Extended Earthmoving with an Autonomous Excavator

Howard N. Cannon

Submitted in partial fulfillment of the
requirements for the degree of Master
of Science in Robotics

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

May, 1999

This research was sponsored by Caterpillar, Inc., and conducted at the National Robotics Engineering Consortium in Pittsburgh, PA.

Acknowledgments

I would like to thank my advisor, Dr. Sanjiv Singh, for his guidance and collaboration on this research. His advice has been invaluable, and it has been both a pleasure and a privilege to work with him.

I would also like to thank Caterpillar Inc., the sponsor of this research, for having the vision to pursue this project, and for educating its engineers in the technologies important to the future of the earthmoving industry.

Thanks also to the other team members on this project: John Bares, Scott Boehmke, Shaun Burnett, Steve Colburn, Lonnie Devier, Jim Frazier, Tim Hegadorn, Herman Herman, Al Kelly, Murali Krishna, Keith Lay, Chris Leger, Oscar Luengo, Bob McCall, Ryan Miller, Richard Moore, Jeff Parker, Jorgen Pederson, Les Rosenberg, Patrick Rowe, Wenfan Shi, Hitesh Soneji, and Tony Stentz. Their developments on the system made this research possible. Particular thanks to Tim Hegadorn and Jim Frazier for logging hundreds of hours in the excavator while I ran my experiments, to Ryan Miller for his collaboration on neural networks, to Oscar Luengo for his collaboration on soil models, and to Keith Lay for all his helpful advice.

Finally, special thanks and appreciation to my wife Jennifer for her love and support throughout this process. Many nights were spent discussing all aspects of the project from technical to personal. Her sympathetic ear and enduring enthusiasm has given me the strength to complete this degree.

Table of Contents

Chapter 1 Introduction

The automation of earthmoving equipment is an endeavor that has the potential for improving the efficiency of the construction and mining industries, removing workers from hazardous situations such as in cleaning up toxic waste, and even enabling extraterrestrial exploration [Singh 97]. In a four year program, we have investigated the automation of truck loading with a hydraulic excavator in a mass excavation scenario. As shown in Figure 1, the excavator sits on top of an elevated bench, removes material from the bench, and deposits it into an awaiting truck.

There are several reasons for investigating the automation of truck loading with an excavator. First, it has the possibility of improving productivity. In a mining situation, even a small fraction of improvement in cycle time can add up over an eight hour shift allowing hundreds of tons of additional material to be excavated. It also requires several years of experience for an operator to be able to run the machine at its full potential. Even then, expert operators cannot maintain peak performance levels due to fatigue. Worker safety is another reason for automation. Every year people are injured while working in or around earthmoving machines. This can be alleviated by removing the worker from the machine, and by appropriate placement of sensors for monitoring the work area.

In this program known as the Autonomous Loading System (ALS), we have demonstrated the ability to completely automate the task of loading trucks with an excavator. The automated excavator is capable of observing the dig face and deciding where to dig. It can then execute the dig in an efficient manner, and capture the material into the bucket. It is capable of observing and recognizing the truck, localizing its position, and deciding where to dump the material into the truck

bed. It can move between the dig and the dump locations in a timely manner while ensuring that it does not hit any obstacles in its path. The system has been demonstrated to accomplish the entire truck loading task at speeds roughly equivalent to an expert human operator. In addition, it has been demonstrated that the system is capable of operating for several hours without any human assistance.

The focus of this thesis is on the development of the system's ability to dig effectively for extended periods of time. This can be broken up into three interrelated problems. The first concern is *how to dig.* We need a method that fills the bucket rapidly, and is robust to unanticipated digging forces. Then there is the problem of *where to dig* so that constraints are not violated and the material is removed from the bench in an optimal fashion. Finally there is the issue of *cleaning up the floor and repositioning the machine* so that excavation can continue after most of the material has been removed. In this thesis, an approach is presented which addresses all three of these issues. Experimental results are also presented which demonstrate the effectiveness of the digging operations with the automated excavator over an extended number of sequences.

Figure 1: 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.

1.1 System Overview

The Autonomous Loading System is a 25 ton commercial excavator that has been modified for the purposes of automation. Figure 2 shows a side view of the system. An excavator is comprised of three planar implements connected through revolute joints known as the boom, stick, and bucket, and one vertical revolute joint known as the swing joint. In addition the excavator has two independently movable tracks. The boom, stick, and bucket are controlled via prismatic hydraulic actuators (also known as hydraulic cylinders) interconnected across the implements, and the swing joint and tracks are controlled with hydraulic motors.

The excavator has been outfitted with electrohydraulic controls, a suite of sensors, and on-board computing. Each implement has a resolver attached to its rotational joint for sensing angular position and velocity. Pressure sensors are located in the hydraulic lines attached to each hydraulic actuator, enabling the measurement of the actuator forces. Two scanning laser range finders are attached at the top of the machine for sensing the surrounding terrain, the truck, and any potential obstacles. All of the decision making processes are conducted on-board the machine with an array of four MIPS processors.

Figure 2: 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.

The overall software architecture is shown in Figure 3. The center of the architecture is a motion planning module which is responsible for guiding the machine through all of its motions. This module receives inputs from several perception modules, selects a path of motion, and then executes the motion by sending commands to a machine control interface. The motion planning module is also responsible for dictating the motion of the range sensors so that they are positioned properly during the work cycle. More details about the motion planning module can be found in [Rowe 99].

The perception modules receive data from the range sensors, and then use this information to accomplish their various tasks. For instance, the truck recognizer module picks out the truck from the range sensor data, and localizes the truck's position. The dump planning module uses the range sensor information to observe the interior of the truck bed and decides where the next bucket of material should be placed. The dig planning module (the focus of this thesis) observes the shape of the terrain, and decides where to dig or where to position the machine so that a suitable dig may be achieved. Finally a work space monitor module uses the range information to look for potential obstacles entering the work area. More information about all of these modules can be found in [Stentz 98]. The outputs of the perception modules correspond to machine configuration goals for the motion planning module, and the motion planning module plans a path for the excavator based on this information.

Finally, the interface to the hardware is accomplished through three separate modules. A left and right sensor interface receives positioning commands from the motion planning module, and communicates this information to the low level range sensor control hardware. The sensor interface modules are also responsible for receiving the range data from the sensors and sending this information to the perception modules. The machine control interface receives commands from the motion planning module and communicates these commands to the low level machine control hardware. The low level hardware can execute closed loop position commands or open loop joint velocity commands. The machine control interface also receives the state of the machine from the low level hardware (such as joint positions, cylinder pressures, etc.) and communicates this back to any of the modules that need the information.

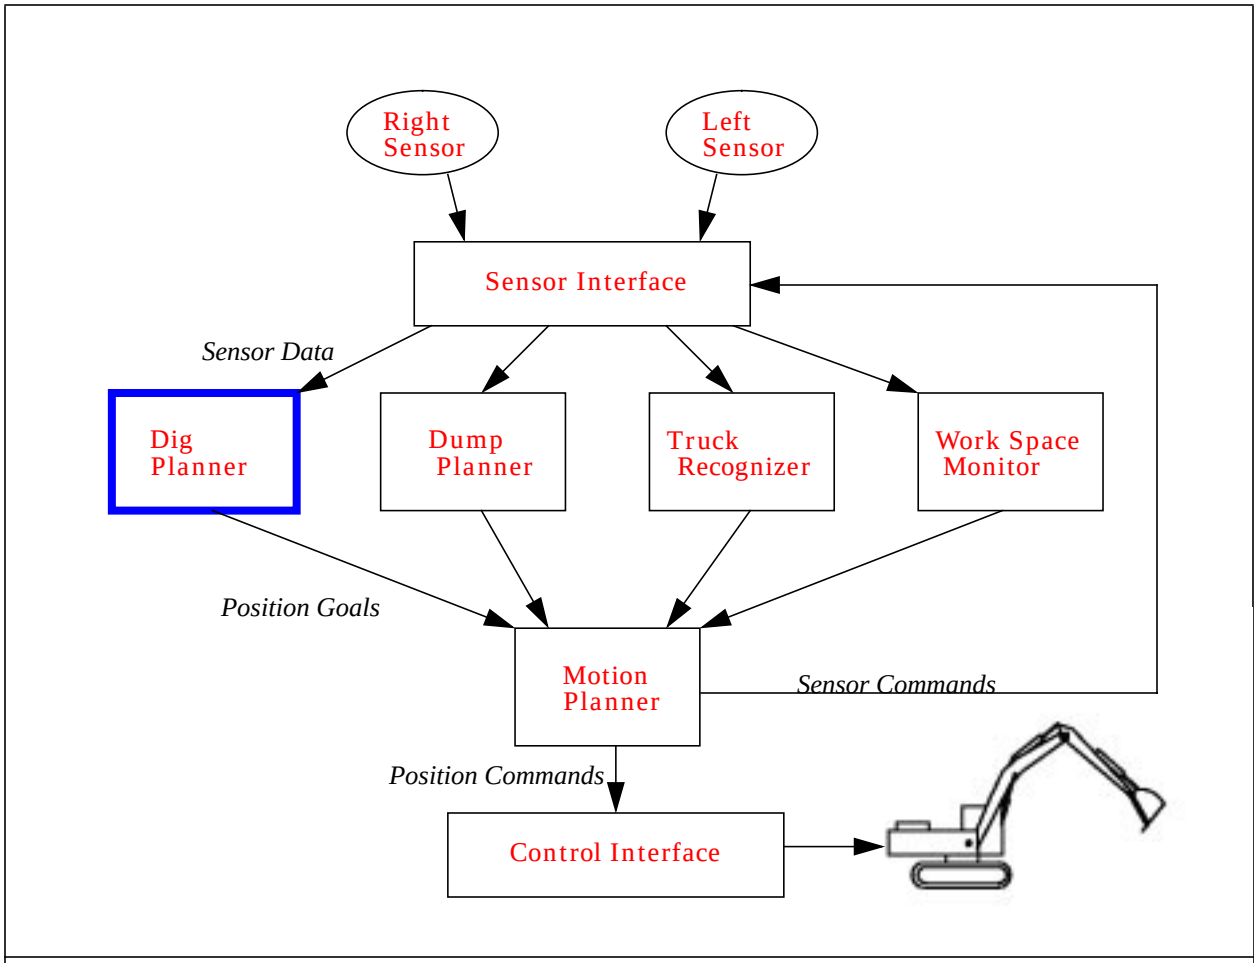

Figure 3: Overall software architecture for the ALS system. A motion planning module is responsible for dictating the motion of the machine. Perception modules are used to select goal points for the excavator based on range sensor information. The hardware interfaces are accomplished through the left and right sensor interfaces and the machine controller interface. The focus of this thesis is on the dig planning module.

A typical working scenario is shown in Figure 4, and this can be used to describe the work cycle. The excavator is situated on top of an elevated bench and the truck is parked at the base of the bench known as the ‘floor’, and situated off to one side. The machine removes a bucket of material from the dig face and begins swinging to the truck. As it is swinging towards the truck, the left sensor is positioned so that it can scan the truck, and the right scanner is positioned so that it can scan the dig face. The truck recognizer module uses the data from the left sensor to localize the position of the truck, and the dump planning module decides where to dump the material. The dig planning module uses the data from the right sensor to decide on the next dig location. After the dump maneuver is executed, the machine begins swinging back to the right to the selected dig location. Meanwhile the dump planning module is using the left sensor data to observe the deposited material in the truck to select the next dump location. This process of digging and dumping is continued until the truck is filled, at which point a new truck arrives, and the process starts over.

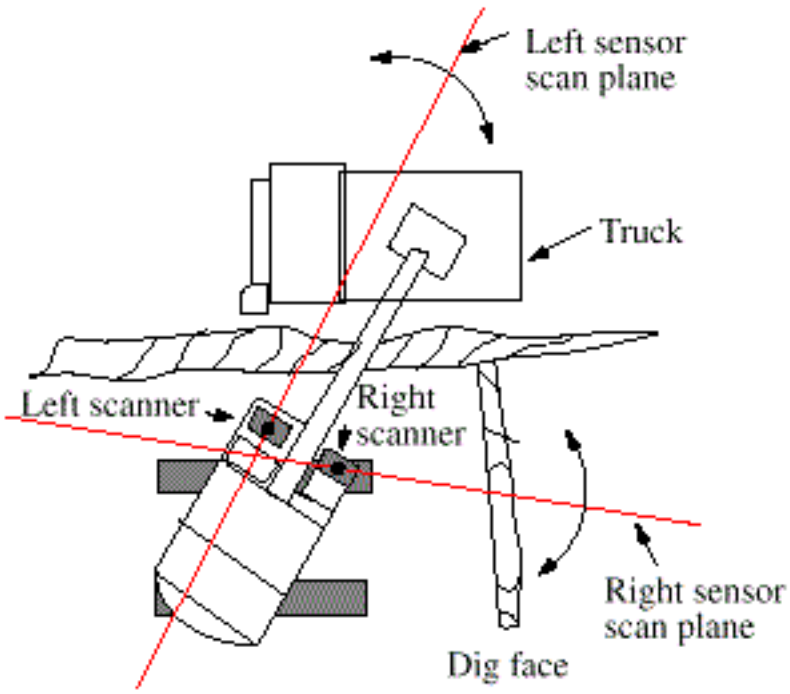

Figure 4: Top view of the ALS system in a typical work configuration [Stentz 98]. The excavator is situated on top of an elevated bench, and the truck is positioned at the base of the bench to the left of the excavator. The orientation of the sensors are shown by the plane of data that is being scanned. The left sensor information is used to decide on the dump location, and the right sensor data is used to select the dig location.

1.2 Summary of Approach

This thesis describes an approach that was implemented for managing the erosion of the bench over long sequences of operation. The approach addresses the problems of how to dig, where to dig, floor cleanup, and repositioning the machine so that excavation can continue. The key to this approach is the ability to sense the shape of the terrain using ranging devices such as laser, radar, or stereo vision. The range data is stored in a data structure we will refer to as a terrain map, which maps terrain elevations to a fixed rectangular grid as shown in Figure 5 [Singh 95].

The problem of “how to dig” is in regards to how the implements should be moved so that the bucket is filled quickly while being robust to extreme variations in digging forces. Excessive loading on any one implement can result in an inefficient and time consuming operation, and a method is needed which prevents this from occurring. To solve this problem, we have utilized a control algorithm known as “Autodig” [Rocke 94, Rocke 95]. Autodig generates commands for the implements based on the pressures that are observed in the hydraulic actuators. The commands are selected from an apriori mapping of pressures to joint angle commands which were generated by observing the way a human operator controls the digging process. One problem with the use of

Autodig in this application is that it requires a human observer to adjust some selectable parameters based on the shape of the terrain and the hardness of the soil. To solve this, we have augmented Autodig with perception to select these parameters automatically. In addition, we have added a feature to Autodig that allows it to be used for floor cleanup.

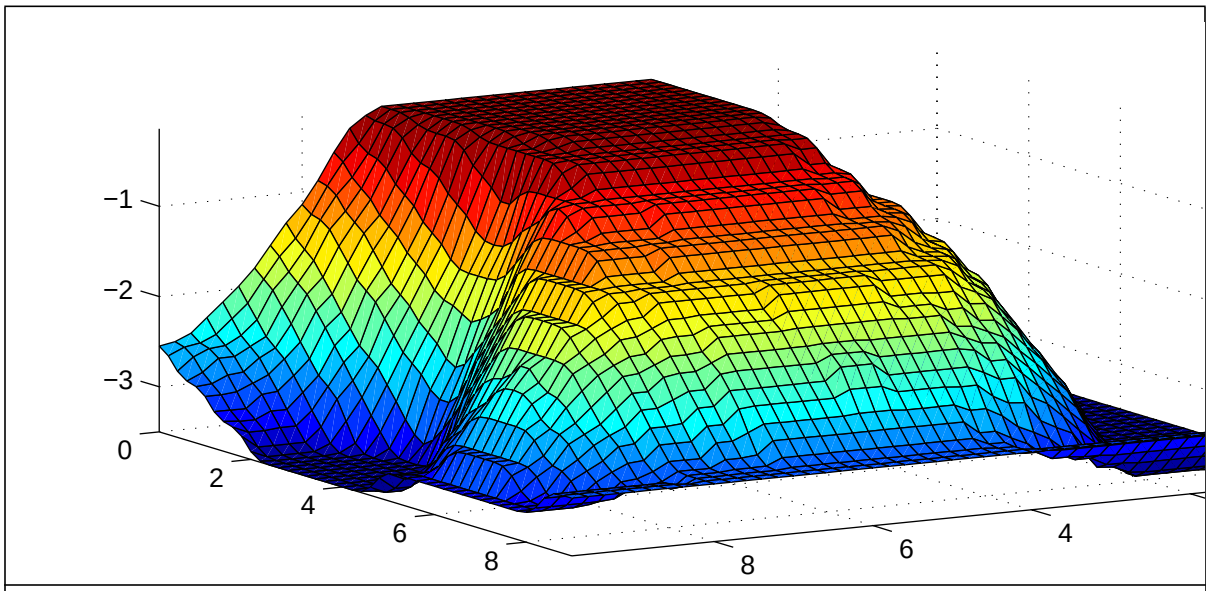

Figure 5: A sample terrain map that was generated by range sensors on the excavator.

