In Proceedings I*nternational Conference on Robotics and Automation,* May 1998, Leuven, Belgium

Multi-Resolution Planning for Earthmoving

Sanjiv Singh and Howard Cannon

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Abstract

We suggest that planning for automated earthmoving operations such as digging a foundation or leveling a mound of soil, be treated at multiple levels. In a system that we have developed, a coarse-level planner is used to tessellate the volume to be excavated into smaller pieces that are sequenced in order to complete the task efficiently. Each of the smaller volumes is treated with a refined planner that selects digging actions based on constraint optimization over the space of prototypical digging actions. We discuss planners and the associated representations for two types of earthmoving machines: an excavator backhoe and a wheel loader. Experimental results from a full-scale automated excavator and simulated wheel loader are presented.

1 Introduction

Automated earthmoving operations such as digging a trench or leveling a mound of soil are distinguished from typical planning problems in two important ways. First, soil is diffuse and therefore a unique description of the world requires a very large number of variables. Second, the interaction between the robot and the world is very complex and only approximate models that are also computationally tractable are available. In previous work we have shown that the large state space and complex robot-world interaction imply that only locally optimal planners can be created [12]. That is, it is not computationally tractable for a planner in this domain to generate an optimal series of digging actions. In order to deal with the practical issues of excavating large volumes of earth in applications such as shown in Fig. 1, we have developed a multi-resolution planning and execution scheme.

At the highest level is a coarse planning scheme that uses geometry of the site and the goal configuration of the terrain to plan a sequence of “dig regions”. In turn, each dig region is searched for the “best” dig that can be executed and finally the selected dig is executed using a force based closed loop control scheme. Treatment of the problem with a layered system meets several objectives. The coarse planner ensures even performance over a large number of digs. The refined planner chooses digs that meet constraints of geometry (reachability and collisions) and optimize a cost function (e.g. volume, energy, time). At the lowest level is a control scheme that is robust to errors in sensing the geometry of the terrain and is effective in capturing material. Empirical data support the contention that the multi-resolution approach can provide near-optimal behavior over an extended sequence of operations.

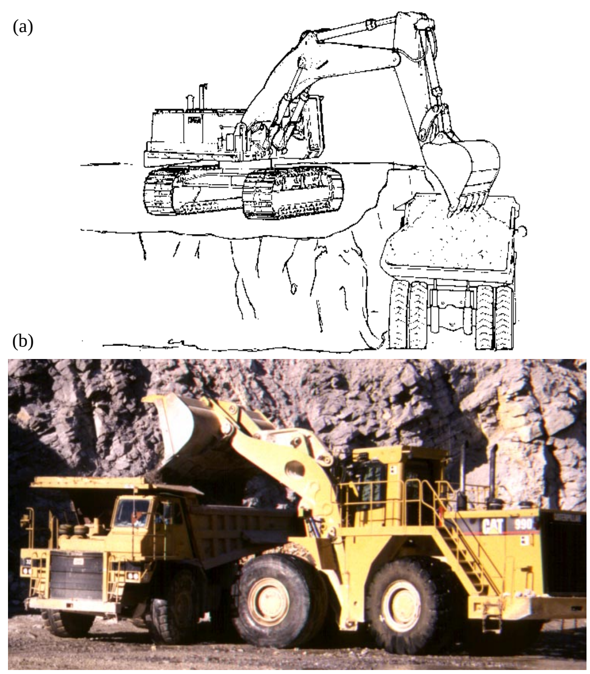

Fig. 1 Truck loading with (a) an excavator backhoe and (b) a wheel loader. We have automated an excavator that performs “bench loading” as in (a).

In this paper we show how the above approach is used to plan and execute earthmoving operations on both an excavator backhoe and a wheel loader. The mechanisms themselves are different and hence the representation used by the coarse and refined planners is quite different, but at an abstracted level, the process is identical.

