$ for all $x, y\}$ under some discretization of x and y. The coupling arises first because both $x-$ and $y-$ derivatives are present; in addition, $\mathcal{R}$ and $\mathcal{R}^*$ are integral operators, and so are not even local. The computational difficulty in solving (1) and related tomographic problems (e.g., emission) has spawned a large body of work in optimization (see [9, 15, 4] and references therein).

3 Decoupled Regularized CT

In this section we formulate the entire regularization problem in a single domain. (We introduced the mathematics of our approach in [1], which we review here; we report on scalability in this paper.). As we shall see, working solely in the projection domain will decouple our large joint optimization problem into many smaller ones. But first we recall the standard technique for inverting the Radon transform: filtered backprojection. Let the Fourier transform of $g = g(t,\theta)$ with respect to t be denoted $(\mathcal{F}_1 g)(\tau,\theta) = (2\pi)^{-\frac{1}{2}}\int g(t,\theta)e^{-i\tau t}dt$, where τ is the spatial frequency along the θ-projection.

Definition 1. *Given the function $h = h(t,\theta)$, the* **Riesz potential** *is the linear operator $\mathcal{I}^\alpha$ satisfying $(\mathcal{F}_1\mathcal{I}^\alpha h)(\tau,\theta) = |\tau|^{-\alpha}(\mathcal{F}_1 h)(\tau,\theta)$.*

Given noise-free observations $h = \mathcal{R}f$, one can solve for the unknown f by directly implementing the following classical formula [16] for the inverse of the Radon transform:

Fact 1 (Filtered Backprojection). $\mathcal{R}^{-1} = \frac{1}{4\pi}\mathcal{R}^*\mathcal{I}^{-1}$.

By defining $\square$ as the Laplacian $\frac{\partial^2}{\partial t^2}$ along each projection, we note that $\mathcal{I}^{-1}$ is the "square root" of $-\square$. This is because $\mathcal{I}^{\alpha_1}\mathcal{I}^{\alpha_2} = \mathcal{I}^{\alpha_1+\alpha_2}$ and the following:

Fact 2. $\mathcal{I}^{-2} = -\square$.

Now recall the Fourier slice theorem, which says that the two-dimensional Fourier transform of f, evaluated at polar coordinates (τ,θ), is just $(\mathcal{F}_1\mathcal{R}f)(\tau,\theta)$. Using this theorem, one can relate the two Laplacians $\triangle$ and $\square$ because the two-dimensional Fourier transform of $-\triangle$ is u^2+v^2 (where u and v are spatial frequencies for x and y, respectively), or $|\tau|^2$ in polar coordinates, which is the one-dimensional Fourier transform of $-\square$. For details and the extension to higher dimensions, see [16] and [11]. This crucial but simple idea is called "**intertwining**".

Fact 3. *The Radon transform $\mathcal{R}$ intertwines $\triangle$ and $\square$, i.e., $\mathcal{R}\triangle = \square\mathcal{R}$.*

Our main contribution is the realization that by applying intertwining to regularized CT*, we can "decouple" this large optimization problem into an equivalent set of much smaller optimization problems.* The idea is to reformulate Problem 1 in terms of $h = \mathcal{R}f$. We need only transfer the smoothness constraint to the projection domain, as the data constraint is already naturally specified there already. Intertwining allows us to analytically transfer the smoothness constraint to the projection domain, as opposed to numerically enforcing it during optimization computations with repeated and expensive forward and backward projections.