The next question is to decide where to dig (i.e. the configuration of the machine for initiating Autodig) so that the bench is eroded as quickly and efficiently as possible while leaving a flat and level floor. Although Autodig will try to fill the bucket from any given dig location, its timeliness, efficiency, and ability to fill the bucket depends greatly on the initial bucket pose. The problem of "where to dig" is distinguished from typical planning problems because of the large state space needed to describe the potential configurations of the terrain, and because of the complexities of the interactions between the bucket and soil. To deal with this, we have developed a multi-resolution planning system. A coarse planning scheme generates a sequence of "dig regions" based on the current geometry and goal configuration of the terrain. A refined planning scheme then searches within a given dig region for the best dig. The search is accomplished by examining a number of candidate digs, and selecting the one that satisfies all constraints and optimizes an objective function.

In order for this planning process to work, we need to be able to predict the outcome of selecting a particular dig candidate. To do this, we have implemented a model of the excavation process that accounts for the behavior of the machine, the soil-tool interaction, and the behavior of Autodig. To span the space of possible candidate configurations, a large number of digs must be predicted in a matter of a few seconds. Therefore the model has been designed to be both computationally fast and reasonably accurate. In addition, since the digging operation is largely dependent on the characteristics of the soil, methods have been developed to adapt the model based on the forces that are encountered during the actual digging process.

A separate planner is used to decide how the floor should be cleaned. In general, the floor cleanup operation is selected based on the distance of the machine to the farthest material. Once floor cleanup has begun, the system begins monitoring the floor to determine when it is appropriate to track the machine backwards so that excavation can continue. The distance the machine can track backwards is based on the ability of the machine to be able to reach all of the material, and to be able to perceive all of the material with the range sensors. When the floor has been cleaned sufficiently so that the machine can move backwards some threshold distance subject to this criteria, then the machine is tracked backwards and the whole process starts over.

1.3 Road-Map to Thesis

Chapter 2 discusses related work conducted by other researchers. Chapter 3 discusses the function of the Autodig algorithm. It also addresses the enhancements for automatically adjusting selectable parameters with the use of perception, and enhancements for cleaning the floor. Chapter 4 introduces the model of the excavation process, which is utilized by the planning algorithm. Chapter 5 discusses the planning algorithms for selecting where to dig, cleaning the floor, and tracking the machine backwards. Chapter 6 discusses the results of our experiments, and how this system compares to an expert human operator. Chapter 7 summarizes the conclusions of this research.

Chapter 2 Related Work

A great deal of research has been performed on automating earthmoving operations because of the potential use in remediation of chemical and nuclear waste sites and extraterrestrial applications. A consolidated summary of the state of the art in this field can be found in [Singh 97]. This chapter focuses on previous work related to automating the excavation process in particular. The work is separated into three categories: automated dig execution, modeling the digging process, and planning digging operations.

2.1 Automated Dig Execution

[Singh 97] describes several systems that have been developed for automatically controlling the machine during the digging process. Since the digging process can involve large forces, simple trajectory control is usually inadequate, and some form of compensation for the forces must be utilized. The simplest methods are to trigger actions based on preset force thresholds [Bullock 89, Bullock 92, Huang 93]. Although these methods are simple, they are probably incapable of handling the large variety of situations that may be encountered, and certainly do not ensure that a full bucket is achieved.

Another control scheme is to use a set of rules to choose between a number of control actions while digging. One example is an automated excavator (LUCIE) that was developed at the University of Lancaster, England [Seward 88, Seward 92, Bradley 93]. Although the excavator tries to follow a predetermined path, a set of rules is used to react to the conditions encountered during excavation. For instance, once the bucket penetrates below a threshold elevation, then the bucket

is rotated. The force compensation is accomplished by monitoring the servo error in the system. It is assumed that higher servo errors are caused by higher forces. Once the servo error exceeds a threshold value, the bucket is raised to compensate.

Researchers at the University of Arizona have developed a means to deal with digging in a heterogeneous materials such as blasted rock [Lever 94, Lever 95, Shi 95, Shi 96]. In this system, the bucket is commanded to follow a specified path, and a fuzzy logic controller is used to guide the machine around immovable obstacles within the path. The inputs to the fuzzy logic controller are the force and torque information at the bucket, and the outputs are the horizontal and vertical step sizes of the bucket, and the bucket speed. Primitive excavation actions are grouped into a hierarchy of behaviors such as *dig-horizontally,* or *go-over-immobile-object,* and neural networks are used to select which behavior should be executed based on environmental information.

Sameshima and Tozawa have also implemented a fuzzy logic controller that guides the bucket through the digging process except that the control scheme specifies the actuation of each degree of freedom versus the motion of the bucket [Sameshima 92]. Three rules are evaluated at every control cycle as shown in Figure 6. The action taken is the weighted output of the three rules which correspond to velocity commands for each joint. In this system the forces are assumed to be reflected in the relative velocities of the stick and bucket. Thus the first rule looks at the relationship between these velocities and adjusts the commands accordingly. For instance, if the velocity of the stick and bucket are both low, then it is assumed that the force acting on the bucket must be large. To compensate the boom is given a larger command which will cause digging to occur at a shallower depth.

	Observations		**Actions**		
	Bucket Vel	Stick Vel	Bucket Vel	Stick Vel	Boom Vel
Rule 1	L	L	B	B	M
	L	H	B	M	S
	H	L	S	M	S
	H	H	S	B	Z
	Bucket Angle		Bucket Vel	Stick Vel	Boom Vel
Rule 2	L		B	S	S
	H		S	B	Z
	Depth of Bucket				Boom Vel
Rule 3	L				Z
	H				S

Figure 6: Fuzzy logic rules used by Sameshima and Tozawa [Singh 97]. *L* and *H* correspond to low and high values for the input observations, and *Z, S, M, B* correspond to zero, small, medium, and big values for the output velocity commands.

The Autodig algorithm [Rocke 94, Rocke 95] which is used in this research, also generates commands for each individual degree of freedom. However the commands are based on actual forces from cylinder pressure measurements versus forces inferred from relative velocities. These commands are generated from a lookup table that was created based on the way a human operator controls the individual joints while digging in various soil conditions. The soil condition must be specified to the routine so that the appropriate mapping is utilized. More information regarding this algorithm will be provided in Chapter 3, and we also discuss how we have augmented the Autodig algorithm with perception.

The Autodig algorithm is a desirable means for controlling the digging process in that it provides at least a piecewise continuous mapping of forces to actuator commands. Thus the motion of the implements can be expected to be somewhat smooth in operation. However, the system is designed to operate in relatively homogenous material. The method probably does not work as well as a rule based method when dealing with a large number of immovable inclusions because there is no means for backing up and trying again. Another disadvantage of this algorithm is that the joint velocity commands are based on the cylinder pressures alone. Therefore the trajectory of the bucket does not follow any selectable path, which makes it is undesirable for shaped excavation.

In contrast to these heuristic methods, a teleoperated mini-excavator developed by Salcudean et al. at the University of British Columbia uses a position-based impedance control to assist a human operator in guiding the bucket during the digging process [Salcudean 97]. The trajectory follows along a desired path specified by the operator until the path is not achievable due to the forces, at which point the control follows the path as closely as possible. In trying to apply this to an automated system, the question is what path should the impedance control try to follow? Bernold also proposed the use of impedance control [Bernold 93], and suggested that the optimal path for the bucket could be ascertained by characterizing the soil-tool interaction. In essence what is needed is a trajectory planning algorithm that generates an optimal path based on the characteristics of the soil.

The use of a trajectory planner alludes to a system that was developed by Singh at Carnegie Mellon University. [Singh 95] reports on a system in which pure position control is used during the digging process. However, he attempts to predict the forces that will be encountered, and rejects any trajectories that cannot be followed due to the limitations of the actuators. This concept will be discussed in more detail in the third section on background work related to planning digging operations. Singh suggested that this system could have been made more robust by the use of stiffness control (a subset of impedance control).

2.2 Modeling the Digging Process

Excavation is a difficult process to model because of the complexities of the machine dynamics and the forceful interactions between the bucket and the soil. Not only are we interested in being able to predict these effects, but we must also be to do so at speeds much faster than real time for use in the planning process. Although there has been some research in each of these individual areas (machine dynamics and soil-tool forces), there has been no work to our knowledge of combining these effects into a single model which is computationally tractable for real time applications.

There has been a significant amount of research related to modeling the dynamic characteristics of an excavator in free space for the purposes of trajectory control. These models are analytically derived from physical principles [Vaha 91, Vaha 93, Lawrence 95]. Although the models characterize the effects of inertial and gravitational forces on the excavator dynamics, they fail to capture the non-linear hydraulic characteristics of the mechanism. Certainly there exist many commercial dynamic modeling packages which can make these predictions, however they are far too slow for real time applications.

Perhaps the most closely related work is a semi-empirical approach taken by Krishna and Bares at Carnegie Mellon University [Krishna 99]. They describe the use of memory based learning to capture the dynamics of the overall machine. Through testing, a map is generated between the space of inputs (cylinder loads and hydraulic valve actuation) to the space of outputs (cylinder velocities). Once the map is created, the process is then to calculate the valve actuation based on control commands, the cylinder loads based on acceleration and gravity forces, and then use the map to predict the cylinder velocities. This method has been shown to predict the motion of the implements approximately 100 times faster than real-time and with reasonable accuracy.

There is also quite a large body of research related to predicting the resistive forces that act on a tool as it moves through the ground. One approach is to use finite element methods (FEM). Yong and Hanna have used this method to predict the forces and the deformation of the terrain due to a flat blade moving through clay soil [Yong 77]. Since this method requires the simultaneous solution of multiple partial differential equations, it is computationally expensive.

Another method is to analytically calculate the forces based on first principle mechanics [Reece 64, Siemens 65, Luth 65, Hettiaratchi 67, Gill 68]. This method was developed for the purpose of estimating tilling forces on agricultural equipment. The idea is that the soil shears away from itself along a failure surface in front of the bucket, and a static analysis is conducted on this "wedge" of material to determine the forces. One of the soil-tool models derived in this thesis is based on a similar static analysis. The equations were modified to account for digging in a sloped terrain. This will be described in more detail in Chapter 4.

This approach requires that some soil-tool properties be measured or estimated, such as the soil-soil friction angle, the soil-tool friction angle, the soil density, and the cohesiveness of the material. These values may be measured in a laboratory [Mckyes 85], or through the use of scaled models [Wadhwa 80]. Alternatively the parameters may be estimated using actual forces that are encountered during the digging process [Luengo 98].

Another approach to estimating the soil forces is to find an empirical relationship between a basis vector and the forces. Singh formed a basis vector that was based on the geometric variables found in the equation developed by Reece [Singh 95]. He attempted several different methods for learning the relationship between the basis vector and the forces, including global regression, memory based learning, and neural networks. A similar approach using global regression is used in this thesis, and the primary difference is the selection of the basis vector.

2.3 Planning Digging Operations

In a totally autonomous system, it necessary to be able to dig automatically, and to select which dig to execute automatically. One alternative is to specify a nominal trajectory for the digging operation. The automated digging routines can then be used to try to follow this trajectory until they must deviate due to the forces. Several researchers have based the nominal dig trajectory on the capacity of the bucket. That is, a trajectory is specified in which the bucket sweeps a volume of material that is equivalent to the bucket capacity [Koivo 92,Bisse 94, Hemami 92, Hemami 94, Sarata 93]. Hemami and Bisse's methods are able to satisfy not only the volume requirement, but also specify the trajectory so that it fills the bucket in some optimal fashion, such as minimizing the path length of the bucket tip.

Our work in planning optimal dig locations is based largely on research that was conducted by Singh at Carnegie Mellon University [Singh 95]. Figure 7 shows an excavation testbed used by Singh in which a bucket is attached to a robot manipulator. A scanning laser range finder is used to scan the shape of the material, and a force sensor on the bucket was used to refine the predictions of the forces that would be encountered during digging. The system automatically planned a trajectory for digging which was executed with closed loop position control.

The planning process was posed as a constrained optimization problem. The constraints included geometric considerations such as kinematic limitations of the machine, shape constraints based on the desired shape of the trench, and the maximum volume of material to sweep based on the capacity of the bucket. It also tried to predict the forces that would be encountered for a given trajectory, and eliminated any trajectories that exceeded the force constraints of the actuators. A search was then conducted over the set of actions that satisfied all of these constraints for the action that optimized a given performance criteria. In the case of trenching, Singh searched for the action that gave at least a 95% full bucket, and minimized the predicted torque to accomplish the motion.

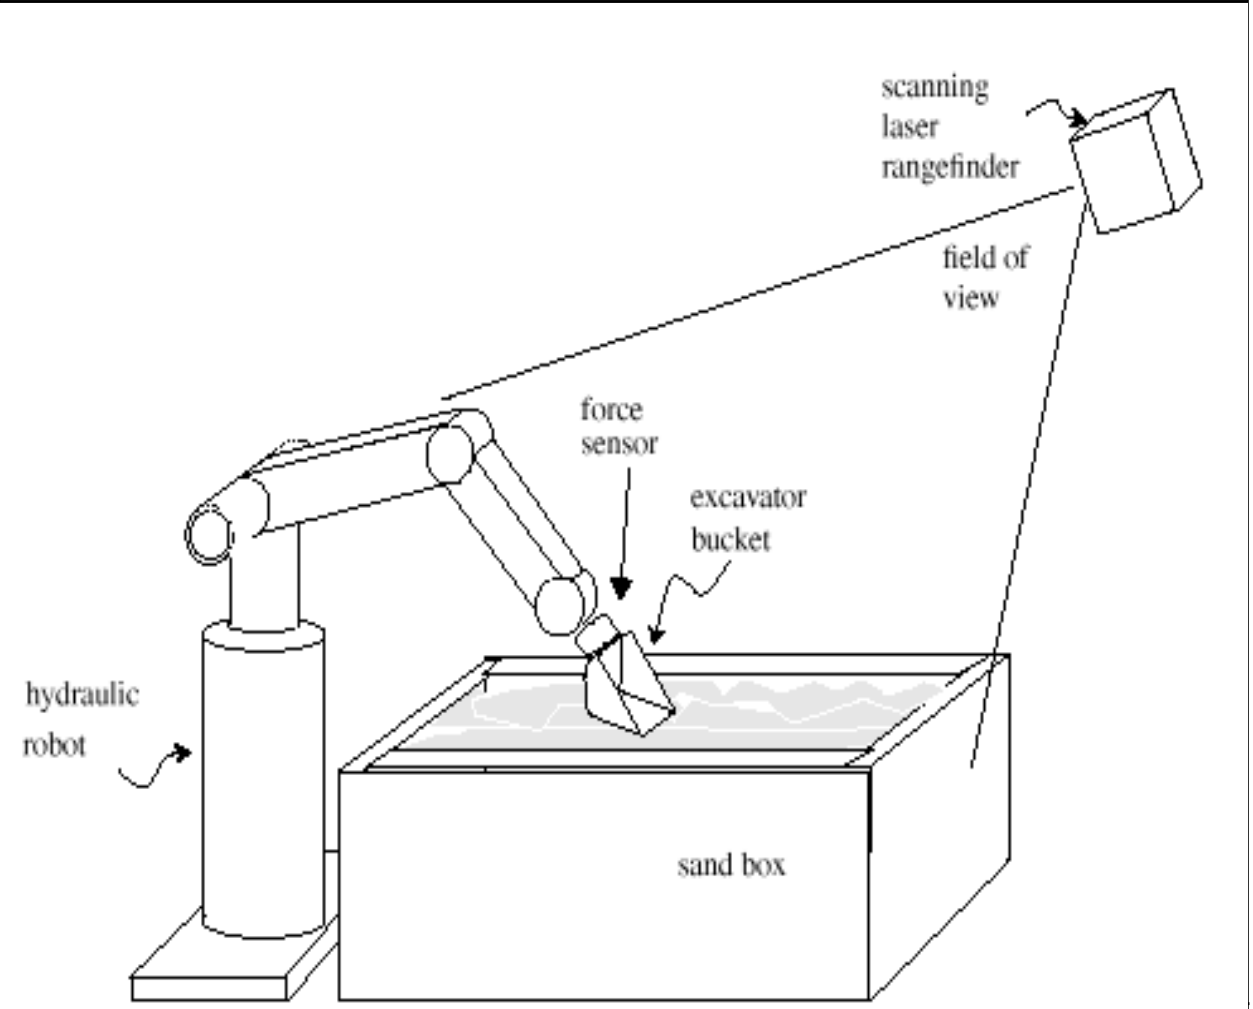

Figure 7: Excavator testbed used by Singh [Singh 95, Singh 97]. A scanning laser range finder is used to map the shape of the terrain prior to each dig. A force sensor is used to refine the force predictions. The system automatically decides on the best dig action.

Singh divided the excavation task into a reduced set of parameters that could describe each dig trajectory. For the trenching task, the set of parameters is shown in Figure 8. After these three variables are specified, the rest of the trajectory can be generated by following a set of rules. Planning therefore takes place in an action space that is spanned by these three variables. Note that even though the entire trajectory is not specified explicitly, the nominal trajectory is defined by the action parameters.

