UIC UNIVERSITY OF ILLINOIS AT CHICAGO

Parallel Processing and Immersive Visualization of Sonar Point Clouds

electronic visualization laboratory

[1]Alessandro Febretti, [2]Kristof Richmond, [3]Peter Doran, [1]Andrew Johnson
[1]Electronic Visualization Lab, UIC [2]Frontier Astronautics LLC [3]Dept. of Environmental Sciences, UIC

Background

The investigation of underwater structures and natural features through Autonomous Underwater Vehicles (AUVs) is an expanding field with applications in archaeology, engineering, environmental sciences and astrobiology.

In other settings, it is possible to use optical methods to analyze the geometry of obstacles and features surrounding an autonomous vehicle, but the optical properties of underwater environments make this impossible. Sonar is still the best technology for underwater location, navigation and scanning.

Sonar point clouds made up of hundreds of millions to billions of points are not uncommon. Highly interactive, immersive visualization is a desirable tool that researchers can use to improve the quality of a final sonar-based data product.

Objective

We present a scalable toolkit for the processing and visualization of sonar point clouds on a cluster-based, large scale immersive visualization environment. The cluster is used simultaneously as a parallel processing platform that performs sonar beam-tracing of the source raw data, and as the rendering driver of the immersive display.

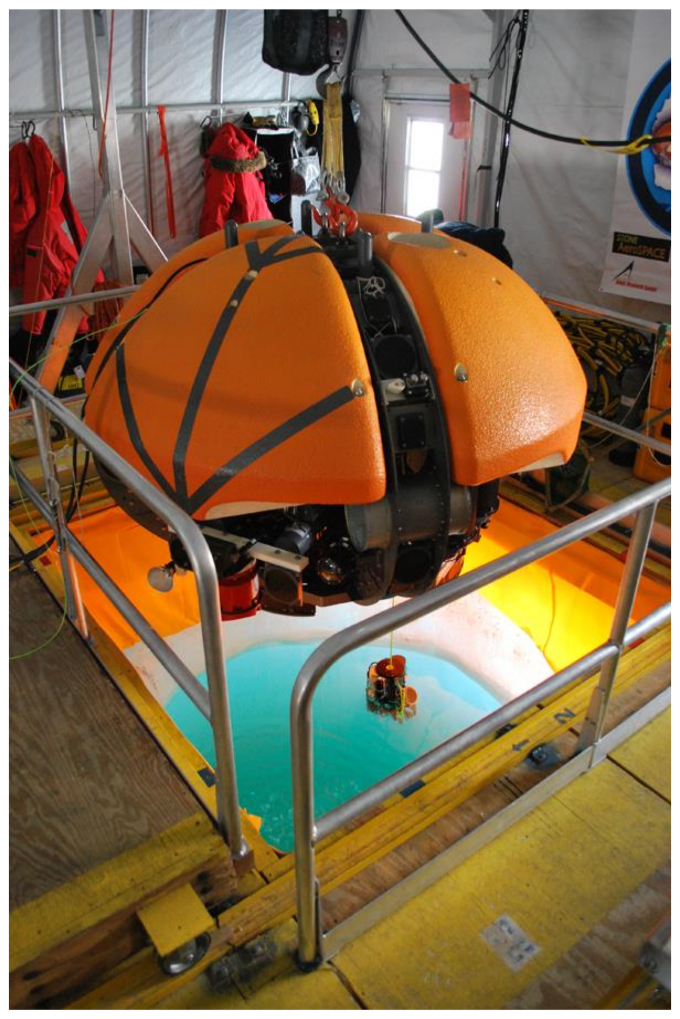

The Environmentally Non Disturbing Underwater Robotic ANTarctic Explorer (ENDURANCE) AUV was deployed in West Lake Bonney, a perennially ice covered lake in Antarctica and operated on 3 science objectives: water chemistry profiling, bathymetry scanning, and glacier exploration. For the purpose of bathymetry reconstruction, the source data consisted of about 350 Million distinct sonar range returns, plus navigation data and AUV attitude information at 0.2 second intervals.

An ideal sonar processing workflow requires continuous feedback between processing (i.e. beam tracing and clustering) and visualization tasks, allowing users to quickly perform corrections on the surveyed data, adjust processing parameters and observe the results of their actions in minutes instead of hours

A team of environmental scientists analyzing sonar data from West Lake Bonney, Antarctica using our visualization tool inside the CAVE2 hybrid immersive environment. The full dataset consists of 350 million sonar points, which can be filtered interactively based on attributes such as ping return time, beam takeoff angle and timestamp. The point cloud is overlaid with a secondary depth data source (the white columns) for cross-checking sonar depth accuracy.

