

Surface Reconstruction of Maltese Cisterns using ROV Sonar Data for Archeological Study

C. Forney, J. Forrester, B. Bagley, W. McVicker, J. White, T. Smith,
J. Batryn, A. Gonzalez, J. Lehr, T. Gambin*, C. M. Clark and Z. J. Wood

California Polytechnic State University, San Luis Obispo, CA, USA and *AURORA
Special Purpose Trust, Malta

Abstract. We present a methodology and algorithm for the reconstruction of three dimensional geometric models of ancient Maltese water storage systems, i.e. cisterns, from sonar data. This project was conducted as a part of a four week expedition on the islands of Malta and Gozo. During this expedition, investigators used underwater robot systems capable of mapping ancient underwater cisterns and tunnels. The mapping included probabilistic algorithms for constructing the maps of the sonar data and computer graphics for surface reconstruction and visualization. This paper presents the general methodology for the data acquisition and the novel application of algorithms from computer graphics for surface reconstruction to this new data setting. In addition to reconstructing the geometry of the cisterns, the visualization system includes methods to enhance the understanding of the data by visualizing water level and texture detail either through the application of real image data via projective textures or by more standard texture mapping techniques. The resulting surface reconstructions and visualizations can be used by archaeologists for educational purposes and to help understand the shape and history of such water receptacles.

1 Introduction

Due to its strategic location in the Mediterranean Sea, Malta's culture has been strongly influenced by neighboring countries in both Europe and Africa. Today, Malta serves as a prime location for exploring archeological artifacts, many of which remain relatively in tact and found beneath the water. The work reported here is interdisciplinary, utilizing emerging robotics and computer science technology to explore previously unexamined, underwater structures. The project concerns investigation of ancient water storage systems, i.e. cisterns, located in most houses, churches, and fortresses of Malta. Cisterns are typically man-made stone underground water storage receptacles, which served a vital role in the survival of the Maltese, especially during sieges. According to [1], their importance was of such significance to residents during medieval times, that they were required by law to maintain access to three years of water supply.

This project was conducted as a part of a four-week expedition on the islands of Malta and Gozo in March 2011. The project focused on obtaining archeological and historical data, as well as validating robot technology and computer

graphics visualizations. This was accomplished through a series of underwater robot system deployments within cisterns. Archaeologists looking to study and document underwater cisterns in Malta have found it too expensive, difficult and dangerous to use people. The goal of this project was to gather data about the underwater cisterns, and enable the creation of geometric models of the cisterns that could later be explored and visualized in a computer graphics application. This work builds from two previous expeditions [2], after which only two-dimensional maps were created.

In all, this project involved the exploration of over twenty cisterns and water features. A small underwater robot, or specifically a VideoRay micro-ROV [Remotely Operated Vehicle] was used, (see Figure 1). The ROV is controlled above the surface using a joystick housed within a control box. A tether connects the ROV to the control box, allowing sensor information (i.e. depth, bearing, and video camera images) to be displayed within the control box. The tether also allows power, and control signals to be sent to the ROV. The ROV is actuated with two co-planar thrusters (motor driven propellers) and one vertical thruster. On board lights are used to illuminate the ROV video camera’s field of view.

Investigators deployed the ROV into cistern access points until it was submerged. The investigators then tele-operated the robot to navigate the underwater environment. Video images were recorded using the ROV’s onboard camera. Two and three-dimensional maps of the cistern are created using data from a SeaSprite scanning sonar mounted on top of the ROV. These sonar measurements were used with multiple algorithms to conduct ROV localization and for the purpose of creating three-dimensional occupancy maps to represent the cisterns. These occupancy maps (or grids) were then used to generate the geometric models of the cisterns.

This paper presents the general methodology for the data acquisition and the novel application of algorithms from computer graphics for surface reconstruction to this new data setting. In addition, we present algorithms for the visualization and texturing mapping of such models (see Figure 6). These visualizations are used by archeologists for understanding the shape, connectivity and history of the cisterns and can be used for educational purposes to illustrate the importance and history of such water features.

2 Related Work

Mapping via Underwater Robot Systems: Creating maps using robots equipped with appropriate onboard sensing has been an active area of research for many years. Laser scanners, stereo-vision, and ultrasonic sensors have all been used to enable a variety of mapping tasks with ground-based, aerial, and underwater robots. This work focuses on mapping with underwater robots (e.g. [3], [4, 5], [6], and [2]), which has seen less attention when compared to mapping in other environments. Regardless, there are a large number of significant real life applications that can benefit from such technology including oceanography [7], marine biology [8], and archaeology [6], as well as for robot navigation (e.g. [9, 10]).