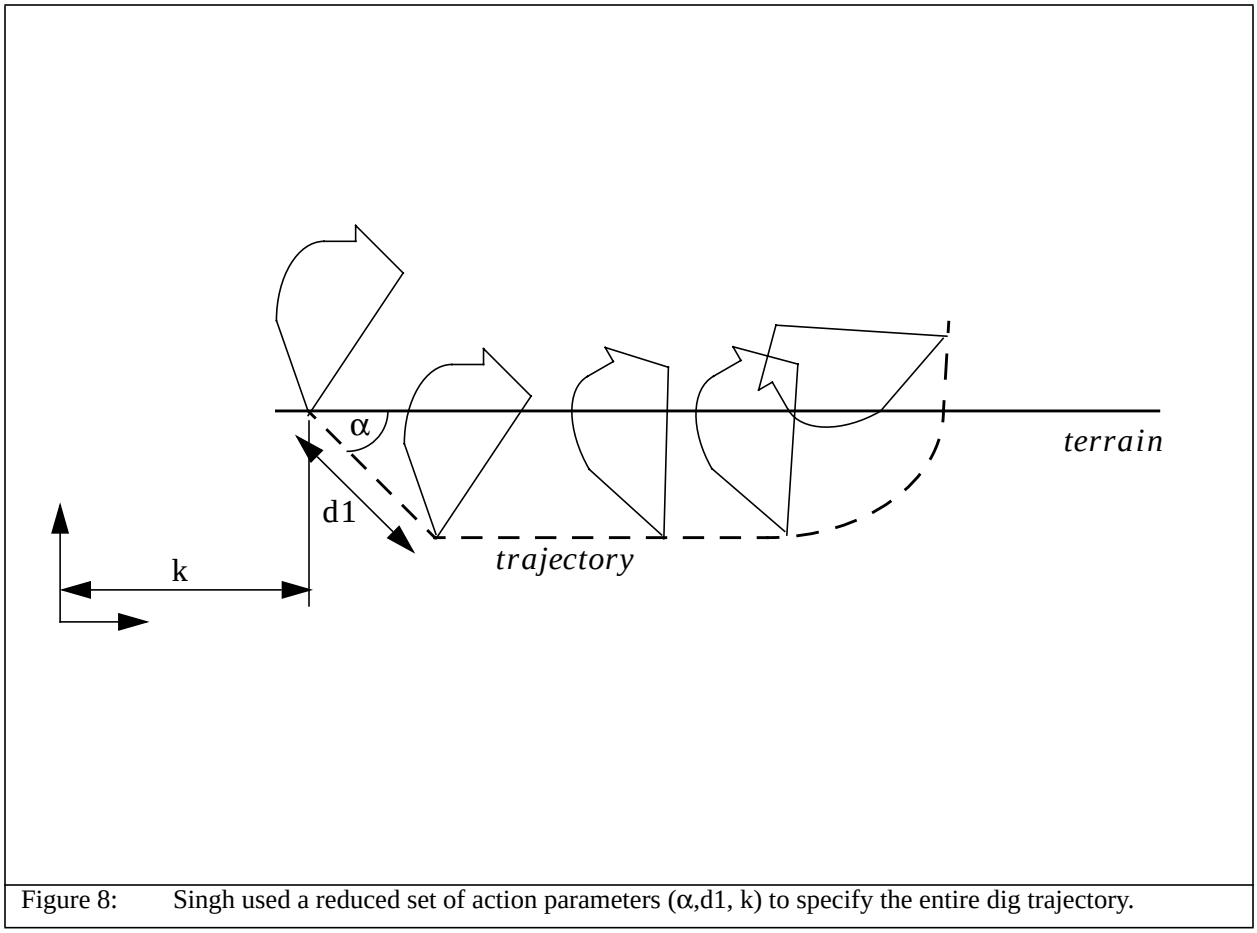

Figure 8: Singh used a reduced set of action parameters (α,d1, k) to specify the entire dig trajectory.

The work reported in this thesis regarding dig location planning is also based on optimizing a function within a set of constraints. Perhaps the main difference between Singh's research and the work reported here is in regards to the specification of the trajectory. By specifying the action parameters and following a set of rules, Singh generated a nominal trajectory for the bucket, which a position control tried to follow. In our system, the trajectory cannot be specified since Autodig is a force based control scheme, and instead the trajectory has to be predicted. Therefore the action space parameters dictate the initial pose of the bucket, and the rest of the trajectory is determined by the forces encountered.

The use of Autodig also impacts the utilization of the predicted forces. Autodig automatically compensates for the forces, therefore a separate force constraint does not apply in our system. The forces are still predicted for each dig, but this is utilized in the prediction of the trajectory itself. Since Singh's work specified the trajectory, a force constraint was needed to ensure that a given trajectory was achievable.

In a lesser sense, some other differences in the research are in the optimization function that is used for selecting the dig, and the implementation of the force models. Note that in contrast to Singh's work, we found that an analytical model of the soil forces produced good results. Perhaps

the reasons for this discrepancy may be due to a reformulation of the equations, better force measurements, and the relative stiffness of the implements compared to the forces.

Finally, this thesis extends the work in planning the excavation task by addressing issues related to digging for extended periods of time. Thus the system considers floor cleanup and tracking the machine backwards when a given area has been excavated. It also considers a three dimensional excavation task such as the removal of a bench, versus a two dimensional task such as trenching.

Chapter 3 Automated Dig Execution

The first question posed in the introduction was how to execute a dig automatically in such a way that the bucket is filled in a timely manner. If a machine is being used to its full potential, the interaction forces between the bucket and the ground should be relatively close to the machine's maximum capabilities so that the bucket is filled quickly. However these large forces can cause the hydraulic actuators to saturate, resulting in a time consuming and inefficient process. Therefore digging effectively requires the use of an algorithm that automatically compensates for the forces that are encountered. The algorithm should limit the forces so that the actuators do not saturate, and at the same time should maximize the rate at which the bucket fills for the sake of productivity.

For this purpose we have used an automatic digging algorithm known as Autodig [Rocke 84, Rocke 85]. Autodig compensates for the digging forces by monitoring the pressures inside the hydraulic actuators, and adjusting the command to the actuators accordingly. The operation is based on a series of apriori maps relating cylinder pressures to actuator commands. These maps were generated based on observing the digging techniques used by expert human operators. Section 3.1 will discuss the operation of the basic Autodig algorithm in more detail.

Autodig was designed to assist a human machine operator by taking over the digging portion of the work cycle. The human operator is still responsible for monitoring the digging performance, and making adjustments to account for the shape of the bench and the hardness of the soil being

excavated. For instance, if the human operator noticed that the bucket was full, Autodig could be terminated early to avoid unnecessary time and energy spent in the digging process. In a totally autonomous machine this correction must be made automatically. Section 3.2 discusses how we have augmented Autodig with perception based enhancements to end the digging process at the appropriate time. Chapter 5 will discuss the ability to automatically compensate for the hardness of the soil.

Finally, the human operator is also responsible for leaving a relatively flat and level floor after the bulk of material has been removed. Autodig itself cannot accomplish this because it is reacting to the pressures in the cylinders which is caused by the digging forces. Section 3.3 discusses a simple addition to Autodig for the purpose of floor cleanup.

3.1 Basic Digging Operation

Autodig uses a paradigm that a dig consists of four basic stages as shown in Figure 9. First, the boom is lowered until contact is made with the ground. Then in a *Pre-Dig* stage the bucket is quickly curled to a favorable angle in which to dig. The *Dig* stage is responsible for pushing the majority of the material into the bucket. Then the last digging phase curls the bucket in order to *Capture* the material, and raises the bucket out of the ground. In general, the implement joint angles are used to determine the digging stage. The digging stage dictates how the cylinder pressures are used to generate the commands.

During the *Boom-Down* stage, the boom is lowered until contact is made with the ground. All of the other implements are not moved. Contact with the ground can be determined by examining the pressures in the cylinders. When contact is made, generally there will be a drop in the head end of the boom cylinder and a rise in the rod end. Any shaking of the machine will cause large pressure oscillations in the boom cylinder, and this can cause a false detection of the ground. For this reason, the boom cylinder pressures have to be at the appropriate levels for a fixed period of time to eliminate any false detection.

During the *Pre-Dig* and *Dig* stages, the commands that Autodig issues are based on a set of prebuilt maps that relate cylinder pressures to actuator commands as shown in Figures 10 and 11. These maps were designed by observing the way an expert human operator digs in a variety of soil conditions. One of the inputs to these maps is a selectable soil hardness index. This value is a qualitative assessment as to the difficulty of digging in the material. In soft soils it is desirable to dig with a quick scooping motion while applying lower forces to the ground. In harder more compact soils, higher forces are required so that the bucket will fill. The soil hardness index basically dictates how much force Autodig applies to the ground. The soil hardness index is a value that is provided externally to Autodig, and an automatic means for obtaining this value is discussed in Chapter 5.

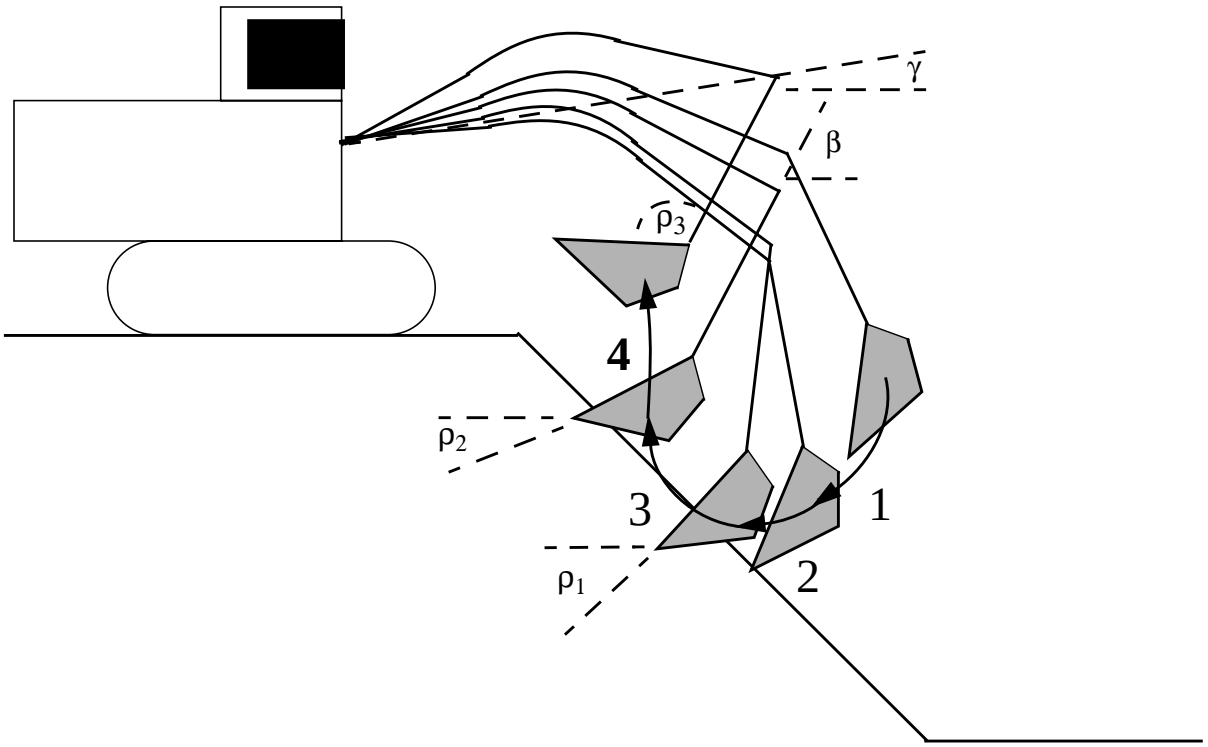

Figure 9: Autodig breaks the digging process down into four stages. 1)*Boom Down* - the boom is lowered until contact is made with the ground. 2)*Pre-Dig* - the bucket is quickly curled to bite into the material. 3) *Dig* - the material is force into the bucket. 4) *Capture* - the implements are positioned for carrying the material to the dump point. β is the stick angle to end *Pre-Dig*, ρ_1 is the bucket angle to end *Pre-Dig*, ρ_2 is the bucket angle to end the *Dig*, ρ_3 is the bucket angle to end *Capture*, and γ is the boom angle to end *Capture*.

The commands and pressures as shown in the curves are in terms of percentages. The percent pressure is based on the maximum pressure achievable in the cylinder as dictated by the hydraulic relief settings. The percent command is based on the maximum velocity of the cylinder. It is important to note here that these commands are “open-loop” in that no cylinder velocity feedback is used to maintain a given velocity. The actual velocity that is attained by the cylinder may be significantly different than the command depending on the forces acting on the cylinder and the available hydraulic pump flow.

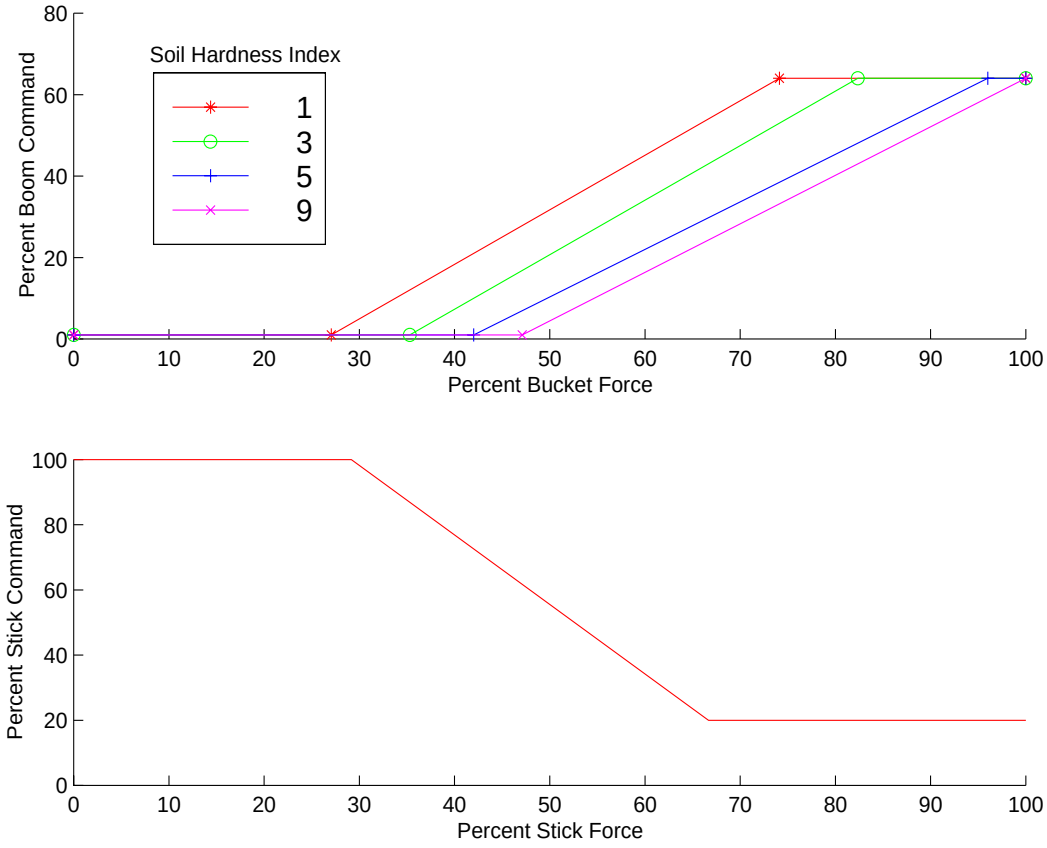

Figure 10: Autodig curves used during the *Pre-Dig* stage. The bucket is given 100% command while the boom and stick commands depend on cylinder pressures as shown. As the soil hardness index is increased, a higher force is required on the bucket cylinder before the boom is raised. The stick command curve is the same for all soil hardness indices.

Again, the purpose of the *Pre-Dig* stage is to bite into the material and quickly curl the bucket to a favorable position in which to begin digging. Therefore during this stage, the bucket is given 100% command. The boom command as shown in Figure 10 is a function of the bucket cylinder pressure, and is designed to regulate the amount of resistance that the bucket encounters. As the pressure in the bucket cylinder increases, the boom command increases to raise the bucket out of the material and hence reduce the bucket forces. Note that for a higher soil hardness index, a larger bucket force is required before the boom is raised. The stick command is designed to regulate its own pressure. When the load on the stick cylinder is relatively light, the stick is given full command so that the bucket tip is forced farther into the material. However as the load on the stick increases, the command is reduced so that the bucket is not forced in too deep. Thus, the *Pre-Dig* stage can be thought of as the boom performing force control on the bucket, and the stick performing force control on itself. The *Pre-Dig* stage ends when the bucket reaches a given angle ρ_1 relative to the horizontal as shown in Figure 9.