Proposition 1 (Smoothness Constraint to Projection Domain).

$$||\nabla f||^2 = (4\pi)^{-1}\langle h, \mathcal{I}^{-3}h\rangle, \quad \textit{where } h = \mathcal{R}f.$$

0-7695-2132-0/04/$17.00 (C) 2004 IEEE

For proof, note that $-||\nabla f||^2 = \langle f, \triangle f\rangle = \langle R^{-1}h, \triangle R^{-1}h\rangle = \langle h, \mathcal{R}^{-1^*}\triangle R^{-1}h\rangle$. Now, note that $\mathcal{R}^{-1^*} = (4\pi)^{-1}\mathcal{I}^{-1}\mathcal{R}$, using Fact 1 and the symmetry of $\mathcal{I}^\alpha$. But then using Fact 3 $-\mathcal{R}^{-1^*}\triangle\mathcal{R}^{-1} = -(4\pi)^{-1}\mathcal{I}^{-1}\mathcal{R}\triangle\mathcal{R}^{-1} = -(4\pi)^{-1}\mathcal{I}^{-1}\Box\mathcal{R}\mathcal{R}^{-1}$. Prop. 1 follows using Fact 2 and because $\mathcal{R}\mathcal{R}^{-1}$ is the identity operator.

Thus we can pose Problem 1 in an equivalent, "easy" form in the projection domain, as follows:

Problem 2 (Decoupled Regularized CT). *Given observed projections g, find projections $h = \mathcal{R}f$ minimizing*

$$||g - h||^2_{\mathcal{W}} + \beta'\langle h, \mathcal{I}^{-3}h\rangle, \qquad \textit{where } \beta' := (4\pi)^{-1}\beta.$$

The corresponding Euler-Lagrange equation,

$$\mathcal{W}h + \beta'\mathcal{I}^{-3}h = \mathcal{W}g, \tag{2}$$

where h is unknown, is easy exactly where (1) is hard. First, the forward and backprojections are eliminated from the optimization; backprojection need only be done once to determine f from solution h. Second, and more importantly, equation (2) is really a decoupled set of systems of equations, where each system corresponds to the unknowns $\{h(t, \theta)$, for all $t\}$, at each *fixed* θ. This follows because operator $\mathcal{W}$ is pointwise multiplication by a scalar and $\mathcal{I}^{-3}$ acts only along t. Thus our "decoupled" approach to regularized CT requires solving the integral equation (2) in the unknown single variable function $h(\cdot, \theta)$, for each fixed θ.

To get an intuitive understanding of (2), recall from the definition of the Riesz potential that the operator $\mathcal{I}^{-3}$ acts along each projection by boosting high spatial frequencies to the third power, much like differentiation. Thus the operator $\mathcal{W}+\beta'\mathcal{I}^{-3}$ is analogous to a shifted derivative operator, and thus its inverse can be expected to have a smoothing effect (integraton). The smoothing will be greatest in those portions of a projection where the weight w is smallest.

3.1 Generalization to Three Dimensions

The same decoupling concept applies in three dimensions as well. To show this, one can argue in a formally analogous way as above, except that the Radon transform is replaced by the X-ray transform (which takes straight line integrals in $\mathbb{R}^3$). One then obtains equations of the same form as (2), and f, g, and h are three-dimensional functions. In future work, we shall implement these equations, assess their scalability (which we anticipate due to the increased computational granularity of processing two-dimensional arrays vs. one-dimensional ones), and make the connection to the cone-beam geometry in addition to the parallel geometry which we have already studied.

3.2 Metal Streak Artifact Reduction

To solve (2), we observe that $\mathcal{W} + \beta'\mathcal{I}^{-3}$ is a positive definite operator (if $w(t, \theta) > 0$), and thus the conjugate gradient method can be applied. To discretize the equations, we sampled in t and θ uniformly. The operator $\mathcal{W}$ was implemented by restriction to the sample locations. The Riesz potential was implemented by taking 1-dimensional FFTs.

We applied our decoupled regularized CT method to the reduction of metal artifacts in X-ray CT (Fig. 1) [20, 12, 14]. Since the actual projection data and scanner parameters were unavailable, we simulated the projections (using "radon" from the Image Processing Toolbox in Matlab) after rescaling the image pixel values (range 0 to 255) by 0.012 (Fig. 3, top left). Alternative and faster methods for computing projections and backprojections include the hierarchical decomposition method of Basu and Bresler [2] and the nonuniform FFT method of Fessler and Sutton [10].

