

Parallel Computing Methods for X-Ray Cone Beam Tomography with Large Array Sizes

David A. Reimann^{1,2,3}, Michael J. Flynn², Vipin Chaudhary⁴, and Ishwar K. Sethi³

1 - Department Mathematics, Albion College, Albion, Michigan

2 - Department of Diagnostic Radiology, Henry Ford Health System, Detroit

3 - Computer Science Department, Wayne State University, Detroit

4 - Electrical and Computer Engineering Department, Wayne State University, Detroit

Abstract—Cone beam geometries are increasingly of interest for x-ray CT applications to improve imaging efficiency. In this paper, we describe our practical experience implementing circular orbit cone beam backprojection on workstation clusters. The reconstruction problem is computationally intensive, particularly for arrays of 512 voxels in each direction. A voxel driven approach is described where the reconstruction volume is partitioned into variable width slabs and each slab given to a workstation. Each projection is filtered by one workstation and then sent to the others for backprojection. While most computation is done in the backprojection step, a significant amount of time must be spent in sending projective data. A method is detailed to further reduce the communication overhead by restricting the amount of projection sent to only what is required by each backprojecting workstation. Furthermore, if the shape of the backprojection slabs is made as square as possible, the total communication requirement can be minimized. By the reduction of communication requirement, an overall improvement in processor utilization was observed, and the crossover point where communications dominates was improved.

Introduction

Cone beam geometries are increasingly of interest for x-ray CT applications. The use of a cone beam geometry allows more efficient x-ray tube heat loading and simpler mechanics (rotate only). The Feldkamp [1] algorithm has been shown to provide superior resolution over both the Grangeat and VOIR methods [2]. Despite the incomplete sampling of the Feldkamp algorithm, it has been shown to have good resolution for small cone angles. Because of the computational efficiency and high quality of reconstructions obtained in our laboratory [3], we have used the Feldkamp algorithm exclusively. Computational requirements are particularly important in our work where we are interested in reconstructing arrays from 256^3 to 512^3 voxels.

We have developed a system for 3D cone beam computed tomography, consisting of a microfocus x-ray source and x-ray image intensifier coupled to a CCD camera [3]. The application of the principles of CT at a microscopic level, or microtomography, allows quantitative investigation of objects in three dimensions. The concept of microtomography has existed since the early development of CT [4, 5], however, only recently have practical systems been developed. Applications for this technology are primarily in the biological and material sciences. Visibly opaque calcified tissues such as bones and teeth are an excellent match for x-ray microtomography systems

and many investigations have been reported [6, 7, 8, 9, 10, 11]. Imaging applications in material microstructure have also been reported [12, 13]. The full width at half maximum resolving power of our system has been experimentally measured to be $70 \mu\text{m}$ when imaging 10 mm diameter objects. The 3D nature of the resulting image data can be used to visualize internal structure and compute parameters such as volume, surface area, and surface/volume orientation.

The large computation time for reconstruction algorithms is the major factor limiting reconstruction size. Dedicated hardware implementations have been proposed [14, 15, 16, 17, 18, 19, 20, 21, 22, 23]. A dedicated hardware chip for backprojection has been recently developed which can backproject a 512×512 image in 4 seconds [16]. Many of these chips can be used in parallel to further reduce reconstruction times. It is doubtful if any of these systems are currently in use, either because of newer hardware solutions, or the hardware is not yet fully integrated. Furthermore, none of these systems was developed for cone beam tomography. Another drawback to a dedicated hardware system is the relatively low market demand and the lack of scalability of such a system. Another active area of research has been in the area of employing commercially available parallel systems to the reconstruction problem. Several investigators have evaluated the Transputer [24, 25, 26, 27, 28]. The MasPar has been used [29, 30, 31]. One investigator used a Cray-1 [32]. The iPSC/2 has been used [33, 34, 35, 36]. Still others have used other specialized multiprocessors [37, 38, 39, 40]. Parallel efficiencies of greater than 88% have been reported [33]. These systems require a large investment (\$50,000 – \$300,00) compared with general purpose workstations.