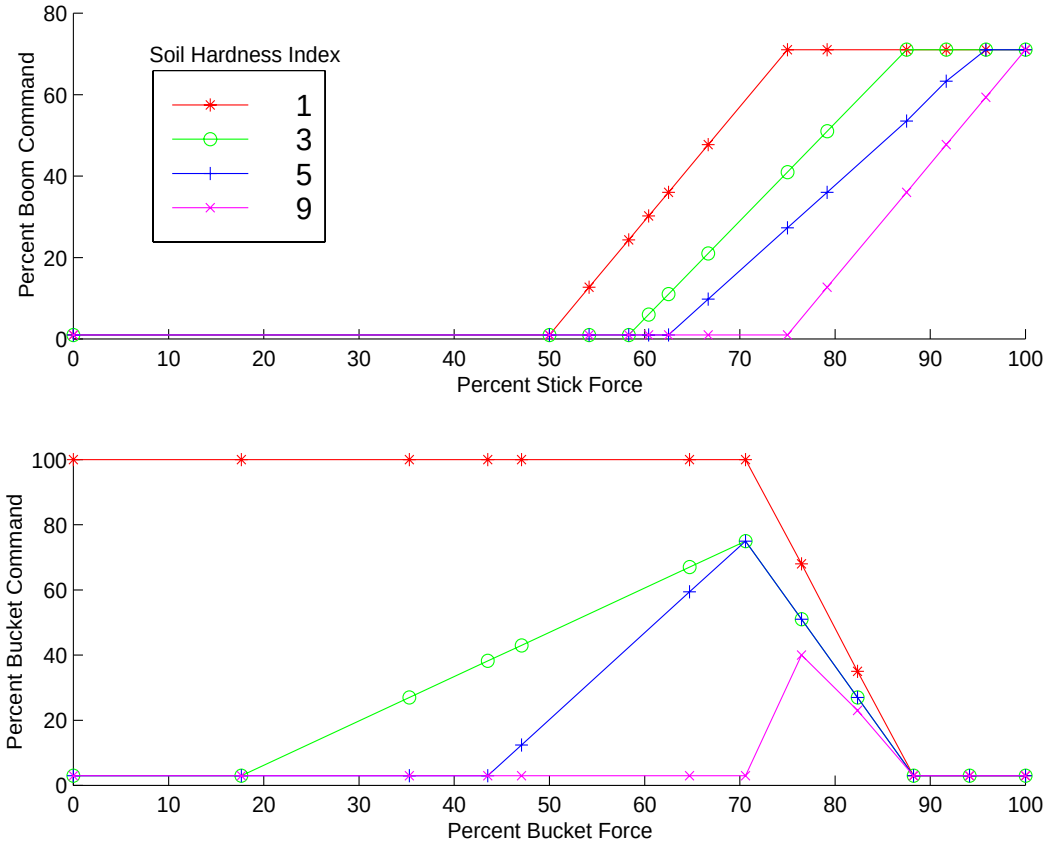

Figure 11: Autodig curves used during the *Dig* Stage. The stick is given 100% command while the boom and bucket commands depend on cylinder pressures as shown. As the soil hardness index is increased, a higher force is required on the stick cylinder before the boom is raised, and the bucket command is reduced.

During the *Dig* stage, the cylinders switch roles. In this stage, the idea is to quickly bring the stick in so that material is forced into the bucket. Therefore the stick command is 100%, and the boom command is used to regulate the load on the stick. As shown in Figure 11, as the stick forces increase, the boom command increases in order to raise the bucket and hence reduce the digging forces. The bucket command is now used to regulate the load on the bucket cylinder. This curve becomes somewhat more complicated depending on the soil hardness index. In soft soils, unless the bucket forces are extreme, the bucket is given full command thus providing a quick scooping motion. In harder soils, the bucket is not curled until an adequate load on the bucket has been achieved. This slows the bucket so that the stick has more time to push the bucket farther into the ground, thus increasing the bucket fill. In all cases, when the bucket loads get too high, the commands are reduced so that the hydraulic relief pressure is not reached. The *Dig* stage ends when the stick angle reaches a given angle β which is usually near vertical, or if the bucket angle reaches angle ρ_2, or if the stick or bucket cylinder limits have been reached.

During the *Capture* stage both the boom and the bucket are given 100% commands, while the stick command is reduced to zero. The idea behind this is to quickly get the material and the

bucket into a position so that it can be carried to the dump location. The boom angle γ at which *Capture* ends needs to guarantee that the bucket is above the ground so that the machine is free to swing to the dump point. The bucket angle ρ_3 at which *Capture* ends is specified so that the material will not fall out of the bucket during the swing.

3.2 Perception Enhancements for Ending the Dig

The beginning of this chapter discussed how a human operator could stop the digging process when he noticed that the bucket was full. Without human oversight, the system must rely on the angles β and ρ_2 as shown in Figure 9 to be properly set to end digging at the appropriate time. If these are not set properly, the efficiency can suffer dramatically. For instance, if the bucket is full, then any additional motion of the bucket through the ground wastes both time and energy. On the other hand, if the bucket is not full enough, then productivity and efficiency suffers for carrying a less than full bucket to the dump point.

In our testing we found that it was *impossible to fix* β *and* ρ_2 *properly for all conditions*. Proper adjustment for these angles is highly dependent on both the shape of the terrain and the hardness of the soil. As an example, suppose in one soil condition it is found that a particular value for β gives a full bucket most of the time. But when encountering a harder material, the bucket will not penetrate as deeply, and hence the buckets will not be as full. Likewise when digging in softer materials, the bucket will fill faster, and effort is wasted in trying to enforce this same value of β. Similarly, if the terrain is steep, less stick motion is required to fill the bucket than with relatively flat terrain profiles.

It became apparent from our testing that it was necessary to determine when the bucket is full so that *Capture* could be initiated. There are several ways in which this might be accomplished. One might suggest that the weight of the material be calculated using the pressures in the hydraulic cylinders. This however would not work because the digging forces can be an order of magnitude higher than the weight of the material, and hence it would be impossible to discern the magnitude of the weight. An alternative method would be to use perception to continuously monitor the material in the bucket to calculate volume, which is similar to what a human operator does. This method could prove difficult however because the material in the front of the bucket could visually occlude the material in the rear of the bucket causing large inaccuracies. Also the shaking and pitching of the machine during digging could cause large errors in the readings from the perception sensor.

The alternative pursued in this system was to store the shape of the terrain prior to digging, and then as the bucket passes through the soil, continually integrate the volume “swept” over the front edge of the bucket. This method also has several sources of inaccuracy. For one, we assume that the soil face does not change from the time that the perceptual image is taken to the time that the bucket edge passes beneath it. For cohesive soils this is a fairly good assumption, but for granular material, the assumption could break down. We also assume that the material that passes over the front edge stays within the bucket. In actuality, some of the material falls off to the side of the bucket as digging progresses. In our experiments we have found that it is sufficient to account for

this spillage by overestimating the capacity of the bucket. Furthermore, visual occlusions caused by undulations in the dig face can cause inaccuracies when integrating the volume. We have found that by careful positioning of the machine, and by interpolation, these inaccuracies can be minimized.

On the other hand, there are several advantages for using this method. First, the perception sensor only needs to take the image of the terrain once, and then the sensor is free to do other things such as monitoring the workspace. This also reduces the computation that is necessary by eliminating the need to continuously update the terrain. Another advantage is that the digging and perception stages are decoupled. The perception stage can be accomplished during relatively smooth machine motions, such as when swinging to the dump point, hence reducing error. This decoupling may also prove to be useful in other applications where the dig face is not visible during digging. For instance, in the wheel loader depicted in Figure 12, the bucket faces away from the machine operator so that the material inside the bucket is not visible during digging. However, since the perception stage is decoupled from digging, the image of the terrain could be obtained when driving to the truck.

Figure 12: A wheel loader digging next to a truck. The operator is incapable of seeing the material in the bucket during the digging process.

One additional enhancement was added to the end of Autodig to speed up the entire loading cycle. Rather than use a fixed angle γ for determining the end of *Capture*, the depth of the bucket below the terrain can be calculated directly from the terrain profile. When the bucket is calculated to be above the terrain, the *Capture* is completed. Calculation of the end of *Capture* on a dig by dig basis allows the swinging motion to the dump point to begin sooner, and hence reduces the overall cycle time.

These perception enhancements have proven to be highly valuable in improving the consistency of digging and in reducing overall cycle time. Table 1 shows a summary of the results for 30 back to back digs where half of the digs were done with the perception enhancements, and half without. The data in the table is arranged in the order of an increasing stick angle, which corresponds to reaching out farther away from the machine. There are two main things to note from the table. First, when using Autodig with the perception enhancements, the weight of the material in the bucket and the time required to dig remained very consistent regardless of where the bucket was initially placed. However without the enhancements the stick angle at which digging ends is fixed. So when the bucket is farther out, more material is obtained at the cost of longer cycle times. Likewise when closer in, a much smaller payload results. The other thing to notice is that even when the payloads were similar, the perception enhancements reduced the cycle time. This is mainly due to being able to stop the dig as soon as the bucket came out of the ground. This discrepancy might be reduced if the angle γ were reduced, except that there is a risk in not raising high enough to clear the bench in some cases.

Table 1: Autodig Performance with Perception Enhancements. Comparison shows the average weight of material captured in the bucket for each dig, and the average time required to execute the dig for five digs at each stick angle. The data is arranged in the order of an increasing stick angle, which corresponds to reaching farther away from the machine.

Initial Stick Angle	Autodig		Autodig with Perception	
	Time(s)	Weight (lbs)	Time(s)	Weight (lbs)
-87 Degrees	6.58	4145	7.13	5331
-70 Degrees	10.35	5172	7.47	5329
-50 Degrees	13.83	6526	8.77	5338

3.3 Modifications for Leaving a Level Floor

Autodig was designed to generate a relatively efficient digging trajectory by reacting to the pressures in the cylinders. Therefore the shape of the dig trajectory is an outcome of the forces that are acting on the bucket. The trajectory itself cannot be specified explicitly. However in many digging applications it is necessary to excavate to a particular shape. Such is the case in mass excavation. In mass excavation a floor elevation is maintained so that the unearthed ground may be traversed by other vehicles. Although a secondary operation may be used to level the floor such as with a wheel loader or bulldozer, the floor should be kept relatively even to reduce this effort.

In our system, two methods are combined to accomplish this goal. At the highest level, the perception based Dig Planner is used to select dig locations with sufficient material coverage so that Autodig will not penetrate the floor. To obtain the material that is just adjacent to the floor, a position based trajectory following routine was added to the beginning of Autodig. We refer to the combination of this position based routine followed by Autodig as a cleanup operation.

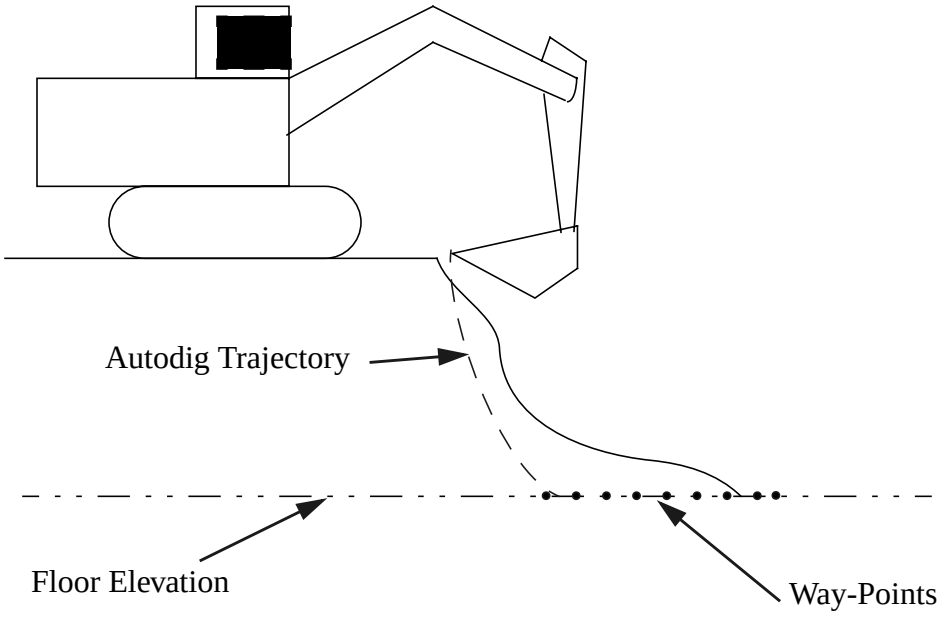

Figure 13: Cleanup operations consist of closed loop position following of way-points, followed by execution of Autodig.

The beginning of the cleanup operation uses the on-board closed loop position control that was provided with our excavator testbed. The routine simply specifies way-points along the floor level that the closed loop control tries to follow. Since the way-points are provided to the closed loop control at a fixed rate, the distance between the way-points dictates the velocity of the bucket, up to the maximum velocity that the control can follow. Unfortunately, as the points are separated farther apart, the machine can take any path from one way-point to the next, and hence the error in the trajectory can increase. We found that it was necessary to start with the way-points positioned closely together and then spread out over time. This effectively causes a ramp up in the velocity of the bucket, and is easier for the closed loop position control to follow than a “step” in the velocity.

The net effect is that the cleanup operation is relatively slow, and can take as much as two to three times the duration of a normal dig.

The way-points along the floor are followed until one of two conditions are met. First, a limit is placed on the stick angle, so that the bucket will not dig under or run into the excavator's tracks. Second, if the pressure in any cylinder exceeds some threshold for a brief period of time, then the trajectory following is ended. The high pressure is a signal that the bucket cannot proceed much further along the floor due to excessive forces. These are generally good indicators that the cleaning is complete, and the bucket is positioned to execute a normal dig. Therefore once one of these conditions are met, Autodig is initiated in the *Pre-Dig* stage.

Although this arrangement for cleaning the floor worked fairly well in our test sites, it should be viewed with some skepticism. The materials that we were working in were relatively soft, and thus the closed loop control did not have much trouble overcoming the resistive forces of the material. Imagine a situation in which hard immovable rocks are buried along the floor. The bucket would run into a rock, not be able to move it, and would therefore cause Autodig to execute prematurely. A better alternative to this method would be to use an impedance control or hybrid force control [Salcludean 97]. These methods effectively soften the position control, such that when high forces are encountered, the bucket can get around the obstruction. This method was not pursued due to lack of time.

Chapter 4 Modeling the Digging Process

The previous chapter described a method for executing a digging action (Autodig) which fills the bucket quickly and is robust to large variations in digging forces. The next question then is where should the dig be executed so that the bench is eroded in an optimal fashion? In the perception based planning algorithm that we have designed, a digging action is selected based on satisfying geometric constraints and optimization of a cost function. This methodology hinges on the ability to accurately model the effect of selecting a candidate action. That is, if a particular bucket pose is specified in which to initiate Autodig, how much material will be swept into the bucket? How long will it take to dig? How much energy is required?

This chapter describes a model of the digging process that is able to predict the outcome of selecting a digging action before it is executed. The model takes into account the function of the Autodig algorithm, machine actuator dynamics, and the soil-tool interaction forces. The soil-tool interaction forces can vary dramatically depending on the characteristics of the soil. Therefore the model was designed to be capable of adapting to the soil encountered at the work site. Also since the planning methodology will examine a number of candidate actions, an emphasis in the design of the model was placed on the computational speed of the predictions.

Section 4.1 discusses the overall structure of the model. Section 4.2 covers the machine actuator dynamics. Section 4.3 discusses two models that were implemented for predicting the soil-tool interaction forces, and methods for adapting the models based on actual digging forces encoun-

tered. Finally, Section 4.4 compares the predictions of the overall dig model to actual digging results.

4.1 Overall Dig Model Structure

The purpose of the Dig Model is to predict the outcome of initiating Autodig from a particular implement configuration. We assume that the trajectory of a dig is influenced by the dynamic characteristics of the hydraulic actuators, the interaction forces between the soil and the bucket, and the closed loop behavior demonstrated by Autodig.

The model of the digging process was set up as shown in Figure 14 as three interlinked predictions. First, a model of the machine's actuators is used to predict the motion of the bucket in response to the actuator commands and forces. The bucket positions define the intersections between the bucket and the terrain, which allows a soil-tool model to predict the soil reaction forces. Using a static analysis of the linkage, these forces are translated into forces at the actuators (see Appendix). The forces on the actuators and the actuator positions are used by the Autodig algorithm to generate actuator commands.

The inputs to the model are the terrain profile and a set of soil-tool properties that dictate the bucket forces. These properties will be described in more detail in Section 4.3. The model is initiated from the candidate start position, and predictions for all three components of the model are made at discrete time steps until the dig is complete. This results in a series of bucket positions which corresponds to the resultant dig trajectory. The dig trajectory and actuator forces can be used to estimate the utility of the candidate dig by estimating the time required to dig, the energy expended during digging, and the volume of material swept into the bucket.

The actuator model and the soil-tool model will be described in more detail in the following sections.