The bright band results from the metal, and is noisier than elsewhere; the noise is obvious in the single projection in Fig. 2 (solid curve). For each fixed projection orientation θ, the decoupled regularized CT equation (2) was solved ($b = 10^9, \beta = 10^3$) to produce the nonhomogeneously smoothed projection shown in Fig. 2 (dotted curve). After smoothing each projection independently, filtered backprojection (using "iradon" in Matlab) was applied to the *set* of smoothed projections (Fig. 3, top right), producing our final result (Fig. 4), which shows reduced streaking artifacts.

4 Parallel Computations

4.1 A Scalable Algorithm

To solve the equation (1) in its decoupled form (2) on a network of processors, we set up a master-slave computation where the master processor performs a single Radon transform (X-ray transform in 3-d) and then sends each of the resulting projections (one per fixed orientation θ) to the next available slave processor. The master also receives the result from each slave processor computation and sends out another projection to said slave if there any left to process or a termination message if there are not. When the resulting solutions for all projections have been received, the master then performs a single filtered backprojection computation (e.g. the Feldkamp method in 3-d) to produce the probabilistically reconstructed image. Each slave processor waits to receive the next projection, solves the decoupled regularization equation (2) and then sends the solution back to the master. In slave then returns to its waiting state until it receives a message from the master to terminate.

As is common for such embarassingly parallel master-slave computations, load balancing and tolerance of latency in the interconnection network are automatically dealt

0-7695-2132-0/04/$17.00 (C) 2004 IEEE

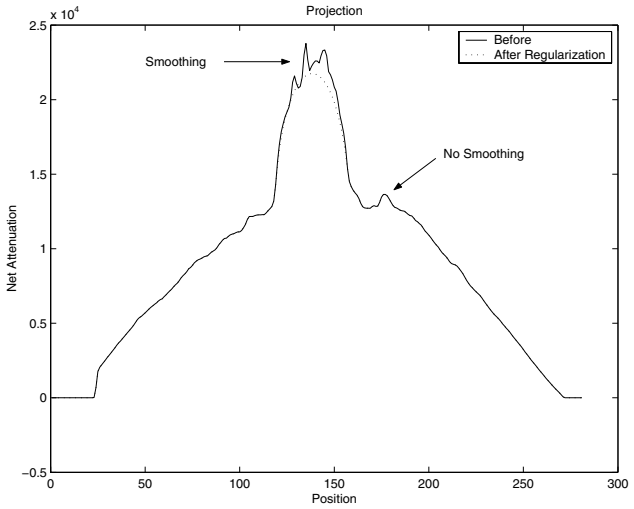

Figure 2. A simulated projection (solid curve, $g(\cdot, \theta)$) oriented at $\theta \approx 45°$ is noisy particularly in the portion due to the metal (large hump in the center). The smoothing due to decoupled regularized CT (dotted curve, $h(\cdot, \theta)$) is greater in the metal portion, where the weights are automatically lowest.

with [19, 6]. In addition, processors of different speeds can be used as slaves: the slower units simply take longer to return their results while the faster ones will process more projections.

4.2 Scalability Study

We also performed a scalability study of our decoupled approach to regularized tomography. Here, the number of processors was increased and the corresponding execution time was measured. The input image was the standard two-dimensional Shepp-Logan phantom (produced by Matlab), and instead of simply solving our decoupled regularization equations once per projection, we used them as the linear step in a Newton's method to fully exploit the (nonlinear) Poisson noise model. Computations were performed in Matlab (version 6.1) and communications were handled by MPITB [7], a MATLAB toolbox for LAM/MPI (version 6.5.7, and built with GNU C compiler version 2.96). The computations were performed on a network of idle 1.5 GHz Pentium 4 nodes running GNU/Linux (Redhat 7.1) and connected by a 100 Mb/s ethernet (with 3Com 3c905C PCI host adapters). Each CPU had 256 kB of L2 cache and 256 MB of main memory with a memory bandwidth of 1.6 GB/s (obtained using Stream "triad" compiled with the GNU C compiler). One node was a dedicated master and 20 slave nodes were available. In Fig. 5 we see the results of this

