Scalable FBP Decomposition for Cone-Beam CT Reconstruction

Abstract

Filtered Back-Projection (FBP) is a fundamental compute intense algorithm used in tomographic image reconstruction. Cone-Beam Computed Tomography (CBCT) devices use a cone-shaped X-ray beam, in comparison to the parallel beam used in older CT generations. Distributed image reconstruction of cone-beam datasets typically relies on dividing batches of images into different nodes. This simple input decomposition, however, introduces limits on input/output sizes and scalability.

We propose a novel decomposition scheme and reconstruction algorithm for distributed FPB. This scheme enables arbitrarily large input/output sizes, eliminates the redundancy arising in the end-to-end pipeline and improves the scalability by replacing two communication collectives with only one segmented reduction. Finally, we implement the proposed decomposition scheme in a framework that is useful for all current-generation CT devices (7^{th} gen). In our experiments using up to 1024 GPUs, our framework can construct 4096^3 volumes, for real-world datasets, in under 16 seconds.

Introduction

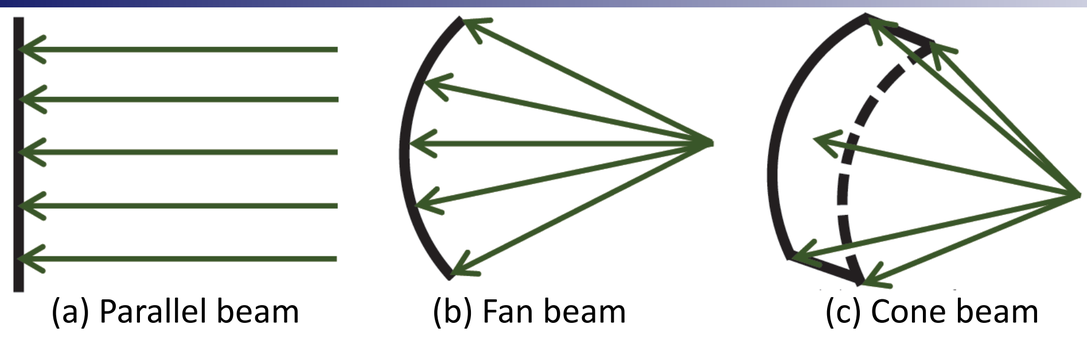

Fig. 1: Different geometries for X-ray sources and detectors. Cone-beam is the geometry used in the latest (7^{th} generation) of CT.

- Filtered Back-projection and computational complexity
  - Filtering computation : $O(N^2 LogN)$
  - Back-projection : $O(N^4)$

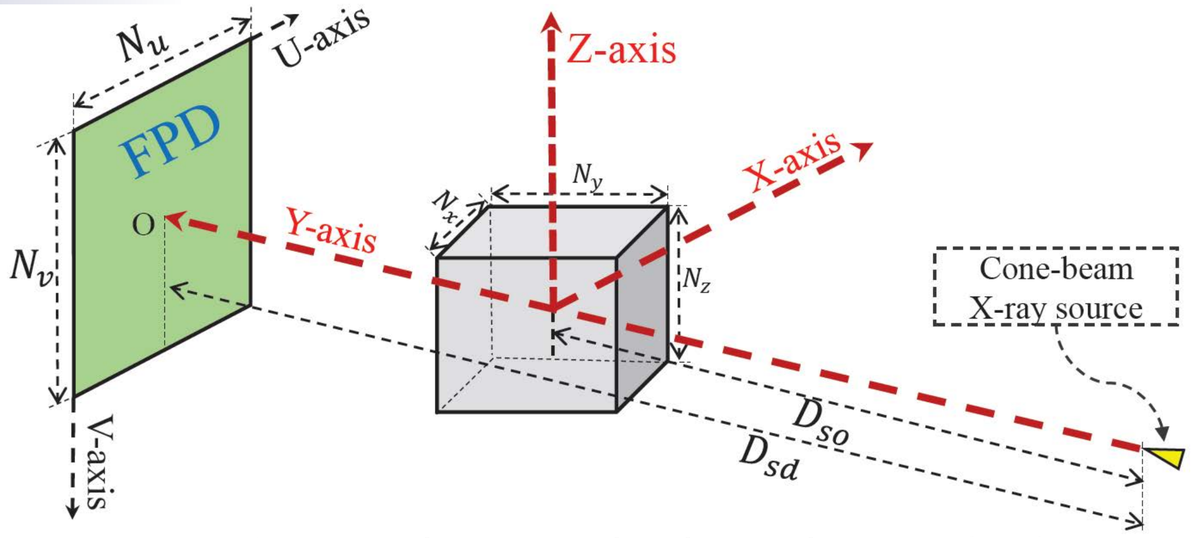

Fig. 2: Cone-beam CT with a Flat Panel Detector (FPD).