2 Related Work

We are not aware of other systems that have attempted the complete automation of earthmoving operations such as truck loading, nor, are there any published results on extended autonomous operation over varying terrain (See [13] for a survey of automation in earthmoving.) Primarily, this is because the actions required of earthmovers change as the terrain is modified, making environmental perception necessary. However, such perception in the conditions that exist at real worksites is a difficult problem. Hence, to date, most automated earthmoving systems have concentrated on the problem of "how to dig" rather than "where to dig". Typically, these systems rely on a human operator to deal with the latter issue [3][5][6][7][9], or, assume that the terrain doesn't change significantly from dig to dig [1]. Once in contact with the terrain, a controller takes over the process of filling the bucket, using force and/or joint position feedback to accomplish the task. A few researchers have proposed methods of selecting digs automatically. Such systems measure the topology of the terrain using ranging sensors [2][15]. Given the profile of the terrain, "optimal" digs, or those that maximize excavated volume while minimizing other criteria such as time and energy, are computed. Previously, we have reported a system that uses a constraint optimization method to select digging actions [11][12] for planar tasks such as excavating a trench. We now report an extension that allows consideration of larger volumes of soil. Our planners are a part of a larger system to deal with the complete task of autonomously loading trucks [14].

In Section 3 we present the representations used by the coarse and refined planning methods for two tasks. Section 4 presents experimental results from the testbed for the excavator and from a simulator for the wheel loader.

3 Multi-resolution Planning

Fig. 3 shows the process of coarse to fine planning for both the excavator and the wheel loader. The cycle starts with determination of the shape of the terrain from data collected by a range sensor that is scanned over the area to be excavated. An example *terrain map* generated by sensors onboard the excavator is shown in Fig. 2. The coarse planner takes as input the terrain map (a 2 D array of height values) and the location of the truck. The output is a sequence of regions, each of which is in turn sent to a refined planner. The refined planner operates on an abstract representation of an atomic action, a single dig. Rather than searching for a bucket trajectory, the refined planner searches for the best action within the bounds specified by the coarse planner. The refined planner then evaluates candidates using a forward model that simulates the result of choosing an action (in our case the starting location of the bucket). An evaluation function scores the trajectory resulting from each action, and the best action is

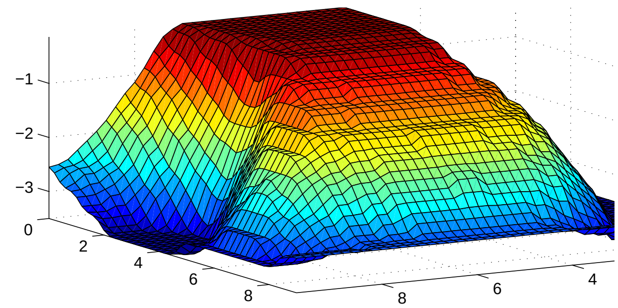

Fig. 2 A sample terrain map generated by sensors onboard the excavator. Such a map is constructed before every dig. The excavator is in the configuration shown in Fig. 5.

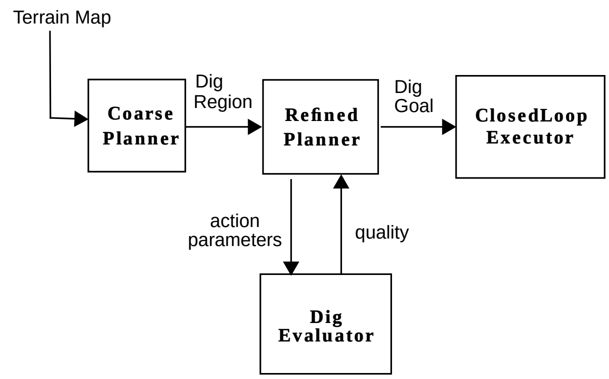

Fig. 3 Coarse to fine planning strategy.

Although in some case backhoes and wheel loaders are used for the same applications, the operation of these machines is very different. Backhoes, by their design are stationary during a digging operation, but they can be used to dig in many different parts of a dig face. In contrast wheel loaders are required to move while digging, but almost always start digging at the base of the pile being excavated. This difference is shown in Fig. 4. It is not sur-

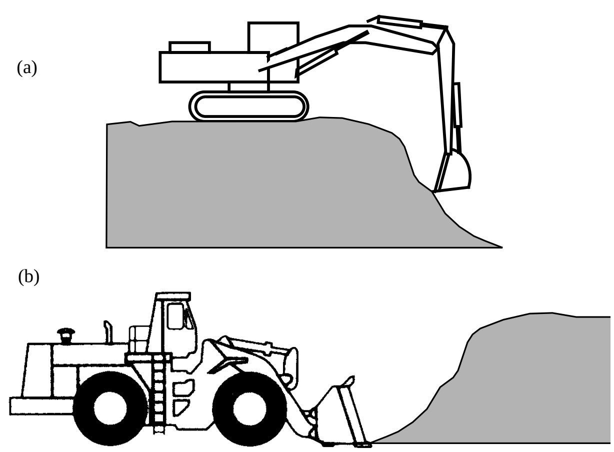

Fig. 4 Nominal operation for (a) backhoe (b) wheel loader. These tasks are encoded with different representations.

prising then that the representations used by the planners for the respective machines are quite different. Below we discuss these representations and the methods used.

