State of the Art in Automation of Earthmoving, 2002

Sanjiv Singh
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
ssingh@ri.cmu.edu

1 Introduction

This paper provides an overview of recent progress in the automation related to earthmoving systems. While a previous survey [28] has comprehensively examined various aspects of earthmoving automation, here we discuss recent representative ideas and systems.

The discussion is divided based on the cycle of operation for any autonomous machine: sensing, planning, and execution. First, an automated machine must sense its own state and the world around it. Next it must use this information along with a description of a goal to be achieved to plan the next action to be taken. In some cases the mapping from sensing to action is direct, and, can take the form of a pre-determined control law. In other cases, deliberation, or the use of models is necessary. Finally, the action must be executed via the mechanism.

This paper starts with a brief summary of sensing technologies as they pertain to earthmoving. Next, it examines the various models used to determine control laws, or, as tools for deliberative planners. Of specific interest are models of machines and their interaction with the earth, an area that has seen some recent activity. Planning and discussed next and in conclusion a three recent systems are discussed.

2 Sensing

Environmental sensing is how earthmoving machines are aware of the world around them. This is particularly important in those cases where the machine must decide where to dig or to evaluate how much material has been transferred. Such sensing is also important to ensure that an earthmover doesn't collide into objects in the environment. Most often environmental sensing seeks to measure shape and this is done outdoors by the use of range imagery from one of *stereo vision*, *ladar* or *radar.*

Stereo vision provides dense range imagery and can be inexpensive to implement. However, range imagery requires good optical texture in images which generally means that even lighting is necessary. Corke et al describe a stereo vision system that operates in underground mines to determine the shape of piles of ore (Figure 1) [6]. Radar is immune to a large range of environmental conditions and works well for detecting large objects but its large beam size produces only coarse localization for the targets that it does detect [16]. A practical scanning radar with a beamwidth narrow enough to be used for imaging remains elusive.

Figure 1 **Depth from stereo (top) left and right images from stereo cameras. (bottom) disparity image computed. Range is computed by triangulation.**

Ladar has the ability to detect small obstacles because of small beam divergence but generally requires mechanical scanning. One solution is to use a 2-D scanning ladar that has a scanning pattern similar to that of a television [12]. Such systems are not only very expensive but require accurate measurement of vehicle motion during the scan to properly register range data. In addition, the raster scan method requires high accelerations and decelerations of a mirror. Instead a popular configuration that has emerged is one that has one continuous rotating axis like a light house. The whole spinning mechanism is rotated around an orthogonal axis to provide a large view of the environment. Figure 2 shows two commercial laser scanners that achieve cm level accuracy. In some cases, a single axis scanner is mounted in such a way that the motion of the earthmover provides the second axis. This approach has been used successfully in [6] and in [30]. Figure 3 shows

Figure 2 **Ranging laser scanners from Riegl Electrooptick. (left) single axis (right) two axis.**

a false colored range image from the dual axis laser range scanner. For earthmoving such range maps can be used to construct high fidelity representations of terrain such as shown in Figure 4. These representations are precise enough to be able to determine volume of soil excavated by a single bucket by differencing terrain maps as shown in the example in Figure 4 [4][8][31].

One reason that laser scanners have been slow to be accepted in earthmoving applications is because most implementations do not address problems of dust penetration or a partially occluded (i.e. dirty) exit window. Recently two different time-of-flight scanning ladar systems have become

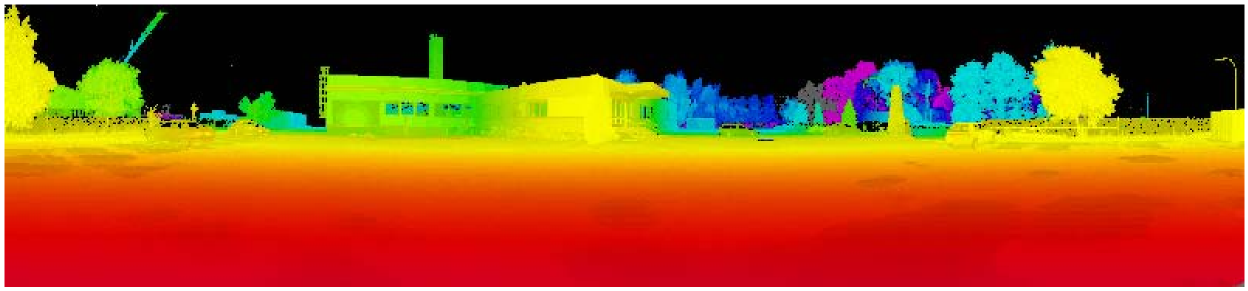

Figure 3 360 degree range panorama from two axis scanning laser range sensor.

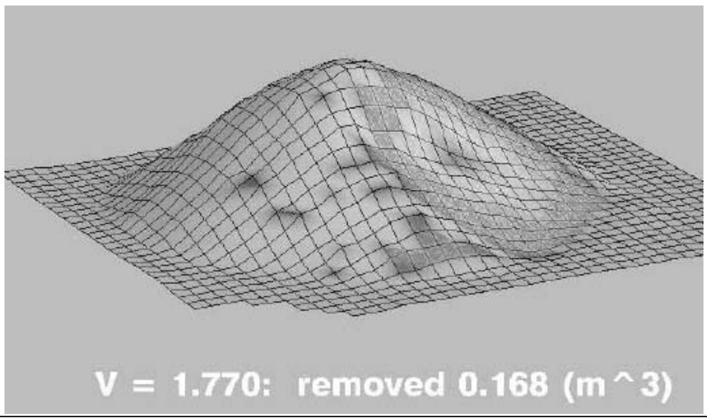

Figure 4 Terrain map of loose soil created from multiple laser scans [31]. Such an elevation map is a two dimensional grid with an elevation value at each cell. Subsequent scans are differenced to provide an estimate of soil removed, in this case by a wheel loader bucket.

available that are impervious to ambient dust conditions [30]. The first uses a "last-pulse" technique that observes the waveform of the returned light and rejects early returns that can arise from internal reflections off of a dirty exit window, or from a dust cloud obscuring the target.In general, the next-to-last pulse returns are due to dust in the scene and are indicative of what a normal "first pulse" rangefinder would see. For instance, in Figure 5 (left), a first pulse rangefinder would detect the dirty exit window and would be unable to "see" the target.

