Poisson reconstruction of extreme submersed environments: The ENDURANCE exploration of an under-ice Antarctic Lake

Alessandro Febretti[2], Kristof Richmond[1] Shilpa Gulati[1], Christopher Flesher[1], Bartholomew P. Hogan[1], Andrew Johnson[2], William C. Stone[1], John Priscu[3], Peter Doran[4]

[1]Stone Aerospace
3511 Caldwell Lane, Del Valle, TX 78617-3006
[2]Electronic Visualization Laboratory, University of Illinois at Chicago
Dept. of Computer Science, 851 S. Morgan St., Chicago, IL 60607-7053
[3]Dept. of Land Resources and Environmental Sciences, Montana State University
334 Leon Johnson Hall, Bozeman, MT 59717
[4]Dept. of Earth and Environmental Sciences, University of Illinois at Chicago
845 W. Taylor St. Chicago, IL 60607

Abstract. We evaluate the use of Poisson reconstruction to generate a 3D bathymetric model of West Lake Bonney, Antarctica. The source sonar dataset has been collected by the ENDURANCE autonomous vehicle in the course of two Antarctic summer missions. The reconstruction workflow involved processing 200 million datapoints to generate a high resolution model of the lake bottom, Narrows region and underwater glacier face. A novel and flexible toolset has been developed to automate the processing of the Bonney data.

1 Introduction

The McMurdo Dry valleys, located within Victoria Land in Antarctica, are one of the world's most extreme deserts, and represent the largest ice-free region in the continent. The unique conditions in the valleys are in part caused by katabatics, winds reaching speeds of 320 kilometers per hour and capable of evaporating all water, ice and snow in the environment.

Some of the lakes in the Dry Valleys rank among the world's most saline lakes. One of them is Lake Bonney, (figure 1), a perennially ice-covered lake at the end of Taylor Glacier. Anaerobic bacteria whose metabolism is based on iron and sulfur live in sub-freezing temperatures under Taylor Glacier: since the Dry Valleys are one of the terrestrial environments closest to Mars and to some of Jupiter's moons, this is considered an important source of insights into possible forms of extraterrestrial life.

For this reason, NASA funded the Environmentally Non-Disturbing Underwater Robotic ANTarctic Explorer (ENDURANCE) project. ENDURANCE is an autonomous underwater vehicle (AUV) designed to map the geometry, geochemistry and biology of Lake Bonney in three dimensions.

ENDURANCE has been specifically designed to minimize impact on the environment it works in. This is primarily to meet strict Antarctic environmental protocols, but will also be a useful feature for planetary protection and improved planetary science in the future: NASA hopes to build upon lessons learned during testing for exploring objects in our solar system known to harbor sizable bodies of water, such as Jupiter's moon, Europa.

Fig. 1. From the left: A satellite image of West Lake Bonney; The ENDURANCE AUV ready to be deployed.

1.1 Challenges

Underwater environments are difficult to navigate, and the simultaneous localization and mapping (SLAM) problem is more challenging to solve, due to the presence of additional sources of uncertainty like water speed, time variant zero-depth levels, and physical properties of the water column influencing the sonar beam geometries.

The sonar mapping at Lake Bonney also required full 3D capabilities. A significant focus of scientific investigation was the degree of influence of the Taylor Glacier at the west end of the lake. To this end, the 3D structure of the glacier face and its interface with the lake bed were to be investigated and mapped in detail. Additionally, the topography of shallow areas around the lake edge, not directly accessible to the vehicle, was of interest.

2 Related work

2.1 Autonomous Underwater Exploration

3D Mapping with underwater robots ([1–3]) has received relatively low attention when compared to 3D mapping in other environments. Even fewer robotic vehicles have been developed for under-ice operation (like the ALTEX [4], Autosub [5] and SeaBED [6]). Some key differences between these vehicles and ENDURANCE are the very restricted operating volume in Lake Bonney (between 3 and 12 meters depth) and a difficult acoustic environment to operate in, due to smooth ice above the operating volume and a severe halocline below it.

2.2 Surface reconstruction / 3D maps

The need for fully 3D, high resolution reconstruction required a careful design of the data processing workflow. The final objective would be to generate a high quality bathymetric and glacier model, while dealing with high noise levels in the data, varying coverage levels, and different sonar configurations from dive to dive.

Delaunay triangulation techniques [7] are affected by data noise so they require pretty aggressive filtering. Also given their plane-fitting nature they do not work well with fully 3D data.