3.1 Planning Backhoe Operations

In a "bench loading" loading scenario (Fig. 1a) an excavator sits on a raised section of the terrain and digs in the face of the bench. The material is then dumped into a nearby truck. For this operation, a coarse planner tessellates the face into smaller regions and provides this region's limits to a refined planner. The refined planner searches within these limits for locally optimal action parameters.

3.1.1 Coarse planner

The coarse planner provides a strategy for removing material as shown in Fig. 5. Material is removed from left to the right and from the top to the bottom of the face.

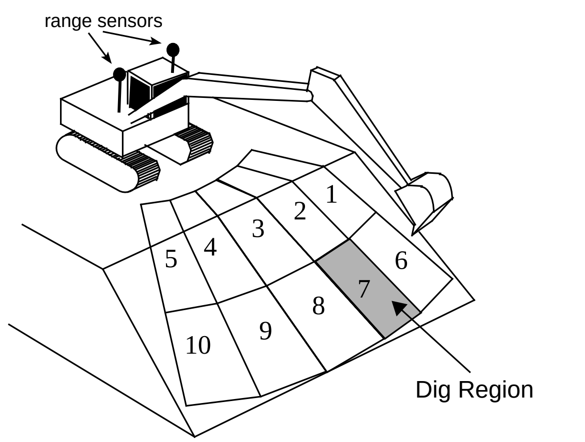

Fig. 5 A coarse plan for an excavator performing bench loading. The strategy encoded is based on recommendations from expert human operators. The dig face is tessellated into regions as a function of the terrain geometry and heuristics that ensure stability of the excavator while digging.

There are several advantages to this strategy. Progressively moving across the bench produces an even erosion. Proceeding from top to bottom reduces the forces required in the lower regions because the weight of the material in the upper regions is eliminated from soil reaction forces. Clearing material away from the upper regions also minimizes areas that are occluded from the field of view of the range sensors that are mounted on the excavator cab. The number of regions selected by the coarse planner and their extents are automatically determined by the shape of the dig-face and the dimensions of the excavator. A new coarse plan is made whenever the excavator is moved, an event that is triggered when no feasible actions can be found.

3.1.2 Refined Planner

The task of the refined planner is to search through the set of feasible actions within the bounds specified and select the locally optimal action. Since control of digging is governed by a closed-loop controller [6] after the bucket enters the earth, the refined planner need only search for the pose of the bucket as it enters the ground. Our experience is that the utility of a dig is sensitive to the starting pose especially as the terrain undulates, hence it is necessary to evaluate a diverse set of starting poses.

Since the process of excavation is planar, the starting pose can be found by searching over the space defined by two action parameters: the distance along the radial line from the excavator, d, and, the approach angle of the bucket, α, as shown in Fig. 6. This pair of parameters is evaluated in two ways. First, a candidate pair is checked for feasibility. That is, does the trajectory implied by the pair satisfy all constraints? Second, the utility of a candidate action is computed to select the action that has the best utility. Both processes require the prediction of the outcome of a selected action, i.e a forward simulation of the trajectory generated by the closed loop controller. For the excavator, we use a model of the closed loop controller which predicts the bucket trajectory given the initial pose of the bucket and the shape of the terrain. To accomplish this, the model predicts the resistive forces that the bucket encounters during digging [4]. In addition the model computes the time and energy required to perform the dig, and the amount of soil that is swept into the bucket.

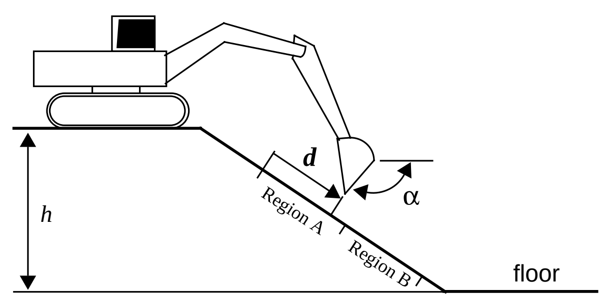

Fig. 6 Action parameters for an excavator. The distance (d) of the bucket from the top of the region, and the bucket angle (α) represent a single action. Bucket trajectory is a function of (d,α) for a given terrain.