Our previous work has shown greater than 80% processor utilization when reconstructing a 256^3 volume from 100 views of 256^2 pixels each by balancing load among 6 workstations [41]. We improved this to greater than 90% processor utilization by performing asynchronous communication [42]. We have used a voxel driven approach where the reconstruction volume is partitioned into variable width slabs and each slab given to a workstation. Each projection is filtered by one workstation and then sent to the others for backprojection.

In this paper, we describe our practical experience implementing circular orbit cone beam backprojection. Results using clusters of workstations and SMP's have been obtained. We have used MPI [43] as the parallel computing library. In particular, we detail a method to further reduce the communication overhead by restricting the amount of projection sent to only what is required by each backprojecting workstation. Furthermore, if the

shape of the backprojection slabs is made as square as possible, the total communication requirement can be minimized.

Methods

The cone beam tomography problem has a high degree of independent operations, and is therefore well suited to parallel implementations. The major nontrivial aspect of the problem is the data sizes needed. An acquisition can be as large as 800 views of 512^2 16 bit integer images. These projectional images are used to create a volumetric data set on the order of 512^3 32 bit floating point voxels. The use of 32 bit floating point voxels is required to provide sufficient accuracy in the result. The memory requirement to store the entire 512^3 32 bit volume is 512 MB, which is at the upper limit of available memory on most tightly coupled parallel systems currently available. However, dividing this memory among 16 workstations requires only 32 MB per workstation. The large amount of memory must be considered in the implementation.

A voxel driven approach taken was used, where the volume is distributed over several workstations and each view is sent to the voxel workstations. Each workstation sees every projection, but only a small subset of the reconstructed voxels. The voxel driven Feldkamp algorithm can be pseudo-coded as

- 1 Initialize each PE
- 2 Read and Broadcast problem specifications
- 3 Partition and allocate memory
- 4 Allocate memory
- 5 Precomputation of weights
- 6 for each θ (N_θ views):
- 7 if (PE is ROOT)
- 8 Read Projection
- 9 Weight and Filter Projection
- 10 Broadcast Projection $P_\theta^*(u, v)$
- 11 Backproject Projection
- 12 Gather Reconstructed Voxels on ROOT PE

where each workstation is represented by a processing element (PE). The total memory required for this implementation is approximately equal to the total number of voxels. One advantage of this method is that the data is acquired in a serial fashion and processing could be done in concert with acquisition.

While most computation is done in the backprojection step, a significant amount of time must be spent in broadcasting the projection. Overlapping computation and communication allows higher utilization for low numbers of processors. Because the communication time only increases and backprojection time decreases as the number of PE's increase, eventually there comes a point where communication time dominates. This is a function of both the communication channel speed and total computational power.

By partitioning the volume into slabs, one really only needs to be concerned with the projection of the slab. There is much wasted communications when a full projection is sent, and only a portion is used. By restricting the communication to send only the portion of the physical projection which is required, this waste can be avoided, as shown in Figure 1. The extent of the projection required to backproject into the subvolume can be computed by taking the minimum and maximum of the horizontal positions where the four corners of the subvolume project.

A further reduction in communication can be made if the slabs are made as close to squares as possible. This minimizes the total projection over all angles, and therefore minimizes the communication requirement. Again, only a piece of the projection is needed in the reconstruction. Partitions can be made squarer by combining adjacent slabs and re-dividing perpendicular to their original cuts as long as increasing the size improves the overall squareness, as shown in Figure 2.




Figure 1: *Reconstruction Column Projection.* A column of data in the reconstruction volume projects only on a subset of each projection. By changing from slabs to columns and sending the required data, a significant reduction in the communications can be made.




Figure 2: *Squarish Partitioning of Reconstruction Volume.* The overall squareness of the partitioned subvolumes is improved, reducing the total communications requirements. An example of repartitioning the xy plane is shown.

To help solve the problem of portability of parallel programs, a group of computer vendors, software engineers, and parallel applications scientists developed the

Message Passing Interface (MPI) Library. The MPI library, or simply MPI, was designed to take the useful feature of other parallel libraries and provide a common programming interface. The portability of this library allows excellent support for comparisons among various hardware architectures and between hardware architectures and theoretic models. The goal of MPI was to define precise semantics and bindings to allow third party implementation of the standard. The extensibility of the standard allows future improvements. The Feldkamp algorithm was implemented using the freely available MPICH implementation of MPI.

