First Results in Autonomous Mapping and Retrieval of Buried Objects

Herman Herman & Sanjiv Singh

Field Robotics Center
Carnegie Mellon University
Pittsburgh, PA 15213-3890 USA

I. INTRODUCTION

We are developing the means to autonomously retrieve buried objects. The most immediate application of this technology is to cleanup operations at sites where toxic and nuclear wastes have been improperly stored. Many cleanup efforts are progressing slowly because they expose human workers to substantial risk. There is also substantial interest in the autonomous retrieval of buried unexploded ordnance at military testing facilities. This paper presents recent results in our development of an autonomous system that detects, locates and retrieves buried objects.

Our system consists of an excavator equipped with surface and subsurface sensors. We use the subsurface sensor to locate buried objects and the surface sensor to model the shape of the terrain. If a buried object is close enough to the surface, we use the excavator to retrieve the object. Otherwise we roughly locate the object and use the excavator to remove a layer of soil. This "sense and dig cycle" is repeated until the object is satisfactorily located. Using this procedure, we do not need very accurate information about a buried object until it is close to the surface. After soil has been removed, the buried objects are closer to the surface and they can be located more accurately than it is possible with strategies that do not modify the terrain.

We use ground penetrating radar (GPR) for subsurface sensing. GPR provides high range resolution to objects underground but has low angular resolution. The result is that it is not easy to spot buried objects from the raw data. We use a series of steps to localize buried objects with high accuracy. Rough location of a buried object provides a prescription of a trench directly above the object. We would like to automatically plan and execute digs to produce such trenches. The proposed approach optimizes the amount of soil excavated at every dig while respecting the constraints imposed by the mechanism, the terrain-mechanism interaction and the goal specified.

II. RELATION TO OTHER WORK

The authors are not aware of any other systems that autonomously perform the task of retrieval of buried objects, though there has been a substantial work in the component technologies. We will discuss how the reported work compares to these various technologies under the categories of Subsurface Sensing and Excavation.

A. Subsurface Sensing

Researchers have used subsurface sensing system for a multitude of tasks such as probing underground caves[1], detecting mines[2], mapping archaeological sites[3], measuring the thickness of ice[4] and locating people buried under snow[5]. The disadvantage of such systems is that they are slow and expose human operators to danger. Further, the output must be manually interpreted and the quality of interpretation depends on the expertise of the operator[1]. All of these systems enhance raw GPR data to make it easier for human interpretation. Numerous methods have been used for enhancement. Fisher[6] used reverse time migration, which is based on the technique used in seismography. Osumi et al. have used microwave holography [7]. Some researchers are looking for the solution to the electromagnetic wave inverse scattering problem, which in general is very difficult to solve. Moghaddam [8] has made a promising advance in this field, but it has been demonstrated only for simulated 2-D objects.

We use some of these standard methods and have developed some new methods to enhance the GPR data. The end result is that we can localize and measure buried objects automatically without human intervention to within a few centimeters.

B. Excavation

There has been some interest in remote controlled excavators for construction and hazardous waste site remediation

[9][10][11]. The aim of this work is to remove the operator from the immediate workspace. All motions of the excavator, however, must be orchestrated by an operator who uses a master manipulator to instruct the excavator, typically by looking at a video monitor. Recently, several manufacturers of construction equipment have developed coordinated motion control for their excavator[19]. Instead of requiring expert operators who manually control several revolute degrees of freedom to produce straight line motion, these machines can dig the bottoms of trenches with minimal assistance. Some researchers have investigated the execution of previously planned trajectories for an automated excavating robot [12][13][14], and others have sought to develop gross plans for digging [15][16][17]. A few systems have shown greater autonomy[18][19][20]. Yoshinada and Otsubo report a bucket loader that has been programmed to load various kinds of materials from a pile into a conveyor system. Whittaker et al. have demonstrated a novel excavation scheme that used a high pressure air knife was shown to automatically uncover buried pipes without contact with the environment.

We have previously suggested a method to plan digging operations [21] and have reported simulation results. In this paper we report results from a system that we have implemented in our laboratory.

III. OVERVIEW OF SYSTEM

We have developed a testbed to conduct experiments in subsurface sensing and excavation. The testbed consists of a sandbox (2.5m x 2.5m x 1m), a Cincinnati Milacron T3 hydraulic robot outfitted with an excavator bucket and GPR antenna, and a Perceptron laser range finder. The setup of our testbed is shown in Fig. 1.

