Distributed Computing on Gadgetron: A new paradigm for MRI reconstruction

Hui Xue[1], Peter Kellman[1], Souheil Inati[2], Thomas Sorensen[3], and Michael Schacht Hansen[1]

[1]Magnetic Resonance Technology Program, National Heart, Lung and Blood Institute, National Institutes of Health, Bethesda, Maryland, United States, [2]National Mental Health Institute, National Institutes of Health, Bethesda, Maryland, United States, [3]Department of Computer Science, Department of Clinical Medicine, Aarhus University, Aarhus, Denmark

Introduction MRI reconstruction has moved beyond the simple Fast Fourier Transform, towards more complicated parallel imaging and non-linear reconstruction algorithms. Often, developers who wish to implement and deploy their advanced algorithms on clinical scanners, find the vendor supported reconstruction hardware is inadequate for advanced algorithms or the provided programming environment is not suitable for clinical integration. The open-source framework, the Gadgetron, which was published recently, aims to address some of these concerns [1]. With this framework, algorithms can be implemented on suitable hardware selected by the user and subsequently connected to the scanner. While this is a valuable step-forward, the original Gadgetron was designed to run the reconstruction task within one process residing on a single computer and does not provide explicit support for cloud computing across multiple computational nodes. Although multiple CPU or GPU cores on one computer can contribute to the computation, single computer may not carry sufficient computing power to achieve clinical usage of non-linear reconstruction. To remove this limitation, we have extended the Gadgetron framework to support cloud computing. With this extension (named "Gadgetron Plus or GT-Plus"), any number of Gadgetron processes can run across multiple computers (thereafter referred as 'nodes') and a dedicated inter-process controlling scheme is implemented to multiple nodes. We demonstrate that with the GT-Plus cloud non-linear reconstruction of real-time cardiac cine imaging can be deployed in a clinical setting with acceptable reconstruction latency. Specifically, a multi slice real-time cine acquisition covering cardiac ventricles and non-linear reconstruction can be completed within 1min.

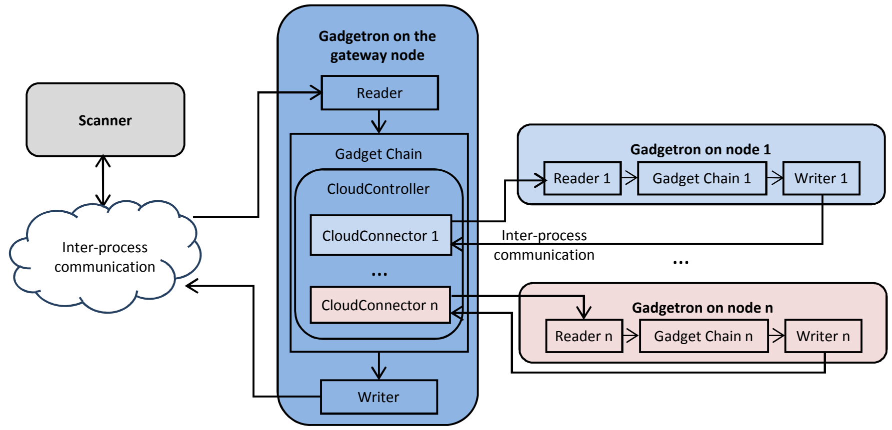

Figure 1. A schematic outline of typical setup of GT-Plus cloud. Here at least one Gadgetron process is running on a node (Multiple Gadgetron processes can run in one node on different ports). The gateway node communicates with the client application (e.g. MRI scanner). It manages the connections to each computing node via the software module `GadgetCloudController` and `GadgetronCloudConnector`. Whenever sufficient data is buffered on the gateway node, the packages can be sent to the computing node for processing. Different computing nodes can run completely different reconstruction chains. The `Reader/Writer` module in each Gadgetron process serializes/deserializes data to/from the Gadget chain for processing.

Architecture and Implementation Schematic outline of GT-Plus is shown in Fig. 1, representing a typical setup of Gadgetron based cloud where at least one Gadgetron running on a specific port at each node. A software module `GadgetCloudController` is implemented to manage the communication between nodes via TCP/IP sockets. Typically, a gateway node receives readout data from the scanner and distributes them through the cloud. For every connected node, a reader thread (`CloudReaderTask`) and a writer thread (`CloudWriterTask`) is spawned as active objects. There is a Gadgetron processing chain running on every connect node and for different cloud nodes, different processing chains can be performed. Whenever the reconstruction results are sent back from a cloud node, the gateway is notified and will take actions defined by the user, e.g. forwarding images back to the scanner or waiting for other jobs to complete.