Approaches based on occupancy grid generation plus moving least squares (MLS) have been used successfully in similar works ([8, 9]) but present three major problems relative to our scenario:

- In presence of non-homogeneous data, occupancy grid + MLS approaches leaves holes in the reconstruction. These holes need to be filled by additional postprocessing steps (for instance by expanding neighboring regions)
- For high resolution data, the amount of required voxels can make memory requirements prohibitive (although there are strategies to alleviate this, i.e. the octree grids used in [8])
- The default sonar beam model used to write data to the grid is a cone with negative log-probability volume voxels and positive log-probability base voxels. This works well for isotropic water volumes. But if physical properties change across the water volume (as is the case for the Bonney environment), the beam model needs to be regenerated for each writing step. Although doable from the technical standpoint, this would require a significant amount of work on existing evidence grid solutions.

These considerations (along with preliminary tests on the data) made the occupancy grid + MLS solution less than ideal for our scenario, and led us to consider a reconstruction approach based an implicit surface technique called Poisson reconstruction [10]. Additional details about this technique will be presented in section 4.3. To the best of our knowledge, this is the first time the Poisson technique has been used in the context of sonar-based surface reconstruction. The rest of this paper will present our workflow implementation and discuss several advantages and disadvantages of the chosen approach.

3 Data collection

ENDURANCE operated during two Antarctic summer seasons (2008 and 2009). Each mission entailed a set of deployments of the AUV, for a total of 45 dives. For each dive, The AUV operated depending on 3 distinct science objectives: Water chemistry profiling, Bathymetry scanning, and glacier exploration. This work concentrates on processing data from bathymetry and glacier dives, but a full coverage of the mission science objective can be found in [11]

The ENDURANCE AUV was equipped with several independent sonar systems. The primary mapping unit was a 480-point multi-beam sonar with a 120° × 3° field of view. This unit could be mounted in both forward-looking and down-looking configurations to ensure full coverage of the lake geometry.

4 Data Processing

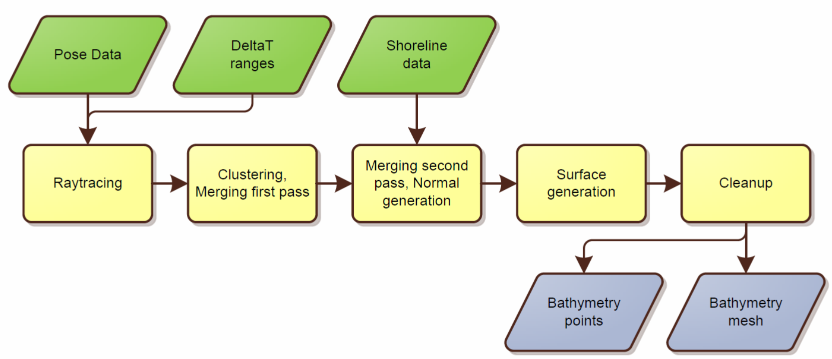

Fig. 2. Overview of the sonar data processing and surface reconstruction workflow. The Pose database contains the AUV position and attitude information. The DeltaT database stores the range return readings from the AUV sonar intrument.

For the purpose of bathymetry reconstruction, the source data consisted of about 200 million distinct sonar range returns, plus navigation data and AUV attitude information at 0.2 second intervals. Figure 2 offers an overview of the data processing steps required to transform the raw range data into the final 3D model of West Lake Bonney. The processing steps required to remove errors from the raw vehicle navigation to arrive at the pose data used for mapping are detailed in [11].

4.1 Raytracing

In section 2.1 we explained how sonar data required corrections dependent on the water column physical properties. The trajectory of sonar beams is influenced

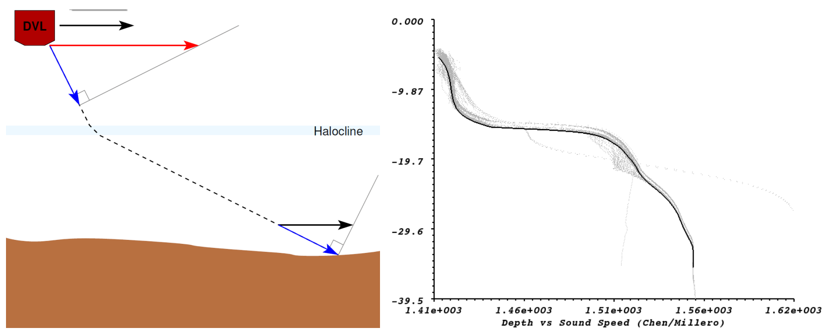

Fig. 3. From the left: geometry of sonar beams in the water column; the sound speed profile for West Lake Bonney: the halocline below the freshwater lens is at a depth of about 16 meters.

by the water density according to Snell's Law. Sound waves get refracted as they travel through the water column. This effect was particularly marked in West Lake Bonney due to the presence of a major halocline at a depth of 16 meters. In order to take this effect into account, we computed a speed-of-sound profile for West Lake Bonney using the Chen-Millero formula [12] (figure 3) To compute this model, we used the measurements obtained from the ENDURANCE science payload (in particular, the temperature and salinity profiles of the water column).