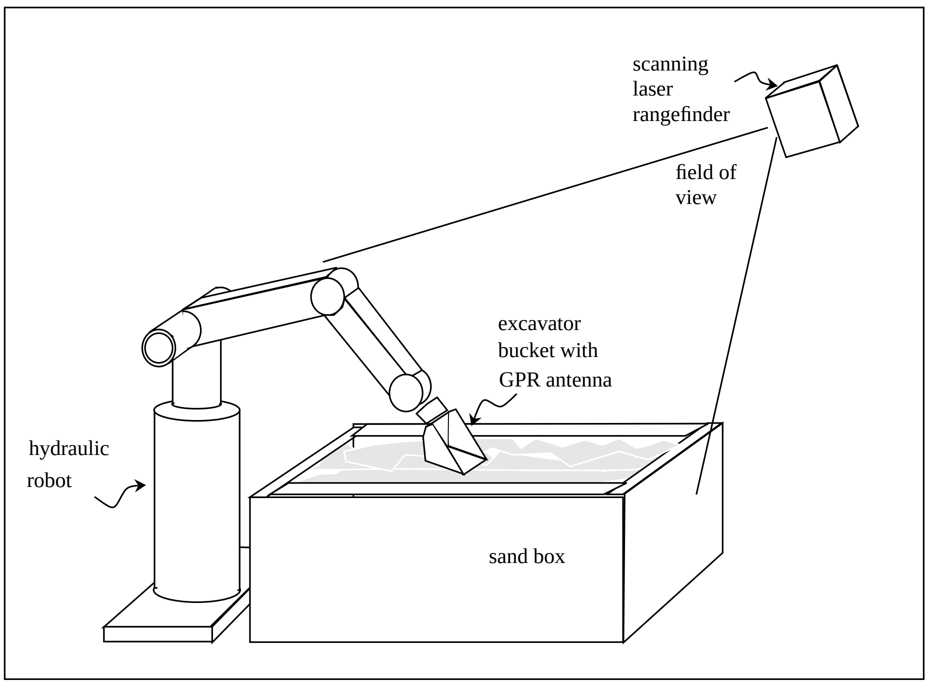

Fig. 1. Testbed

The robot we use is a large industrial manipulator with an end effector payload of approximately 125 lbs. We have attached a small excavator bucket with a volume of 0.02 m^3 as an end effector. We have mounted a GPR antenna underneath the bucket for subsurface sensing.

The laser range scanner produces an image of the terrain such that the value of each pixel in the image is based on the distance from the scanner to the world along a ray that sweeps in a raster fashion. The scanner has a 60 degree horizontal field of view and a 45 degree vertical field of view. We have adapted a perception and mapping system developed at CMU [23] to produce terrain elevation maps for our task. An example

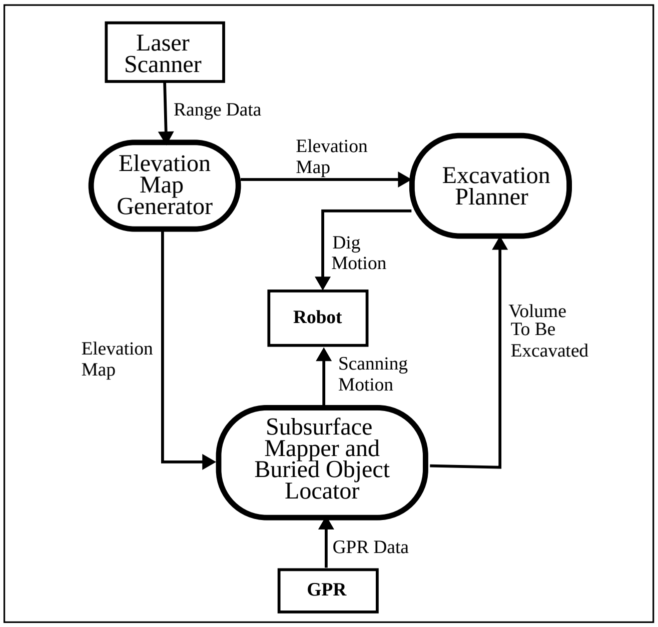

Fig. 3. System diagram

terrain map of the testbed is shown in Fig. 2.

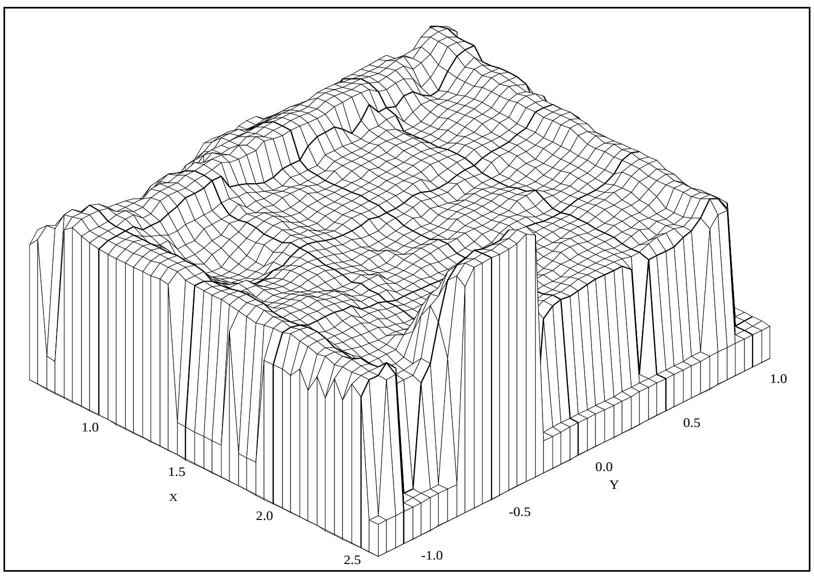

Fig. 2. Terrain map of testbed

Fig. 3. shows how the various modules of our system interact.

