Data Processing Methods for a High Throughput Brain Imaging PET Research Center

Judson P. Jones, *Member, IEEE*, Arman Rahmim, *Member, IEEE*, Merence Sibomana, Andrew Crabb, Ziad Burbar, *Member, IEEE*, Charles B. Cavanaugh, Christian Michel, *Member, IEEE*, Dean F. Wong

***Abstract*– We describe a computer system designed to meet the data processing needs of a high-volume brain PET research center based on the High Resolution Research Tomograph (HRRT). Listmode data are collected by an acquisition computer and stored on a high-speed disk. A workflow management program transfers the data through a gigabit network, rebins events into sinograms, and calculates correction factors. Reconstruction jobs are performed on a 64 processor cluster. We developed methods for dynamically allocating subclusters from the pool of available nodes, and reconstructing multiple images on multiple subclusters simultaneously. We also studied overall workflow. In our initial plan, scatter and randoms calculation unexpectedly became a bottleneck. We therefore adjusted our plan so that scatter estimation was performed initially in low resolution, and later expanded to high resolution.**

I. Introduction

THE scientific roles of PET in neuroscience and neuropsychiatry are exponentially increasing with the global use of neuroreceptor and other functional brain PET for the study of normal brain physiology [1] and neuropathological conditions [2]–[4]. Also PET is now well accepted as an element of the infrastructure for drug discovery [5]–[7], and has attracted the attention of major pharmaceutical/industrial investments in PET imaging. Due to the special anatomical properties of the brain, high spatial resolution coupled with a large axial field of view is of central importance for a high-performance brain PET camera. These considerations led to the development of the High Resolution Research Tomograph (HRRT), currently in use at Johns Hopkins and other centers.

Here we describe a data processing system designed for a high-volume brain PET research center based on the HRRT. Each typical study consists of 30 frames of high resolution data, which, including estimated scatter and randoms sinograms, represents approximately 140 GB per study (2.1 TB per week at 15 studies/wk) of sinogram data for reconstruction. Fig. 1 illustrates horizontal, coronal, and sagittal slices from a typical continuous infusion C11-Raclopride study of a small (4 kg) monkey conducted on the HRRT. Scan time for frame 18 was 5 minutes (55-60min) within a 90-minute dynamic 30-frame reconstruction protocol. Clear separation of the caudate and putamen nuclei in this study demonstrates the high spatial resolution of the HRRT.

Fig. 1. Horizontal, coronal, and saggital, sections from a typical C-11 raclopride study of a small (4 kg) monkey. Clear separation of the caudate and putamen nuclei illustrates the resolving power of the HRRT. Images courtesy of M. Weed.

II. Data Processing for the HRRT

The HRRT gantry contains 8 rectangular "panel array" planar heads, arranged in an octagon surrounding a 312 mm in-plane and 252 mm axial field of view. Each head bears an array of 9x13 blocks of 8x8x2 LSO/LYSO crystals spaced on a 2.1 mm pitch utilizing quadrant-shared 19 mm photomultipliers; there are a total of 119,808 scintilllator crystals in the camera. For details, see [8]. Fig. 2 illustrates an outline of the methods used for processing HRRT data.

A. Normalization

A normalization scan is performed approximately once every three months using a rotating rod "direct normalization" method [9]. Line of response (LOR) coincidence events, stored in list mode format, are initially rebinned into a 256x288x6367 "span-3" sinogram and further processed into a 256x288x2209 "span-9" sinogram. These sinograms are modulated by both the dwell of the axial rebinning (span 3 or 9, max ring difference 67) and by the rebinning from LOR to equidistant parallel projection space using nearest neighbor approximation. The two normalization corrections (span 3 and 9) are obtained after inverting those sinograms corrected for the dwell of the rotating rod (16 cm radius). These two normalizations (floating point sinograms) are used in all

Manuscript received November 17, 2006. We thank M. Weed for the images illustrated in Fig. 1. This work was supported in part by the National Institutes of Health under grants S10RR017219, RO1AA12839, K24DA00412, NS38927, and U01MH075378.

J. P. Jones, M. Sibomana, Z. Burbar, and C. Michel are with Siemens Medical Solutions, Molecular Imaging, Knoxville, TN 37932 USA (telephone. 865-218-2389, e-mail Judson.Jones@Siemens.com).