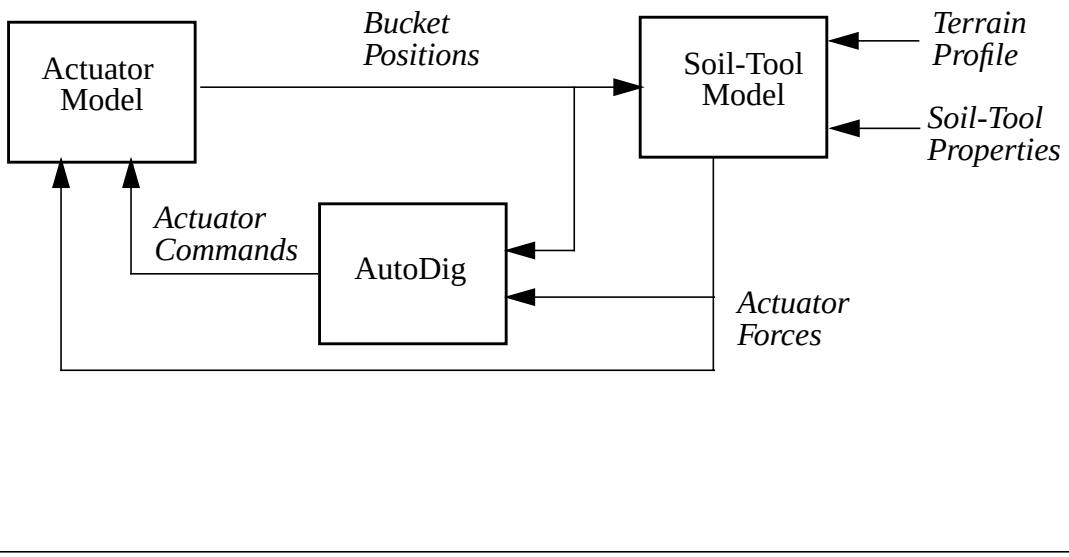

Figure 14: Composite model of the digging process. The model consists of 3 main components: Actuator Model, Soil-Tool Model, and Autodig. Based on the actuator commands and forces, the Actuator model predicts the next actuator position which defines the position of the bucket. The Soil Force model intersects the bucket with the terrain to predict the force acting on the bucket, and converts these into actuator forces. Autodig uses the actuator positions and forces to generate commands.

4.2 The Actuator Model

This section describes the hydraulic actuator model. The model predicts the motion of the hydraulic actuators which in turn defines the motion of the bucket. The motion of the actuators is dependent on the operation of the entire hydraulic system on board the excavator. Therefore Section 4.2.1 discusses the excavator's hydraulic system. Then the model itself is described in detail in Section 4.2.2. Finally the predictions of the actuator model are analyzed in Section 4.2.3.

4.2.1 Hydraulic System Description

The purpose of the actuator model is to predict the motion of the implements given the machine's current state, and the command being issued by Autodig. A simplified schematic of the system that is being modeled is shown in Figure 15. Since this system is not modeled explicitly, it is not necessary to go into great detail into how each of the sub-systems function. However it is important to note the complexity of the system, sources of non-linearity, and nature of the dynamics so that there is a foundation for approximating the system.

A hydraulic excavator is comprised of four revolute joints: swing, boom, stick, and bucket. Since digging is typically a planar motion, we model only the last three degrees of freedom. The boom, stick, and bucket are controlled by extending or retracting the hydraulic actuators across each joint. The velocity of the actuator extension is proportional to the hydraulic oil flow into the actuator which is dictated by the main implement valves. These valves are controlled by a low pressure pilot system. A controller sends electrical current to the solenoids in the pilot system for generating the pressures. The controller is responsible for converting the Autodig commands to appropriate current level in the solenoids. The source of hydraulic oil flow in the system is two variable displacement pumps which are directly coupled to the engine. The boom and bucket are controlled by one pump, while the stick and swing are controlled by the other.

There are several factors which contribute to the difficulty of modeling this system. First, the system is complex, and highly non-linear in several regards. To start with, the control valve arrangement in the implement valve stack is known as an “open-center” system. Detailed descriptions of this type of system can be found in [Merritt 67, Krishna 99]. What this means is that when the control valves are near their neutral positions, then some of the flow is leaked to tank while the remainder is used to move the actuator. The control valve can be approximated as two orifices in parallel with variable flow areas. As the area in the orifice providing flow to the cylinder is increased, the area of the orifice leading to tank is reduced. The equation governing flow in an orifice is given by:

$$Q = C_d A \sqrt{\Delta P} \quad \textbf{(1)}$$

where Q is the flow rate, C_d is the discharge coefficient, A is the area of the orifice opening, and ΔP is the pressure drop across the opening. The relative flow rates therefore are dependent on the pressure drops across the valves which in turn are dependent on the forces acting on the actuators.

Unlike many typical robot systems, the digging forces that are exerted by the excavator are extremely high, and are a significant contributor to system non-linearity. For instance, high pressures and flow rates in the hydraulic system are a significant drain on available engine power. To ensure that the engine does not die, the pumps are designed to destroke and limit flow when the available engine power is being approached. Also if the pressure in any one actuator exceeds a threshold value, then a pressure relief valve opens up which essentially stops any further motion of that implement.

We also expect there to be some coupling between the actuators due to flow limitations of the hydraulic pumps. Since the boom and the bucket get their flow from the same pump the velocity of these two implements are related. When the bucket is moving at a rapid rate, less flow is available for the boom, and hence the boom velocity is reduced. This effect is mitigated in some degree by the use of crossover valves. If the boom is given a large enough command, then a crossover valve opens allowing flow that is unused by the swing-stick pump to be used by the boom.

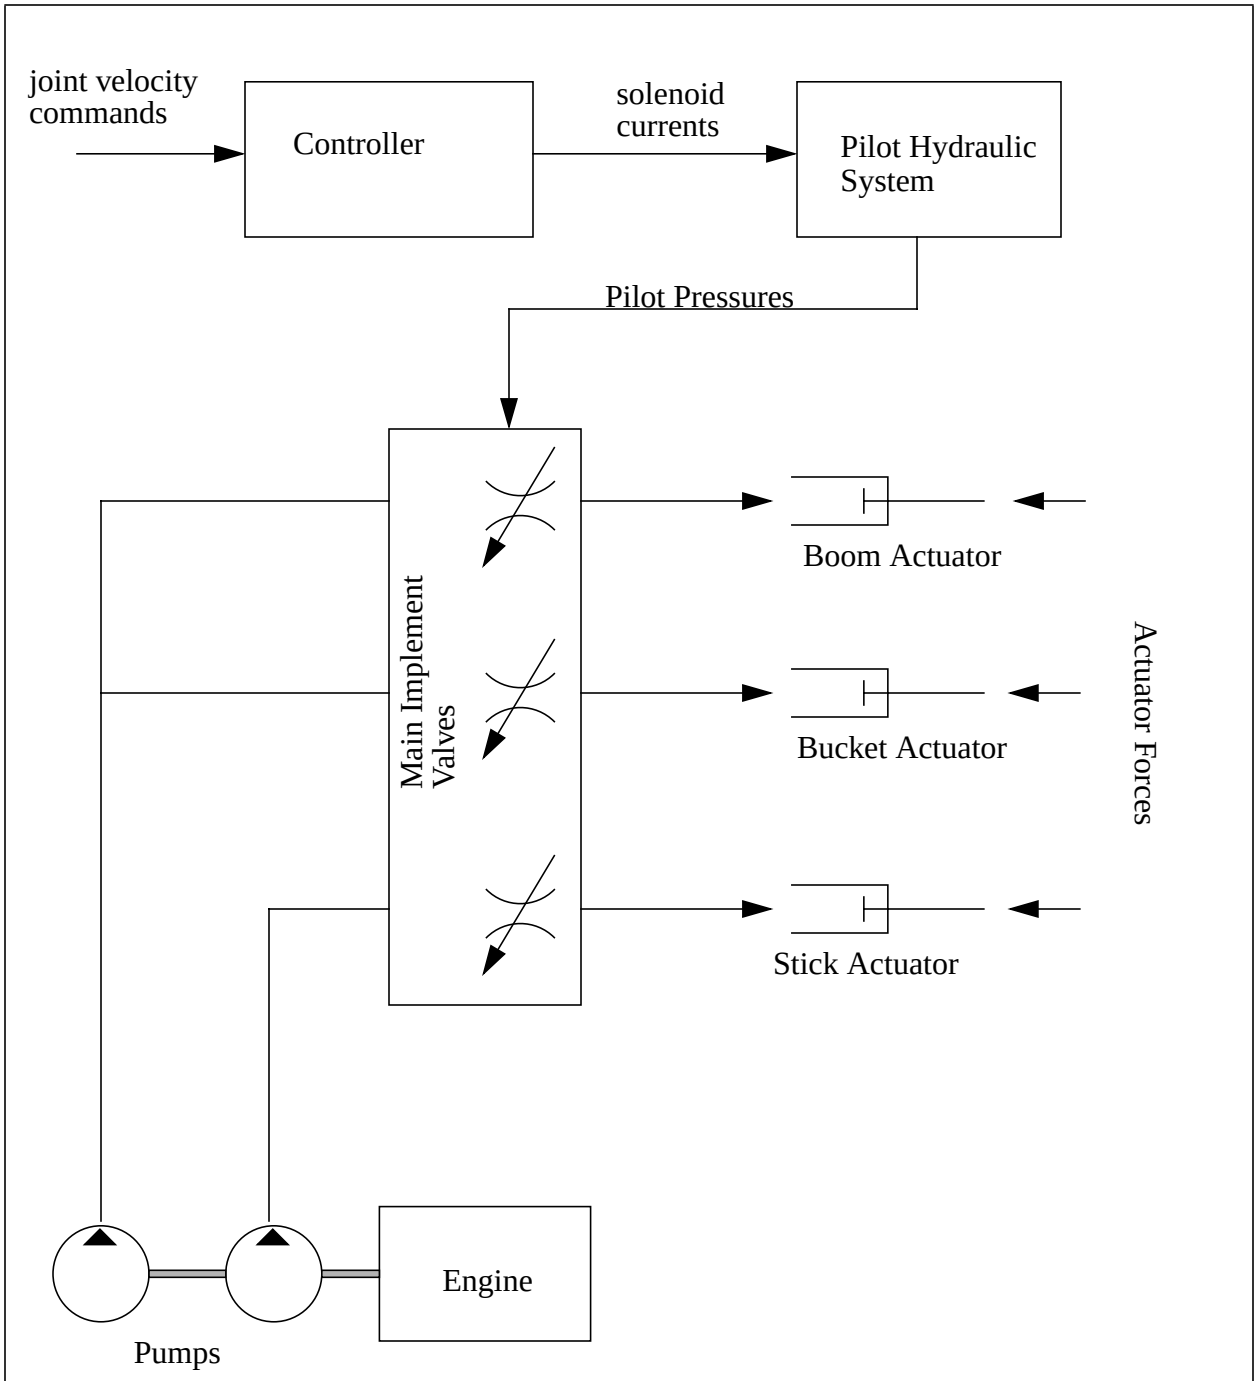

Figure 15: Simplified illustration of the hydraulic system of the excavator. A computer controller interprets commands and sends current to a low pressure pilot hydraulic system. The pressures from this system are used to control valves in the main implement valve stack. These valves control the flow from two variable displacement pumps to the cylinders.

There are many sources of system delays which also need to be addressed in the model. Some of these can be considered to be fixed delays such as the communication delays inside the controller, sensor delays, and perhaps the time required to actuate a valve in the pilot system. Probably the principle time lags however are due to the main hydraulic system. First, there is the time required for the pump to stroke in order to match the demand. Then there is the compressibility of the fluid which limits the rate at which the pressure in the system can rise. In a closed control volume, the pressure rise rate is given by:

$$\frac{dP}{dt} = -\beta\frac{Q}{V} \tag{2}$$

where P is the pressure, β is the bulk modulus, Q is the flow rate, and V is the volume of oil. Air entrainment in the hydraulic system results in a low bulk modulus, thus causing a low pressure rise rate. Finally, there is the dynamic response of the cylinder itself, which is given by:

$$M\frac{d^2x}{dt} = P_1A_1 - P_2A_2 - F \tag{3}$$

where x is the actuator position, M is the mass of the actuator rod, P and A are the pressures and areas of the two ends of the actuator, and F is the external forces. These forces include the forces due to digging, the gravitational forces caused by the implements, and any damping and inertial forces due to the motion of the implements.

Based on this understanding of the system, we have developed a model of the vehicle's actuators. As shown in Figure 14, we assume that the actuator velocities and hence the bucket motions are dependent on the command signals from the control and the actuator forces. We expect the system to be highly non-linear, and for the boom and bucket joints to be coupled. Finally there are both fixed delays in the system, and larger transient time lags.

4.2.2 Actuator Model Implementation

Due to the complexity and non-linear nature of the system, we decided that an analytically based model would be insufficient. During the planning cycle we want to be able to analyze perhaps a hundred different digs in just a few seconds. It was felt that an analytical model sufficient to capture the complex dynamics would require too much computational time. In lieu of an analytical model, we chose to investigate the use of neural networks.

A neural network is a method for mapping a non-linear relationship between a set of outputs and a set of known basis functions. Inspired by the observation of biological learning systems, it consists of a complex web of interconnecting units called perceptrons. A perceptron takes a number of inputs, combines them linearly through a set of weights, and then thresholds the result so that if the sum exceeds this threshold, the output is one. Otherwise the output is zero. The thresholding can be done by a variety of methods, but we used a function called a 'sigmoid'. A sigmoid provides a continuous means for closely approximating a unit step at the threshold value. A perceptron is shown in Figure 16.

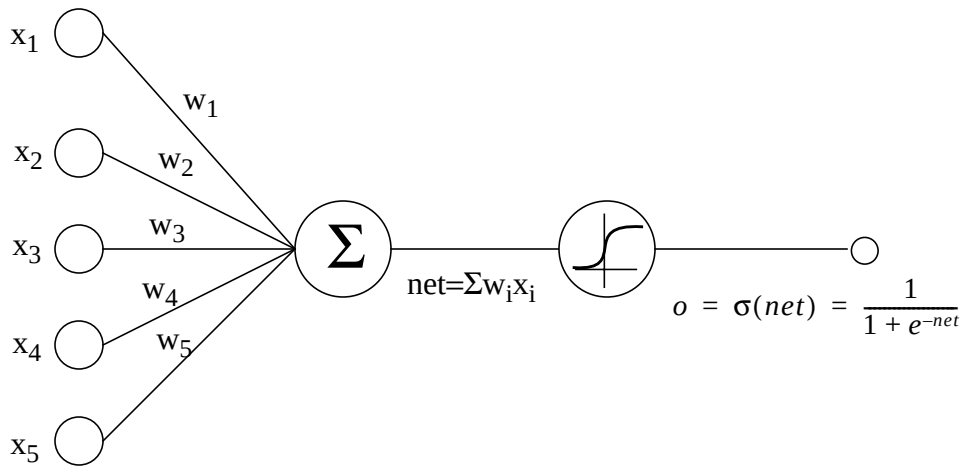

Figure 16: A perceptron used in neural networks. The perceptron sums the weighted inputs, which is then used in a thresholding function called a sigmoid.

To characterize a complex function, several perceptrons are combined together to form a neural network. The inputs to the neural network are a set of basis functions that are assumed to be the major causes of the system's outputs. For instance, in modeling the actuators, we expect that the Autodig commands and the actuator forces would be two major components that dictate the actuator motions, and thus would be included in the inputs. The mapping between the system inputs and outputs is encompassed within the weights that connect the perceptrons together. A method called "back-propagation" can be used to adjust the weights, and hence train the network. More details about back-propagation can be found in [Mitchell 97]. Training a network consists of repeatedly showing the network sample inputs and the true system outputs. This training process generally can take a long period of time as it requires numerous samples to span the possible system configurations and for distinguishing between true system fluctuations and noise in the training data. Once the network is trained however, it can generate predictions extremely fast.

There are several motivations for the use of the neural networks in this application. To start with, the system is highly non-linear and could not be reasonably captured with a globally linear regression. The neural network is also a compact form for representing the data in contrast to memory based learning methods. Finally, although the neural network is slow to train, the predictions are extremely fast, which is desirable for our application. The characteristics of the machine generally do not change significantly for long periods of time. Therefore the neural network can be trained off-line, with as many data points and as much time as needed. Fast predictions are used on-line to estimate the vehicle's motions during the digging process.

By examining the structure of the system, and through trial and error experimentation, we designed the actuator model as shown in Figure 17. The core of the model consists of three neural networks, one for each actuator. The neural networks predict the velocity of the actuators at the

next time step based on delayed commands from the controller, and actuator forces. The velocity of the actuator is then integrated in order to obtain position. Using simple trigonometry, the actuator positions are transformed into implement joint angles, which can subsequently be converted into a bucket tip position using forward kinematics [Singh 95].