IV. Subsurface Sensing

This section discusses the characteristics of Ground Penetrating Radar and our processing methods.

C. Ground Penetrating Radar

We use Ground Penetrating Radar (GPR) as the main sensing modality. GPR works in a fashion similar to conventional radar except that GPR is used to detect objects underground. GPR is a much more capable and flexible sensor than metal detectors and magnetometers. In addition to detecting buried objects, GPR could also be used to compute the depth and dimensions of buried objects. GPR works by transmitting a short pulse from its antenna and then recording the return signal. Several factors influence the strength of the return signal. They include the distance to the object, the type of material the object is made of, and the orientation of the surface that causes the reflection. Surfaces that are normal to the direction of the transmitted signal create the strongest reflections.

Our GPR unit is equipped with a 1GHz antenna. Its high frequency provides sub-centimeter range resolution. However, the beam width (60 degrees) is large; when a single reflection is obtained, it is only possible to localize the reflecting object to an arc that subtends 60 degrees. As a result, raw GPR images appeared blurred and it is not at all simple to localize the reflecting buried objects by inspection.

D. Processing Method

Our subsurface sensing system is unique in that the position and orientation of the GPR antenna can be controlled automatically by positioning the end effector of the robot. This enables us to quickly scan any area of our testbed.

The first processing method that we use is correlation. We correlate the returned signal with the transmitted signal. The resulting signal contains peaks at the locations where the similarity between the transmitted signal and the returned signal is high. This step is necessary because the peaks in the returned signal does not indicate the beginning of the reflections, but the peaks in the correlated signal do. Without this step, the range information could be off by about half the wavelength of the transmitted pulse. In sand this error is 7.5cm. This processing is only relevant to objects that are buried in shallow soil. If an object is deeply buried, this processing step is omitted since typically soil transmits different frequency EM waves with different velocity. As a result the returned signal has a very different shape compared to the transmitted signal.

The second step is a common synthetic aperture radar technique[24][25]. This technique uses an array of antennae to improve the angular resolution of GPR. This is illustrated in Fig. 4. To detect if there is a target in certain location, the array must be focused to that particular location. This is done by inserting delay *D1,D2* and *D3* (Fig. 4.) before the signals from the individual antenna are summed together. These delays correspond to the different distances that the signal from each antenna has to traverse to the target. For example, in Fig. 4. the reflection signal from the target arrives at antenna A2 before it arrives at antenna A1 and A3 since antenna A2 is closer to the target than antenna A1 and A3. Therefore the output signal from antenna A2 must be delayed so that at the summation process the reflection signals from the same target appear at the same time. The summation is constructive if there is an object in the location in question and destructive otherwise. We have to do this for every possible target location that we are interested in. The resulting data contain a much sharper boundary for the objects due to the focusing effect of the antenna. In our system we simulate the effect of the antenna array by moving a single antenna over a 2-D grid.

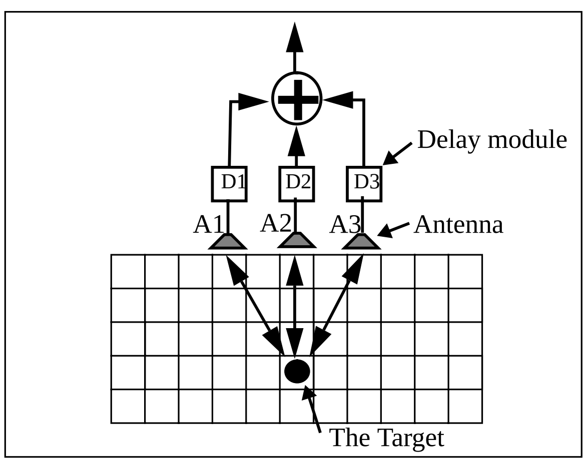

Fig. 4. Synthetic Aperture Antenna Array and its focusing method.

The next step is thresholding the data based on the percentage of positive votes. For example, in Fig. 4.

every antenna position gives a positive vote if it can detect the target. This thresholding method emphasizes consistency rather than strength of the reflection signal. As a result we are able to detect a small object located closed to a large object. If we use a thresholding method based on the strength of the reflected signal, a large object could create a strong enough signal to cover up the reflection caused by a smaller object.

The thresholded data, which represents a 3-D volume are segmented into several connected volumes or blobs. We then determine the most probable target by examining the top surface area and the reflection strength of each blob. This step is sufficient if the targets create a stronger reflection than the noise caused by the heterogeneity in the soil (clutter noise). A more sophisticated model based recognition technique is required to be able to recognize weak targets in presence of substantial clutter noise.