A. Rahmim, A. Crabb, and D. F. Wong are with the Johns Hopkins Medical Institutions, Baltimore, MD 21287 USA (telephone 410-955-8433, e-mail ARahmim1@jhmi.edu).

C. B. Cavanaugh is with IBM Corp. Dallas, TX 75234 USA (telephone 214-233-3429, e-mail cbc@us.ibm.com)

subsequent reconstructions. Presently the dead-time correction is accounted for post-reconstruction.

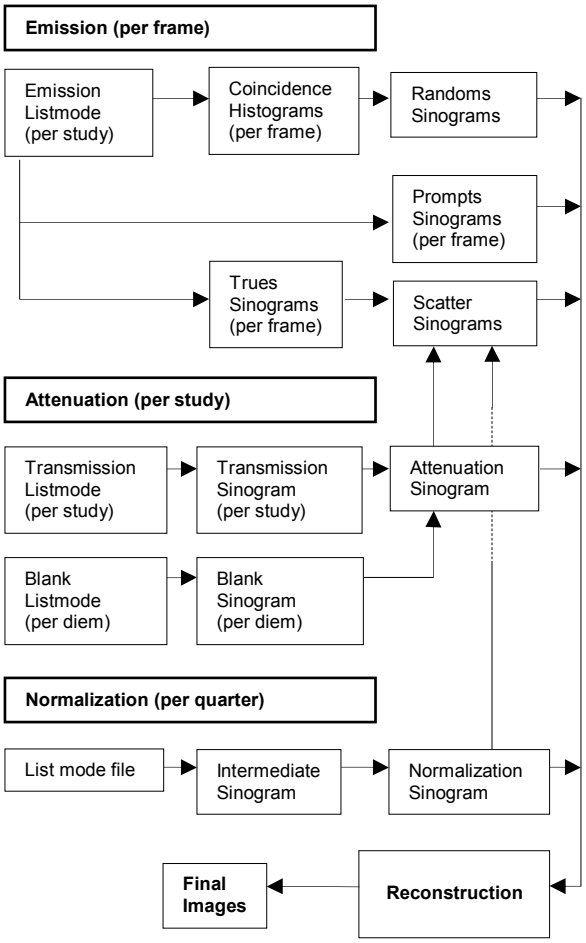

Fig. 2 Data processing chart for HRRT image reconstruction.

B. Attenuation

The processing chain for attenuation correction consists of a 2D blank scan performed once per day using a scanning 137 Cs point source. For each study, a 2D transmission scan is performed using the same point source. Blank and transmission scans are further processed by a Maximum Likelihood dedicated transmission algorithm including smoothness and intensity priors to produce the attenuation image (mu map) at 511 keV [10]. A mock scan from a virtual shifted source is automatically subtracted from the transmission scans when created which corrects for contaminating emission and scatter [11].

C. Emission

Emission data are initially collected into a listmode file and processed into time frames. Three sinograms are associated with each frame: prompt, estimated random and true coincidences sinograms. The estimated (variance reduced) random sinogram is obtained from histogramming the measured delayed coincidences into singles per detector, estimating the total singles per detector and applying the random rate equation [12]. The true sinogram is made from prompt and measured random and is used only to estimate the scatter correction term [13].

D. Reconstruction

The five sinograms which result from the above preprocessing steps (normalization, attenuation, scatter, prompts, and randoms) are reconstructed using the Ordinary Poisson variant of 3D OSEM (OP-OSEM) [14], [15]. The resulting images constitute the final product of the data processing chain; typically no post-reconstruction spatial filters are used. Decay, deadtime and quantitation information are added subsequently.