The sound speed profile was used in conjunction with the sonde telemetry data and sonar configuration to re-trace the beam paths as they traveled across the water column. A modified version of the raytracing algorithm implemented in the MB system software has been used in this step.[1]

Some minor noise filtering was performed during this step to eliminate clear outliers (i.e. using range return timing and beam angle thresholds)

4.2 Clustering

In order to reduce the noise in the raytraced point cloud, we integrated a fairly standard clustering step in our processing pipeline. Range returns were collected in 3D bins, using an octree grid implementation similar to the one presented in [1]. If the number of points in a bin fell under a certain threshold, the entire bin was discarded as noise. Otherwise, points in the bin were averaged in order to generate a single 3D point for each bin.

To process the Bonney dataset, we used a 3D grid with a resolution of 1 meter and a 16 point rejection threshold.

[1] MB System website: `http://www.ldeo.columbia.edu/res/pi/MB-System/html/mbsystem_home.html`

4.3 Surface Reconstruction

As mentioned in section 2.2, we decided to employ an implicit surface technique in our reconstruction step. In [13], Kazhdan et al. show how surface reconstruction from oriented points can be formulated as a Poisson problem. The advantage of this formulation is that it considers all the points at once and is quite resilient to noise in the source data.

Like other solutions, Poison reconstruction is based on computing a 3D indicator function χ defined as 1 for points inside the model and 0 otherwise. The gradient of the indicator function $\bigtriangledown\chi$ is a zero vector everywhere but for points near the surface, where it equals the inward surface normal. The oriented point samples can therefore be viewed as samples of $\bigtriangledown\chi$: the indicator function χ can be defined as the scalar function whose gradient best approximates a vector field $\boldsymbol{V}$ defined by the samples:

$$min_{\chi} \| \bigtriangledown \chi - \boldsymbol{V} \|.$$

By applying the divergence operator, this problem can be transformed into a standard Poisson problem: compute the scalar function whose Laplacian (divergence of the gradient) equals the divergence of a specified vector field:

$$\triangle\chi \equiv \bigtriangledown \cdot \bigtriangledown\chi = \bigtriangledown \cdot \boldsymbol{V}.$$

Due to their interdisciplinary interest, A number of efficient and robust methods have been developed to solve Poisson problems. One additional insight in [13] is that, since we are interested in an accurate solution only near the reconstructed surface, we can use adaptive solvers in order to keep memory and time constraints proportional to the size of the reconstructed surface. Once we have an estimate of the indicator function χ, we can generate the reconstructed surface by extracting an appropriate isosurface. The isovalue is chosen so that the corresponding surface approximates the positions of the input samples. Extraction is then performed using an adapted Marching Cubes technique [14, 15]

Since the Poisson reconstruction technique works on oriented point sets, we used the normal estimation algorithm described in [16] to add normal information to the source points. This method constructs a K-nearest-neighbor-points graph over the input data (a Riemannian graph), and propagates a seed normal orientation using a minimum spanning tree over this graph.

The Poisson reconstruction step was run on the point cloud using a solver depth of 11, in order to produce a final surface with an average vertex distance of 1 meter.

4.4 Postprocessing and improvements

Cutting Since the Poisson algorithm solution is a watertight surface, we needed to filter out the surface components that were lying above-water: these were just a low resolution reconstruction artifact generated to close the mesh. This was easily done as a postprocessing step, by placing a cutting plane at the zero water depth level, and eliminating all the geometry above it.

Cleanup At this point, The reconstructed surface still presented a significant amount of noise, in the form of spherical clusters or blobs disconnected from the main surface. This noise was due to the ENDURANCE science payload deployments, which caused spurious sonar returns every time the AUV stopped to collect data. These undesired sonar returns were coherent enough to pass the cleanup threshold during the clustering step. To address this issue, a second postprocessing identified all the connected components in the mesh below a certain size threshold, and eliminated them from the final surface.