Results

A volume of 256^3 was reconstructed from 200 views of 680×572 pixels each. A slab of memory is allocated to each processor as shown in Figure 3. The root node provides all the filtering of the projections and writing of the final result. Computations were done on a DEC AlphaServer 2000 4/233. This machine is an SMP with 2 processors, 256 MB of RAM, and a 2 GB disk. The serial version ran in 1713.3 seconds and the parallel version ran in 1063.2 seconds. The speedup is therefore 1.61 and the efficiency is 80.6%. Using this architecture, it is feasible to reconstruct large volumes in reasonable times.

Discussion

Making the slabs as square as possible further reduces the overhead compared to long and narrow slabs. However, optimal packing of squares is related to other packing problems which have been shown to be unsolvable in polynomial time. Our partitioning method should be useful in other 3D imaging applications such as volume rendering. The partitioning presented is also applicable to pixel driven planar tomography.

A future goal is to fully develop a theoretic model of the voxel and ray driven backprojection algorithms. A comparison of theory with actual timings would give insight into architecture specific problems one may encounter on various systems.

Acknowledgements

Partial support of this work was provided by NIH Grant RO1-AR42101.

References

- [1] L. A. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1:612–19, June 1984.
- [2] Nicolas J. Dusaussay. VOIR: A volumetric image reconstruction algorithm base on fourier techniques for inversion of the 3-D radon transform. *IEEE Transactions on Image Processing*, 5(1):121–131, January 1996.
- [3] David A. Reimann, Michael J. Flynn, and Sean M. Hames. A flexible laboratory system for 3D x-ray microtomography of 3–50 mm specimens. In *3D Microscopy: Image Acquisition and Processing 1995*, Proceedings of the SPIE 2412, pages 186–195, San Jose, California, 5–10 February 1995.
- [4] Steve Webb. *From the Watching of the Shadows: The Origins of Radiological Tomography*. Adam Hilger, Bristol, England, 1990.
- [5] William H. Oldendorf. Some possible applications of computerized tomography in pathology. *Journal of Computer Assisted Tomography*, 4(2):141–144, April 1980.
- [6] Ulrich Bonse, Frank Busch, Olaf Günnewig, Felix Beckmann, Reinhard Pahl, Günter Delling, Michael Hahn,

and Walter Graeff. 3D computed x-ray tomography of human cancellous bone at $8 \mu\text{m}$ spatial and 10^{-4} energy resolution. *Bone and Mineral*, 25:25–38, 1994.