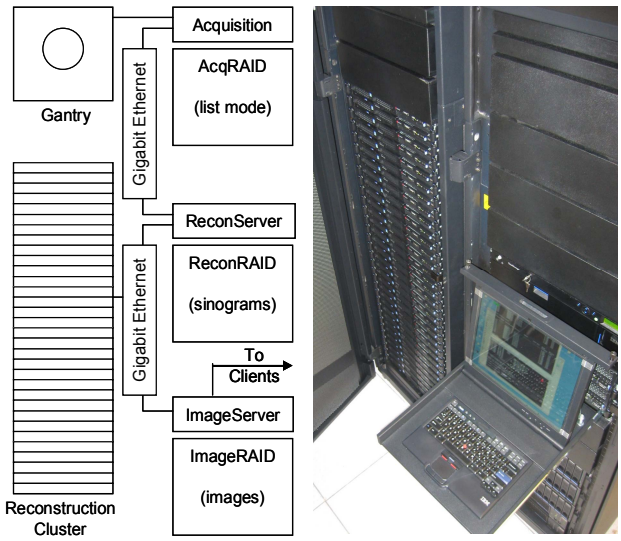

Fig. 3 (A) General computer architecture of the HRRT data processing system. (B) Photograph of the installed system. The 32-node (64 processor) reconstruction cluster is contained in the left rack; the ReconServer, the ImageServer, and the mass storage RAIDs are contained in the right rack. The acquisition computer is not illustrated.

III. Computer Architecture & Clustering

A. General Computer Architecture

The HRRT data processing system consists of three principal computers used for data acquisition, preprocessing and process control, and serving images. These principal computers are assisted by 32 compute nodes, used exclusively for calculation. The *acquisition computer* is connected directly to the HRRT gantry and contains a single 2.4 GHz Intel Xeon, 2 Gbytes memory, gigabit Ethernet (GBE), and a high-speed 1.0 Terabyte Level 0 RAID. The acquisition computer collects listmode events from the gantry and commits them to disk. Once the acquisition is complete, a workflow management program transfers listmode files through the GBE to the *ReconServer*, an IBM xSeries 346 with dual 3.4 GHz Intel Xeons, 3 GB memory, and two 1.3 TB RAIDs. The ReconServer rebins the events into the appropriate sinograms, and executes the relevant preprocessing steps. When the preprocessing is completed, reconstruction jobs are submitted to a queue, and are processed by a 32 node (64 processor) reconstruction cluster composed of IBM xSeries 336 1U

servers with dual 3.4 GHz Intel Xeon processors, 2.0 GB memory, GBE, and 40 GB local disks. Reconstructed images are placed on the *ImageServer*, which has the same specifications as the ReconServer, and made available to the center's clientele via the hospital network.

B. Clustering

Image reconstruction represents the most computationally challenging component of the data processing chain. Our approach to parallelism for dynamic PET in this system is based on a two-level model of parallelism – at a lower level, each reconstruction is performed on multiple compute nodes in parallel, and at a higher level, many frames are reconstructed in parallel on multiple "subclusters" within the larger cluster.

Our approach to parallelizing reconstruction for individual frames is described in [16]: identical programs run on all participating nodes, data are distributed so that the computation is balanced, and internode synchronization and data sharing are accomplished through message passing. Sinogram data are requested from the ReconServer when they are needed, and reconstructed images are deposited onto the server when they are ready. Load balancing is accomplished by assigning "views" of the sinogram data onto nodes using a round-robin distribution scheme. This insures that regardless of the number of nodes or the number of subsets chosen, the maximum computational imbalance will be no more than the load associated with the smallest view.

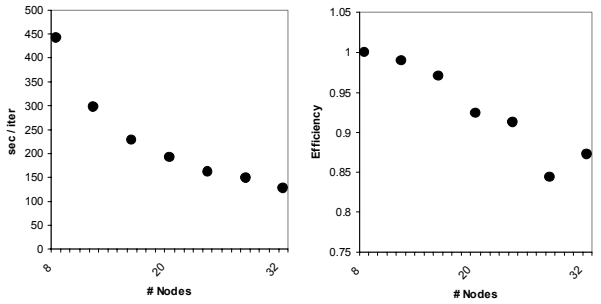

A **B**

Fig. 4 (A) Ordinary Poisson OSEM3D reconstruction time per iteration for HRRT span-3 data as a function of the number of nodes in a cluster. Speedup is obtained through 32 nodes. (B) Efficiency of the calculation relative to the 8-node case. High efficiency can be retained by reconstructing multiple frames in parallel using *multicluster* methods.