A candidate pair that requires the bucket to exceed the workspace is rejected without need for further evaluation. Now the bucket trajectory is analyzed for other constraints. For instance, it is desirable to produce an even "floor" as the excavator digs and hence any trajectories that penetrate the floor depth, h, can be rejected.

The set of actions that meet these constraints are evaluated using a quality metric that is a multiple of three functions

(Fig. 7). These functions are dependent on *volume swept, time* and *energy* required for digging, respectively. When the bucket sweeps less than V_{min}, $\xi = 0$, that is, all actions that sweep less than V_{min} are vetoed. Above V_{max} ξ is constant because V_{max} represents bucket capacity and no additional value is attached to sweeping more material. Similarly, the ψ and ζ functions decrease linearly as the time and energy required to dig increases and have similar cut-offs beyond which the overall quality drops to zero.

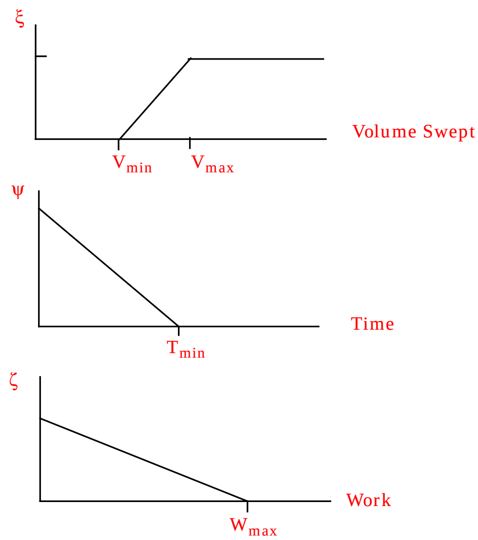

Fig. 7 Evaluation criteria for selecting a dig for the excavator. Overall quality is a function of volume swept, ξ, the energy required, ψ, and time, ζ, required to dig.

3.2 Planning Wheel-loader Operations

Another scenario for autonomous operation is an application in which a wheel loader selects dig locations along a wall of material and the material is dumped into a waiting truck as shown in Fig. 8. Since an automated wheel loader

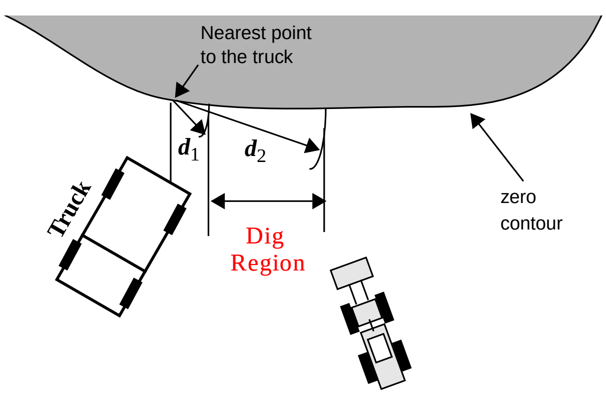

Fig. 8 Coarse Planning methodology for a wheel loader. The wheel loader is constrained to start within the dig region.

was not available, we tested our algorithms in a simulated world. The wheel loader is assumed to have environmental sensors similar to the excavator backhoe that enable it to perceive the shape of the terrain to be excavated to locate the truck.

As in the case of the excavator, a coarse planner is used to reduce the search space. The refined planner considers a number of possible actions each of which is defined by the initial pose of the wheel loader as it enters the terrain.

3.2.1 Coarse Planner

Coarse planning enforces a strategy that minimizes distance traveled between the dig-face and the truck. As with the excavator, the coarse planner finds the boundaries of the material, specifically the contour of the pile at ground level (*zero contour)*. The point on zero contour closest to the truck is located and the coarse plan is defined as the section of the zero contour that lies within a range of distances, d_1 and d_2 (Fig. 8).