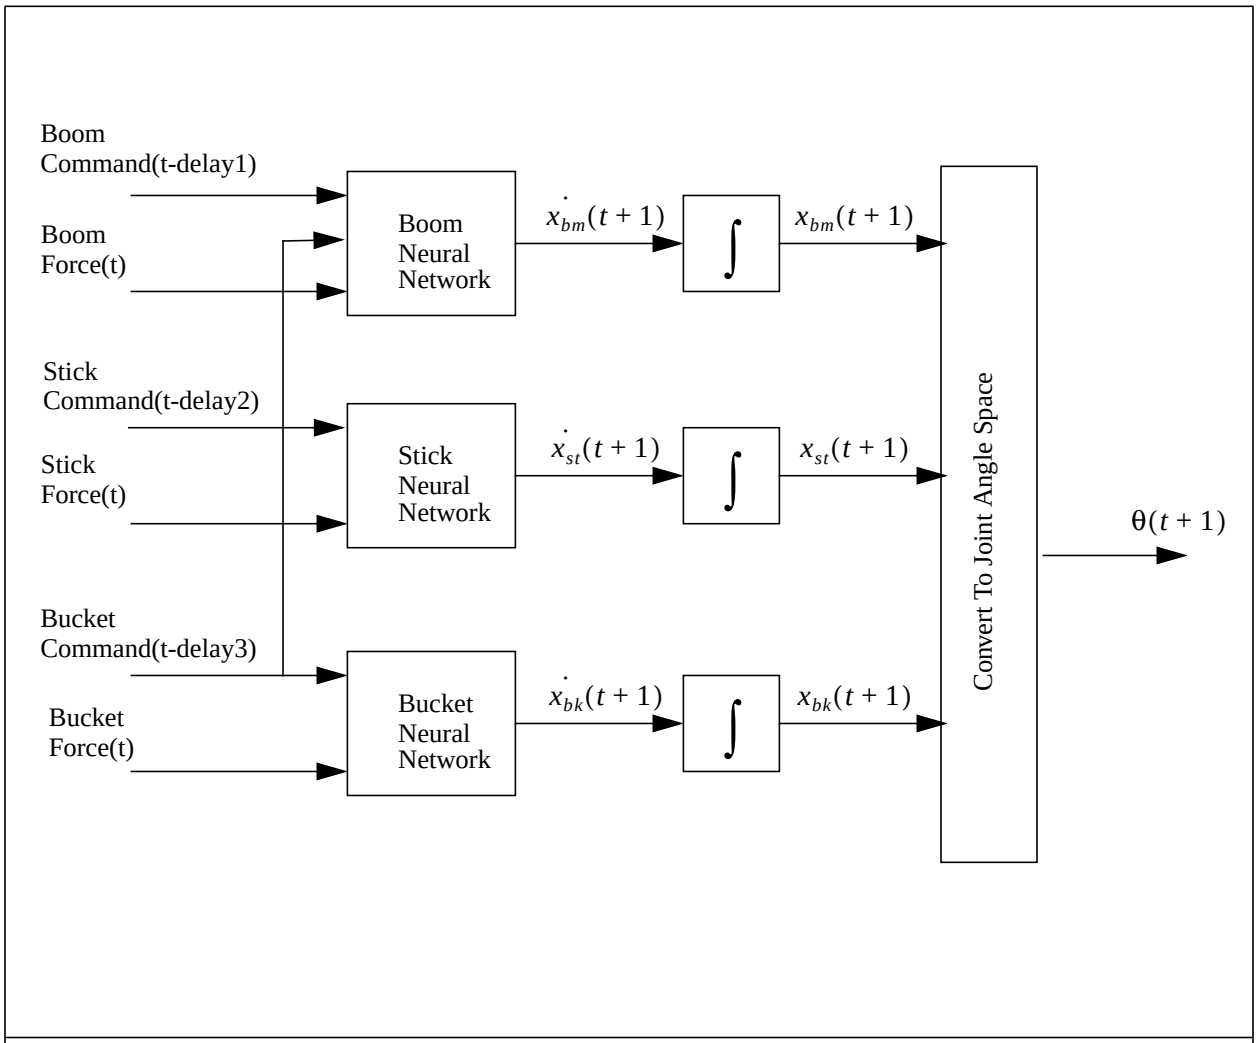

Figure 17: The vehicle model consists of three neural networks, one for each actuator. The inputs to networks consist of delayed commands, and the actuator forces. The networks predict the velocity of the actuator at the next point in time, which is then integrated to find position. The cylinder positions are then used to calculate implement joint angles. To account for coupling, the bucket command is fed to the boom network.

There are several reasons for selecting this model structure. First, the predictions are made in actuator space versus joint angle space to reduce the amount of non-linearity that the neural networks have to handle. That is, the actuator velocity is always proportional to the hydraulic flow into the cylinder, whereas the angular velocity is a highly non-linear function dependent on actuator position.

The fixed delays in the system are taken into account by using a delayed actuator command. The commands are delayed by accumulating the commands in a buffer, and then sending the com-

mands to the neural network after the delay time has elapsed. The forces are assumed to act immediately upon the actuator.

In order to take into account the large transient lags in the system, a recurrent neural network is used. This means that the output velocity of the network is used as one of the inputs. The use of the recurrent network accounts for the dynamic nature of the system. In other words, a change in velocity cannot happen instantaneously, it is dependent on the previous velocity. In equation form we can write this as:

$$v(t+1) = w_v v(t) + u(t) \quad \textbf{(4)}$$

which can be recognized as the difference form of a second order dynamic equation. The recurrent networks for the stick and bucket are shown in Figure 18. Through trial and error testing we found it sufficient to use five hidden nodes in each network.

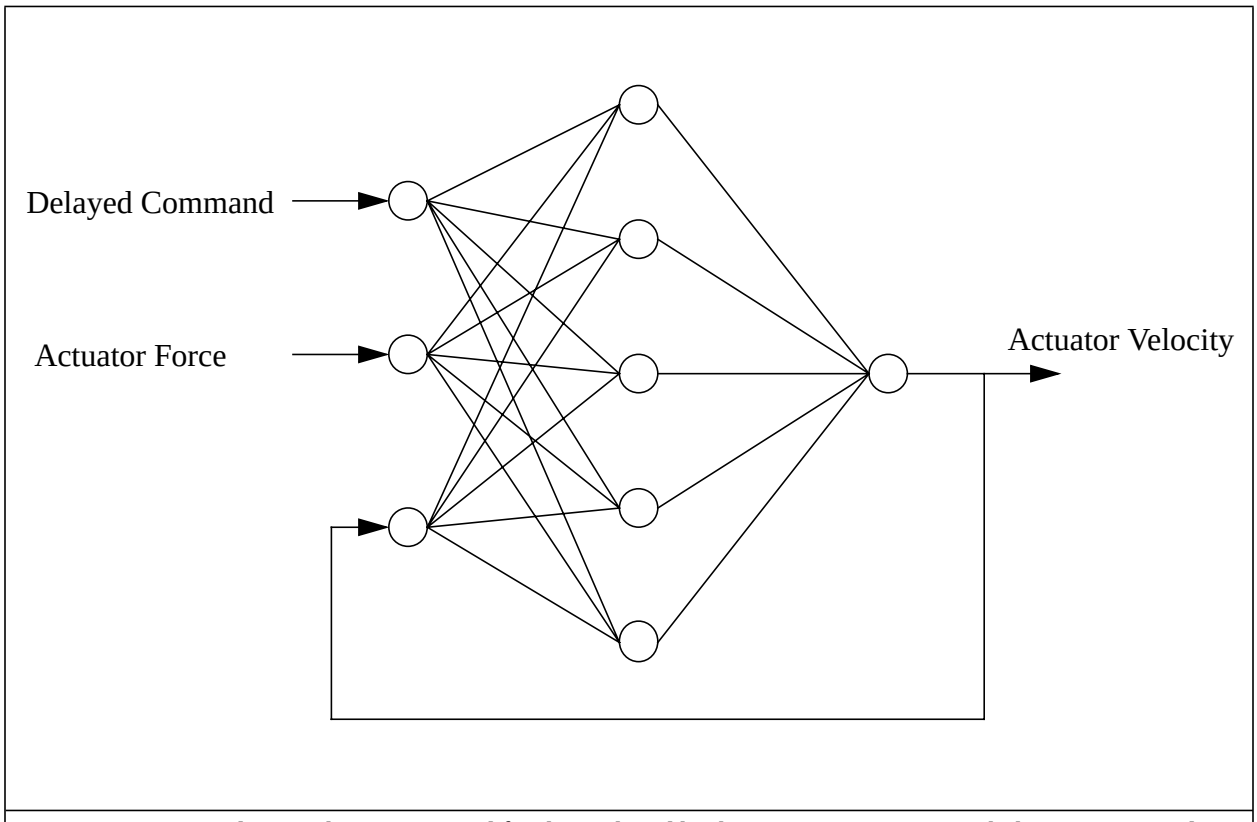

Figure 18: Neural network structure used for the stick and bucket. In a recurrent network the output is used as one of the inputs.

As previously mentioned we expected some degree of interdependency between the joints. This would suggest that a single combined network would be needed to adequately describe the dynamics of the system. However it turns out that the coupling between the joints during digging is fairly minimal. Perhaps this is because Autodig ensures that the power limitations of the machine are not reached, and also because of the crossover valves. As expected we found there

was minor coupling between the boom and bucket joint due to sharing one flow source. This was adequately accounted for by feeding the bucket command to the boom neural network as one of the inputs.

4.2.3 Actuator Model Results

The networks were trained using 30 digs from various terrain profiles. This provided approximately 1700 data points. The input data for training corresponded to the actual commands, forces, and velocities observed on the machine. Training required approximately ten minutes on a Sparc 20 workstation. After training, the models were tested on a separate set of ten digs for observing the accuracy of the predictions. Comparisons showing the predicted versus measured actuator velocities are shown in Figures 19, 20, and 21. The mean absolute error for these predictions range between 5% to 8% of the peak velocities observed in the test.

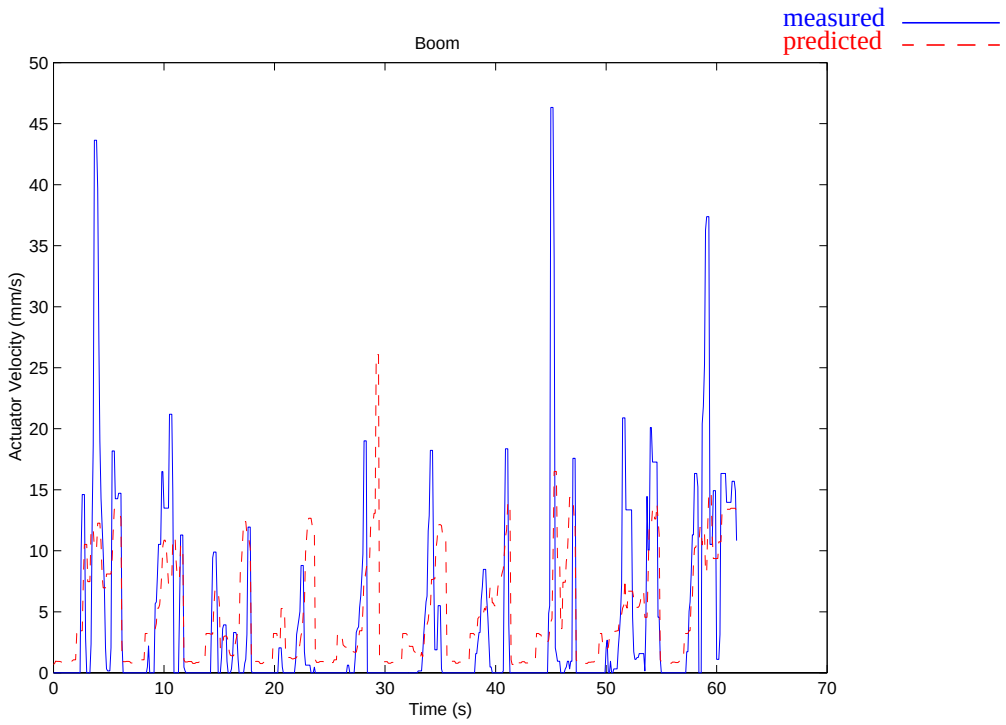

Figure 19: Comparison of predicted versus measured boom actuator velocity over 10 dig cycles. The mean absolute error is 4.3 mm/s. Peak velocity observed in testing was 87 mm/s. The measured velocity has been filtered for clarity.

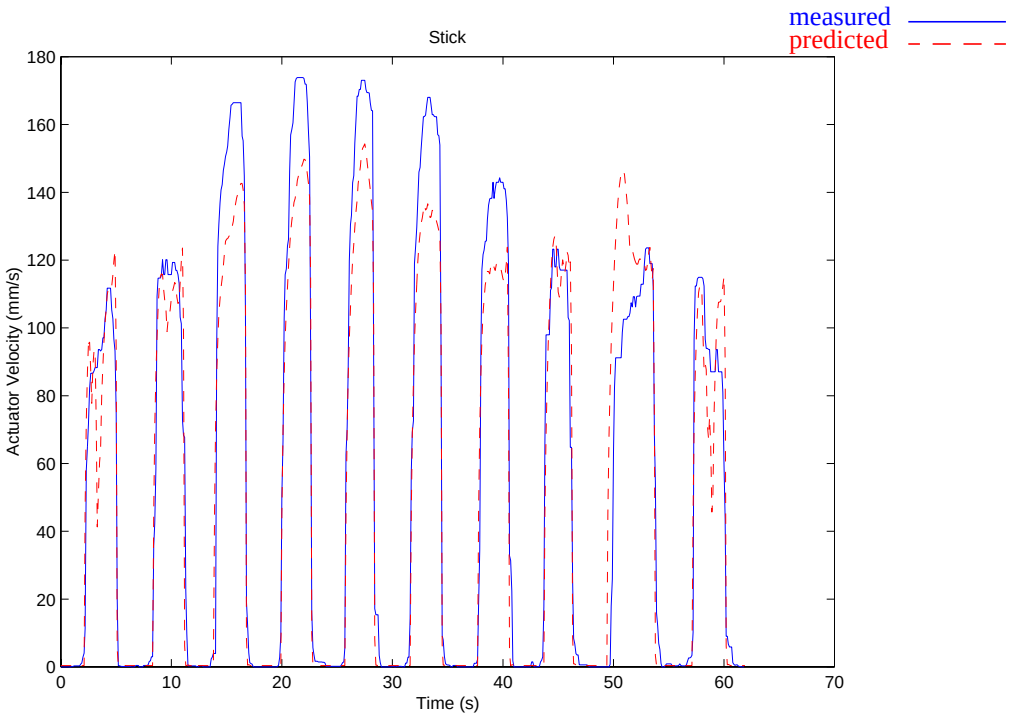

Figure 20: Comparison of predicted versus measured stick actuator velocity over 10 dig cycles. The mean absolute error is 11.8 mm/s. Peak velocity observed in testing was 243 mm/s. Measured velocity has been filtered.

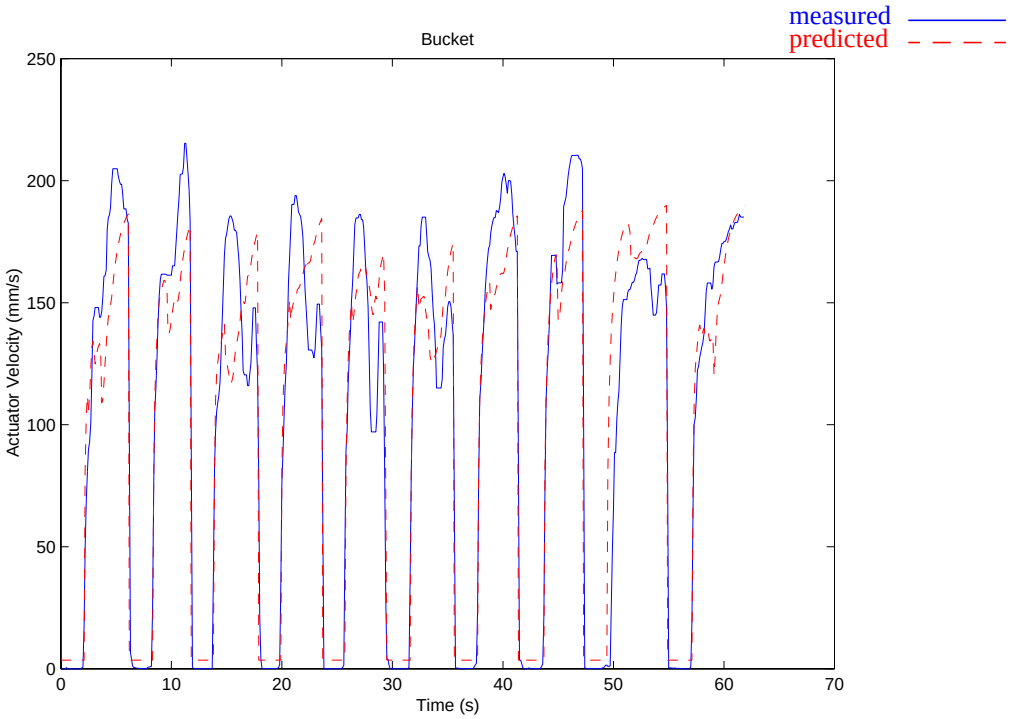

Figure 21: Comparison of predicted versus measured bucket actuator velocity over 10 dig cycles. The mean absolute error is 18.6 mm/s. Peak velocity observed in testing was 250 mm/s. Measured velocity has been filtered.

4.3 The Soil-Tool Interaction Model

This section describes two different methods for modeling the resistive force of the soil that acts against the bucket during digging. The models are based on the well known “Fundamental Earthmoving Equation” (FEE) in soil mechanics as described by [Reece 64]. This equation was developed for estimating the cutting forces of tilling implements in agricultural engineering. For our purpose, we have reformulated the equation to account for digging in a sloped terrain. Both of the soil-tool models discussed in this section are adaptive, in that the resistive forces that are encountered during digging may be used to improve future predictions. This is accomplished by estimating a set of soil-tool properties. Section 4.3.1 gives a basic description of the FEE. Section 4.3.2 discusses modifications that were made to the FEE to account for digging in a sloped terrain. Section 4.3.3 describes the two modeling methods that were employed for predicting the resistive forces. Section 4.3.4 discusses estimation of the soil-tool properties for both methods. Finally section 4.3.5 compares the results of the two models to measured data.

4.3.1 Fundamental Earthmoving Equation

The FEE predicts the resistive forces of the soil acting against a flat blade moving horizontally through the soil as shown in Figure 22. When the blade moves forward, the soil is sheared away from itself in front of the blade, creating a “wedge” of material that slides along the failure surface. The FEE predicts the static force required to shear the material based on all of the forces that are acting on the wedge.