E. Results

TOur subsurface sensing subsystem has reliably detected, located and measured small metal objects (12cmx12cm). The position of the objects is located to within 2 cm while size is located to within 5 cm. Fig. 5. shows one of the configuration of the buried objects that we used to test the subsurface sensing. The first object is a hollow metal cylinder (length=15cm, diameter=12cm) and the second object is a metal plate (length=12cm, width=12cm, thickness=1cm).

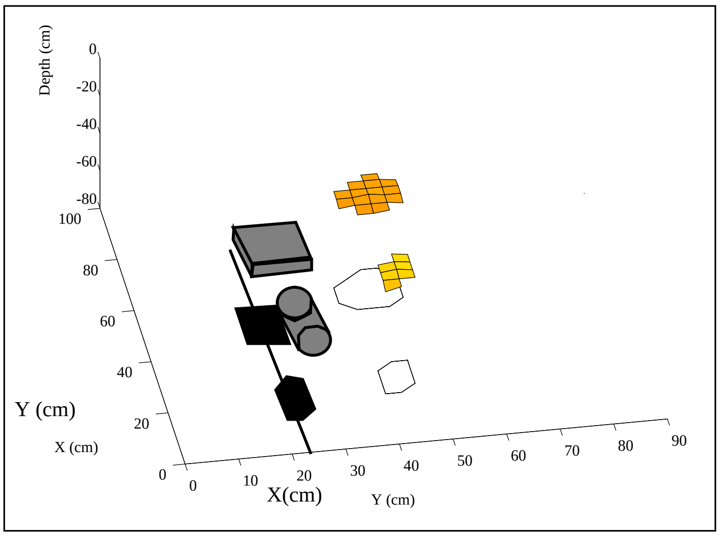

Fig. 5. The configuration of the buried objects

We took radar scans over an area 90 cm x 90 cm by positioning the GPR antenna above the soil, along a

3 cm grid pattern. Fig. 6. shows a vertical cross section of the raw data. The cross section is a plane of con-

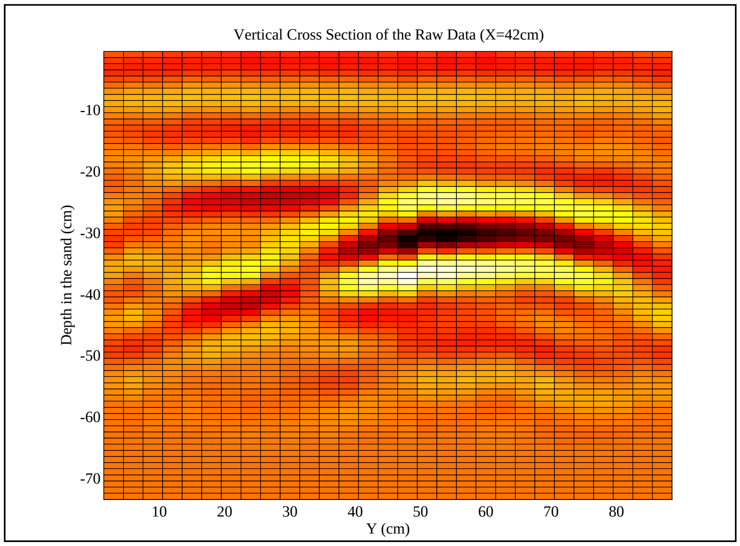

Fig. 6. Raw GPR image of a cylinder and a plate

stant x (x=42cm). The horizontal projection of this plane is shown in Fig. 5. as a line at x=42cm.The same vertical cross section of the processed data is shown in Fig. 7. It clearly shows two buried objects. The resulting clusters are segmented and thresholded. From experimentation we set the threshold so we only select the objects that have a top surface area between half and twice the known top surface area of the target we are looking for. We only use the top surface area because the thickness of the objects could not be measured accurately. This method seems to work well for many small objects if their sizes are known. We are currently working on model based recognition techniques to recognize objects of any size.

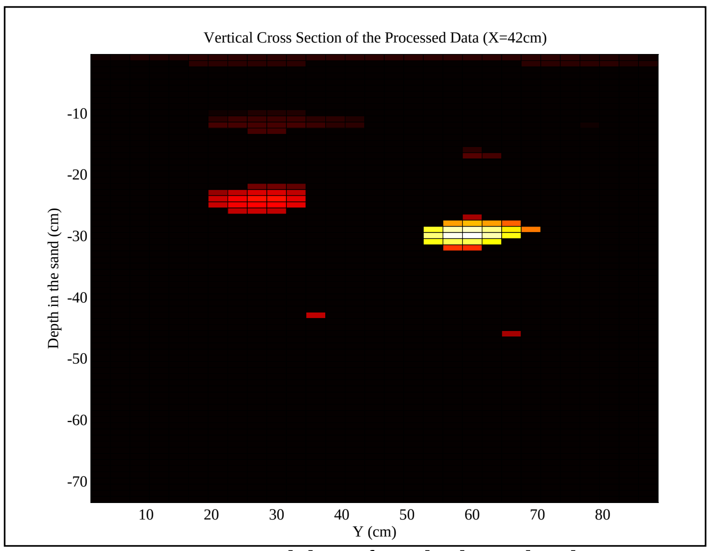

Fig. 7. Processed data of a cylinder and a plate

Fig. 8. shows the two objects that pass the threshold. The measured positions of these two buried objects

are correct, while the sizes are slightly off.