Fig. 1. On the left are images of the data acquisition process, while the right shows acquired (and mosaiced) sonar data. The bottom left photo is the VideoRay micro-ROV. Top left is a photo of investigators flying the ROV through the Mdina Archives cistern to gather sonar and video data. The middle photo shows the deployment of the ROV via lowering it into the cistern access point.

In the last 20 years, research has been conducted in developing algorithms capable of fusing data from multiple sensors to create maps. In particular, Simultaneous Localization and Mapping (SLAM) algorithms have been developed that can simultaneously construct an estimate of a map and calculate an estimate of the state of the robot (e.g. position) with respect to the map. A good survey of the core techniques including both Kalman Filtering and Particle Filtering based techniques can be found in [11].

Other relevant work includes research in underwater robot SLAM. One of the first instances includes the work done in [3] where sonar scans were used to map and track features of the environment. Scan-matching techniques have been used by researchers conducting mapping and localizing with robots in harbor environments [12]. More recently, successful 3D tunnel mapping in underwater environments was demonstrated in [4, 5]. Unlike the related work, this research requires mapping in tunnels with passages and access points that are relatively small (e.g. 0.30m diameter at some points). This requires a smaller robot with low payload capacity and minimal sensors (i.e. scanning sonar, depth sensor, and compass). White et. al. demonstrated 2-dimensional maps of cisterns could be constructed with this limited sensor suite [2], and here the work in [2] is extended to construct 3-dimensional occupancy grid maps, which can later be surfaced and texture-mapped for high quality visualization.

Surface Reconstruction and Visualization: The problem of digital surface reconstruction is well known in the computer graphics community, (e.g. [13–15]), due to the importance of creating three dimensional computer models with which users can interact. Once such a geometric model has been constructed, a user

can then examine it from any view in order to gain insight about its nature and areas of interest. Creating such a tool for this setting assists archaeologists in their examination of the underwater structures and helps with the planning of future robot missions. Surface reconstruction in the maritime setting is a relatively new area of research [7]. This project involves the novel application of well known surface reconstruction algorithms [16] and data visualization [17] to this new setting and data.

3 Methodology and Algorithms




Fig. 2. A diagram of the major pipeline stages of this project.

As illustrated in Figure 2, this project generally involves a two stage process, data acquisition in the field and off line data post-processing to construct the maps and visualizations. The focus of this paper is the later stages in the pipeline, but we include here a brief mention of the process by which sonar and image/video data is acquired.

3.1 Data Acquisition

At each cistern site, the ROV is lowered into the cistern for data acquisition. During deployment, video images of the cistern are recorded, and stationary sonar scans are obtained. A portion of each horizontal plane scan will be taken while the ROV rests on the bottom of the cistern. Due to scale or occlusions, the ROV must be re-positioned several times to acquire complete coverage of the shape of each cistern. For each horizontal plane, the ROV is positioned such that scans would overlap each other to facilitate mosaicing. Sonar data in three dimensions can be acquired by flying the ROV to different depths and continuing to scan the cistern horizontally. Alternatively, in certain cases, the sonar can be configured to scan in the vertical plane, and the ROV can then be rotated to different headings while resting on the bottom for scans. Figure 1 shows versions of mosaiced horizontal and vertical scans. Control signals, depth and heading measurements are recorded for use in producing the occupancy map. A more complete overview of the data acquisition system can be found in [18].




Fig. 3. The cistern located at Wignacourt Museum College Garden. This was a fairly large and complex cistern, including several pillars and stairs down into the cistern. On the left is shown the visualization system’s GUI and a wire frame version of the mesh extracted from the occupancy grid. On the right is a texture mapped model.

3.2 Occupancy Grid Construction

Once data has been acquired from the deployment, it is fused with a SLAM algorithm to construct 3-dimensional maps. Since little information regarding the cisterns is available a-priori, a feature-based mapping was not used. Instead, an occupancy grid mapping approach was used to represent the map as done in [19]. To be specific, the space being mapped was discretized into a three dimensional grid of cells, and each cell was given a likelihood $p_{i,j,k} \in [0, 1]$ of being occupied.