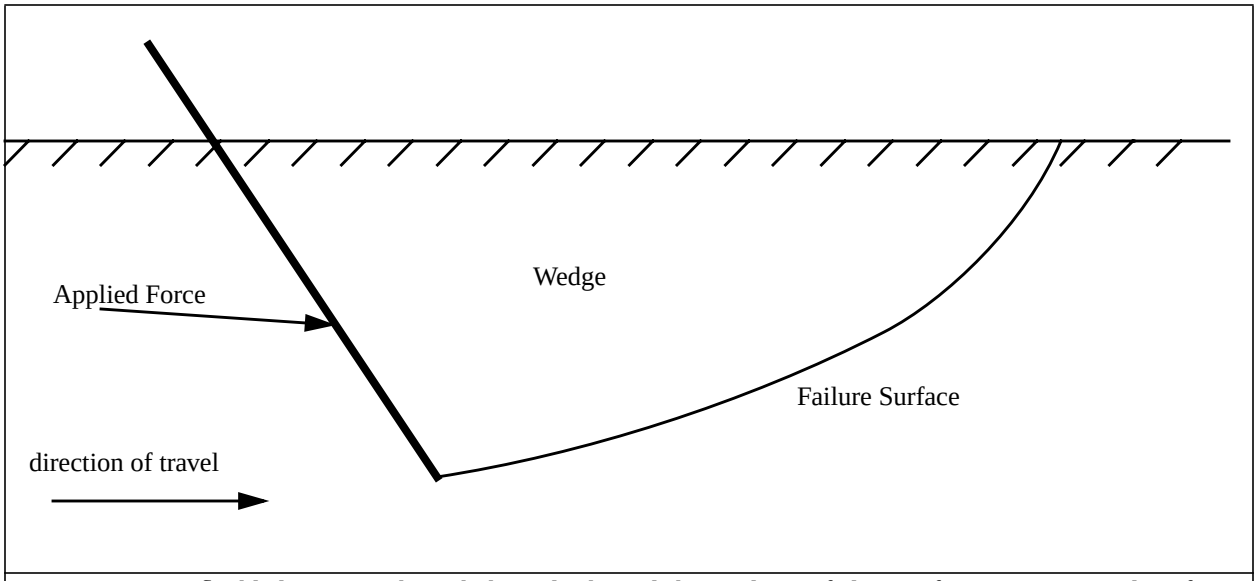

Figure 22: A flat blade moving through the soil. The soil shears along a failure surface creating a wedge of material. By analyzing the forces on the wedge, the force applied by the blade can be predicted.

Assuming that the failure surface is a plane, the wedge can be represented as shown in Figure 23.

The forces that are acting on the wedge consist of:

- The shear force of the material away from itself which is a function of the cohesiveness of the material.
- The reaction force of the soil against the sliding wedge.
- The weight of the material in the wedge, and the weight of previously dug material known as the surcharge.
- The adhesion of the soil to the tool.
- The force of the tool against the wedge.

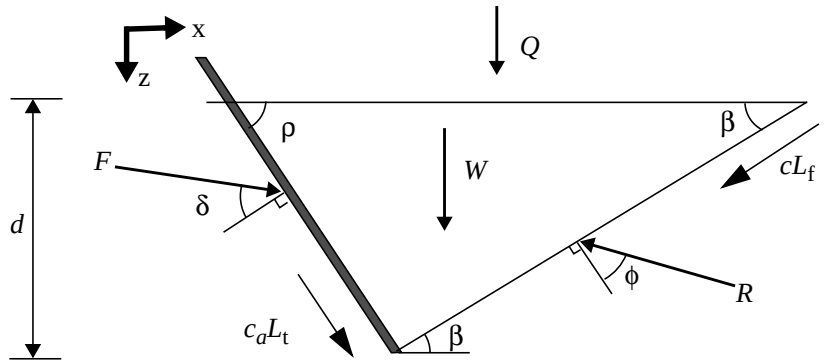

Figure 23: Static analysis of wedge model. W is the weight of the wedge, L_t is the length of the tool, L_f is the length of the failure surface, Q is the weight of the surcharge, ϕ is the soil-soil friction angle, c is the cohesiveness of the soil, c_a is the adhesion between the soil and blade, δ is the soil-tool friction angle, β is the failure surface angle, ρ is the rake angle, d is the depth of the tool in the soil, R is the force of the soil resisting the moving of the wedge, and F is the force exerted by the tool on the wedge. [McKeys85]

Writing the force equilibrium equations for a blade of unit width:

$$\begin{aligned} \sum F_x &= F\sin(\rho+\delta) + c_a L_t \cos\rho - R\sin(\beta+\phi) - cL_f\cos\beta = 0 \\ \sum F_z &= -F\cos(\rho+\delta) + c_a L_r \sin\rho + cL_f\sin\beta - R\cos(\beta+\phi) + W + Q = 0 \end{aligned} \quad (5)$$

and then solving the equations for F:

$$F = \frac{W + Q + cd[1 + \cot\beta\cot(\beta+\phi)] + c_a d[1 - \cot\rho\cot(\beta+\phi)]}{\cos(\rho+\phi) + \sin(\rho+\phi)\cot(\beta+\phi)} \quad (6)$$

The weight of the soil in the wedge is given by:

$$W = \gamma g \frac{d^2}{2}(\cot\rho + \cot\beta) \tag{7}$$

where γ is the density of the material. The weight of the surcharge is given by:

$$Q = qd(\cot\rho + \cot\beta) \tag{8}$$

where q is the surcharge pressure.

Since the force due to adhesion is small compared to the other forces, it will be ignored. Rearranging the force equation, and accounting for the width of the bucket w, we obtain:

$$F = (\gamma g d^2 N_\gamma + cdN_c + qdN_q)w$$

$$\begin{aligned} N_\gamma &= \frac{\cot\rho + \cot\beta}{2[\cos(\rho+\delta) + \sin(\rho+\delta)\cot(\beta+\phi)]} \\ N_c &= \frac{1 + \cot\beta\cot(\beta+\phi)}{\cos(\rho+\delta) + \sin(\rho+\delta)\cot(\beta+\phi)} \\ N_q &= \frac{\cot\rho + \cot\beta}{\cos(\rho+\delta) + \sin(\rho+\delta)\cot(\beta+\phi)} \end{aligned} \tag{9}$$

The three factors N_γ, N_c, and N_q, can be calculated based on the geometry of the wedge. The three factors dictate the force due to the weight of material in the wedge, the force due to cohesion of the soil, and the force due to the surcharge pressure.

4.3.2 Modifications to the FEE

As previously noted, the FEE was developed for agricultural tools, and thus it was assumed that the terrain profile would be relatively flat. In mass excavation however, the ground profile is usually sloped, and the material that passes over the blade is captured and retained by the bucket.

Figure 24 shows a cross section of the wedge model that compensates for accumulating the material in the bucket. Note that in this model, the surcharge is assumed to accumulate behind the bucket tip versus being evenly distributed over the wedge. The material that is shaded in gray accounts for all of the material that has passed over the bucket tip, and it is assumed that all of this material stays inside the bucket. We refer to this material as the “swept volume”, V_s, and it is calculated by continuously integrating the volume of material that passes over the tip during digging.

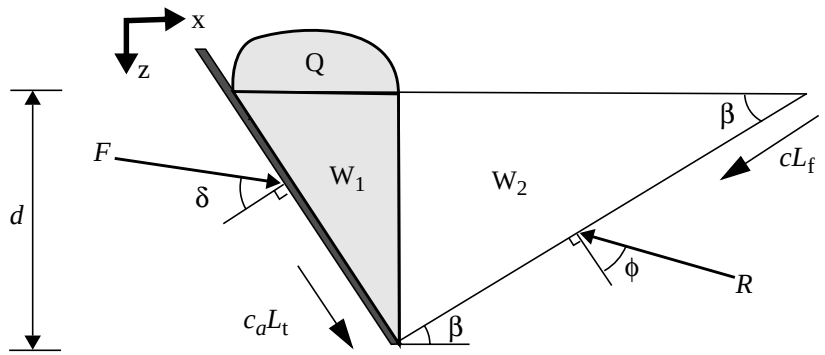

Figure 24: Wedge model that accounts for the material being retained in the bucket. The material in the shaded region corresponds to the swept volume V_s. Q is the surcharge, W_1 is the weight of the material above the bucket, W_2 is the weight of the rest of the material in the wedge. L_t is the length of the tool, L_f is the length of the failure surface, ϕ is the soil-soil friction angle, c is the cohesiveness of the soil, c_a is the adhesion between the soil and blade, δ is the soil-tool friction angle, β is the failure surface angle, ρ is the rake angle, d is the depth of the tool in the soil, R is the force of the soil resisting the moving of the wedge, and F is the force exerted by the tool on the wedge.

The weight of the material in the shaded region for a unit bucket width can now be calculated by:

$$W_1 = V_s \gamma g \tag{10}$$

and the weight of the remaining material in the wedge is given by:

$$W_2 = \frac{1}{2}\gamma d^2 \cot\beta \tag{11}$$

The force equation for a given bucket width now becomes:

$$F = \frac{1}{2}d^2 w g N_w + c w d N_c + V_s \gamma g N_q$$

$$\begin{aligned} N_w &= \frac{\cot\beta}{2[\cos(\rho+\delta)+\sin(\rho+\delta)\cot(\beta+\phi)]} \\ N_c &= \frac{1+\cot\beta\cot(\beta+\phi)}{\cos(\rho+\delta)+\sin(\rho+\delta)\cot(\beta+\phi)} \\ N_q &= \frac{1}{\cos(\rho+\delta)+\sin(\rho+\delta)\cot(\beta+\phi)} \end{aligned} \tag{12}$$

Finally to compensate for the slope of the terrain, the wedge model as shown in Figure 25 is used.

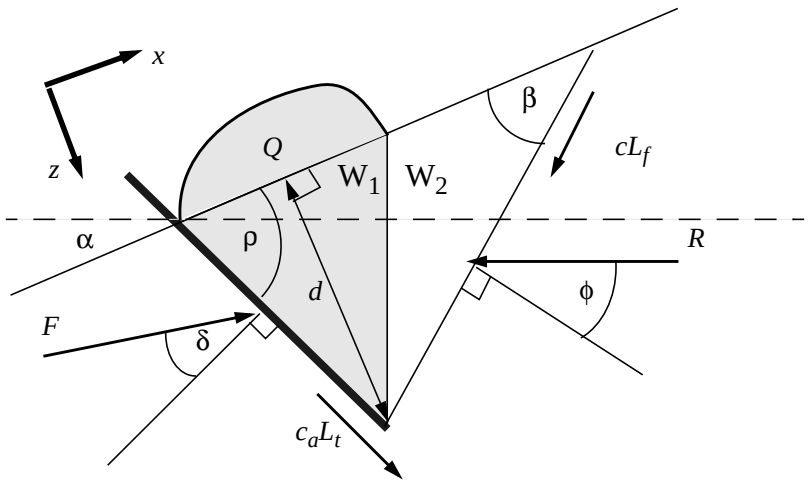

Figure 25: Wedge model that accounts for the material being retained in the bucket, and for the slope of the terrain.The material in the shaded region corresponds to the swept volume V_s. Q is the surcharge, W_1 is the weight of the material above the bucket, W_2 is the weight of the rest of the material in the wedge. L_t is the length of the tool, L_f is the length of the failure surface, ϕ is the soil-soil friction angle, c is the cohesiveness of the soil, c_a is the adhesion between the soil and blade, δ is the soil-tool friction angle, β is the failure surface angle, ρ is the rake angle relative to the soil surface, d is the depth of the tool perpendicular to the soil surface, R is the force of the soil resisting the moving of the wedge, F is the force exerted by the tool on the wedge, and α is the terrain slope.

Note how the depth is now measured perpendicular to the terrain, and the rake angle ρ is measured between the surface of the terrain and the blade. Again the material that is shaded corresponds to the swept volume. The *x* axis of the coordinate system has been oriented parallel to the terrain, so that the equilibrium equations become:

$$\sum F_x = F(\sin(\rho+\delta)) - V_s g \sin\alpha - R\sin(\beta+\phi) - cL_f w\cos\beta - W_2\sin\alpha = 0$$

$$\sum F_z = F(\cos(\rho+\delta)) - V_s g \cos\alpha - R\cos(\beta+\phi) - cL_f w\sin\beta - W_2\cos\alpha = 0 \quad \textbf{(13)}$$

where α is the terrain angle. Removing the soil reaction force R from the equation, we obtain:

$$F = \frac{V_s g(\cos\alpha + \sin\alpha\cot(\beta+\phi)) + cL_f w(\cos\beta\cot((\beta+\phi)+\sin\beta)) + W(\sin\alpha\cot(\beta+\phi)+\cos\alpha)}{\sin(\rho+\delta)\cot(\beta+\phi)+\cos(\rho+\delta)} \quad \textbf{(14)}$$

The weight of the material in the unshaded region is given by:

$$W_2 = \frac{1}{2}d^2(\cot\beta - \tan\alpha)w\gamma g \quad \textbf{(15)}$$

and the length of the failure surface:

$$L_f = \frac{d}{\sin\beta} \tag{16}$$

Rearranging the equation, we get:

$$F = d^2 w \gamma g N_w + c w d N_c + V_s \gamma g N_q$$

$$\begin{aligned} N_w &= \frac{(\cot\beta - \tan\alpha)(\cos\alpha + \sin\alpha\cot(\beta+\phi))}{2[\cos(\rho+\delta) + \sin(\rho+\delta)\cot(\beta+\phi)]} \\ N_c &= \frac{1 + \cot\beta\cot(\beta+\phi)}{\cos(\rho+\delta) + \sin(\rho+\delta)\cot(\beta+\phi)} \\ N_q &= \frac{\cos\alpha + \sin\alpha\cot(\beta+\phi)}{\cos(\rho+\delta) + \sin(\rho+\delta)\cot(\beta+\phi)} \end{aligned} \tag{17}$$

Note that when the terrain angle is zero, this equation is identical to Equation 12.

The coordinate system in which this model was derived had the x axis parallel to the terrain. However for the rest of this paper the soil forces will be transformed to a coordinate system attached to the machine called the base frame. The base frame is shown in Figure 26. Given the new coordinate system, the forces acting on the soil by the bucket are given by:

$$\begin{aligned} F_x &= F(\cos\delta\sin(\rho-\alpha) - \sin\delta\cos(\rho-\alpha)) \\ F_z &= F(\cos\delta\cos(\rho-\alpha) + \sin\delta\sin(\rho-\alpha)) \end{aligned} \tag{18}$$

Finally, we assume that the soil-tool forces are applied at the cutting edge of the bucket. In general this appears to be a good assumption. Obviously however the gravitational force of the material acts closer to the centroid of the bucket. To account for this discrepancy, a moment is applied to the bucket which compensates for the moment created by offsetting the force. The moment is given by:

$$M_y = F_g x_c \tag{19}$$

where F_g is the force due to the weight of the material in the bucket, and x_c is the distance from the centroid of the bucket to the cutting edge along the x axis.

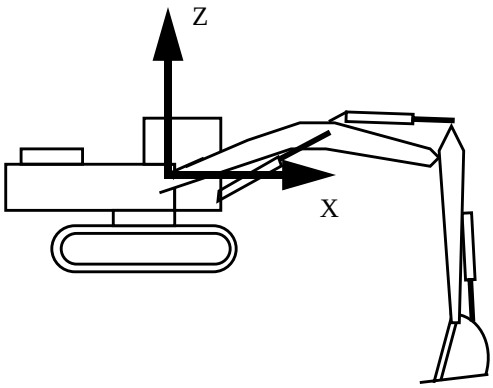

Figure 26: The forces are represented in the base coordinate frame.

4.3.3 Soil Modeling Methods

The objective of the soil model is to predict the force that the soil is imparting to the bucket given the geometry of the intersection between the bucket and the ground, and the soil characteristics. In equation form:

$$F = f(\Gamma, \Psi) \tag{20}$$

where F is a vector representing the soil forces and moments acting on the bucket, f is the model, Γ is the geometry of the intersection between the bucket and the ground, and Ψ is a set of parameters that characterize the soil. Γ can be calculated by knowing the shape of the terrain and the position and orientation of the bucket. Estimation of the soil-tool properties Ψ will be discussed in the next section.

Two methods were pursued for modeling the soil reaction forces. The first method is to use the modified FEE analytical model as was developed in the previous section. When using the analytical model, the geometry can be specified by $\Gamma=(d, \alpha, \rho, V_s)$, and the soil characteristics $\Psi=(\gamma, \beta, \phi, \delta, c)$.

The second method was to find an empirical relationship between Γ and F. In this method we assume that we can find a relationship:

$$f = \sum_i \Gamma_i \Psi_i = \Gamma_1\Psi_1 + \Gamma_2\Psi_2 + \Gamma_3\Psi_3 + \ldots \tag{21}$$

In this model it is necessary to find an empirical relationship for all of the force components, (F_x, F_z, and M_y), whereas with the FEE model the magnitudes of the force components are given by Equation 18. Using the terms in the FEE as inspiration for deriving *f*, we selected $\Gamma=(d^2, \cos(\rho), \alpha, V_s)$. Although it is unnecessary for the terms of Γ to be identical for each of the force components, we found that it was sufficient to do so. If any one term Γ_i is unnecessary for calculating a component of the force, then the corresponding Ψ_i should reflect this by having a smaller magnitude compared to the other parameters.