C. Multiclustering

Due to the requirement for global communication, parallel tomography does not scale onto large numbers of nodes as well as problems that require only local communication. Fig. 4 illustrates the time required per OSEM iteration as a function of the number of nodes, and the efficiency (relative to the 8 node case) of the calculation on clusters of that size. While we obtain some speedup on clusters of all sizes, clearly there is a significant loss of efficiency at larger scales. To retain efficiency in the current system we have therefore developed methods for dynamically allocating *subclusters* from the pool of available nodes, and processing reconstruction jobs on multiple subclusters simultaneously. The choice for the size of the subcluster allocated for each reconstruction job is bounded from above by efficiency constraints and the number of available nodes, and from below by the memory requirements of the job.

Building on previously established infrastructure [17], the main challenges for multicluster computing were (1) managing the pool of available nodes; (2) adaptive and dynamic allocation of subclusters; (3) managing resources in the event of failures; and (4) managing resources contention during parallel distributed I/O. These problems were addressed through the development of a queuing subsystem specialized for multiclustering which includes two thread-safe databases, one for compute nodes, and one for I/O ports. The queuing system counts the number of frames to be reconstructed, allocates subclusters from the available pool of nodes, and manages the execution of the reconstructions. The queue also provides *fault tolerance* by detecting and avoiding single node failures and subcluster failures.

IV. Workflow Analysis

To understand the performance of the system as a whole, we performed a workflow analysis for a typical study, which consists of 30 frames of span-3 data (see Fig. 5). Following the initial assembly of the system, random and scatter corrections were performed in span-3 resolution using a single-processor, resulting in 28 min. of preprocessing per frame. We found that dedicating 8 nodes per job resulted in optimum performance of 15.9 hours/study of 30-frames, or roughly 31 minutes per frame. Thus, the preprocessing steps became a bottleneck in the whole scheme. We therefore adjusted our approach so that scatter correction is calculated in span 9 resolution, and subsequently expanded to span 3, as justified due to the low frequency content of the scattering.

The performance of the remaining steps is as follows:

1) Histogramming. List-mode data are read from the acquisition RAID through the gigabit ethernet and histogrammed into a sequence of dynamic frames. Most common framing sequences typically produce 30 frames. Currently, for a typical 40 GB dataset, histogramming requires ~1.5 min/frame corresponding to a total of 30 frames x 1.5 min/frame = 45 minutes.

2) Attenuation estimation: This step involves estimation of the amount of attenuation of emitted events due to interaction with patient tissue and bones. This step needs to be performed only once for the entire study, and therefore is not time-limiting at all (processing time: 13 minutes).

3) Estimation of scattered and random events. Scatter and randoms estimation require 28 minutes per frame and are not currently parallelized.

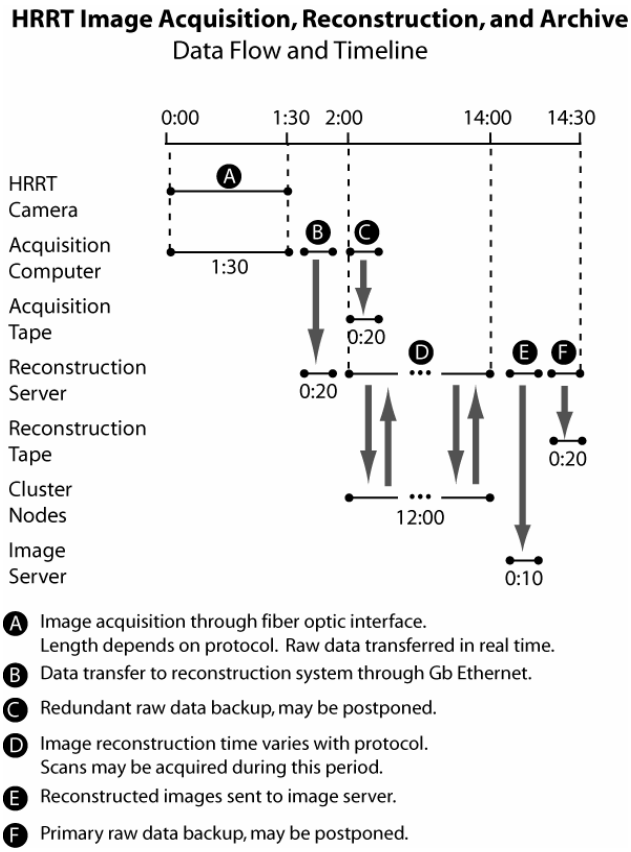

Fig. 5. Workflow analysis. This chart depicts time in hh:mm for a typical 30 frame dynamic study.