Methodology

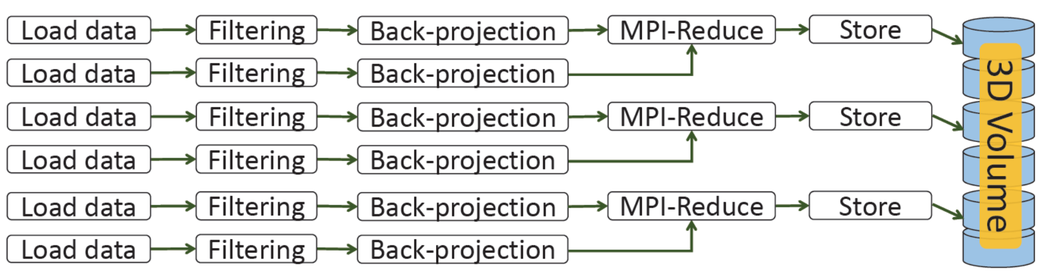

Fig. 3: Overview of the proposed framework.

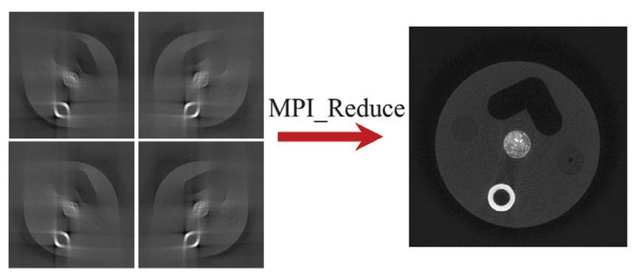

Fig. 4: An example of MPI-Reduce.

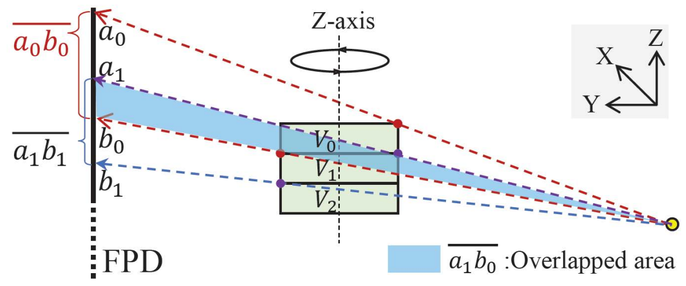

Fig. 5: Schematic YZ-view of overlapped projections.

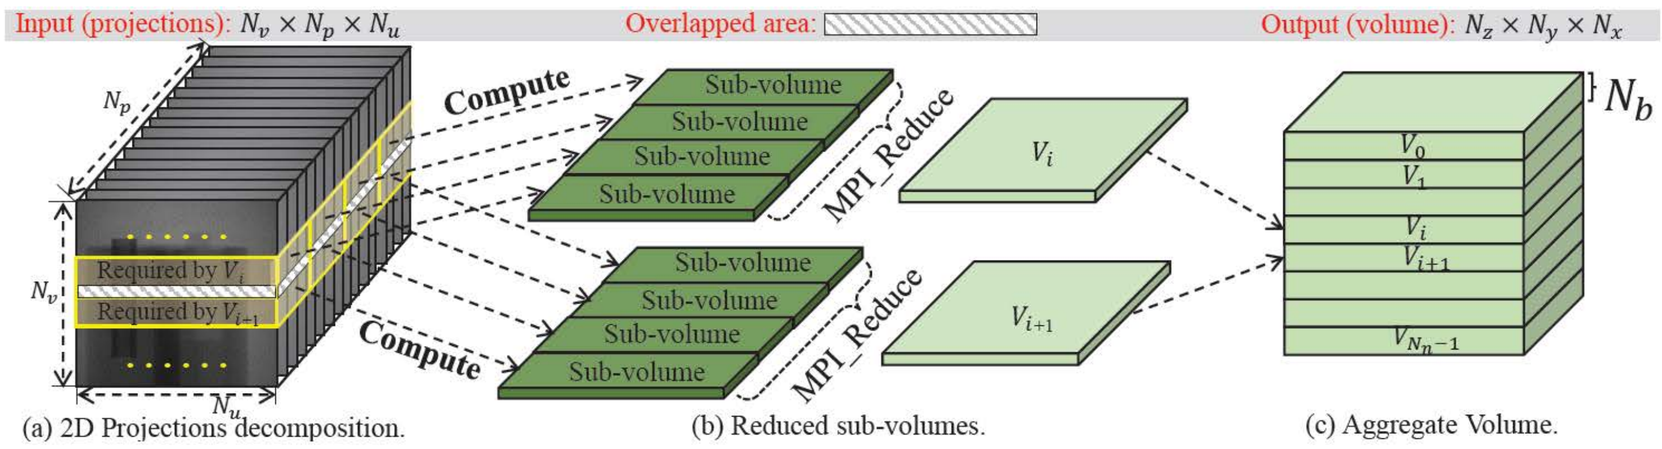

Fig. 6: Overview of the projection and volume decomposition. An example of four MPI ranks working as a group.