Point Cloud Visualization

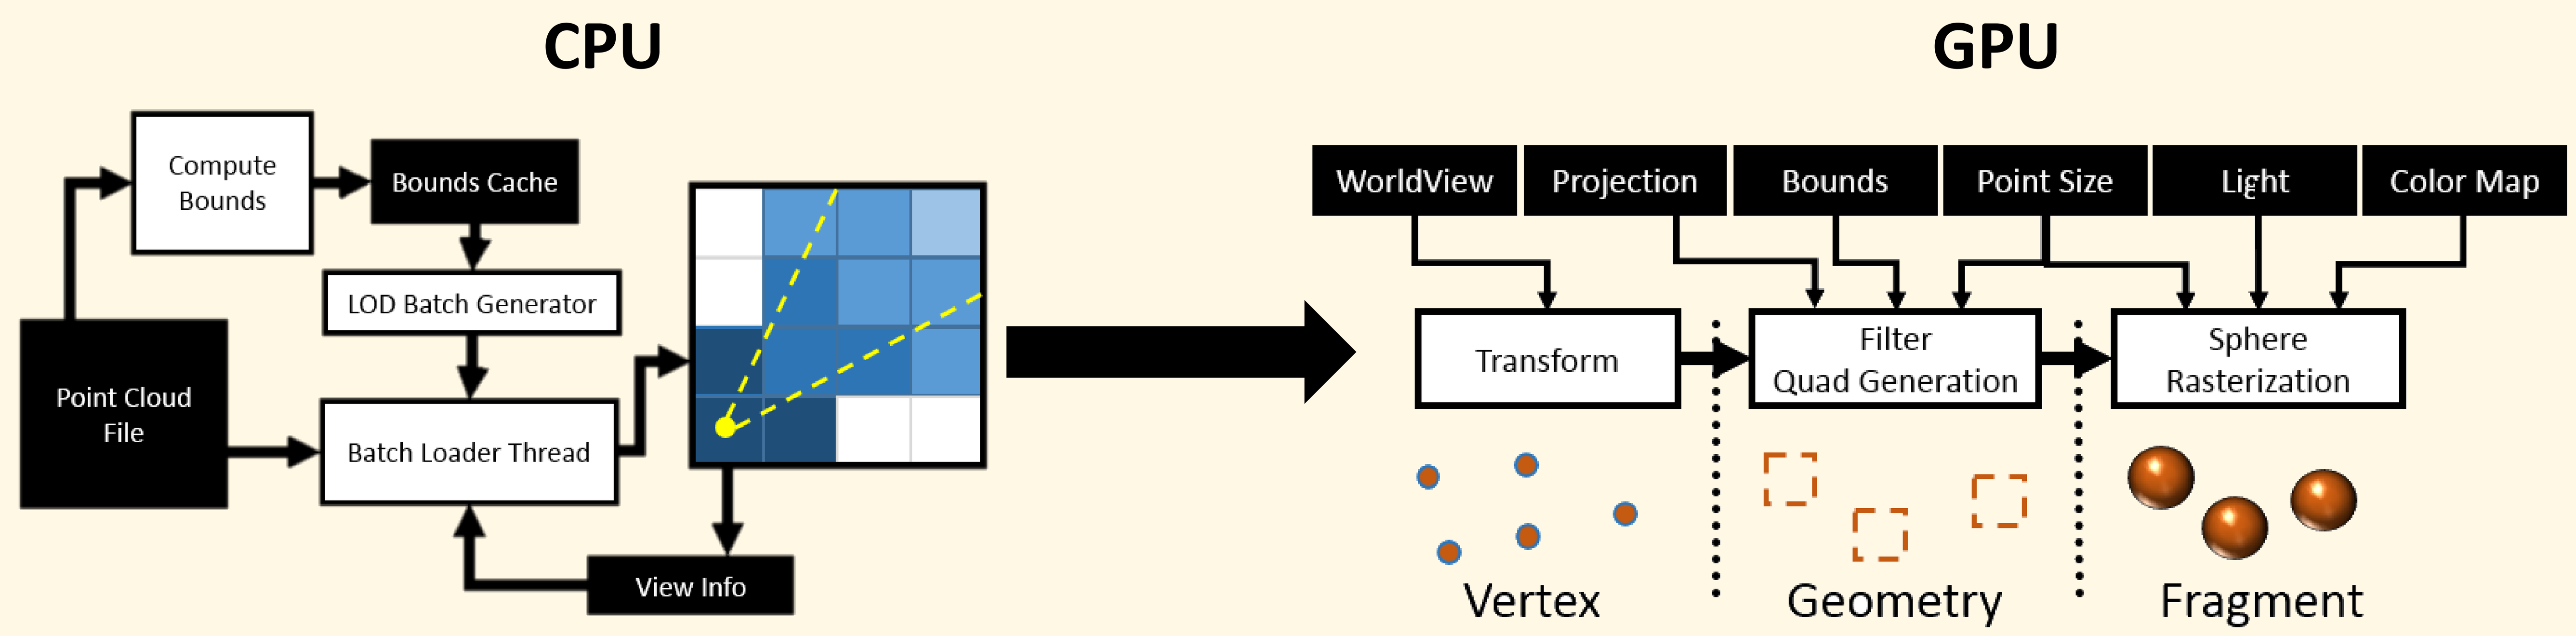

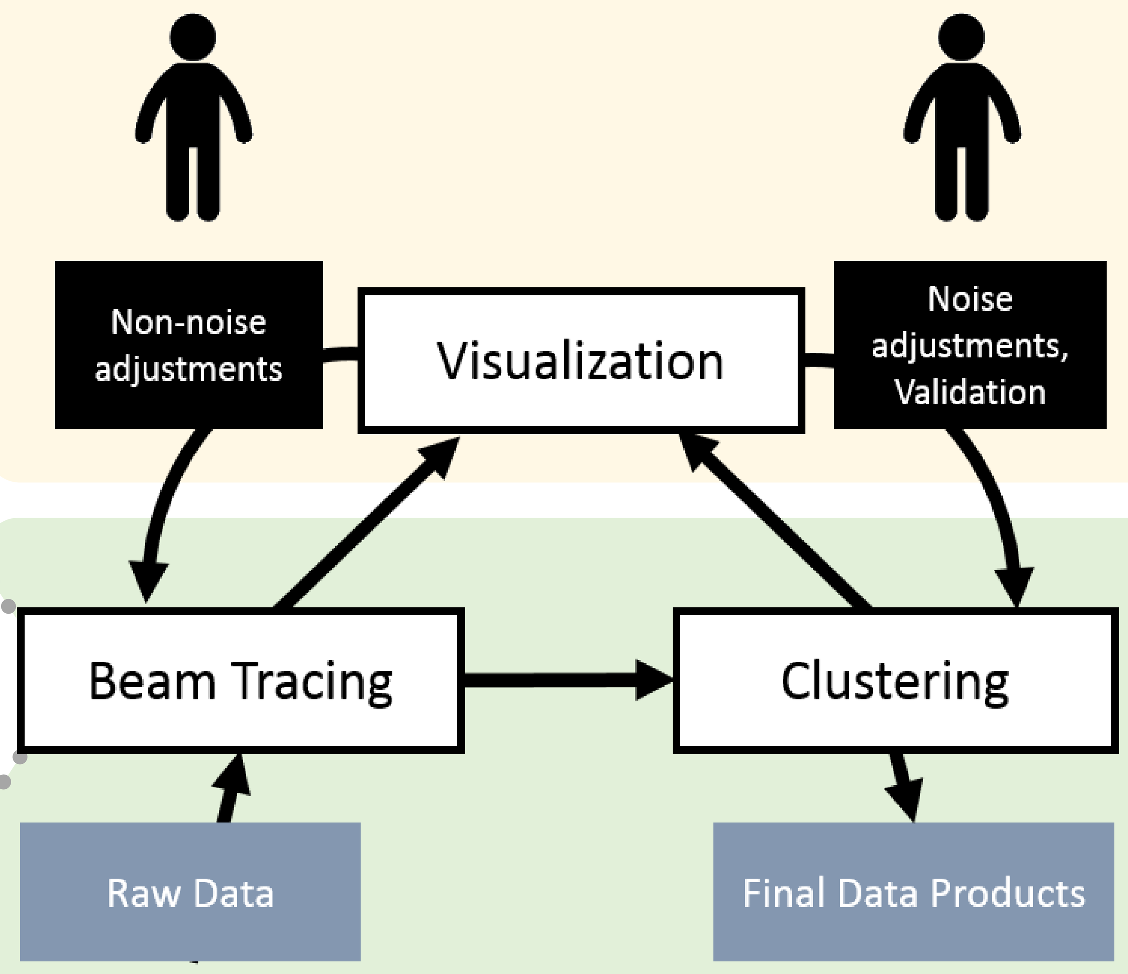

GPU Point Filtering and Rendering

The pipeline processes a vertex buffer containing point positions decorated with attribute data for each point. Conditional primitive emission inside a geometry shader is used to discard points that are not within desired attribute bounds. The geometry shader emits a screen-space quadrilateral, which is then rendered as a high quality sphere.

CPU LOD Batch Loading

Upon loading, data from a point cloud file is split in batches of a target size. Based on observer distance, each batch is loaded at a different decimation level. Batches are queued for loading on a separate thread, so the visualization remains interactive regardless of the speed at which LODs update.

Sonar Beam Processing

The speed of sound in water varies with changes in temperature, salinity and pressure, causing the trajectory of sound waves to bend due to refraction. To correct for beam bending, sonar data is processed by a ray-tracing algorithm that adjusts the beam trajectory as it goes through water layers with different sound speed characteristics. After beam tracing, the 3D points generated from multiple missions are merged into a single point cloud: overlapping data is used to filter noise and to generate final output points at the desired resolution level.