Both the cutting plane and the connected components filter have been implemented as automated scripts running in MeshLab.

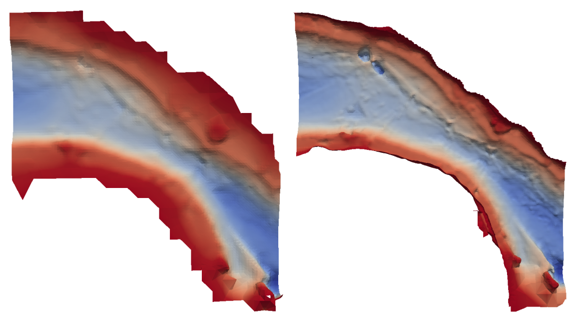

Fig. 4. From the left: detail of Narrows region reconstruction without shore line information; the same region reconstructed with shore line information.

Shore line data integration In order to increase the reconstruction quality along the lake shores, we integrated shore line information generated by satellite imagery to the sonar point cloud. This dramatically reduced reconstruction artifacts for the lake zero-depth contour, as seen in figure 4.

4.5 dttools

The entire workflow has been implemented as a C++ toolkit, consisting of a set of programs implementing the raytracing, clustering and reconstruction steps of the data processing pipeline. All the programs use configuration files to specify their inputs, outputs and parameters. This allowed us to create scripts to run the entire pipeline easily for multiple configurations, and compare results in order to identify the optimal processing settings. The software source code is available online and can easily be adapted to different source datasets. [2]

[2] dttools website: `http://code.google.com/p/dttools/`

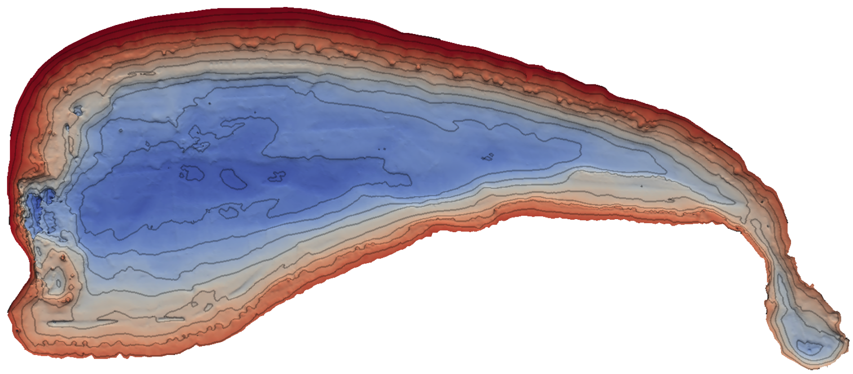

Fig. 5. The new 3D West Lake Bonney bathymetry.

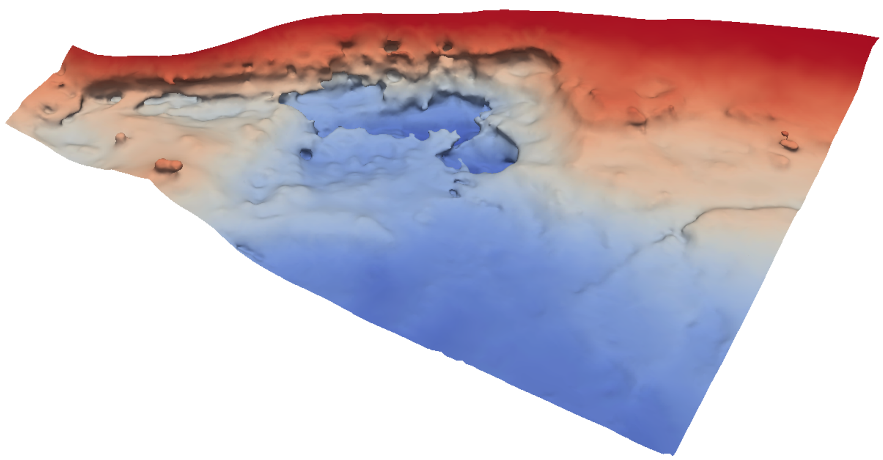

Fig. 6. A detail of the glacier face reconstruction.

5 Results

The final reconstruction of West Lake bonney can be seen in figures 5 and 6. The high resolution mesh consists of about 200k vertices and 400k faces. As the detail view of the glacier face shows, the reconstruction is fully 3D, while previous bathymetric models from the lake (generated by manually taken samples of the lake depth) were limited to approximate depth contours.