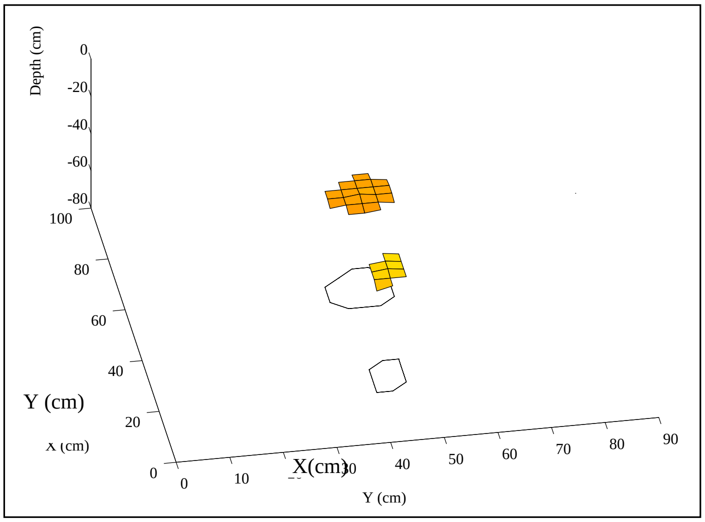

Fig. 8. The two objects that are found after processing

The cylinder is detected as having a diameter of 9.0cm and length of 15.0cm. The plate is detected as having a width of 15.0cm and length of 18.0cm. We suspect that these errors are due to the fact that the antenna size is not infinitely small, an assumption of our processing methods.

The scanned area is 90cm by 90cm and the resulting 3-D data has a voxel size of 3cm by 3cm by 1cm (the depth resolution). The processing time is about 30 minutes on a Sparc 2. We have also developed a parallel version that can be run concurrently on several workstations connected to a network. This is possible since many of the processing steps are local in nature and need to share only the raw GPR data. After the raw GPR data is distributed, the processing can be executed concurrently. Running on 15 Sparc workstations, we are able to reduce processing time to approximately 3 minutes.

V. AUTOMATED EXCAVATION

We have customized our general excavation scheme for trench digging. Excavation proceeds in two steps. First, we use a planner to select digging motions that optimize the amount of soil excavated at each dig. Informally, the one-dig problem can be stated as follows: *Sense the terrain. Of all possible digs, consider a subset that the robot can feasibly perform and choose one that satisfies a bound on a cost criterion. Execute the dig.* Complete excavation is accomplished by concatenating a sequence of such plans. When the amount of soil that can be obtained by a single dig falls below a threshold, the excavator bucket is made to follow the outline of the trench specified. The first stage optimally fills the bucket while the second ensures that the desired shape is achieved.

In this section we will develop the notion of an “action space”, a tool that is used to plan digging motions. We will then illustrate the method with results from experiments conducted on our testbed.

F. Posing Excavation in an Action Space

We will consider digging actions that are parameterized by a small set of variables. For example, consider the variables necessary to describe the dig shown in Fig. 9. This dig can be parameterized by the variables (α, d, k) where α is the angle at which the excavator bucket approaches the soil, k is the distance from some common coordinate frame to the point where the bucket enters the soil and d is the distance for which the bucket drives into the soil. At the end of this distance the bucket is made to curl up out of the soil.

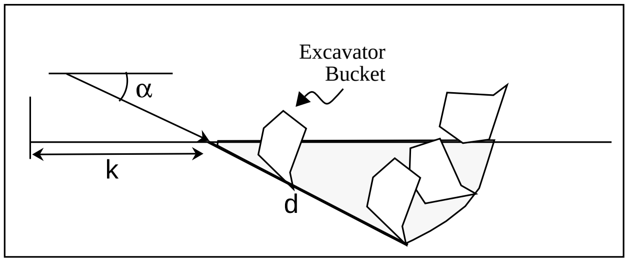

Fig. 9. A prototypical dig.

The values that these variables (α, *d*, *k*) can attain define a multi-dimensional space of possible actions and represent the set of all digs under consideration, independent of the terrain and the excavator. We call this an *action space*. From this space we will want to exclude certain digs because they are not feasible. A dig may not be feasible if the excavator is required to reach outside its workspace, excavates more soil than the bucket can hold, or, if the bucket is required to intrude past specified geometric boundaries. Among the set that satisfies all the constraints, it remains to find one dig that satisfies a threshold on a cost function, for example, the amount of soil excavated. Fig. 10. shows the action space spanned by the variables (α, *d*, *k*) and a hypothetical surface that constrains the set of feasible digs for a particular excavating robot and a particular terrain. Here, the constrained volume lies beneath the surface shown. Note that the constraint surface changes as the excavation proceeds and in general, the constrained volume need not be a single, connected space.