Since we have encoded the task as the determination of a digging location relative to the truck whose pose is independently controlled, only one dig region is necessary. Without this constraint the wheel loader could choose to dig anywhere along the wall. In this case it would be necessary to define multiple regions, and the order of region selection would be based on a global strategy to evenly erode the entire pile as with the excavator.

3.2.2 Refined Planner

Once again we assume that a closed loop controller is used to control the bucket while the bucket is in contact with the terrain. Hence, as with the excavator, the action parameters need specify only the starting pose of the machine. Other heuristics may be used to reduce the size of the action space. For instance, to reduce tire damage from loose rocks, wheel loaders typically start digging with the bucket flat and on the ground. Another heuristic is that the wheel loader should dig perpendicular to the wall to prevent uneven loading of the bucket. This is ensured by requiring that both front corners of the bucket touch the zero contour and thus constraining the heading. Hence, in this case, the action space is a set of one dimensional distances along the dig region (Fig. 9).

Since at the moment no closed-loop controller exists for wheel-loader digging at the moment, evaluation of actions is based only on the starting location from there the wheel loader dig, as opposed to the predicted closed-loop trajectory through the soil. Three criteria are used: *side loading*, *concavity* and *location*. First, the selected dig location should minimize side loading of the bucket. The side loading criteria is calculated by examining the volume of material inside and in front of the bucket. The function ξ has a greater value when the volume is more equally distributed.

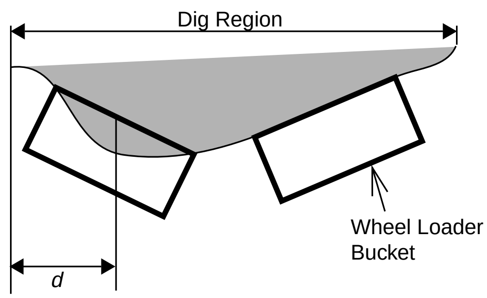

Fig. 9 Wheel loader loading requires only one action parameter, *d*. The heading or angle of attack is dictated by the terrain profile.

Second, digging should occur at locations where the terrain protrudes away from the rest of the dig face, and avoid areas that are recessed. This improves the efficiency of the dig because the force applied by the wheel loader is directed to the cutting edge of the bucket, and not to the side walls. If a surface is highly curved or concave, there will be more volume in the bucket at its starting location than if the surface were flat or recessed inwards. Therefore ψ is simply a ratio of the volume of material in the bucket to the maximum bucket capacity. And third, dig locations should be chosen as close as possible to the truck for the sake of productivity. The function ζ decreases as the travel distance required increases. Fig. 10 shows the evaluation functions which are used to calculate the overall quality of a given dig location.

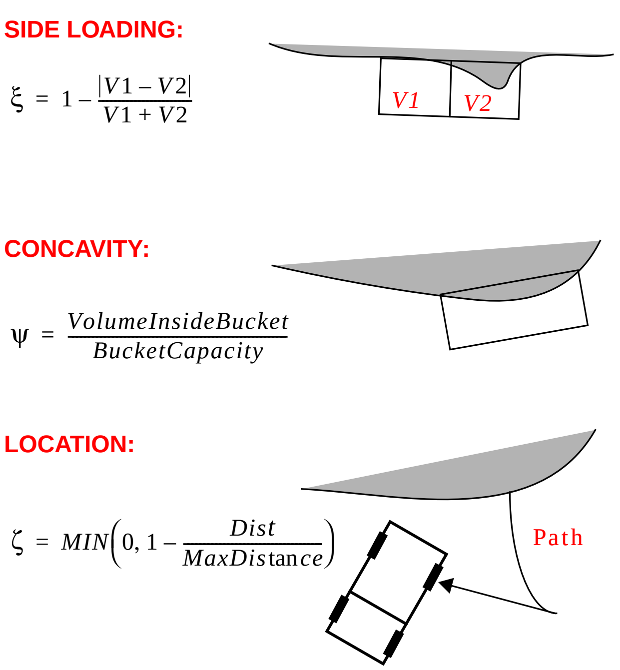

Fig. 10 Wheel loader dig evaluation criteria. The overall evaluation of an action is a product of three functions that measure the utility of choosing a particular location to start digging.

Apart from the evaluation functions above, two hard constraints are placed on the trajectories of the wheel loader that are dictated by the choice of the starting location. First, it should be possible for the trajectory to be represented by two arcs with a turning radius greater than the minimum turning radius that the wheel loader is able to execute. Second, the trajectory of the wheel loader must not cause a collision with the truck or the pile. Any actions that violate these constraints are rejected.