Results

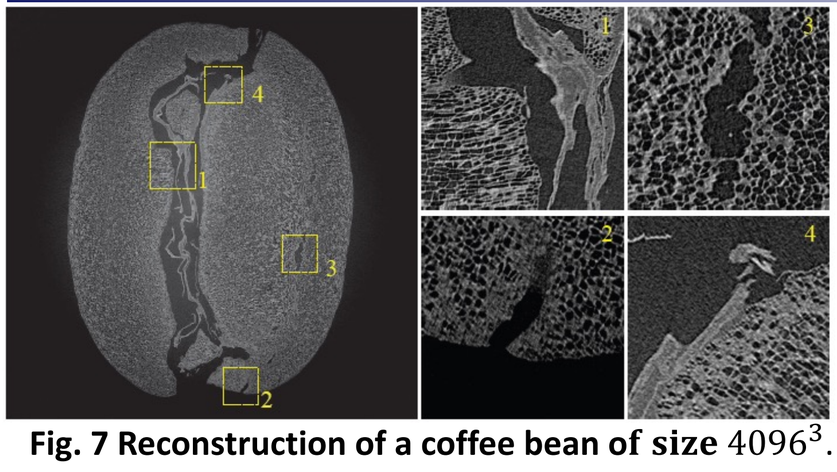

Fig. 7 Reconstruction of a coffee bean of size 4096^3.

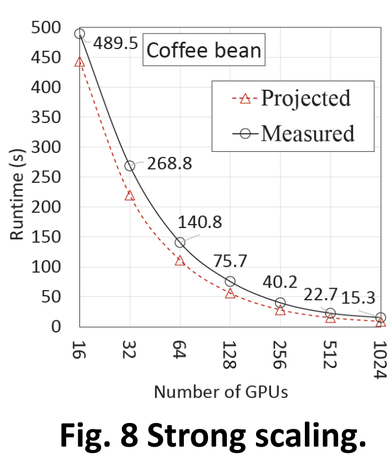

Fig. 8 Strong scaling.

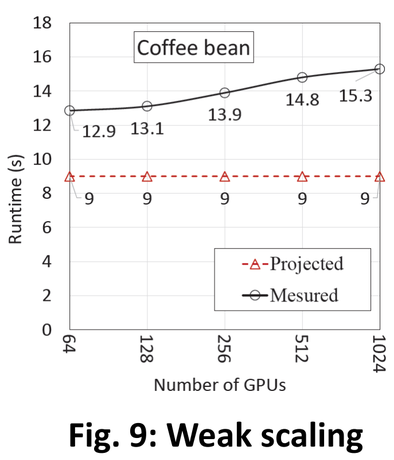

Fig. 9: Weak scaling

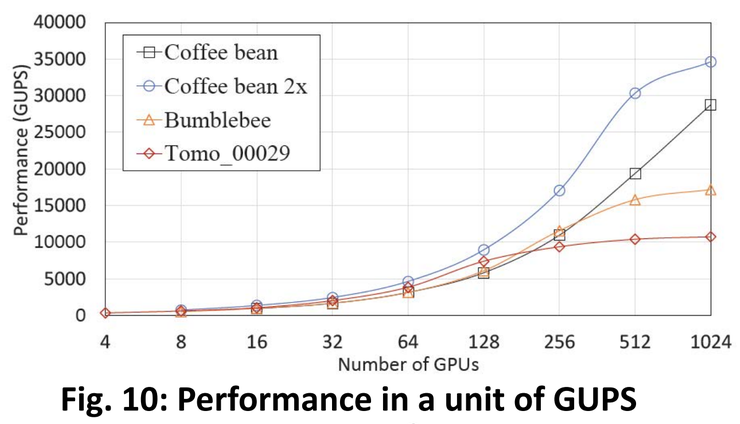

Fig. 10: Performance in a unit of GUPS when generating 4096^3 volumes.

Conclusion

- We propose a novel algorithm to decompose the image reconstruction problem for current generation cone-beam CT devices.
- Different computations were orchestrated on CPUs and GPUs to take advantage of the heterogeneous architecture of the GPU-accelerated supercomputers.
- To our knowledge, this is the first framework that supports out-of-core capability for cone-beam CT.