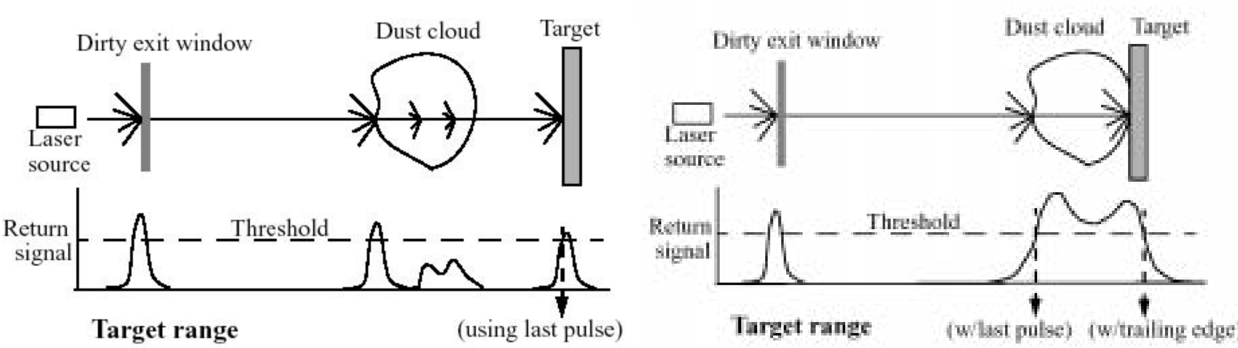

Figure 5 Last pulse (left) and trailing edge (trailing) detection of target when target is obscured in dust cloud.

Even if the window were clean, the first pulse unit would still “see” the dust cloud instead of the target. Since reflections off the exit window are rejected with the last pulse technique, the unit can be environmentally sealed using an inexpensive transparent cover that does not have to be optically perfect or clean. Another advantage is that the laser system can also report when multiple returns occur, giving a warning that dust is present. This is important because overall ranging reliability and accuracy is decreased in dusty conditions, so an autonomous machine might need to adopt a slower, more conservative motion strategy. There is, however, a limitation to last-pulse rangefinding. When the target is within the dust cloud, the receiver electronics can have difficulty separating the dust and target returns. A second type of dust penetrating scanner system identifies targets by locating the “trailing edge” of the last return signal as is shown in Figure 5 (right). Like the last pulse system, this device is also robust to occlusions on the exit window making it ideal for construction and mining environments. Though the trailing edge detection technique forgoes some range accuracy, it is a superior approach for environments where dust may frequently surround the target. Note that the last pulse device is unable to separate the dust cloud from the truck and reports the front of the cloud. The trailing edge device correctly reports range to the truck regardless of the presence of dust.

Two dual-axis laser scanners were used on the autonomous backhoe developed at Carnegie Mellon for automated truck loading. Mounted on either side of the boom (Figure 6), they are used to sense the dig face, truck, and obstacles the workspace [30]. Two scanners are needed for full cov-

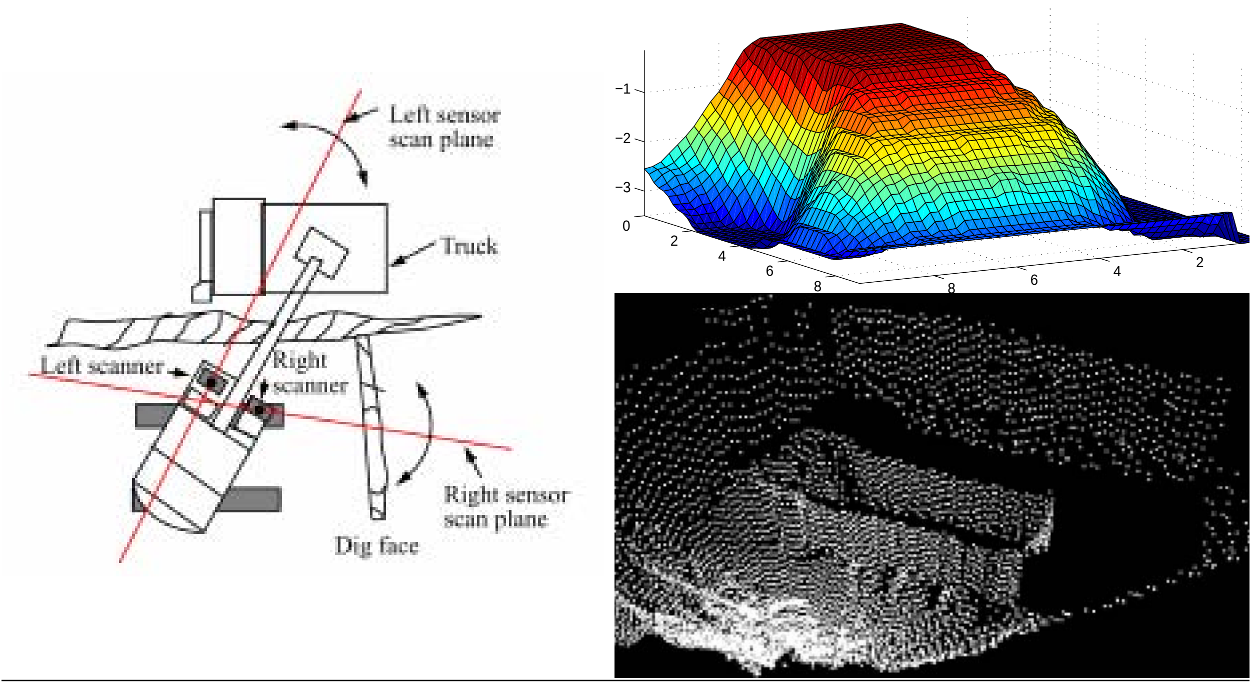

Figure 6 (left) Two single-axis laser scanners are mounted on either side of the boom on an automated backhoe. The motion of the backhoe as it swings between the digface and the truck allows the backhoe to build a terrain map of the bench (top right) as well as a model of the truck that is being loaded.

erage of the workspace and to enable concurrent sensing operations. Each sensor has a sample rate of 12 kHz, and a motorized mirror sweeps the beam circularly in a vertical plane. Additionally, each scanner can pan at a rate of up to 120 degrees per second, enabling this circle to be rotated about the azimuth. The scanner positioned over the operator’s cab is called the “left scanner", and it is responsible for sensing the workspace on the left side of the excavator. That is, this scanner is used to detect obstacles and locate the truck. The "right scanner", which is located a symmetric position on the right side of the boom, is responsible for sensing the workspace on the right side of the excavator which includes detecting obstacles and mapping the dig face.

3 Modeling

An important issue in the automation of earthmovers is the modeling of actuators and of the interaction between tools and the earth. Such modeling can be used to improve the design of tools and also actively to control automated earthmoving. This section focuses on the latter case.

Canon and Singh have demonstrated a system on an automated backhoe that select digs based on a method that scores digging actions by evaluating, the resultant trajectory for a given action or the starting conditions of a dig [3]. This trajectory not only varies with the control parameters of an action, but also with the shape of the terrain as well as soil properties. Terrain shape is measured directly, but soil properties must be estimated. They developed a feedforward model of the excavation process as shown in Figure 7 A model of the machine's actuators is used to predict the motion of the bucket in response to the actuator commands and reaction forces. The resultant motion is used by a soil-tool model to predict reaction forces on the bucket. Reaction forces and actuator positions are used by a control law to dictate actuator commands. This cycle is started with initial conditions (starting location and orientation of the bucket) and continued until the bucket is out of the ground. Predicted trajectories are scored by the use of a utility function composed of factors such as the time spent during digging, the energy expended, and the volume of soil captured.