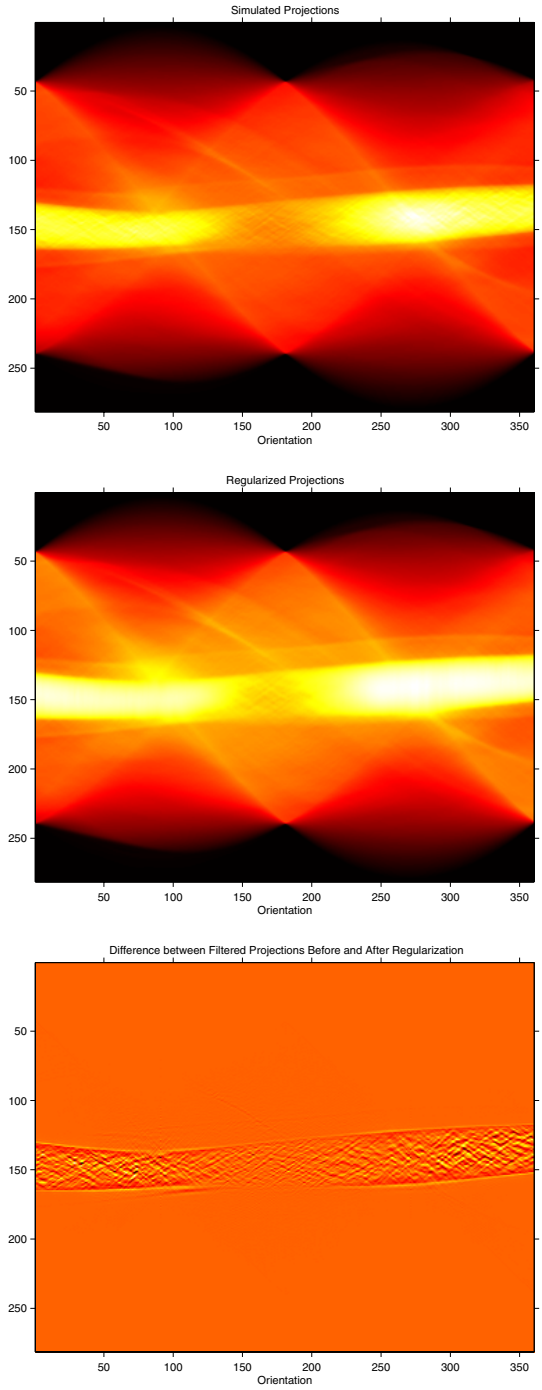

Figure 3. The simulated projection data $g(t, \theta)$ (top), or Radon transform, of Fig. 1 (bottom) shows a bright band due to the metal. The key smoothing action of decoupled regularized CT ($h(t, \theta)$, middle) is localized on the metal band, as emphasized (bottom) in the result of applying $\mathcal{I}^{-1}$ (a kind of differentiation and the first step in filtered backprojection) to the difference $g - h$.