For a given robot’s position, the cells in an occupancy grid map can be updated via the log likelihood approach [11]. Specifically, the map was updated using a FastSLAM algorithm [11] that fuses a series of time-stamped sensor measurements including: a 360 degree sonar scan, depth measurement, compass measurement, and robot state predicted by either a) a dynamic model of the robot, b) a manual match of overlapping sonar scans, or c) an algorithmic match of overlapping sonar scans. Important to this sensor fusion is the experimentally modeled noise in both the sensors and motion model. Further details about this mapping algorithm will appear in a manuscript under development.

An example of a three dimensional occupancy grid produced from this algorithm is shown in Figure 4. Note that in certain cases where vertical flight proved difficult due to water depth (i.e. less than 2 m) or depth sensor failure, three dimensional maps were constructed by extruding two dimensional maps in the vertical direction.

3.3 Isosurface Extraction

The occupancy grid built during the 3D mapping is used as input data for our surface reconstruction pipeline. The goal of this stage is to produce a geometric

model of the surface of the cistern. To accomplish this, we treat the occupancy grid as a volume of probabilities and extract an isosurface from the occupancy grid using marching cubes [16]. Figure 3 shows the wireframe structure of a mesh extracted from the corresponding occupancy volume. For the surfaces shown, a probability value greater than 0.51 was used to extract the surfaces representing the most likely location of the walls of the cistern.

In order to construct visually appealing and more meaningful geometric models, our surface reconstruction algorithm uses two additional algorithms to enhance the data, volumetric smoothing and diffusion.

Volume Smoothing: To counteract noise in the data, a smoothing function was used as a pre-process on the volume data. The smoothing function is applied locally in a one voxel neighborhood in the volume, and averages the probability values of the surrounding six neighbors and the current grid point's probability. Specifically, for all i, j, k : $p_{i,j,k} = (p_{i,j,k} + p_{i+1,j,k} + p_{i-1,j,k} + p_{i,j+1,k} + p_{i,j-1,k} + p_{i,j,k+1} + p_{i,j,k-1})/7$. All of the images shown include one step of smoothing. Note that applying the smoothing function removes small extraneous noise from the occupancy grid (which can arise from small occlusions to the sonar, e.g. sticks in the water, etc.) which would translate into small disconnected components in the data. Examples of the small disconnected surfaces can be seen on the right side of Figure 7, but have been removed from the same data in Figure 6. Alternate filters, such as a Gaussian will be explored in the future. An additional step is used to remove other small disconnected components from the final geometric model by starting from a random seed face and counting the number of connected faces, f . Any component with $f < threshold$ can be eliminated from the final geometric model.




Fig. 4. Two images of the acquired data of the Mdina Archive. On the left is the occupancy grid produced from the SLAM algorithm. On the right is a visualization of the texture mapped surface reconstruction. This cistern included a main shaft and two bell like structures branching off the shaft.

Volume Diffusion: In more complex cisterns, due to occlusions or incomplete data, gaps in the occupancy volume result in holes in the reconstructed surface. A specific scenario that creates gaps in the occupancy data is when sonar readings are composited from multiple horizontal layers of data with gaps between horizontal slices. Our goal is to construct a single connected surface for this sequence of horizontal scans. In order to form a solid surface, we experimented with a very simplified diffusion function to distribute known surface data to vertical neighboring voxels when assumptions about shape could be made. Our algorithm is a very simplified version, related to work done by Davis et. al in [20]. Future work will include investigations into alternative diffusion functions and hole filling.

In our setting, we can make a number of simplifying assumptions when tackling hole filling due to the fact that the occupancy grid was created with the knowledge of the ROV's orientation and the volume is written in planar sections where 'height' (traversing in j) directly maps to the vertical orientation in the world. We also know that the shape of the exterior of the cistern should be continuous. Thus, in our setting, the function diffuses volume data through the volume only for surface patches of a particular orientation. Current experiments have focused on applying diffusion for any voxel adjacent to an isosurface polygon with a surface normal, \hat{n} , pointing in a vertical direction. Specifically, each polygonal facet of the isosurface is evaluated and if the facet's normal, \hat{n} is within 45 degrees from vertical, then every $p_{i,j,k}$ adjacent to the polygon will diffuse its probability value to vertical neighbors in the volume. Specifically, each $p_{i,j+n,k}$, within some neighborhood of size n , is set to the probability value of $p_{i,j,k}$. Figure 5 shows two steps of diffusion using a step size, n , of ten.