- [7] J.C. Elliott, P. Anderson, X.J. Gao, F.S.L. Wong, G.R. Davis, and S.E.P. Dowker. Application of scanning microradiography and x-ray microtomography to studies of bones and teeth. *Journal of X-Ray Science and Technology*, 4(2):102–117, 1994.
- [8] K. Engelke, W. Graeff, L. Meiss, M. Hahn, and G. Delling. High spatial resolution imaging of bone mineral using computed microtomography. *Invest. Rad.*, 28(4):341–349, 1993.
- [9] Lee A. Feldkamp, Steven A. Goldstein, A. Michael Parfitt, Gerald Jession, and Michael Kleerekoper. The direct examination of three-dimensional bone architecture in vitro by computed tomography. *Journal of Bone and Mineral Research*, 4(1), 1989.
- [10] Steven A. Goldstein, Robert Goulet, and Doris McCubrey. Measurement and significance of three-dimensional architecture to the mechanical integrity of trabecular bone. *Calcified Tissue International*, 53(Suppl 1):S127–S133, 1993.
- [11] S. Yasumura, K. Jones, P. Spanne, G. Schidlovsky, L. Wielopolski, X. Ren, D. Glaros, and Y. Xatzikonstantinou. In vivo animal models of body composition in aging. *Journal of Nutrition*, 123(2 Suppl.):459–464, February 1993.
- [12] J.H. Kinney, N. Lane, S. Majumdar, S.J. Marshall, and G.W. Marshall, Jr. Noninvasive three-dimensional histomorphometry using x-ray tomographic microscopy. *Journal of Bone and Mineral Research*, 7(Supplement 1):S136, August 1992.
- [13] J.C. Russ. Imaging the three-dimensional structure of materials. In et al. C.S. Barrett, editor, *Advances in X-ray Analysis*, volume 35, pages 1219–1227. Plenum Press, 1992.
- [14] B.K. Gilbert, R.A. Robb, and L.M. Krueger. Ultra high-speed reconstruction processors for x-ray computed tomography of the heart and circulation. In Morio Onoe, Kendall Preston, Jr., and Azriel Rosenfeld, editors, *Real-Time Medical Image Processing*, pages 23–40. Plenum Press, New York, 1980.
- [15] E. Di Sciascio, R. Guzzardi, and D. Marino. Proposal of a real-time reconstruction processor for 3D positron emission tomography. In *Conference Record of the 1992 Nuclear Science Symposium and Medical Imaging Conference*, volume 2, pages 921–923, Orlando, Florida USA, 25–31 October 1992.
- [16] Iskender Agi, Paul J. Hurst, and K. Wayne Current. An image processing IC for backprojection and spatial histogramming in a pipelined array. *IEEE Journal of Solid-State Circuits*, 28(3):210–221, March 1993.
- [17] Smil Ruhman and Isaac Scherson. Associative processor for tomographic image reconstruction. In *MEDCOMP '82 Proceedings*, pages 353–358, Philadelphia, Pennsylvania USA, 23–25 September 1982. IEEE.
- [18] H. Wani and H. Ishihara. Real-time image processing in CT—convolver and back projector. In Morio Onoe, Kendall Preston, Jr., and Azriel Rosenfeld, editors, *Real-Time Medical Image Processing*, pages 99–106. Plenum Press, New York, 1980.
- [19] Sally L. Wood. Efficient MVE image reconstruction for arbitrary measurement geometries. In *ICASSP 82 Proceedings*, volume 2, pages 1158–1161, Paris, France, 3–5 May 1982. IEEE.