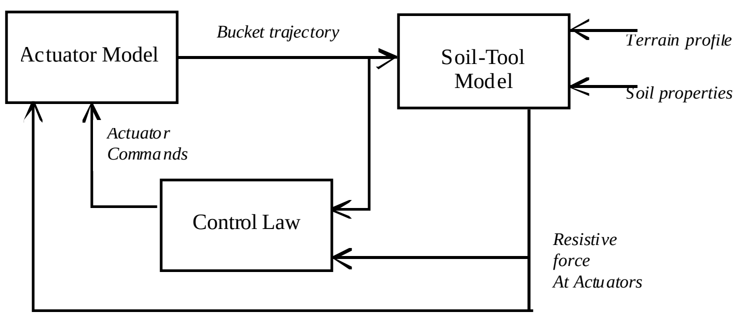

Figure 7 A forward model of the digging process proposed by Canon and Singh [3]. Given models of the actuators and the interaction between the earthmover and the terrain, the model predicts the trajectory of the implements in the soil. The advantage of this model is that it runs much faster than real-time, allowing its use in selecting between trajectories to execute.

3.1 Actuator Models

Dynamical models of hydraulic actuators specific to earthmovers are well studied (e.g. [1], [2], [15]) but often cannot be processed with requisite speed when they must run much faster than real-time. Several factors contribute to the difficulty of modeling this system. Typically, open-center valves are used to control the hydraulic actuators. Use of these valves requires that some of the flow is leaked back to a reservoir while the remainder is used to move the actuator. The leakage is partly dependent on the load on the actuator. The implements are also coupled in motion because a single hydraulic pump is used to power two joints and is limited in the amount of flow it can produce. There are also many transient effects, such as the time required for a pump to stroke up to meet demand, and the compressibility of the hydraulic fluid. Finally, unlike many typical robot systems, contact forces at the tool can be very high, and significantly contribute to system non-linearity. For instance, when the pressure in an actuator exceeds a threshold, a pressure relief valve opens stopping the actuator. Models of such complex mechanisms tend to be computationally expensive.

Cannon and Singh developed a neural net to model actuator response as shown in Figure 8 [3]. Neural nets are particularly suited to this application because system performance does not change significantly over time and a large set of training data (operation in various geometric configurations as well as soil conditions) is available prior to prediction. The advantage of neural nets is that they can capture the non-linearity not possible with globally linear methods and are much faster than locally linear methods. The disadvantage is that good performance requires trial and error experimentation with various configurations of neural nets.

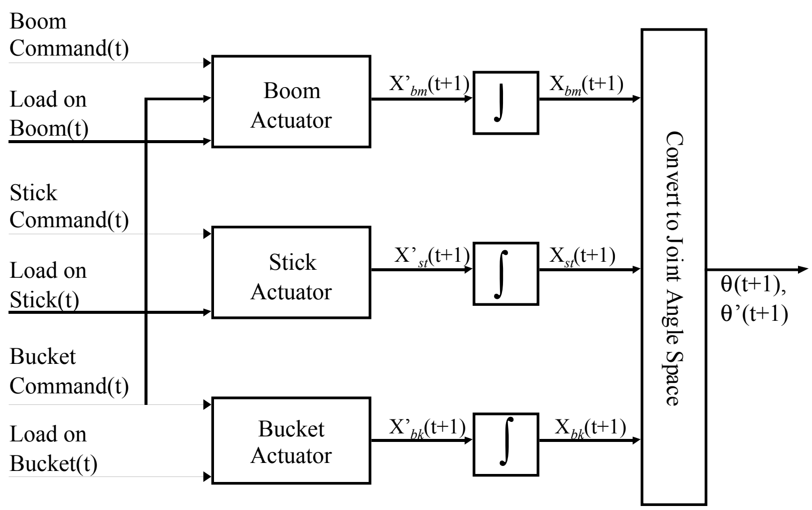

Figure 8 The actuator model for three links of a backhoe predicts velocities given commands and loads (From [3]). Note the coupling between the bucket and boom due to the fact that both are powered by a single source. Each actuator is modeled with a neural net.

Such a model was found to be effective in predicting velocities based on commands and loads as shown in Figure 9 Accuracy is typically within a few percent and the actuator model can be executed within a few milliseconds.

3.2 Soil-Tool Models

Another model is necessary to predict contact forces at the bucket due to resistance from the soil. Canon and Luengo implemented two methods for estimating reaction forces on the bucket [3][4][20]. The first method uses an analytical model based on the well known Fundamental Earthmoving Equation (FEE) in soil mechanics [21]. The second method uses a lumped parameter empirical model that uses terms from the FEE as a basis set. Reaction forces can vary dramatically depending on terrain geometry, soil compaction, and even climate. For instance even at one test site, soil varies from hard and frozen in the winter, wet and soupy in the rainy seasons, and dry and powdery during the summers. Therefore it is necessary that the soil reaction force models reflect the changing soil characteristics encountered. For each of the models a method for extract-

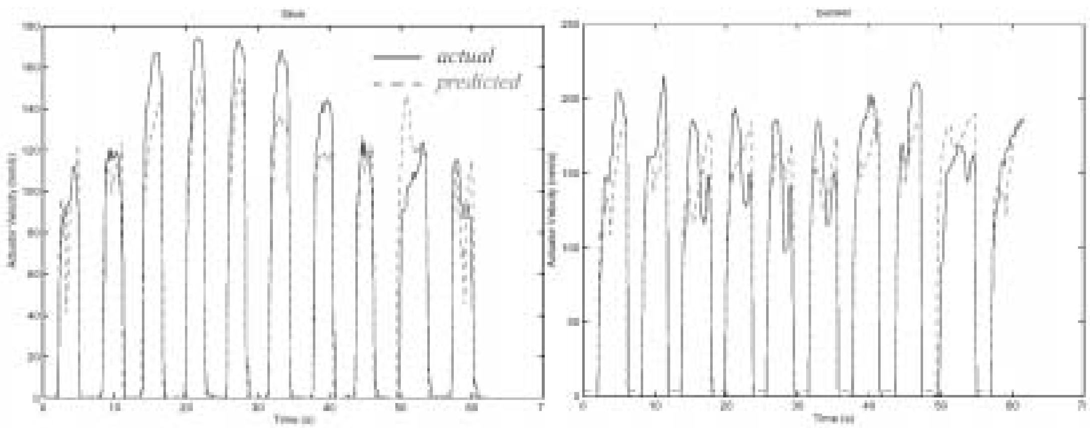

Figure 9 Comparison of predicted versus measured actuator velocities on an automated backhoe excavator for the stick (left) and the bucket (right) and stick over 10 digging cycles. Mean absolute velocity errors for the boom, stick and bucket were 4.3 mm/s, 11.8mm/s and 18.6 mm/s respectively. Peak velocities for these actuators were 87 mm/s, 243 mm/s and 250 mm/s. (From [3]).

ing soil properties based on the forces encountered during digging was implemented. Reaction forces are estimated based on the measurement of pressure in the hydraulic actuators. These pressures can be transformed into joint torques, which can subsequently be used to estimate the reaction forces. A comparison between the measured force and both the force predicted by the analytical and empirical method is shown in Figure 10.