V. Discussion

The HRRT produces unprecedented quantities of emission data and requires a complicated processing scheme to produce high-resolution, quantitatively accurate images. The system described here for processing these data consists of 69 general purpose processors housed in 35 separate computers, with almost 8 Terabytes of available disk space. Currently, 32 of these computers (64 processors) are used exclusively for OSEM-3D; other calculations are performed primarily on the ReconServer. The resulting calculation is nearly balanced.

However, anticipated improvements in the OSEM-3D algorithm [18] as well as a need to reduce disk I/O time and network contention will drive further optimization and a redistribution of tasks within the data processing system. We have already developed a prototype parallel randoms smoothing algorithm and are currently developing a parallel method for creating attenuation sinograms from mu-maps; these will be integrated into the OSEM-3D application at a later date. We can also further reduce the I/O burden on the ReconServer by communicating only scatter segment 0 and axial scatter scaling factors, rather than a complete 3D scatter sinogram. These gains in efficiency are likely to be matched by new demands on performance: a new rate dependent component-based normalization technique [19] will require calculation and communication of unique normalization data for each frame, and a new image-based resolution recovery technique [20] will require more OSEM-3D iterations [21] to be performed for each frame. Since our approach to high-throughput data processing employs general purpose computing technology rather than special-purpose hardware, adapting to innovative new algorithms and reassigning tasks to different machines is not as difficult as it would otherwise be.

References

[1] C. Munro, M. McCaul, D. F. Wong, L. Oswald, Y. Zhou, H. Kuwabara, L. Choi, J. Brasic, G. Wand, "Sex differences in striatal dopamine release in healthy adults." *Biological Psychiatry*, 2006. *In press.*

[2] D. F. Wong, "In vivo imaging of D2 dopamine receptors in schizophrenia: The ups and downs of neuroimaging receptors." *Arch. Gen. Psychiatry*, 59: 31-34, 2002.

[3] D. F. Wong, "Neurotoxin effects of drug abuse: Imaging and mechanisms." In: *Cell Biology of Addiction*, B.K. Madras, C. M. Colvis, J. D. Pollock, J. L. Rutter, D. Shurtleff, M. von Zastrow (eds), Cold Spring Harbor Laboratory Press, 2006.

[4] D. F. Wong, H. Kuwabara, D. J. Schretlen, K. R. Bonson, Y. Zhou, A. Nandi, J. R. Brai, A. S. Kimes, M. A. Maris, A. Kumar, C. Contoreggi, J. Links, M. Ernst, O. Rousset, S. Zukin, A. A. Grace, C. Rohde, D. R. Jasinski, A. Gjdde, E. D. London, "Increased occupancy of dopamine receptors in human striatum during cue-induced cocaine craving.", *Neuropsychopharmacology*, 31: 2716-2727, 2006.

[5] D. F. Wong, M. G. Pomper, "Predicting the success of a radiopharmaceutical for in vivo imaging of central nervous system neuroreceptor systems." *Molecular Imaging and Biology* 5(6): 350-362, 2003.

[6] D. F. Wong, W. Z. Potter, J. R. Brasic, "Proof of concept: Functional models for drug development in humans." In: *Neuropsychopharmacology: The Fifth Generation of Progress*, K. L. Davis, D. Charney, J. T. Coyle, C. Nemeroff (eds), Lippincott Williams & Wilkins, Baltimore, MD, pp. 457-473, 2002.

[7] G. Grunder, A. Carlsson, D. F. Wong, "Mechanism of new antipsychotic medications: Occupancy is not just antagonism." *Archives of General Psychiatry* 60: 974-977, 2003.