Deployment The GT-Plus package can be deployed on various platforms. Two cloud setups were tested here. ***Amazon EC2 based cloud*** 19 nodes were launched to build up a virtual cloud on the Amazon Elastic Compute Cloud (Amazon EC2). Every node had two eight-core Intel Xeon E5-2670 2.60GHz processers and 60GB RAM, running Ubuntu Server 13.10. ***NIH Biowulf cluster*** The Biowulf system (http://biowulf.nih.gov) is a GNU/Linux parallel processing system built at the National Institutes of Health, USA**.** A total of 38 nodes were requested from the Biowulf. The gateway node had two eight-core Intel Xeon E5-2670 2.60GHz processers and 72GB RAM. All other nodes have two six-core Intel Xeon X5660 2.80GHz processers and 24GB RAM. For both setups, the "online" reconstruction was achieved.

In-vivo test Cloud based reconstruction was performed for multi-slice free-breathing myocardial cine imaging with non-linear l1-SPIRiT reconstruction [3]. A healthy volunteer (female, 23.8yrs) was scanned on a 1.5T clinical MRI system (MAGNETOM Area, Siemens, 32-channel surface coil). Acquisition parameters were: balanced SSFP readout, TR = 2.53/TE = 1.04ms, acquired matrix 192×100, flip angle 60°, FOV 320×240mm^2, slice thickness 8mm with a gap of 2mm, bandwidth 723 Hz/pixel, interleaved acquisition pattern with acceleration factor R=5. The ventricles of the heart were covered by 9 slices and for every slice the acquisition lasted 1.5s with one dummy heartbeat between slices.

Results Figure 2 shows the reconstruction results generated on the GT-Plus based cloud, illustrating noticeable improvement in image quality using non-linear reconstruction. The scan time (defined as the time to perform data acquisition) for this test was 22.6s. On the described Amazon EC2 cloud, the total imaging time (defined as the time from the start of data acquisition to the moment when images of all slices were sent back to scanner) was 52.6s. The computing time (defined as the time used to perform the reconstruction computation) was 48.9s. Note the reconstruction started once the first slice was acquired, rather than waiting for the completion of all 9 slices. On the NIH's Biowulf cloud, imaging and computing times were 62.1s and 58.2s. If only the single gateway node was used, the computing time was 427.9s and 558.1s for two cloud setups respectively.

Conclusions The GT-Plus extension for the Gadgetron framework was developed to support cloud computing over multiple computing nodes. As a demo, the increased computational power significantly speeds up the l1-SPIRiT reconstruction for multi-slice cine imaging, enabling the 1min imaging time for whole left ventricular coverage with significant improvement in image quality. On the other hand, the framework is independent from specific use cases and can be applied to other MRI applications.

References [1] Hansen MS, et al., MRM 69:1768-1776 (2013) [2] Lustig M, et al., MRM 64:457-471

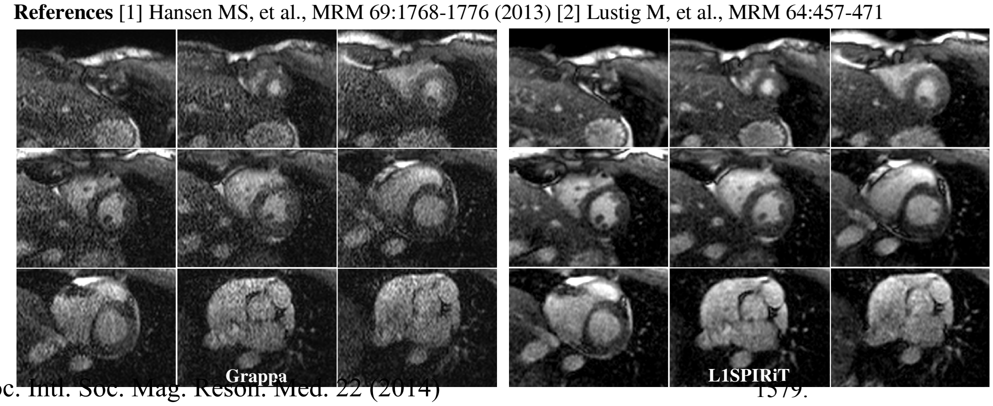

Figure 2. Reconstruction results of multi-slice myocardial cine imaging on the GT-Plus cloud. Compared to the Grappa linear reconstruction shown on the left, non-linear reconstruction on the right gives noticeable improvement in image quality. For this 9-slice cine protocol with entire ventricular coverage, the cloud version of non-linear reconstruction was completed within 30s after the end of data acquisition, while the single-node reconstruction took ~9mins. This significant reduction of computing time could be crucial for clinical usage of advanced image reconstruction techniques.