4 Results

4.1 Excavator

Fig. 11 shows the results of one of our experiments on the excavator. Each graph shows the profile of the terrain

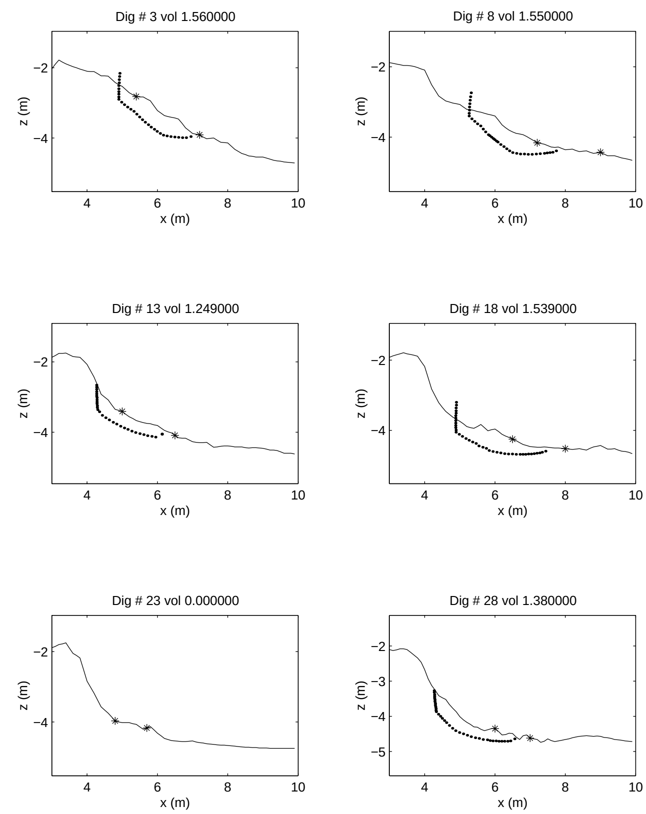

Fig. 11 Side profiles of the face of terrain being excavated along a long section of the bench (fixed swing angle). Dotted line shows the predicted trajectory of the bucket. The *s show the bounds of the coarse plan. A total of 28 digs were done on the bench although only five digs were done along the shown long section.

along a long section of the bench during a sequence of digs. The profile was constructed from range data taken from sensors on the excavator. Note that by the fifth dig along the long section (23rd dig on the bench) most of the material had been removed. The refined planner did not select a dig within this coarse plan because none of the actions evaluated met the requirements for a suitable dig.

However the material was eventually removed by the next dig along the long section because the coarse plan limits allowed the machine to reach farther out. The coarse to fine planning methodology has allowed the excavator to remove most of the material from a given excavator location while still obtaining a full bucket with every dig.

This method has been used as a part of a system that automatically loads trucks at our test site. A complete loading cycle can be accomplished in less than a minute on average and is about as fast as the operation by an expert operator. Dig selection via evaluation of approximately 200 starting poses, typically takes 2-3 seconds of a cycle and can be completed without pausing the machine.

4.2 Wheel Loader

Fig. 12 illustrates the performance of the wheel loader dig planning over the first eight digs during a run with simulated terrain. Each graph shows a top view of the terrain profile as it is being eroded by a wheel loader bucket. It is possible to observe the wheel loader digging perpendicular to the terrain profile, and at the center of protrusions.