[8] K. Wienhard, M. Schmand, M. E. Casey, K. Baker, J. Bao, L. Eriksson, W.F. Jones, C. Knoess, M. Lenox, M. Lercher, P. Luk, C. Michel, J. H. Reed, N. Richerzhagen, J. Treffert, S. Vollmar, J. W. Young, W. D. Heiss, R. Nutt, "The ECAT HRRT: Performance and First Clinical Application of the New High Resolution Research Tomograph." *IEEE Trans. Nuc. Sci.* vol. 49, pp. 104-110, 2002.

[9] D. L. Bailey, D. W. Townsend, P. E. Valk, M. N. Maisey, *Positron Emission Tomography*. Springer-Verlag, London, 2005.

[10] J. Nuyts, P. Dupont, S. Stroobants, A. Maes, L. Mortelmans, P. Suetens; "Evaluation of Maximum-Likelihood based attenuation correction in positron emission tomography," *IEEE Trans. Nucl. Sci.,* vol. 46, pp. 1136-1141, Aug. 1999.

[11] M. Sibomana, L. Byars, V. Panin, M. Lenox, F. Kehren, J. Rist, Z. Burbar, C. Michel, C. Knoess, H. W. A. M. de Jong. Simultaneous measurement of transmission and emission contamination using a collimated 137Cs point source for the HRRT, 2004 IEEE Nuclear Science Symposium Conference Record, M2-317.

[12] L. G. Byars, M. Sibomana, Z. Burbar, J. Jones, V. Panin, W. C. Barker, J.-S. Liow, R. E. Carson, C. Michel, "Variance reduction on randoms from delayed coincidence histograms for the HRRT." 2005 IEEE Nuclear Science Symposium Conference Record, pp. 2622-2626.

[13] C. C. Watson "New, faster, image-based scatter correction for 3D PET." *IEEE Trans. Nuc. Sci.* vol. 47, pp. 1587-1594, 2000.

[14] X. Liu, C. Comtat, C. Michel, P. Kinahan, M. Defrise, D. Townsend, "Comparison of 3-D reconstruction with 3D-OSEM and with FORE+OSEM for PET." *IEEE Trans. Med. Imag.* vol. 20, pp. 804-814, 2001.

[15] C. Comtat, F. Bataille, C. Michel, J. P. Jones, M. Sibomana, L. Janeiro, R. Trebossen, OSEM-3D Reconstruction Strategies for the ECAT HRRT", IEEE Nuclear Science Symposium Conference Record, vol. 6, pp. 3492-3496, 2004.

[16] J. P. Jones, W. F. Jones, F. Kehren, D. F. Newport, J. H. Reed, M. W. Lenox, K. Baker, L. G. Byars, C. Michel, M. E. Casey, "SPMD cluster-based parallel 3-D OSEM", *IEEE Trans. Nuc. Sci.* vol. 50, pp. 1498-1502, 2003.

[17] J. P. Jones, W. F. Jones, F. Kehren, Z. Burbar, J. H. Reed, M. W. Lenox, K. Baker, L. G. Byars, C. Michel, M. E. Casey, "Clinical Time OSEM3D: Infrastructure Issues", IEEE Nuclear Science Symposium Conference Record, 2003.

[18] I. K. Hong, S. T. Chung, Y. D. Son, H. G. Lee, H. K. Kim, Z. H. Cho, "Fast forward projection and backprojection algorithm using SIMD", IEEE Nuclear Science Symposium Conference Record, 2006.

[19] M. Rodriguez, J.-S. Liow, S. Thada, M. Sibomana, S. Chelikani, T. Mulnix, C. A. Johnson, C. Michel, W. C. Barker, R. E. Carson, "Count-rate dependent component-based normalization for the HRRT", *submitted*.

[20] A. J. Reader, P. J. Julyan, H. Williams, D. L. Hastings, J. Zweit, "EM algorithm system modeling by image-space techniques for PET reconstruction." *IEEE Trans. Nucl. Sci.,* vol. 50, pp. 1392-1397, Oct. 2003.

[21] F. C. Sureau, C. Comtat, A. J. Reader, C. Leroy, M.-J. Santiago-Ribeiro, I. Buvat, R. Trébossen, "Improved clinical parametric imagaing using list-mode reconstruction via resolution system modeling", IEEE Nuclear Science Symposium Conference Record, 2006.