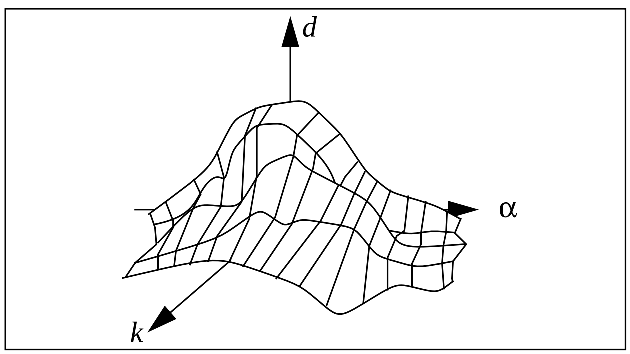

Fig. 10. A three variable action space.

If the action space is small it is possible to find the "optimal" dig by enumerating all the options of the discretized space. Otherwise, a numerical method is necessary. Since there is no guarantee of a unique extremum in the cost function, a method like *simulated annealing* can be used to optimize the cost function [22]. Once the optimization procedure has selected a dig, it can be mapped back to the joints of the excavator.

G. Posing Geometric Constraints

For now we will use the following constraints for the robot:

Reachability Constraint: This constraint separates the action space into digs that are kinematically feasible and those that are not. Given a candidate dig, that is a trajectory for the bucket tip to follow, a standard inverse kinematics method is used to find the corresponding joint displacements (θ_1- θ_6) along the trajectory. A candidate dig may fail this constraint if it requires the robot to reach outside its workspace or if in the course of the dig, any of the links are required to interpenetrate the terrain.

Volume Constraint: Since the excavator bucket can only hold a volume V_{max}, then an (α, *d*, *k*) triplet should not excavate more than this amount of soil. This gives us a further basis on which we can limit the set of feasible digs.

Shaping Constraint: This constraint is given by the goal state of the terrain which may be an arbitrary stable configuration of the earth. We specify this as a elevation profile of the terrain of the completed trench. Any dig that intrudes below this elevation profile is considered inadmissable.

The union of these three constraints delimits the set of digs that are geometrically feasible. It remains to select a plan that optimizes the amount of soil excavated. Although in the past we have used more efficient numerical methods for this optimization, currently a simple exhaustive search of the action space suffices.

H. Results

We have buried a several small metallic objects in our sandbox. The subsurface sensing system estimates the location of an object to be retrieved and prescribes a trench to be excavated. After the trench has been excavated, subsurface sensing confirms the location of the object and a single dig is selected to retrieve the buried object.

In a recent experiment, the excavation module was required to produce a trench 30 cm deep, 20 cm wide and 50 cm long. To produce such a free volume, extra soil must be excavated because the soil in our testbed consists of dry, fine sand. Hence the walls of the trench collapse until the they are angled at less than the natural angle of repose of the soil. The planner continues to dig until the volume prescribed has been completely excavated.

The terrain starts in an approximately level state as shown in Fig. 2. Fig. 11. shows the action space before the first dig given the goal stated above, using the kinematics constraints of the T3 robot.

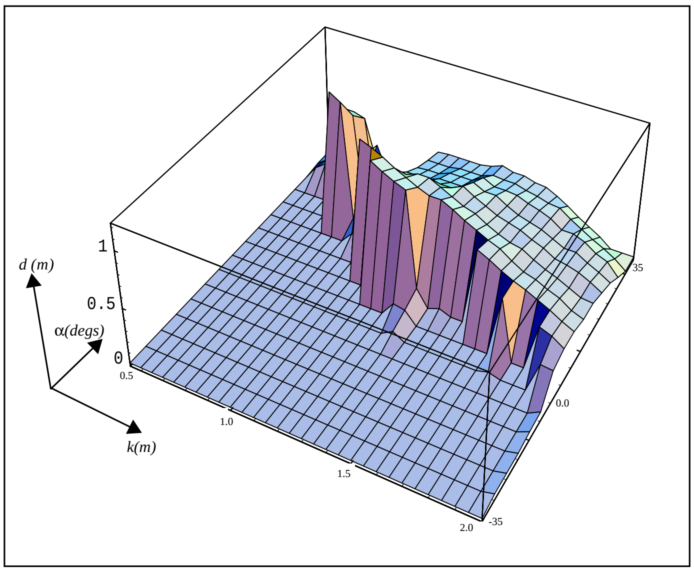

Fig. 11. Action space for the first dig.

Fig. 12. shows a cross-section of the sandbox along the long axis of the trench, as the task progresses. The best dig selected at each step is shown as a triangle that connects the points where the bucket enters the terrain, the point to which it digs to and then the point to which the bucket curls.

