Accelerated Real-time MRI with Distributed Computing Using Standard Scanner Hardware

E. K. Brodsky[1,2], and W. F. Block[1,3]

[1]Biomedical Engineering, University of Wisconsin - Madison, Madison, WI, United States, [2]Radiology, University of Wisconsin - Madison, Madison, WI, United States, [3]Medical Physics, University of Wisconsin - Madison, Madison, WI, United States

INTRODUCTION

We previously demonstrated a 3DPR-based real-time acquisition, reconstruction, and interactive visualization system capable of limited resolution, phased array 3D scanning [1]. The system utilized only the unused processing resources of the manufacturer's operator console host workstation to provide real-time reconstruction of non-Cartesian data at limited resolutions of 48x48x48 with 8 receivers at a rate of one frame per second. While this resolution was adequate for fluorotriggering in comprehensive MR angiography exams of the thorax and abdomen, higher spatial and temporal resolution is necessary for most interventional applications.

We have extended the system to support distributed computing and here present a version that utilizes a cluster of computers to accelerate the reconstruction. The system can be run using an off-line cluster of servers or on-line using the computing nodes in the manufacturer's standard product data acquisition and reconstruction server. This distributed system has allowed resolution to be increased to 96x96x96 with 8 receivers, an eight-fold increase in processing requirements, while maintaining the frame rate of one frame per second, using 4-8 nodes.

MATERIALS AND METHODS

The distributed real-time system, shown in Figure 1, is built upon our previous system and maintains the data acquisition and visualization interfaces from the host-only real-time system described earlier. The reconstruction module running on the master node distributes the raw data onto the cluster's slave nodes on a per-receiver basis. The slave nodes perform the regridding and inverse Fourier Transform and send compressed receiver magnitude volume images back to the master for combination and display. Communication between nodes was achieved using the Message Passing Interface (MPI), a standard API for distributed computing. The mpich2 1.0.4 implementation developed by Argonne National Laboratories was used [2].

The cluster used for off-line reconstruction consisted of eight rack-mounted servers, each with two AMD Opteron 270 2.0 GHz dual-core processors and 8 GB of memory. The cluster was interconnected with 1 GB/s Ethernet using a private switch (Linksys SRW2024), and connected to the scanner over the hospital's standard 100 MB/s Ethernet. On-line reconstruction was performed using the General Electric Excite HDx 3T Volume Reconstruction Engine (VRE) with four computing nodes (each node processing two receivers), using the operator console host computer as the master.

RESULTS AND DISCUSSION

The eight-node off-line cluster reconstructs 96x96x96 images in real-time at a rate of one frame per second. Each slave node spends 480 ms regridding, 30 ms on the inverse FFT, and 100 ms on other processing (chopping and frame combination). Additionally, it spends 20 ms compressing the image data and 15-200 ms sending it (as the master receives data synchronously, some nodes must wait while the master receives data from others). For each frame, the master spends 100-120 ms distributing raw data, 150-200 ms receiving image data, and 180 ms on receiver combination and normalization. Total reconstruction time is 625-675 ms, with latency from acquisition completion to image display of 500 ms. For online real-time reconstruction on the VRE using four nodes, total recon time is 700-750 ms, with latency of 580 ms. Achieving comparable performance on the manufacturer's smaller cluster is attributed to better network utilization and faster processing cores in the VRE.

Dividing the workload on a per-receiver basis is well-suited to inexpensive cluster solutions, which offer high processing performance but low network bandwidth. The most computationally intensive operations, regridding and the inverse FFT, are receiver-independent, so network utilization is limited to transmitting the raw data and receiver volume images. A substantial portion of the speedup is currently lost to network transmissions delays, but reducing the transmission of unused raw data and compressing data before transmission will significantly improve this. Compression ratios of 2:1 for raw data and 3:1 for image data are achievable using LZ77 and dynamic Huffman encoding.

CONCLUSIONS

Using the distributed computing cluster has allowed us to accelerate real-time reconstruction sufficiently to reduce the voxel volume by a factor of eight. The per-receiver distribution optimizes utilization of slow networks and scales well with the trend towards higher receiver channel counts. Optimizations to the gridding algorithm, increased CPU speeds, and higher performance network architectures will enable resolutions of 128^3 and higher. Reconstruction speed can also be further improved by using fewer receivers or reconstructing onto an asymmetric matrix, as is often possible in interventional imaging.

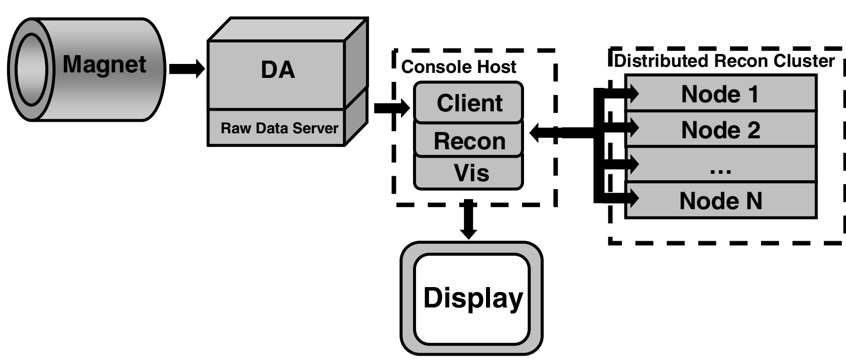

Figure 1: The distributed implementation of the real-time 3D fluoroscopy system adds a networked cluster of high-performance servers to accelerate image reconstruction. The improved performance allowed for reconstruction of 96x96x96 volumes with 8 receivers at one frame per second.

REFERENCES AND ACKNOWLEDGEMENTS

1. E Brodsky, *et al.*, MRM 56(2):247-254('06)
2. W Gropp, Euro PVM/MPI LNCS 2474:7('02)

Research supported by NIH NCI 1R01CA116380-01A1. The authors appreciate the assistance of GE Healthcare.