Fig. 5. Three side view images of the Mdina archive, on the right is the archive with no diffusion, then one step, then two. Note that Figure 4 illustrates the final shape of the archive geometric model.

3.4 Texturing and Visualization

Once a satisfactory geometric model has been constructed from the occupancy grid, textured models can be displayed using two different methods; either by traditional texturing using input images of stone (for example from the cistern walls) or by using projective texture mapping from the video obtained during deployment. For traditional texturing, given the enclosed shape of cisterns, we found that both cylindrical and spherical texture coordinates were sufficient when texturing the models.

Projective Textures: One of the goals of this project is the replication of an actual cistern environment using both acquired geometry and textures. To complete this goal, images from the real world acquired with the onboard camera were used for texturing. A completely accurate representation requires that acquired images be placed in the location in the virtual world that corresponds with the precise location in the real world where they were acquired. In addition, these multiple images must be stitched and blended together. In our current implementation, localization of the virtual camera is not computed automatically, and instead an interactive system allows the user to select and place the acquired images in the virtual environment while navigating through the scene. These images are then displayed via projective texture mapping [17].

Due to the fact that numerous projections are necessary to cover the acquired geometry, we sought solutions that did not require rendering the entire scene multiple times per projection. We optimize rendering by only displaying the viewable triangles for each individual projection. Upon creation, all of the faces of the geometry are compared against the current projection's frustum and only those within this boundary are rendered for that projection. As the program iterates, it displays the combination of all of the saved projections and their viewable geometry in the virtual world with each vertex rendered once. This makes runtime smooth and interactive-friendly. During runtime, each projection is oriented using inverse matrix multiplication of the current view and matrix multiplication of the projection's view. Figure 7 shows the results of texturing the geometric data with real images acquired from the on board camera. Note that we found most cisterns to contain fairly murky water, resulting in poor video quality as seen in Figure 7. Future work includes automatic virtual camera placement, image enhancements and better blending of multiple images.

General visualization In order to help the user see and understand the shape of the cistern, our visualization system has the ability to explore and manipulate the surface data. A simple GUI was created to allow the user to rotate, translate, view as a wireframe mesh, and change the probability value for the reconstructed surface (see Figure 3). This enables users to have a good representation of the cistern, while allowing them to examine the surface and construction of the walls. This is especially important to archaeologists who wish to analyze the shape of the cistern. In addition for visualization purposes, a water surface is created

to convey the water level within the cistern (for these visualizations the water height is mapped for appearance only).




Fig. 6. On the left is a geometric model of the cistern located at the Mдина Cathedral Sacristine, Malta. On the right is the Gozo ‘key hole’ cistern. Both models are visualized with representational water and using standard texture mapping.




Fig. 7. Images showing the ‘keyhole cistern’ with projective textures acquired from the video camera on the ROV during flight. On the left is a close-up view within the cistern showing the textured walls. On the right is an overview showing the textured cistern. Note that the water in the cisterns tended to be murky and thus images pulled from the video are likewise murky.

4 Results and Future Work

We have presented the methodology and algorithms used for the acquisition, surface reconstruction and visualization of cisterns located on Malta. For the March

2011 trip to Malta and Gozo, over twenty different water features were explored using the ROV. In this current paper, we show surface reconstruction results from four different cisterns. For a complete database of the cisterns visited, see “<http://users.csc.calpoly.edu/~cmclark/MaltaMapping/sites.html>”.

The physical scale of all grids used for this paper is 0.1m per cell. The four cisterns shown include:

- Figure 6, shows the cistern located at Mdina Cathedral Sacristine (Malta). This figure shows the surface reconstructed from the occupancy grid, textured and with a water visualization included. This model was generated from an occupancy grid of 120x80x20 cells.
- Figure 6 and 7 shows the cistern located at House Dar Ta’Anna - Upper courtyard, Gozo, which is a keyhole shaped cistern. These models were generated from an occupancy grid of 122x117x11 cells.
- Figure 3 is the cistern at Wignacourt Museum College Garden (Malta). It was situated in a courtyard with the cistern access in a centrally located stone structure. The cistern itself is fairly complex and connects to another cistern access site on the same grounds in the Wignacourt Museum hallway. It includes stairs entering the cistern at one end and several pillars. This model was extracted from an occupancy grid of size 60x150x15.
- Figure 4 shows results from sonar scans of the Mdina Archives. The archive’s cistern is located in a small courtyard directly facing the main entrance to the building. The ROV was lowered down a shaft which opened up to two bell shaped chambers. The occupancy grid for the Archive data is 135x145x69.