0-7695-2132-0/04/$17.00 (C) 2004 IEEE

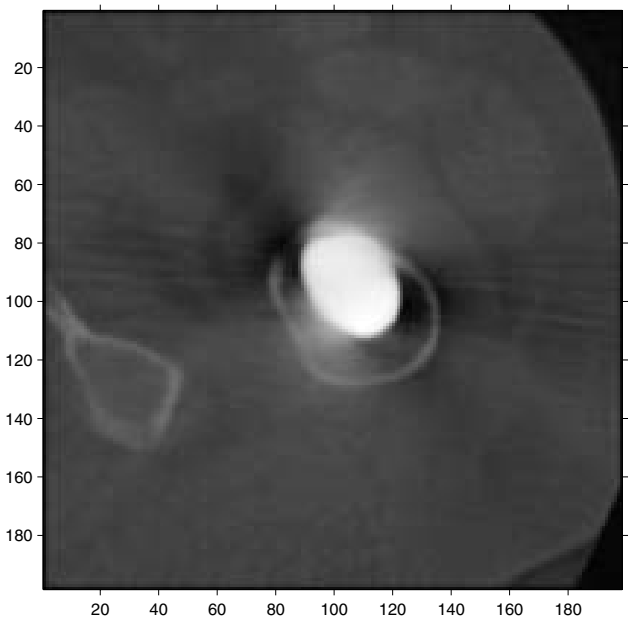

Figure 4. The result of decoupled regularized CT for reducing streaking artifacts (compare to Fig. 1). Observe that the smoothing effect is primarily near the metal; the rest of the image is still sharp.

study, where the scalability is essentially linear for the realistic image sizes of 256x256 or 1024x1024. Anecdotally we note that master processor was mostly idle at these large image sizes, suggesting our algorithm is granular.

5 Conclusion

In this paper we have presented the decoupling concept for simplifying regularized tomography. Our approach is fundamentally unlike most current approaches to the statistical estimation of tomographic images as it eliminates the most expensive steps of those techniques, viz., the repeated projections and backprojections. In addition, the decoupling idea makes regularized tomography embarassingly parallel and our results show how this leads to scalable computations. As discussed at length in the introduction, the impacts of this approach include the many practical applications of faster regularized tomography.

Acknowledgements: We thank Branko Jaramaz for introducing us to the metal artifact problem. The National Science Foundation provided funding under the ACR program.

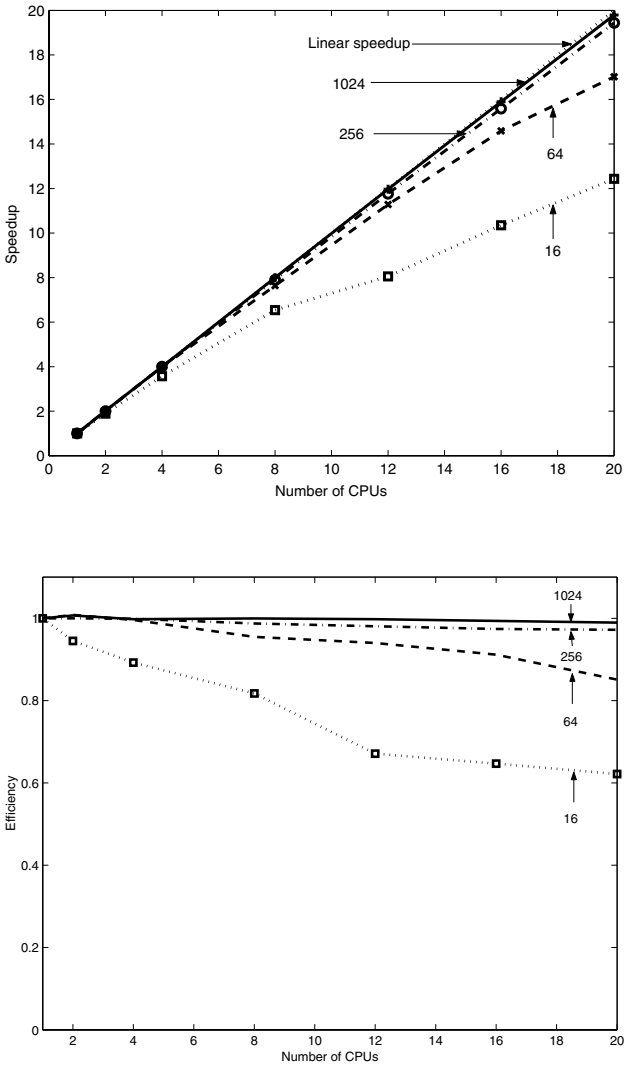

Figure 5. (top) Speedup (the run-time for one slave processor / the run-time for several slave processors) of decoupled regularized tomography with increased number of processors. From bottom to top are the speedup results for images of size 16x16, 64x64, 256x256, and 1024x1024. Observe that while speedup is poor for the smaller, unrealistic, image sizes, it is almost linear (i.e., ideal) for typical image sizes and improves with increased resolution, which is very encouraging. (bottom) Efficiency of processor use for decoupled regularized tomography with increased number of processors. Efficiency is defined as speedup / number of processors. Observe that for the larger (and typical) image sizes, the scalability is essentially optimal.