- [20] N.A. Wilkinson, M.S. Atkins, and J.G. Rogers. A tomograph VMEbus parallel processing data acquisition system. *IEEE Transactions on Nuclear Science*, 36(1):1047–1051, February 1989.
- [21] W.F. Jones, L.G. Byars, and M.E. Casey. Design of a super fast three-dimensional projection system for positron emission tomography. *IEEE Transactions on Nuclear Science*, 37(2):800–804, April 1990.
- [22] C.J. Thompson and T.M. Peters. A fractional address accumulator for fast back-projection. *IEEE Transactions on Nuclear Science*, 28(4):3648–3650, August 1981.
- [23] Eric Sheih, K. Wayne Current, Paul J. Hurst, and Iskendar Agi. High-speed computation of the radon transform and backprojection using an expandable multiprocessor architecture. *IEEE Transactions on Circuits and Systems for Video Technology*, 2(4):347–360, December 1992.
- [24] M. Stella Atkins, Donald Murray, and Ronald Harrop. Use of transputers in a 3-D positron emission tomograph. *IEEE Transactions on Medical Imaging*, 10(3):276–283, September 1991.
- [25] N. Kingswood, E.L. Dagless, R.M. Belchamber, D. Betteridge, T. Lilley, and J.D.M. Roberts. Image reconstruction using the transputer. *IEE Proceedings*, 133, Part E(3):139–144, May 1986.
- [26] S. Barresi, D. Bollini, and A. Del Guerra. Use of a transputer system for fast 3-D image reconstruction in 3-D PET. *IEEE Transactions on Nuclear Science*, 37(2):812–816, April 1990.
- [27] C. Comtat, C. Morel, M. Defrise, and D.W. Townsend. The FAVOR algorithm for 3D PET data and its implementation using a network of transputers. *Physics in Medicine and Biology*, 38:929–944, 1993.
- [28] K.A. Girodias, H.H. Barrett, and R.L. Shoemaker. Parallel simulated annealing for emission tomography. *Physics in Medicine and Biology*, 36(7):921–938, 1991.
- [29] C.S. Butler and M.I. Miller. Maximum *a posteriori* estimation for SPECT using regularization techniques on massively parallel computers. *IEEE Transactions on Medical Imaging*, 12(1):84–89, March 1993.
- [30] C.S. Butler, M.I. Miller, T.R. Miller, and J.W. Wallis. Massively parallel computers for 3D single-photon-emission computed tomography. *Physics in Medicine and Biology*, 39:575–582, 1994.
- [31] M.I. Miller and C.S. Butler. 3-D maximum *a posteriori* estimation for single photon emission computed tomography on massively-parallel computers. *IEEE Transactions on Medical Imaging*, 12(3):560–565, September 1993.
- [32] Linda Kaufman. Implementing and accelerating the EM algorithm for positron emission tomography. *IEEE Transactions on Medical Imaging*, 6(1):37–51, March 1987.
- [33] C.M. Chen, S.-Y. Lee, and Z.H. Cho. A parallel implementation of 3-D CT image reconstruction on hypercube multiprocessor. *IEEE Transactions on Nuclear Science*, 37(3):1333–1346, June 1990.
- [34] C.M. Chen, S.-Y. Lee, and Z.H. Cho. Parallelization of the EM algorithm for 3-D PET image reconstruction. *IEEE Transactions on Medical Imaging*, 10(4):513–522, December 1991.
- [35] Chung-Ming Chen and Soo-Young Lee. On parallelizing the EM algorithm for PET image reconstruction. *IEEE Transactions on Parallel and Distributed Systems*, 5(8):860–873, August 1994.
- [36] Henri-Pierre Charles, Jian-Jin Li, and Serge Miguet. 3D image processing on distributed memory parallel computers. In Raj S. Acharya and Dmitry B. Goldgof, editors, *Biomedical Image Processing and Biomedical Visualization*, Proceedings of the SPIE 1905, pages 379–390, San Jose, California USA, 1–4 February 1993.
- [37] M. Stella Atkins, M. Zastre, K. Buckley, and L. Byars. Evaluation of the use of the i860 supercard in a 3-D PET tomograph. In *Conference Record of the 1992 IEEE Nuclear Science Symposium and Medical Imaging Conference*, volume 2, pages 913–914, Orlando, Florida USA, 25–31 October 1992.
- [38] T.M. Guerrero, S.R. Cherry, M. Dahlbom, A.R. Ricci, and E.J. Hoffman. Fast implementations of 3D PET reconstruction using vector and parallel programming techniques. In *Conference Record of the 1992 IEEE Nuclear Science Symposium and Medical Imaging Conference*, volume 2, pages 969–971, Orlando, Florida USA, 25–31 October 1992.
- [39] T.M. Guerrero, A.R. Ricci, M. Dahlbom, S.R. Cherry, and E.J. Hoffman. Parallel image reconstruction for 3D positron emission tomography from incomplete 2D projection data. In Raj S. Acharya and Dmitry B. Goldgof, editors, *Biomedical Image Processing and Biomedical Visualization*, Proceedings of the SPIE 1905, pages 978–986, San Jose, California USA, 1–4 February 1993.
- [40] Michael I. Miller and Badrinath Roysam. 3-D maximum *a posteriori* estimation for single photon emission computed tomography on massively-parallel computers. *Proc. Natl. Acad. Sci USA*, 88:3223–3227, 1991.
- [41] David A. Reimann, Vipin Chaudhary, Michael J. Flynn, and Ishwar K. Sethi. Parallel implementation of cone beam tomography. In A. Bojanczyk, editor, *Proceedings of the 1996 International Conference on Parallel Processing*, volume II, pages 170–173, Bloomington, Illinois, 12–16 August 1996.
- [42] David A. Reimann, Vipin Chaudhary, Michael J. Flynn, and Ishwar K. Sethi. Cone beam tomography using MPI on heterogeneous workstation clusters. In *Proceedings, Second MPI Developer's Conference*, pages 142–148, University of Notre Dame, Notre Dame, Indiana, 1–2 July 1996. IEEE Computer Society Press.
- [43] Message Passing Interface Forum. MPI: A message-passing interface standard. Technical report, 12 June 1995. <http://www.mcs.anl.gov/mpi/mpi-report-1/mpi-report.html>.