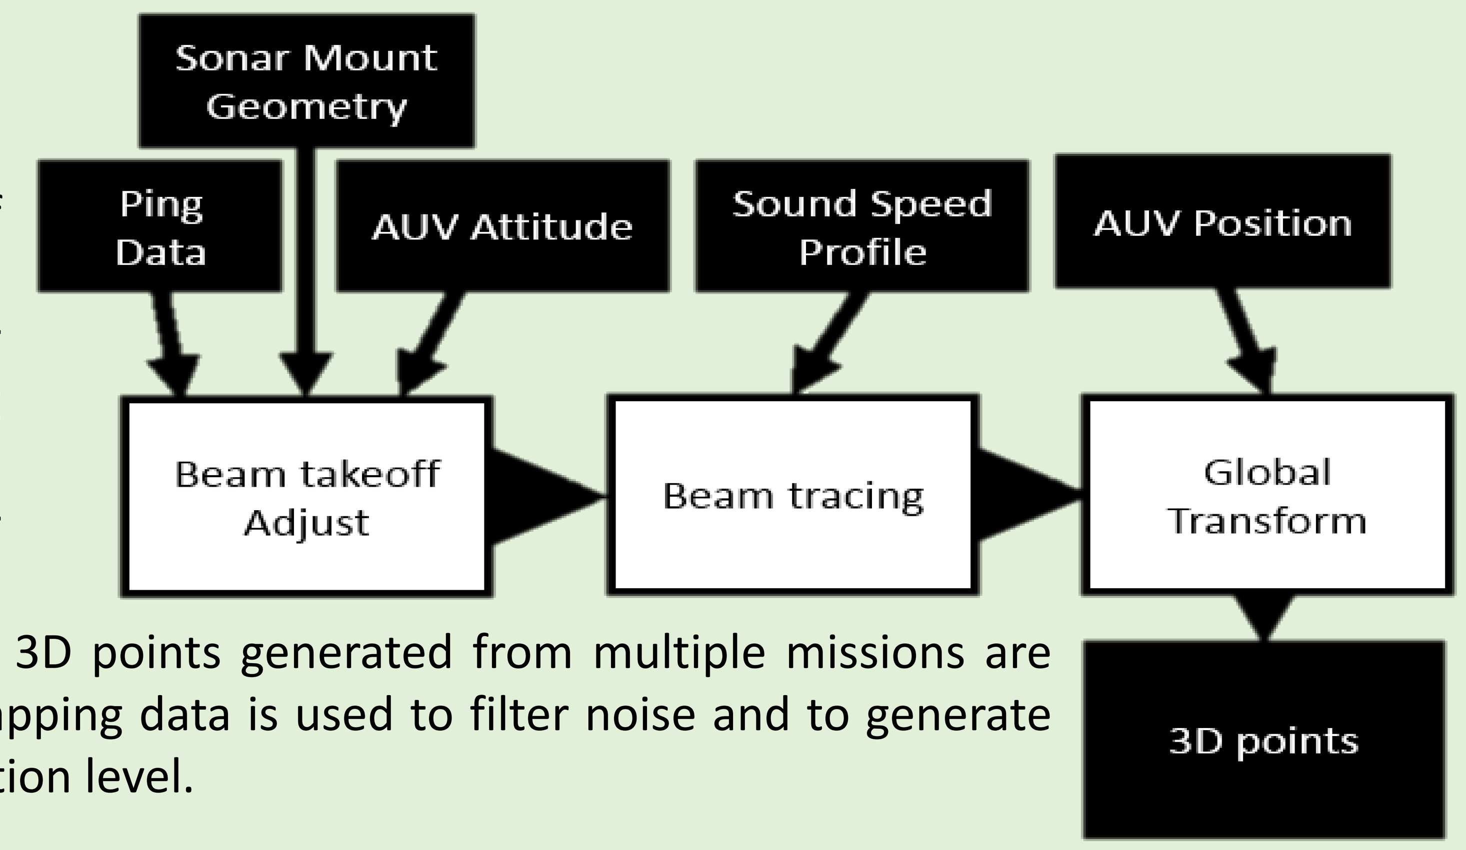

Parallel Performance

We evaluated the performance of our parallel tools on the 36-node CAVE2 cluster, varying the number of used cores. At 128 cores, beam tracing the 350 million points of the ENDURANCE dataset takes approximately 4 minutes, compared to 1 hour needed by a sequential run. For the merge step, we measure the speedup for two clustering resolutions: 1 meter and 2 meters. At 128 cores, merging all the ENDURANCE dives at 2-meter resolution requires about 5 seconds.