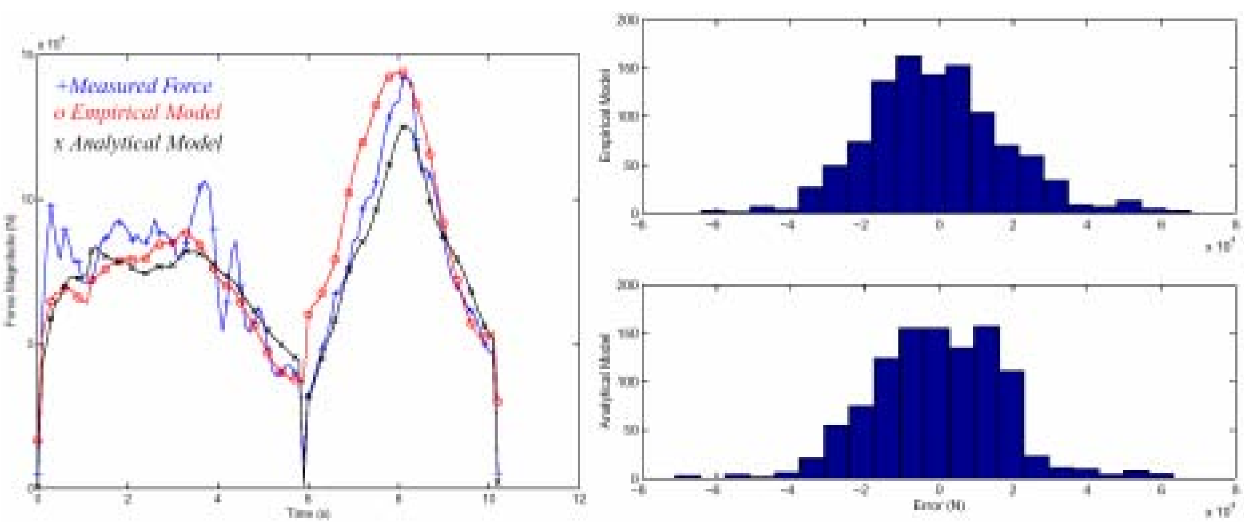

Figure 10 Cross-validation test for resistive force using the analytical and empirical models. The left plot shows a comparison of the force magnitudes for two digs.The right plots show histograms of errors. Using the analytical model, the mean absolute error in force magnitude is 13,980 N vs. 14,790 N when the empirical model is used. Standard deviations are 10,970 N and 11,880 N respectively. Peak forces are approximately 142,000N.

Somewhat surprisingly, the difference in error between the two models is small. However, while the time required to predict reaction forces using the empirical model is only slightly faster than the analytical model (13 ms vs. 16 ms), extraction of the soil properties with the empirical model is significantly faster (8 ms vs. 3400 ms)[1]. Soil properties estimated via the analytical model appear to make physical sense, and are relatively consistent during an extended sequence of operations. The empirical model parameters however vary a great deal since they are not tied to any physical relationship [4].

4 Planning

Automated earthmoving operations such as 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. The large state space and complex robot-world interaction imply that only locally optimal planners (i.e. per dig) can be created. In order to deal with the practical issues of excavating large volumes of earth in applications, a multi-resolution planning and execution scheme is typically necessary. Below, the discussion is separated into planning for digging and for dumping.

4.1 Dig Planning

At the highest level of the system developed by Canon and Singh is a coarse planning scheme that uses the geometry of the site and the goal configuration of the terrain to plan a sequence of "dig regions" [29]. In turn, each dig region is searched for the "best" dig that can be executed in that region. Finally, the selected dig is executed by a force-based closed loop control scheme. Treatment of the problem at three levels meets different objectives. The coarse planner ensures even performance over a large number of digs. The refined planner chooses digs that meet geometric constraints (reachability and collisions) and which locally 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. Fig. 10 shows the process of coarse to fine planning for the excavator.

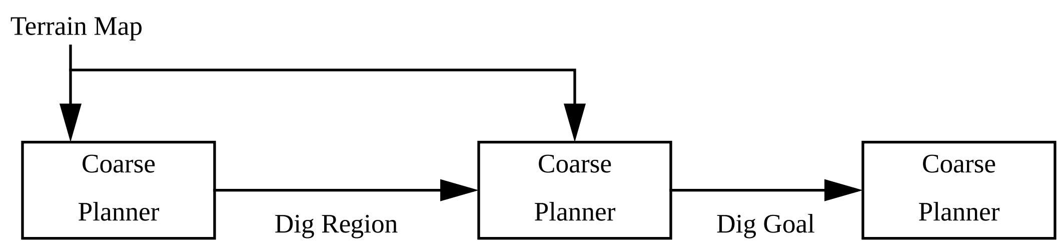

Figure 11 Coarse to fine planning strategy

The coarse planner takes as input processed sensor data which it places in a terrain map (a 2-D grid of height values). The output is a sequence of dig regions, each of which is in turn sent to a refined planner. Fig. 12 shows a strategy for removing material that was recommended by an expert excavator operator. Each box indicates a region, and the number within the box indicates the order in which the region is provided to the refined planner. In this strategy, material is removed from left to right, and from the top of the face to the bottom. There are several reasons for choosing this strategy. In most cases, the truck is parked on the operator's left hand side so that

1. Tests were conducted with an SGI R10000 processor.

the operator has an unobstructed view of it. By digging from left to right, the implements do not need to be raised as high to clear material when swinging to the truck. In digging from top to bottom, less force is required from the implements because it is not necessary to work against the weight of the material up above. In addition, clearing material away from the top minimizes the range shadows cast on the face of the terrain given a scanning range sensor that is mounted on the cab. The refined planner operates on an abstract representation of an atomic action (i.e. a single dig). Rather than searching for a bucket trajectory, the refined planner searches through compact task parameters within the bounds specified by the coarse planner. In order to select the best digging action, the refined planner evaluates candidates through the use of a forward model that simulates the result of choosing an action. An evaluation function based on energy and time scores the trajectory resulting from each action, and the action that meets all constraints and optimizes the cost function is chosen.