These visualizations will be used by archeologists in order to explore the shape of the cisterns and in educational material about these hidden underwater structure in Malta.

Future work on this project includes further enhancements on surface reconstruction, such as improved handling of holes in the data or the addition of ‘caps’ on top of the cisterns and improvements to projective texturing including registering video data for more accurate texturing of the extracted surface. In addition, we would like to include uncertainty visualization to better depict regions of uncertainty or error based on acquired measurements.

5 Acknowledements

This material is based upon work supported by the National Science Foundation under Grant No. 0966608.

References

1. Blouet, B.: The story of malta. In: Allied Publications. (2007)
2. White, C., Hiranandani, D., Olstad, C., Buhagiar, K., Gabmin, T., Clark, C.: The malta cistern mapping project: Underwater robot mapping and localization within ancient tunnel systems. In: Journal of Field Robotics. (2010)

3. Williams, S., Newman, P., Dissanayake, G., Durrant-Whyte, H.: Autonomous underwater simultaneous localization and map building. In: Proceedings of the 2000 IEEE International Conference. (2000)
4. Fairfield, N., Kantor, G., Wettergreen, D.: Three dimensional evidence grids for SLAM in complex underwater environments. In: Proceedings of the 14th International Symposium of Unmanned Untethered Submersible Technology (UUST). (2005)
5. Fairfield, N., Kantor, G., Wettergreen, D.: Real-time SLAM with octree evidence grids for exploration in underwater tunnels. In: Journal of Field Robotics, Vol 24, Issue 1-2, pp. 03-21. (2006)
6. Clark, C.M., Olstad, C., Buhagiar, K., Gambin, T.: Archaeology via underwater robots: Mapping and localization within maltese cistern systems. In: Proc. of the 10th International Conference on Control, Automation, Robotics and Vision (ICARCV 08). (2008)
7. Pizarro, O., Eustice, R.M., Singh, H.: Large area 3D reconstructions from underwater optical surveys. IEEE Journal of Oceanic Engineering (2009) In Press.
8. Williams, S.B., Pizarro, O., Jakuba, M., Barrett, N.: AUV benthic habitat mapping in South Eastern Tasmania. In: 7th International Conference on Field and Service Robotics. (2009)
9. Ribas, D., Ridao, P., Neira, J., Tardos, J.: SLAM using an imaging sonar for partially structured underwater environments. In: IEEE/RSJ International Conference on Intelligent Robots and Systems - IROS. (2006)
10. Conte, G., Zanoli, S., Gambella, L.: Acoustic mapping and localization of an ROV. In: 14th Mediterranean Conference on Control and Automation, 1-6. (2006)
11. Thrun, S., Burgard, W., Fox, D.: Probabilistic robotics. In: MIT Press. (2005)
12. Hernandez, E., Ridao, P., Ribas, D., Batlle, J.: Mispic: A probabilistic scan matching algorithm using a mechanical scanned imaging sonar. In: Journal of Physical Agents 3:311. (2009)
13. Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: Proceedings of the 23rd Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH '96. ACM Press (1996) 303-312
14. Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L., Ginzton, M., Anderson, S., Davis, J., Ginsberg, J., Shade, J., Fulk, D.: The digital Michelangelo project: 3D scanning of large statues. In: SIGGRAPH 2000. (2000)
15. Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W.: Surface reconstruction from unorganized points. In: SIGGRAPH 92. Volume 26. (1992) 7178
16. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. In: Proceedings of the 14th Annual Conference on Computer Graphics and interactive Techniques. (1987) 163-169
17. Segal, M., Korobkin, C., van Widenfelt, R., Foran, J., Haeberli, P.: Fast shadows and lighting effects using texture mapping. In: In Computer Graphics (SIGGRAPH 92 Proceedings). (1992)
18. Hiranandani, D., White, C., Clark, C., Gambin, T., Buhagiar, K.: Underwater robots with sonar and smart tether for underground cistern mapping and exploration. In: The 10th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST. (2009)
19. Moravec, H.: Sensor fusion in certainty grids for mobile robots. In: AI Magazine, Vol 9, pp. 61-74. (1988)
20. Davis, J., Marschner, S., Garr, M., Levoy, M.: Filling holes in complex surfaces using volumetric diffusion. In: 1st International Conference on 3D Data Processing, Visualization, and Transmission. (2002)