This improved model will allow for precise calculations of the lake water volume, and together with the data collected by the ENDURANCE science payload will allow a better understanding of the chemical and biological features of this extreme environment.

A drawback of the current reconstruction approach is its dependence on clustering and normal estimation parameters. Slight variations of those parameters could lead to incorrectly estimated normals for some regions, which would in turn cause errors in the final reconstruction. Thus, multiple iterations of the reconstruction process had to be run, varying parameters and visually inspecting the result in order to determine the best pipeline configuration. In particular, it was important to balance the parameters in order to get a good reconstruction in all three of the lake regions (the Glacier face, lake bottom, and Narrows). Reconstruction with slight differences had to be compared side by side. Given the high resolution of the models, the use of tiled displays (figure 7) helped in the process. In the future it may be beneficial to expand dttools, integrating tools that help in the iterative refinement of reconstruction parameters and comparison of corresponding results.

Fig. 7. Side-by-side analysis of West Lake Bonney reconstruction results on a high resolution display wall.

References

1. Forney, C., Forrester, J., Bagley, B., McVicker, W., White, J., Smith, T., Batryn, J., Gonzalez, A., Lehr, J., Gambin, T., Others: Surface reconstruction of Maltese cisterns using ROV sonar data for archeological study. Advances in Visual Computing (2011) 461–471
2. Fairfield, N., Kantor, G., Wettergreen, D.: Real-Time SLAM with Octree Evidence Grids for Exploration in Underwater Tunnels. Journal of Field Robotics **24** (2007) 03–21
3. White, C., Hiranandani, D., Olstad, C.S., Buhagiar, K., Gambin, T., Clark, C.M.: The malta cistern mapping project: Underwater robot mapping and localization within ancient tunnel systems. J. Field Robot. **27** (2010) 399–411
4. McEwen, R., Thomas, H., Weber, D., Psota, F.: (Performance of an AUV navigation system at arctic latitudes)
5. McPhail, S.: Autosub operations in the Arctic and Antarctic. In Griffiths, G., Collins, K., eds.: Proceedings of the Masterclass in AUV Technology for Polar Sci-

ence, National Oceanography Centre, Southampton, UK, Society for Underwater Technology (2006) 27–38
6. Jakuba, M.V., Roman, C.N., Singh, H., Murphy, C., Kunz, C., Willis, C., Sato, T., Sohn, R.A.: Long-baseline acoustic navigation for under-ice autonomous underwater vehicle operations. Journal of Field Robotics (2008) 861–879
7. Okabe, A., Boots, B., Sugihara, K.: Spatial Tesselations, Concepts and Applications of Voronoi Diagrams. Wiley series in probability and mathematical statistics. John Wiley & Sons, Chichester, England (1992)
8. Fairfield, N., Jonak, D., Kantor, G.A., Wettergreen, D.: Field results of the control, navigation and mapping system of a hovering AUV. In: Proceedings of the Unmanned Untethered Submersible Technology Conference (UUST) 2007, Durham, New Hampshire, AUSI (2007)
9. Papadopoulos, G., Kurniawati, H., Bin Mohd Shariff, A.S., Patrikalakis, N.M.: 3D-surface reconstruction for partially submerged marine structures using an Autonomous Surface Vehicle. 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (2011) 3551–3557
10. Kazhdan, M., Bolitho, M., Hoppe, H., Burns, R.: Poisson surface reconstruction. In: Proceedings of the fourth Eurographics symposium on Geometry processing, Eurographics Association (2006) 61–70
11. Richmond, K., Febretti, A., Gulati, S., Flesher, C., Hogan, B.P., Murarka, A., Kuhlman, G., Sridharan, M., Johnson, A., Stone, W.C., Priscu, J., Doran, P.: Sub-Ice exploration of an antarctic lake: results from the ENDURANCE project. Arctic (2011)
12. Dushaw, B.D., Worcester, P.F., Cornuelle, B.D., Howe, B.M.: On equations for the speed of sound in seawater. Acoustical Society of America Journal **93** (1993) 255–275
13. Hoppe, H.: Poisson surface reconstruction and its applications. Proceedings of the 2008 ACM symposium on Solid and physical modeling - SPM '08 (2008) 10
14. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. SIGGRAPH Comput. Graph. **21** (1987) 163–169
15. Wilhelms, J., Van Gelder, A.: Octrees for faster isosurface generation. ACM Trans. Graph. **11** (1992) 201–227
16. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points. ACM SIGGRAPH Computer Graphics **26** (1992) 71–78