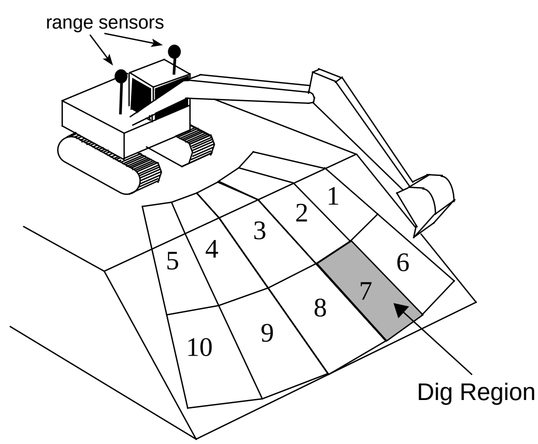

Figure 12 A coarse plan for an excavator performing bench loading. The strategy encoded is based on recommendations from expert human operators. The dig face is tessellated into regions as a function of the terrain geometry and heuristics that ensure stability of the excavator while digging. From[29].

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 Figure 13 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.

The 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). Once again it is assumed 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

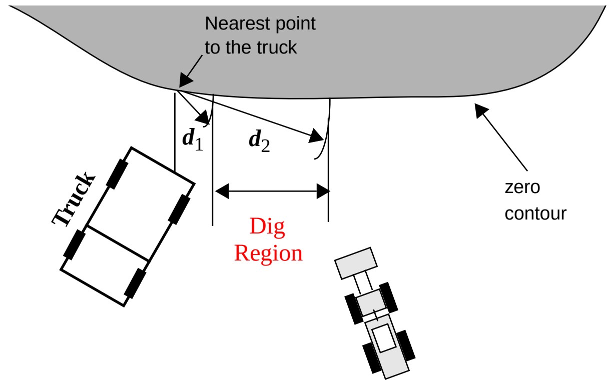

Figure 13 Coarse planning methodology for a wheel loader. The wheel loader is constrained to start within the dig region.

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.

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. 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. Second, 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. And third, dig locations should be chosen as close as possible to the truck for the sake of productivity.

4.2 Dump Planning

Apart from choosing where to dig, earthmovers that load trucks must determine where to dump the excavated material. The truck must be loaded evenly and completely. Because of uncertainty in soil settlement, the dumping strategy may need to be revised for each successive bucket load. The dump point planner applies a template-based approach to robustly find the low regions of soil distribution in the truck bed. Sensor data are gathered after each bucket of soil has been dumped in the truck. Like with the dig point planner, the sensor data are placed in a 2-D terrain map. Occlusion of the deposited soil by the truck bed walls can be a serious problem. Rather than assuming that nothing is in the unseen regions of the truck bed, the dump point planner fills in any unknown grid cells with the average elevation of the known grid cells. Finally, a specific terrain shape template is convolved over the entire truck bed terrain map to produce a score for each grid cell. This template looks for a certain profile of the material in the truck bed, such as a slope or a hole. Simple templates of constant elevations can be used to find the lowest elevation in the truck bed terrain map as well. The convolution operator produces a score which represents how well the template matched the particular region in the truck bed, and the location of the cell with the best score is returned as the desired dump location.

5 Control

This section examines free-space motion control as well as control of the implements while in contact with the terrain.

5.1 Free-Space Motion Control

Since excavators use highly non-linear actuators, efficient control of the implements in free space requires special consideration. While motion between any two points in space can be optimized in an offline fashion [15], it is important to be able to plan motion of an earthmover in free space online and preferably in an adaptive fashion that compensates for changing characteristics of the actuators over long periods of time.

Rowe developed a planner to plan motions of an excavator in free space, that is to avoid all known obstacles in the workspace and execute each loading cycle as quickly as possible [25]. Because of power constraints and joint coupling effects of the excavator's hydraulic system, as well as the difficulty in accurately modeling the dynamics of such a machine, more traditional optimal trajectory generation schemes do not work well. Instead, recognizing the fact that the excavator's motions are highly repetitive and very similar from loading cycle to loading cycle, and that it operates in a relatively small portion of its total workspace, a script based approach to motion planning was adopted. A script is a set of rules which define the general motions of the excavator's joints for a certain task, in this case loading trucks. These rules contain a number of variables, known as script parameters, which are instantiated on every loading cycle.The script rules were designed with the input of an expert human excavator operator and implicitly constrain what the excavator is and is not allowed to do. For example, if it was advised that moving two particular joints simultaneously was a bad idea, then the rules of the script make that motion impossible. Figure 14 shows the script rules for the truck loading task.

A new set of script parameters are computed right after digging is complete and before the free motion starts using the information about the truck's location and the desired dig and dump points. There are two types of script parameters, those which appear in the left hand side of the script rules and affect which commands are sent by the planner, and the joint commands themselves which appear on the right hand side of the rules. The command script parameters in the right hand side of the rules are primarily computed by geometric and kinematic means. The script parameters in the left hand side of the rules are found through a combination of simple excavator dynamic models and heuristics. These simple dynamic models capture first order effects of the excavator's closed loop behavior when given desired angular position commands. These models provide information about the velocities, accelerations, and command latencies for each joint, which are used to intelligently coordinate the different joint motions, resulting in fast loading times. As an example, consider the case when the excavator has finished digging, and the bucket is raising up out of the ground. The excavator does not need to wait until the bucket has raised to its full clearance height before swinging to the truck. Instead, it can begin swinging at some earlier point as the bucket is still raising, but it must have the knowledge provided by the dynamic models about how much time it will take to swing to the truck and to raise the bucket so it can safely couple the two motions to avoid a collision. Rowe also showed how critical parameters could be learned from actual experimentation and that such a parameterized script could be used to create smooth and fast motion of a complex, coupled non-linear system.

Joint 1: Swing	Command
1) When digging finishes, wait	$\theta 1 = 5^\circ$
2) If $\theta 2 > 14^\circ$, swing to truck	$\theta 1 = 101^\circ$
3) If $\theta 4 > 10^\circ$, swing to dig	$\theta 1 = 0^\circ$
4) If $\theta 1 = 0^\circ$, stop and execute dig	

Joint 2: Boom	Command
1) When digging finishes, raise	$\theta 2 = 18^\circ$
2) If $\theta 1 < 60^\circ$, lower to dig	$\theta 2 = 6^\circ$

Joint 3: Stick	Command
1) When digging finishes, wait	$\theta 3 = -100^\circ$
2) If $\theta 1 > 31^\circ$, move to spill point	$\theta 3 = -76^\circ$
3) If $\theta 4 > -30^\circ$, move to dump point	$\theta 3 = -92^\circ$
4) If $\theta 1 < 65^\circ$, move to dig	$\theta 3 = -75^\circ$

Joint 4: Bucket	Command
1) When digging finishes, curl	$\theta 4 = -90^\circ$
2) If $\theta 1 > 60^\circ$ and $q3 > -89^\circ$, open	$\theta 4 = 30^\circ$

Figure 14 Truck loading script for an automated excavator [25]. The left hand side of the rules are functions of the excavator's state, and the right hand side of the rules are the commands which the planner sends to the excavator's low level joint controllers. Thus, when the left hand side of a particular rule evaluates to true, its corresponding command gets sent to the excavator. The rules get re-evaluated at a 10 Hz during the execution of the excavator's motion.