In this case, the planner stopped after 11 digs and the second phase was initiated where the excavator bucket follows the profile of the desired trench. This step was repeated three times before the bottom of the trench was clear of soil collapsing from the trench walls. The final shape of the terrain is shown in Fig. 13.

The volume of soil required to be excavated is approximately 0.1 m^3. Given a maximum volume of the bucket of 0.02 m^3, the overall efficiency of our digging scheme is 35.7%, i.e on average, 35% of the bucket was filled at every dig.

As a final step in our experiment, a single dig was chosen such that the flat plate described in the section on subsurface ended up in the bucket.

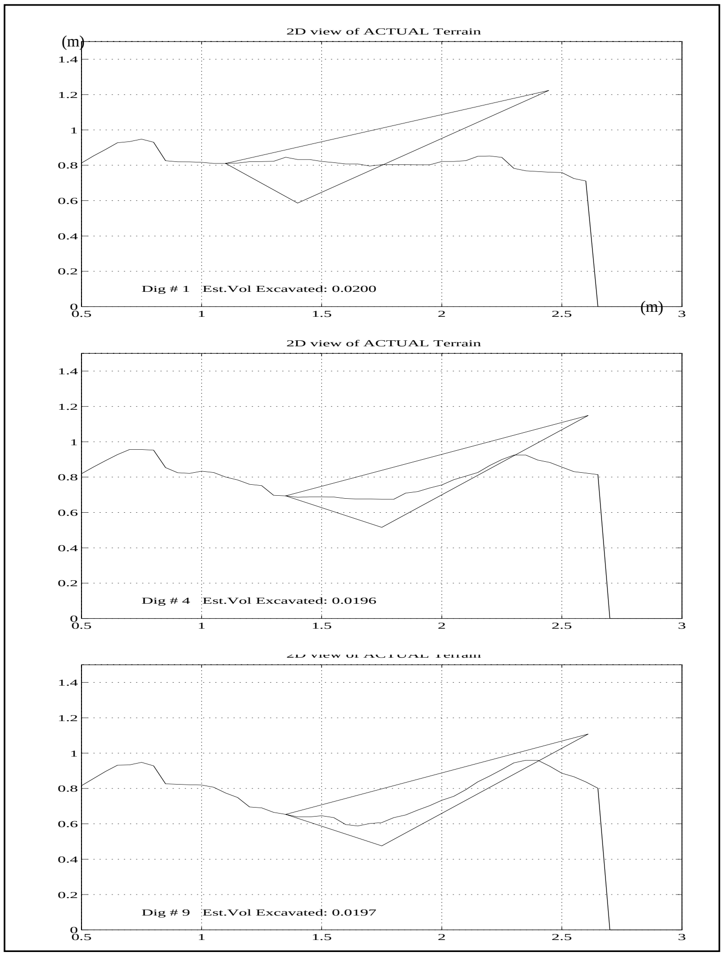

Fig. 12. Progression of terrain.

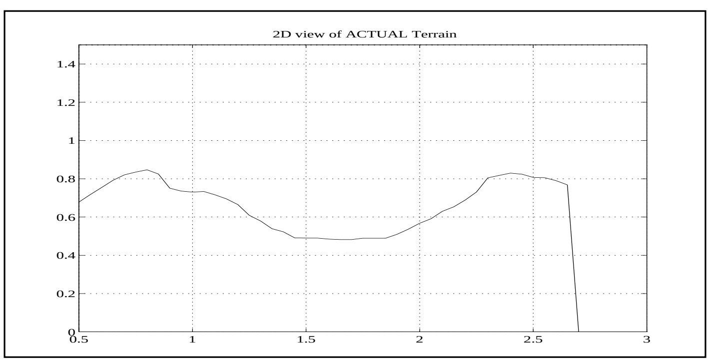

Fig. 13. Final shape of terrain.

VI. Summary

At present we have the capability to locate and retrieve small buried objects. So far, our system can only deal with simple objects— we have some problems with long, thin objects. Our method sometimes splits such objects into several blobs. We are currently working on several areas to improve the performance of the subsurface sensing system. One of them is to use the elevation map obtained from the laser rangefinder to correct the error that is caused by ground surfaces that are not flat. Another area for improvement involves scanning from several different orientations so we are able to measure the thickness of buried objects as well as their width and length. Ultimately we would like to build a model based recognition system which would also automatically position the antenna for best viewing angle.

Our excavation scheme suffers in efficiency because at present our prototypical digs are very simple. In the future we will use more variables (at the risk of a larger search space) to describe the trajectory of bucket. Further, we hope to incorporate force constraints that limit the digging motions based on the force required to perform a dig. Although the force constraint is difficult to specify accurately, it will allow for a more efficient strategy; currently we artificially limit the maximum volume that a dig might excavate, to keep the robot from encountering trajectories that it cannot execute due to force limitations.

ACKNOWLEDGEMENTS

The authors would like to thank Regis Hoffman for his help with perception and mapping. This work was sponsored by the Geomechanical, Geotechnical and Geo-Environmental Systems program of the National Science Foundation (number 9114674), and by the Department of the Air Force (number F08635-92-C-0019).