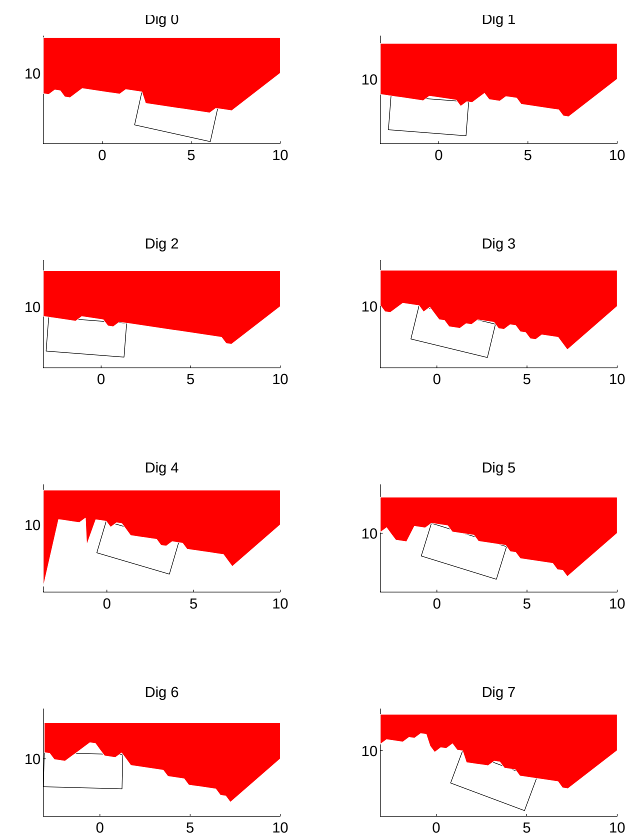

Fig. 12 Top view of the terrain profile being excavated by a wheel loader bucket. The wheel loader digs perpendicular to the wall, and the bucket is centered on protrusions

5 Conclusions

We have used a layered system as a means of producing plans for earthmovers during extended operations. In this methodology, a coarse planner is used provide overall strategy by encoding heuristics used by expert operators. A refined planner selects locally optimal actions. Experimental results to date are encouraging. Future work will investigate operation over longer sequences and in varying soil conditions.

References

[1] Bullock, Darcy M. and Oppenheim, Irving J., “A Laboratory Study of Force-Cognitive Excavation,” In *Proc. Sixth International Symposium on Automation and Robotics in Construction*. June, 1989.

[2] Feng, P. and Yang, Y. and Qi, Z. and Sun, S., “Research on Control Method of Planning Level for Excavation Robot,” In *Proc. 9th International Symposium on Automation and Robotics in Construction*. Tokyo, 1992.

[3] Huang, X. D. and Bernold, L. E., “Control Model for Robotic Backhoe Excavation and Obstacle Handling,” In *Proc. Robotics for Challenging Environments*. Albuquerque, 1994.

[4] Luengo, O. Singh, S., Cannon, H. “Modeling and Identification of Soil-tool Interaction in Automated Excavation,” Submitted to *IROS 98*.

[5] Lever, P. and Wang, F. and Chen, D., “Intelligent Excavator Control for a Lunar Mining System,” In *Proc. ASCE Conference on Robotics for Challenging Environments*. Albuquerque, February, 1994.

[6] Rocke, D. J., “Control system for automatically controlling a work implement of an earthmoving machine to capture material,” US Patent 5528843, 1994.

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

[8] S.E. Salcudean, S. Tafazoli, P.D. Lawrence and I. Chau “Impedance control of a teleoperated mini excavator,” In *Proc. of the 8th IEEE International Conference on Advanced Robotics* (ICAR), Monterey, California, July 1997.

[9] Sameshima, M. and Tozawa, S., “Development of Auto Digging Controller for Construction Machine by Fuzzy Logic Control,” In *Proc. of Conf. Japanese Society of Mechanical Engineers*, 1992.

[10] Seward, D. and Bradley, D. and Mann, J. and Goodwin M., “Controlling an Intelligent Excavator for Autonomous Digging in Difficult Ground,” In *Proc. 9th International Symposium on Automation and Construction*, June 1992, Tokyo.

[11] Singh, S. *Synthesis of Tactical Plans for Robotic Excavation*. Ph.D. Thesis, Robotics Institute, Carnegie Mellon University. 1995

[12] Singh, S., Kelly, A., “Robot Planning in the Space of Feasible Actions: Two Examples,” In *Proc. of the International Conference on Robotics and Automation*, Minneapolis, April 1996.

[13] Singh, S., “State of the Art in Automation of Earthmoving, ASCE Journal of Aerospace Engineering,” Vol 10, Number 4, Oct 1997.

[14] Stentz, A., Bares, J., Singh, S., Rowe, P., “A Robotic Excavator for Autonomous Truck Loading,” Submitted to IROS 98.

[15] Takahashi, H., Kamata, H., Masuyama, T., Sarata, S., “Autonomous shoveling of rocks by using image vision system on LHD,” In *Proc., International Symposium on Mine Mechanization and Automation*, June 1995, Golden, Colorado.