5.2 Digging control

Digging control or the control of the implements while the excavator is in contact with the terrain is complicated essentially by the fact that the resistive force experienced can vary significantly. Further this application requires both the control of position ('do not dig beyond a particular boundary') as well as force ('regulate contact force to not exceed a preset threshold'). Several researchers have considered the use of impedance control to deal with this issue [10][11][26]. Such methods have been shown to work well when encountering a stiff boundary such as an impenetrable layer of rock. However, the issue of efficiently digging has not been addressed very much. Rocke has proposed a rule based system (AutoDig) to control the implements that is similar to the parameterized scripts used for free-space motion control [24].

AutoDig uses the position of the implements and cylinder pressures to directly control cylinder velocities. The behavior of AutoDig can be modified by changing a small set of parameters some of which have to do with implement position, while others are used to map cylinder pressures to actuator velocities. The latter can be thought of as an adjustment of end-effector stiffness or conversely a means of indicating soil hardness. Experiments show that the control law is effective in most cases where it is possible to capture a full bucket. When a stiff inclusion such as a boulder is encountered, AutoDig guides the bucket along the boundary of the object without stalling the excavator. AutoDig requires prior specification of certain joint angles at which the digging should end. In some cases, this means that digging is terminated early, while in others, the bucket sweeps a volume of soil greater than bucket capacity.

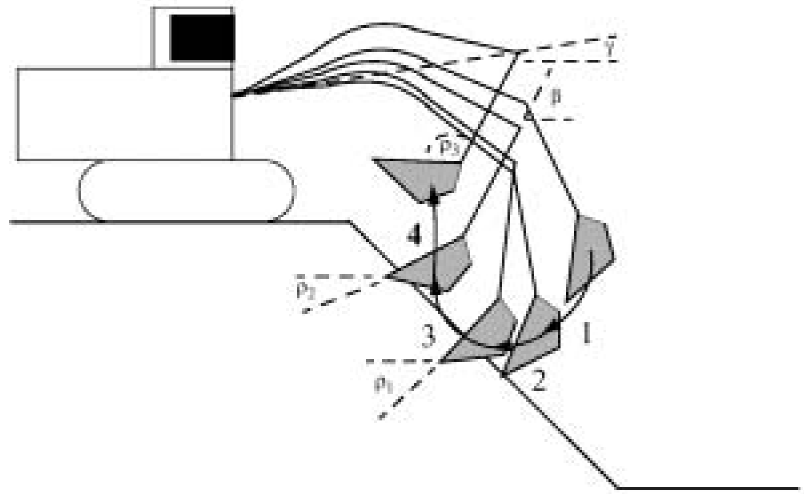

Figure 15 Autodig breaks the digging process down into four stages. 1)*Boom Down*: the boom is lowered until contact is made with the ground. 2)*Pre-Dig*: the bucket is quickly curled to bite into the material. 3) *Dig*: the material is force into the bucket. 4) *Capture*: the implements are positioned for carrying the material to the dump point. β is the stick angle to end Pre-Dig, ρ1 is the bucket angle to end Pre-Dig, ρ2 is the bucket angle to end the Dig, ρ3 is the bucket angle to end Capture, and γ is the boom angle to end Capture.

Canon extended AutoDig to deal with this issue by continuously monitoring the amount of soil in the bucket during the execution of a dig [4]. The amount of soil estimated to be swept into the bucket was determined by differencing the bucket trajectory from the terrain profile. In case the bucket filled up sooner than expected, capture is initiated. If the bucket has not filled up as much as expected, the implements are kept moving until the bucket is at capacity or a joint limit is encountered. In practice this simple control law was very useful is producing an effective digging strategy.

6 Systems

Automated earthmoving systems have been used for a variety of tasks recently. These include systems for laying pipes and working around buried utilities and ordnance[13][14][17][18] and transporting materials underground with a Load Haul Dump machine [22]. An impressive rule based method was used to effectively load from ore piles using a wheel loader [27]. This section examines three recent systems in the automation of earthmoving.

6.1 Dragline Automation

In surface mining, reduction of overburden removal costs has been identified as an important means of improving the economic performance of open cut coal mines. Draglines are key components in the overburden removal process, so improvement in dragline productivity realizes a significant reduction in costs. At a cost of $50 million to $100 million, a dragline is a major investment. It has been estimated that increasing the productivity of a dragline by around 4 per cent, would save the typical coal mine $1-2 million a year. A project at CSIRO in Australia has developed a swing automation system to improve the productivity of the swing-to-dump and dump-to-dig phases of a draglines operation. This project has automated a dragline and is presently conducting field testing [5][7][32]

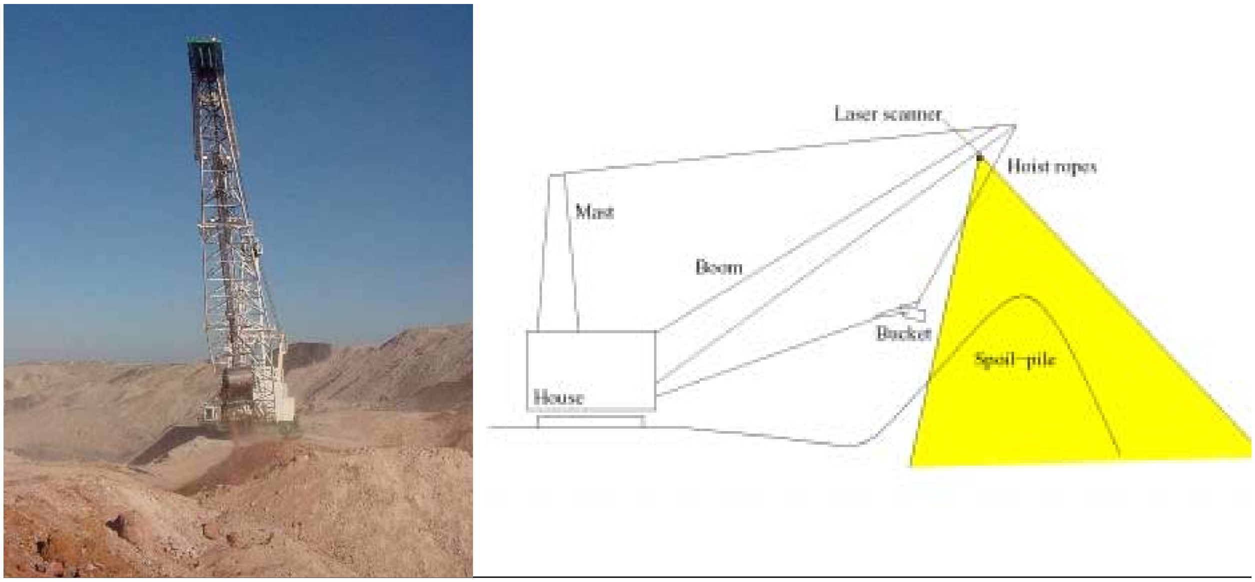

Figure 16 Automated Dragline (left) in strip mine (right) configuration showing the location of the laser scanner.