0-7695-2132-0/04/$17.00 (C) 2004 IEEE

References

[1] J. August. Decoupling the equations of regularized tomography. In *International Symposium on Biomedical Imaging*, pages 653–656, 2002.

[2] S. Basu and Y. Bresler. $O(n^2 \log_2 n)$ filtered backprojection reconstruction algorithm for tomography. *IEEE Transactions on Image Processing*, 9(10):1760–1773, 2000.

[3] M. Bertero and P. Boccacci. Application of the os-em method to the restoration of lbt images. *Astronomy and Astrophysics, Supplemental Series*, 144:181–186, 2000.

[4] J. Browne and A. R. D. Pierro. A row-action alternative to the em algorithm for maximizing likelihoods in emission tomography. *IEEE Trans. on Medical Imaging*, 15(5):687–699, 1996.

[5] C. M. Chen, S.-Y. Lee, and Z. H. Cho. Parallelization of the em algorithm for 3-d pet image reconstruction. *IEEE Trans. on Med. Imaging*, 10(4):513–521, 1991.

[6] K. Dowd and C. R. Severance. *High Performance Computing*. O'Reilly, Sebastopol, CA, 1998.

[7] J. Fernandez, A. Canas, A. F. Dias, J. Gonzalez, J. Ortega, and A. Prieto. Performance of message-passing matlab toolboxes. In J. P. et al, editor, *VECPAR 2002*, pages 228–242, 2003.

[8] J. A. Fessler and S. D. Booth. Conjugate-gradient preconditioning methods for shift-variant pet image reconstruction. *IEEE Trans. on Image Proc.*, 8(5):688–699, 1999.

[9] J. A. Fessler, E. P. Ficaro, N. H. Clinthorne, and K. Lange. Grouped-coordinate ascent algorithms for penalized-likelihood transmission reconstruction. *IEEE Trans. on Medical Imaging*, 16(2):166–175, 1997.

[10] J. A. Fessler and B. P. Sutton. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing*, 51(2):560–574, 2003.

[11] S. Helgason. *The Radon Transform*. Birkhauser, Boston, 2nd edition, 1999.

[12] J. Hsieh. Adaptive streak artifact reduction in computed tomography resulting from excessive x-ray photon noise. *Med. Phys.*, 25(11):2139–2147, 1998.

[13] C. Johnson and A. Sofer. A data-parallel algorithm for iterative tomographic image reconstruction. In *Proceedings of the 7th Symposium on the Frontiers of Massively Parallel Computation*, pages 126–137. IEEE Computer Society Press, 1999.

[14] B. D. Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. Reduction of metal streak artifacts in x-ray computed tomography using a transmission maximum a posteriori algorithm. *IEEE Trans. on Nuclear Science*, 47(3):977–981, 2000.

[15] E. U. Mumcouglu, R. M. Leahy, and S. R. Cherry. Bayesian reconstruction of pet images: methodology and performance analysis. *Phys. Med. Biol.*, 41:1777–1807, 1996.

[16] F. Natterer. *The Mathematics of Computerized Tomography*. Wiley, Chirchester, 1986.

[17] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Transactions on Signal Processing*, 41(2):534–548, 1993.

[18] S. Smallen. On-line parallel tomography. Master's thesis, University of California, San Diego, 2001.

[19] T. Sterling, editor. *Beowulf Cluster Computing with Linux*. MIT Press, Cambridge, MA, 2002.

[20] G. Wang, T. Frei, and M. W. Vannier. Fast iterative algorithm for metal artifact reduction in x-ray ct. *Acad. Radiol.*, 7:607–614, 2000.