REFERENCES

[1] Peters, L. Jr., and Young, J.D., “Application of Subsurface Transient Radar” Chapter 9 from Time Domain Measurements in Electromagnetic, Ed. By Edmund K. Miller.

[2] Ozdemir, T., Roy, S.,Berkowitz, R.S., “Imaging of Shallow Subsurface Objects: An Experimental Investigation”, IEEE Transactions on Geoscience and Remote Sensing, Vol 30, No3, May 1992, pp. 472-481.

[3] Imai, T., Sakayama T., Kanemori, T., “Use of Ground-probing Radar and Resistivity Surveys for Archaeological Investigations”, Geophysics, Vol 52, February 1987, pp. 137-150.

[4] Ulricksen, C.P.F., “Application of Impulse Radar to Civil Engineering”, Doctoral Thesis, Department of Engineering Geology, Lund University of Technology, Sweden, 1982.

[5] Yamaguchi, Y., Maruyama, Y., Kawakami, A., Sengoku, M., Abe, T., “Detection of Objects Buried in Wet Snowpack by an FM-CW Radar”, IEEE Transactions on Geoscience and Remote Sensing, Vol 29, No 2, March 1991, pp. 201-208.

[6] Fisher, E., McMechan, G.A., Annan, P.A., Cosway, S.W., “Reverse-time migration of single-channel ground penetrating radar profiles”. The University of Texas at Dallas, Center for Lithospheric Studies.

[7] Osumi, N. and Ueno, K.,”Microwave Holographics Imaging of Underground Objects”, IEEE Transactions on Antennas and Propagation, Vol AP-33, No 2, February 1985, pp. 152-159.

[8] Moghaddam, M. and Chew, W.C., “Nonlinear Two-Dimensional Velocity Profile Inversion Using Time Domain Data”, IEEE Transactions on GEoscience and Remote Sensing, Vol 30, No 1, January 1992, pp. 147-156.

[9] Muramatsu, T., Kenji, C., “A Concept of Control System For Construction Robot,” in *Proc. IEEE/RSJ International Workshop on Intelligent Robots and Systems*, September, 1989.

[10] Ostaja-Starzewski, M., Skibiniewski, M., “A Master Slave Manipulator for Excavation and Construction Tasks,” Robotics and Autonomous Systems, North Holland, Vol 4, 1988/89.

[11] ”Force Feedback Excavator and Material Handling System,” Kraft TeleRobotics, Inc. Overland Park, Kansas, July 1990.

[12] Bernhold, L., “Experimental Studies on Mechanics of Lunar Excavation,” ACSE Journal of Aerospace Engineering, Vol 4 (1), January, 1991.

[13] Bullock, D M., Oppenheim, I. J., “A Laboratory Study of Force-Cognitive Excavation”, in *Proc. Six International Symposium on Automation and Robotics in Construction*, June 1989.

[14] Vaha, P.K., Skibiniewski, M. J., Koivo, A. J., “Kinematics and Trajectory Planning for Robotic Excavation,” in *Proc. ASCE Construction Congress II*, Cambridge, MA 1991.

[15] Apte, S. M., “A Representation for Spatial Constraints in Material Removal,” Ph.D. thesis, Civil Engineering Department, Carnegie-Mellon University, July 1992

[16] Baker, N., “Simulation of a Monitor/Controller for a Robotic Excavator,” Master’s Thesis, Department of Civil Engineering, Carnegie Mellon University, Pittsburgh, PA, 1986.

[17] Romero-Lois, H., Hendrickson, C., Oppenheim, I., “A Strategic Planner for Robot Excavation,” in *Proc.*

Sixth International Symposium on Automation and Robotics in Construction, June 1989.

[18]Sakai, T., Cho, K., "Operation System for Hydraulic Excavator for Deep Trench Works," *In Proc. of the 5th International Symposium on Robotics in Construction*, Tokyo, June 1988.

[19]Yoshinada H., Otsubo K, Komatsu Corporation, Personal Communication, August 1992.

[20]Whittaker, W., Turkiyyah, G., Bitz, F., Balash, J., Guzikowski, R., Montgomery, B., Akdogan, R., and Swetz, S., "First Results in Automated Pipe Excavation," in *Proc. of the Second International Conference on Robotics in Construction*, Pittsburgh, PA, May 1985.

[21]Singh, S., Simmons, R., "Task Planning for Robotic Excavation," In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, Raleigh, 1992.

[22]Singh, S., "An Operation Space Approach to Robotic Excavation," in *Proc. IEEE Symposium on Intelligent Control*, Alexandria, August 1991.

[23]Hoffman, R., and Krotkov, E., "Terrain Mapping for Long Duration Autonomous Walking", In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, Raleigh, 1992.

[24]Fitch, "Synthetic Aperture Radar", Springer-Verlag, 1988.

[25]Eaves and Reedy, "Principle of Modern Radar", Van Nostrand Reinhold, 1987.