While a previous system requires an operator to enter so-called “no go” points in to the system indicating where the bucket must not travel, the current system is seeks a higher level of automation by allowing the dragline to “see” the local terrain, and thus automatically generate via points in real-time as the dragline swings from dig to dump. This would also allow the automation system to “see” the top of the spoil pile as it grows with each swing and hence allow the automatic adjustment of the dump height. This is expected to reduce hoist time since the current system contains a dump height safety margin to reduce the risk of collision bucket to spoil pile collision.

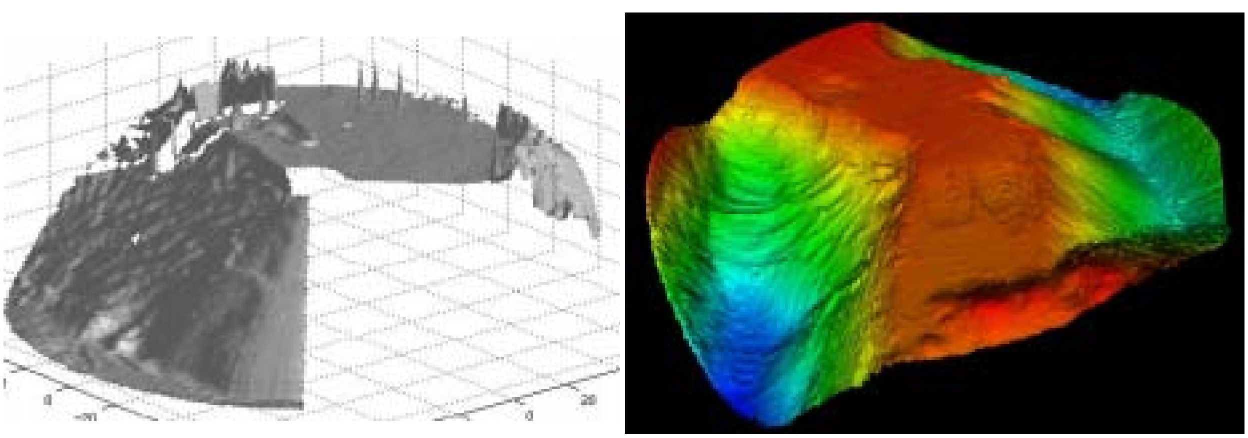

Figure 17 Digital terrain maps created from the laser scanner mounted on the dragline. From [23].

6.2 Volume measurement of truck haulage loads

Haulage trucks are large and costly machines used in open-pit mining to remove large volumes of overburden and ore. There are a number of design and scheduling issues that need to be resolved to optimize the performance of haul trucks. In particular, whilst haul-trucks are designed to carry weight, they are used to move a volume of dirt. Mine operators are primarily interested in removing a volume of material from the ground, irrespective of its weight. Historically, the volume carried by the truck has been inferred from a volumetric survey of the mine (divided by the number of loads) or a weight survey of each truck (bulk density of the material divided by weight). Neither technique actually measures the load in the tray and does not take into account changes in packing, or swelling. Laser scanning makes it possible to measure the in-situ volume of each load directly. CSIRO in Australia has developed an automated system that can automatically measure the volume of material in the tray of haul trucks. Currently, it is able to estimate volumes off-line (with human assistance) but is intended measure the volumes in real-time in the future [9].

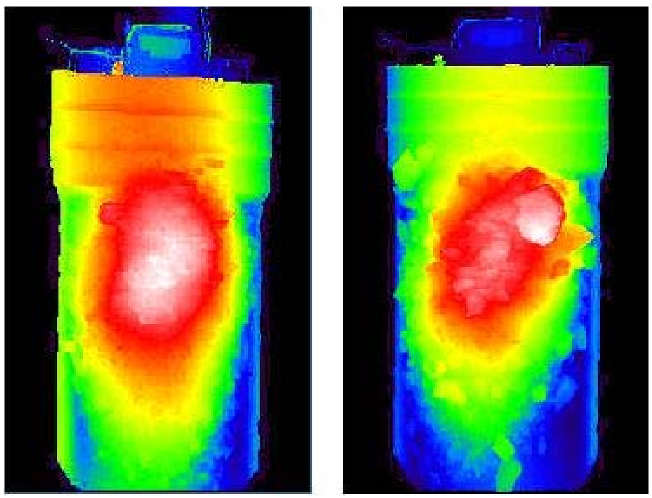

Figure 18 Depth images of an overhead view showing load profile and fragmentation.

The hardware for the current system consists of a field computer and two scanning lasers. The lasers are inexpensive and eye-safe. They can operate in rain and dusty environments. The lasers are placed several metres above the truck. As the truck is driven under the lasers, a 3D profile is built up and the volume can be estimated by subtracting the 3D surface of the empty tray from the 3D surface of the full tray. Although this is conceptually simple, the process is computationally complex. It requires knowing the precise location and orientation of each tray, and a 3D representation of the empty tray. In addition to estimating the volume of material in each tray, the 3D profile of the load can be used to measure load distribution, load profile and fragmentation as shown in Figure 18

6.3 Automated Truck Loading

The surface mining of metals, quarrying of rock, and construction of highways require the rapid removal and handling of massive quantities of soil, ore, and rock. Typically, explosive or mechanical techniques are used to pulverize the material, and digging machines such as excavators load the material into trucks for haulage to landfills, storage areas, or processing plants.

As shown in Figure 19, an excavator sits atop a bench and loads material into trucks that queue up to its side. The operator is responsible for designating where the truck should park, digging material from the face and depositing it in the truck bed, and stopping for people and obstacles in the loading zone. The opportunities for automation are immense. Typically, loading a truck requires several passes, each of which takes 15 to 20 seconds. Reducing the time of each loading pass by even a second translates into an enormous gain across the entire job. The operator’s performance

Figure 19 Excavator loading a truck in a mass excavation scenario. The excavator sits on top of an elevated bench, removes the material from the bench, and deposits it into the back of the truck.

peaks early in the work shift and degrades as the shift wears on. Scheduled idle times, such as lunch and other breaks, also diminish average production across a shift. All of these factors are areas where automation can improve productivity. Safety is another opportunity. Excavator operators are most likely to be injured when mounting and dismounting the machine. Operators tend to focus on the task at hand and may fail to notice other site personnel or equipment entering the loading zone. Automation can improve safety by removing the operator from the machine and by providing complete sensor coverage to watch for potential hazards entering the work area.

The Automated Loading System developed at Carnegie Mellon completely automates the truck loading task (Figure 20). The excavator uses two scanning laser rangefinders to recognize and localize the truck, measure the soil face, and detect obstacles. The excavator's software decides where to dig in the soil, where to dump in the truck, and how to quickly move between these points while detecting and stopping for obstacles. The system was fully implemented and was demonstrated to load trucks as fast as human operators on extended term operations that have lasted for several hours at a time.

References

[1] Bilodeau, G. and Papadopoulos, E., Modeling, Identification and Experimental Validation of a Hydraulic Manipulator Joint for Control, In Proc. International Conf. on Intelligent Robots and Systems, Grenoble, France, September 1997.

[2] Bonchis A, Corke P I, Rye D C and Ha Q P. 1999. Robust position tracking in hydraulic servo systems with asymmetric cylinders using sliding mode control, In Proceedings of the International Conference on Field and Service Robotics, CMU Pittsburgh, USA, August, pp 316-321.

[3] Cannon, H. and Singh, S., Models for Automated Earthmoving, In P. Corke and J. Trevelyan, editors, Experimental Robotics VI, Lecture Notes in Control and Information Sciences, pages 183-192. Springer Verlag, Sydney, June 2000.

Figure 20 A side view of the Autonomous Loading System (ALS). The ALS system is a commercially available 25 ton excavator that has been outfitted with a suite of sensors and on-board computing for the purpose of automation.

[4] Cannon, H., Extended Earthmoving with an Autonomous Excavator, Masters Thesis, Carnegie Mellon Robotics Institute, CMU-RI-TR-99-10, 1999.

[5] Corke, P., Hainsworth, D., Winstanley, G., Li, Y., and Gurgenci, H., Automated control of a dragline using machine vision. In Proc. Electrical Engineering Congress, pages 597-600, Sydney, November 1994. Institute of Engineers Australia.

[6] Corke, P., Roberts, J., Winstanley, G., 3D perception for mining robotics. In Proceedings Field and Service Robotics, pages 41-47. Springer-Verlag, 1999.

[7] Corke, P., Roberts, J., Winstanley, G, Experiments and experiences in developing a mining robot system. In P.Corke and J.Trevelyan, editors, Experimental Robotics VI, Lecture Notes in Control and Information Sciences, pages 183-192. Springer Verlag, Sydney, June 2000.

[8] Cheok, G., Lipman, R., Witzgall, C., Bernal, J., Stone, W., Field Demonstration of Laser Scanning for Excavation Measurement. In Proceedings of ISARC, 2000.

[9] Duff, E., Automated volume estimation of haul-truck loads. In Proceedings of the Australian Conference on Robotics and Automation, pages 179-184, Melbourne, Australia, 2000.

[10] Ha Q. P., Nguyen H. Q., Rye D. C. and Durrant-Whyte H, Robust impedance control of excavator dynamics, in Proceedings of the International Conference on Field and Service Robotics, CMU Pittsburgh, USA, August, pp 226-231. 1999

[11] Ha Q. P., Nguyen H. Q., Rye D. C. and Durrant-Whyte H, Force and position control of electrohydraulic position systems of a robotic excavator, in Proceedings of the IAARC/IFAC/IEEE Int. Symposium on Automation and Robotics in Construction (ISARC 99),

Madrid, Spain, September, pp 483-489.

[12] Hancock, J. et al., Active laser radar for high performance measurements, Proceedings of the IEEE International Conference on Robotics and Automation, 1998.

[13] Huang, X. and Bernold, L.E., CAD Integrated Excavation and Pipe-Laying, ASCE Journal of Construction Engineering and Management, Vol. 123, No. 3, September, pp. 318-323. 1997

[14] Huang, X., Bernd, D. and Bernold, L.E., Innovative Technology Development for Safe Excavation, ASCE, J. Construction Engineering and Management, Vol. 122, No. 1. March, 91-97. 1996

[15] Krishna, M., and Bares. J., Hydraulic System Modeling through Memory-based learning, In Proc. International Conf. on Intelligent Robots and Systems, Victoria, B.C, Canada. October 1998.

[16] Langer, D. and Jochem, T., Fusing Radar and Vision for Detecting, Classifying, and Avoiding Roadway Obstacles, Proceedings of the IEEE Symposium on Intelligent Vehicles, 1996.

[17] Lee, J., Lorenc, S.J., and Bernold, L.E., Teleoperated Pipe Manipulation, In Proceedings The Third International Conference and Exposition on Robotics for Challenging Environments, Albuquerque, NM, Apr.26-30, pp.188-194, 1998

[18] Lorenc, S.J. and Bernold, L.E., Equipment Mounted Buried Utility Detection System, Seventh Topical Meeting on Robotics and Remote Systems, Augusta, GA, Apr.27-May 1, pp.169-175, 1997.

[19] Lorenc, S.J., and Bernold, L.E., Smart Attchment for Utility Damage Prevention, The Third International Conference and Exposition on Robotics for Challenging Environments, Albuquerque, NM, Apr.26-30, pp.140-146, 1998.

[20] Luengo, O., Singh, S., Cannon, H, Modeling and Identification of Soil-tool Interaction in Automated Excavation, In Proceedings, IEEE/RSJ International Conference on Intelligent Robotic Systems, October 1998 Victoria, B.C., Canada.

[21] Reece, A. R., The Fundamental Equation of Earthmoving Mechanics, In Proc. of the Institution of Mechanical Engineers, 1964.

[22] Roberts, J., Duff, E., Corke, P., Sikka, P., Winstanley, G., and Cunningham, J., Autonomous control of underground mining vehicles using reactive navigation. In Proceedings of IEEE Int. Conf. on Robotics and Automation, pages 3790-3795, San Francisco, USA, 2000.

[23] Roberts, J., Winstanley, G., Corke, C., 3D imaging for a very large excavator, In Proceedings of 3rd International Conference on Field and Service Robotics, Helsinki, Finland, June 2001.

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

[25] Rowe, P. and Stentz, A., Parameterized Scripts for Motion Planning. In Proc. International Conference on Intelligent Robots and Systems (IROS). Grenoble, September 1997.

[26] Salcudean, T., Tafazoli, S., Lawrence, P. and Chau, I., Impedance control of a teleoperated mini excavator, In Proc. of the 8th IEEE International Conference on Advanced Robotics

(ICAR), Monterey, California, July 1997.

[27] Shi, X. Wang, F., and Lever, P., *Autonomous Rock Excavation: Intelligent Control Techniques and Experimentation*, World Scientific Pub Co; 1st edition, 1998.

[28] Singh, S., The State of the Art in Automation of Earthmoving, ASCE Journal of Aerospace Engineering, Vol 10, Number 4, October 1997.

[29] Singh, S. and Cannon, H., Multi-Resolution Planning for Earthmoving, In Proc. Int'l Conference on Robotics and Automation, May 1998, Leuven, Belgium.

[30] Stentz, A., Bares, J., Singh, S., Rowe, P., A Robotic Excavator for Autonomous Truck Loading, with A. Stentz, J. Bares, and P. Rowe, Autonomous Robots, 7(2):175-186, September 1999.

[31] Stone, W., Cheok, G., Lipman, R., Automated Earthmoving Status Determination, In Proceedings Robotics 2000, ASCE conference on Robotics for Challenging Environments, Albuquerque 2000.

[32] Winstanley G J, Corke P I and Roberts J M., Dragline swing automation, in Mineral Resources Engineering 8(3): 301-312. 1999.