4.3.4 Soil-Tool Property Estimation

Again, the objective of this work is to predict the forces that would be experienced by the bucket during digging, so that the dig model can estimate the trajectory that the bucket would take through the soil. At each time step the dig model predicts the state of the machine, and hence the pose of the bucket is known. By intersecting the bucket with the terrain at each time step, we can calculate all of the parameters of Γ, such as depth and rake angle, and by integrating the intersection over time, we can calculate V_s. But to predict the force, we still need to have some estimation for the soil-tool properties Ψ.

We know that the soil characteristics can vary dramatically from one work site to another depending on the type of material that is being excavated. Even within one type of material the soil characteristics can be quite different due to compaction or seasonal changes. In our test environment, the soil was quite hard and difficult to dig in winter months when the ground was frozen. During rainy months, the soil was soft and sticky. And during the summer the soil was dry and powdery. Therefore to accommodate the multitude of different soil characteristics that could be encountered, it is necessary for the soil model to be adaptable. That is, we need to modify the soil-tool properties Ψ based on the forces that are encountered during digging.

Suppose therefore that we have collected some data during several digging operations, and that this data includes the poses of the bucket during the digs, and the shape of the terrain. Using this information, we can find $\hat{\Gamma}$ which is a matrix containing the geometry of the actual intersection of the bucket with the terrain for all of the digs. Suppose also that our data includes the actual forces $\hat{F}$ that are encountered during digging which can be derived by monitoring the pressures in the hydraulic actuators. Our task then is to find an estimate of Ψ based on $\hat{\Gamma}$ and $\hat{F}$. Figure 27 illustrates this process.

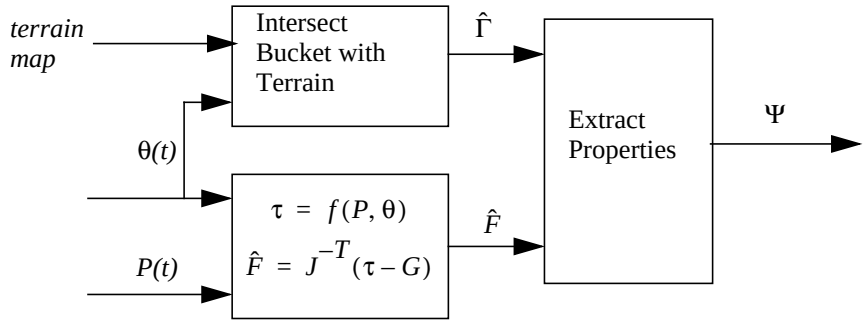

Figure 27: Extraction of the soil-tool properties. The path of the bucket during digging is intersected with the terrain map to form a matrix containing the intersection geometries. The pressures during digging are converted to forces at the bucket tip using the Jacobian of the mechanism. (See Appendix for details). The forces and geometries are then used to extract the soil-tool properties.

Since we have formulated the empirical model as a linear combination of terms, we can use least squares regression analysis to solve for Ψ. We assume:

$$\hat{F} = \hat{\Gamma}\Psi \tag{22}$$

We can then find Ψ by forming the pseudo-inverse of the geometry matrix:

$$\Psi = (\hat{\Gamma}^T\hat{\Gamma})^{-1}\hat{\Gamma}^T\hat{F} \tag{23}$$

For the analytical model however, the equation is not invertible. Therefore a search is conducted for a set of parameters Ψ that minimizes the discrepancy between modeled forces and actual forces. That is, find Ψ that minimizes $J(\hat{F}, f(\hat{\Gamma},\Psi))$, where J is some objective function. For this work we have selected J to be the root mean square error of the force in the x and z directions. It would be equally valid to include the moment within the calculation or to use a linear combination of the errors in each direction.

There are numerous methods that could be used for conducting this search. One alternative is an exhaustive search. This simply means that the soil parameters are discretized within their possible ranges, and then each possible combination of discretized parameters is tested to determine if it is the combination that minimizes J. This method has two major problems. First, it will probably not provide an exact solution since the parameters must be discretized, and the accuracy of the result is limited by the resolution of the discretization. Second, the computational requirements can be enormous since every possible combination is explored.

Alternatively we have selected a method that combines a stochastic search method called simulated annealing with an efficient form of gradient descent [Luengo 98]. The simulated annealing algorithm works by conducting a pseudo-random search for a set of parameters that minimizes *J*. Each trial of simulated annealing begins by selecting a random set of parameters and then calculating *J*. Then a set of parameters is randomly selected that is "near" the first set, and again calculating *J*. If the objective function for the second set is better than the first set, then the second set becomes the base for the next move. Otherwise the second set will become the base set based on the probability function:

$$e^{\frac{-(J_2 - J_1)}{kT}} \tag{24}$$

where T is the current number of trials, k is the cooling factor, and J_1 and J_2 correspond to the objective functions of the first and second sets. As the number of trials increase, the probability of moving to a set of parameters with a worse objective function is reduced. The cooling factor can be adjusted to dictate how quickly the probability decreases. By moving the base set towards the parameters with better values of *J* most of the time, it is hoped that eventually a reasonable estimate of the global minimum can be found. The purpose of moving towards worse objective functions some of the time is to avoid becoming trapped in local minima.

In gradient descent, the idea is again to move from parameter sets with higher values of *J* to parameter sets with lower values of *J*. Except instead of moving randomly, the gradient of the function is followed. There are several different methods for doing gradient descent, however we have chosen one called Powell's method because it is computationally less taxing [Press 88]. In Powell's method the gradient is calculated using the conjugate directions starting from the basis vectors of the search space.

To combine these methods, several trials of simulated annealing are conducted, and the best solution is stored from all of the trials. This serves as the starting point for conducting gradient descent. The idea is that the use of simulated annealing will find a set of parameters close to the global minima, and that gradient descent will find the true minimum within some finer accuracy. The combined methods work well together because simulated annealing can conduct a broad search over the workspace quickly, while the gradient descent method provides the needed accuracy.

4.3.5 Soil Model Results

We performed a cross validation test using data from 23 separate digs (approximately 1100 data points) that were taken back to back on the same day. Through trial and error testing, we found that 300 data points usually provided an adequate amount of data for estimating the soil-tool properties. Since the duration of each dig is five seconds on average (50 data points), then this corresponds to a window size of approximately six digs.

Figures 28 and 29 show a comparison of predicted versus measured force components for the analytical model and the empirical models respectively. For each dig, 300 data points adjacent to the dig were used to estimate the soil-tool properties. In addition, the last digs in the test were used to

estimate the properties for the first digs. Qualitatively speaking both models appear to fit the data fairly well for all of the separate force components. Figure 30 overplots the predictions of the force magnitudes for two of these digs and Figure 31 shows histograms of the error in the predicted force magnitudes for all twenty three digs. The mean absolute error for both models is roughly 10% of the peak force observed in the test.

Somewhat surprisingly, there is not a significant difference between the analytical and empirical model in terms of accuracy. However there are other factors to consider. First, in terms of computational speed, we must examine the time required to make a prediction and the time required to extract the soil-tool properties. On an SGI R10000 processor, the difference in prediction times was negligible (16 ms analytical versus 13 ms empirical on average for each dig). The difference in computational time required to extract the soil-tool properties however is dramatic. The analytical model required 3400 ms to extract the soil-tool properties from a set of 300 data points, whereas the empirical model required only 8 ms.

Another factor to consider is the ability of a model to extrapolate to new soil conditions. The soil parameters in the empirical model do not correspond to any identifiable soil characteristic. They are simply coefficients for each geometry term. This can be seen by examining the soil-tool properties that were extracted for both types of models in Figures 32 and 33. Note that with the analytical model a relatively constant estimate of the soil properties is maintained, whereas the empirical model coefficients are more erratic. That is, the analytical model seems to have honed in on the intrinsic characteristics of the soil. Thus the soil-tool properties in the analytical model are physically based, and the model may be used with some confidence to extrapolate to new soil conditions. With the empirical model, it would be unclear as to how the soil-tool properties should be modified.

Since the extraction of the soil-tool properties for the empirical model was much faster than the analytical model, we used the empirical model in the perception based dig planning system. The analytical model however has proven to be useful in our off-line simulator allowing us to extrapolate to much harder soil conditions.

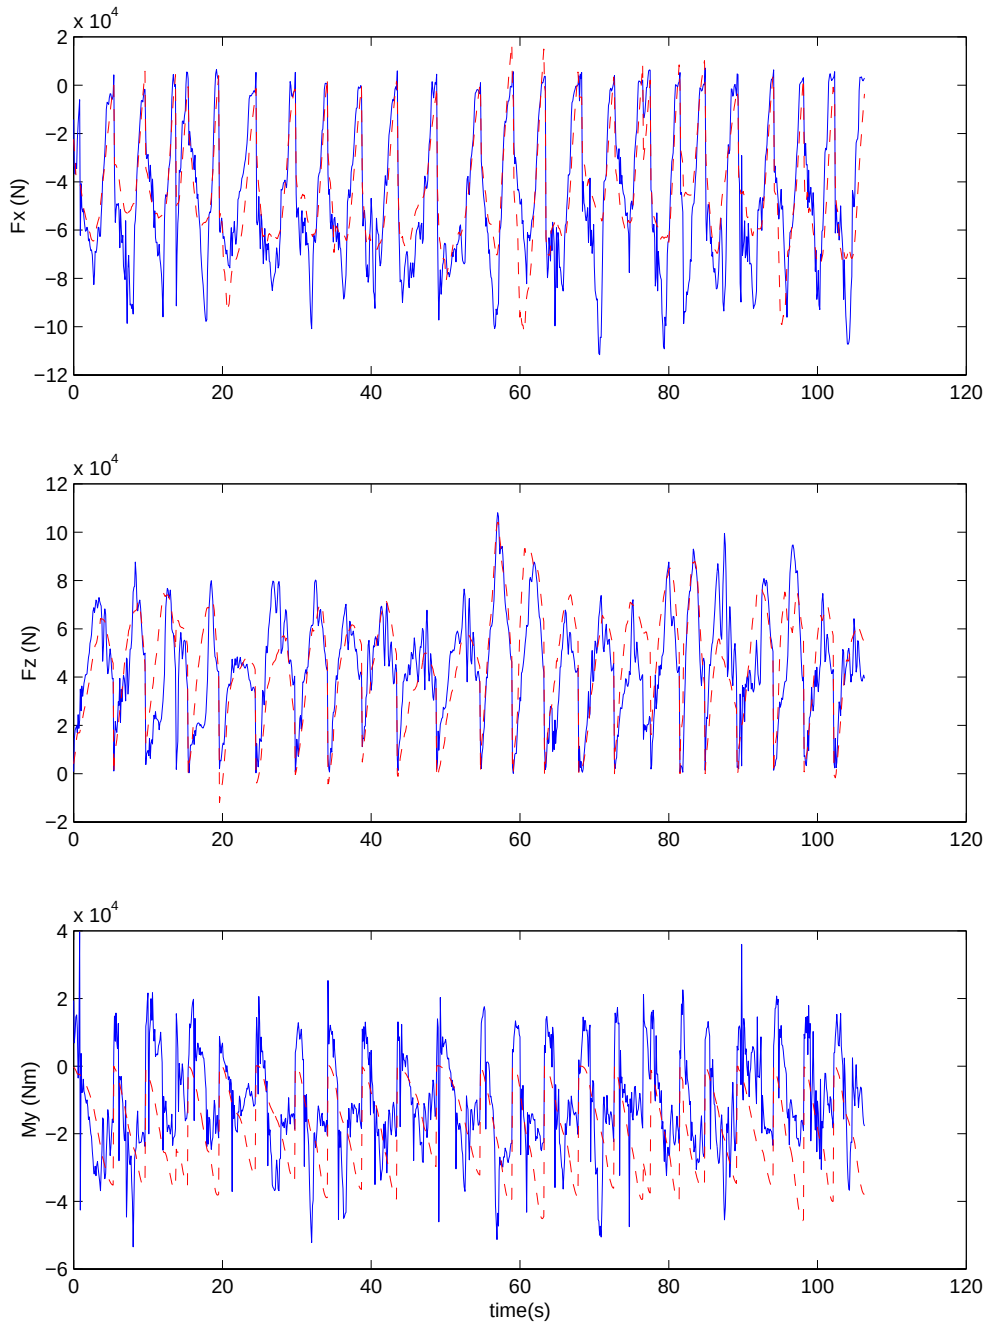

Figure 28: Comparison of measured force components to predicted force components for the analytical model for a set of 23 separate digs. 300 adjacent data points are used to estimate the soil-tool properties before each dig.

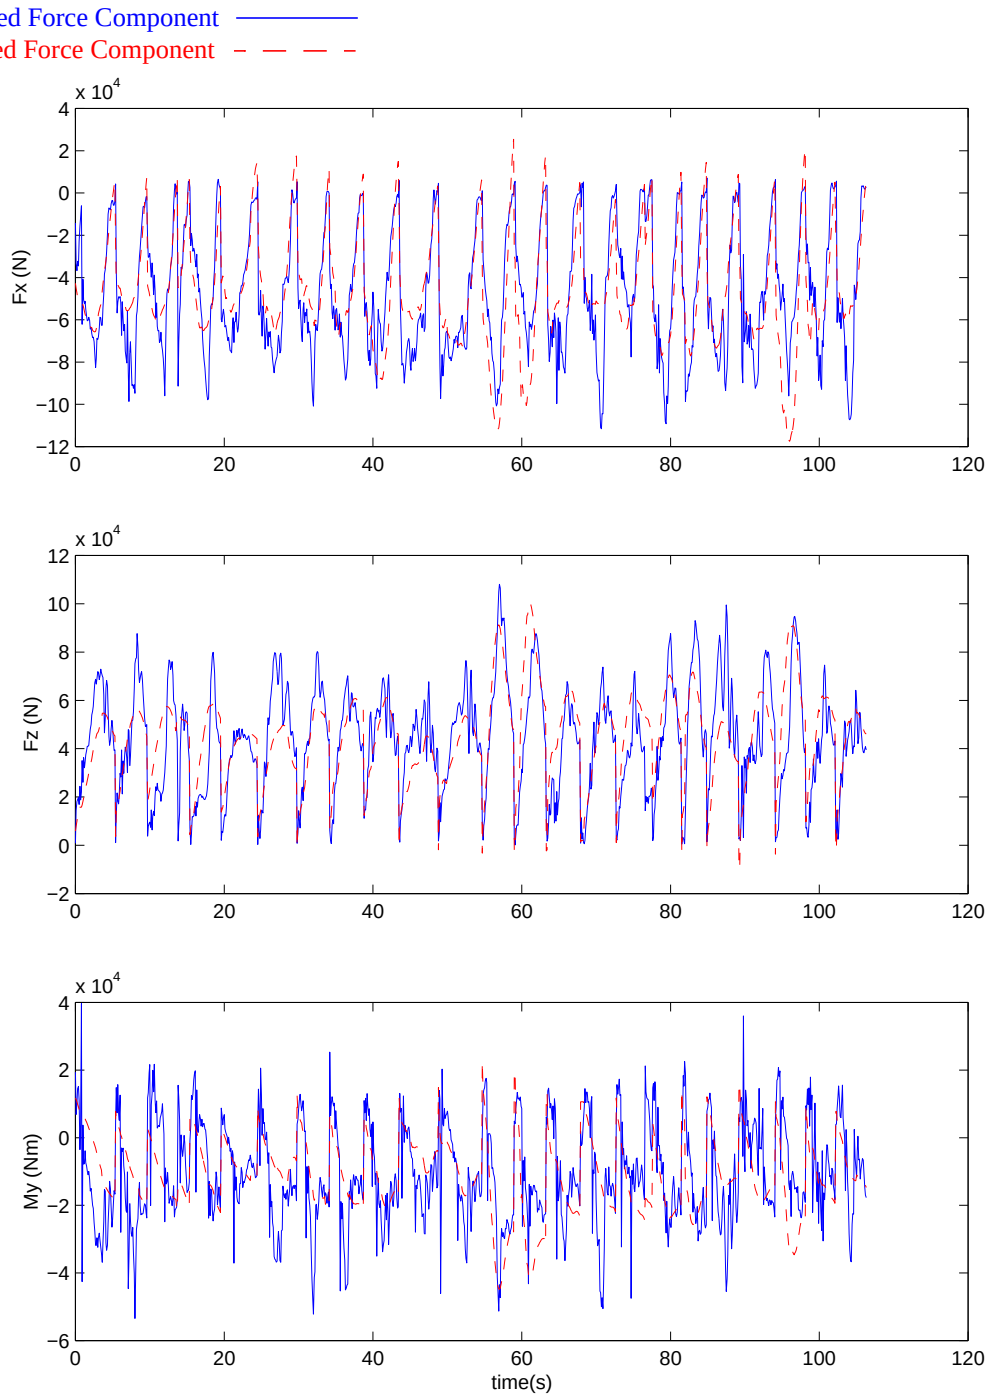

Figure 29: Comparison of measured force components to predicted force components for the empirical model for a set of 23 separate digs. 300 adjacent data points are used to estimate the soil-tool properties before each dig.

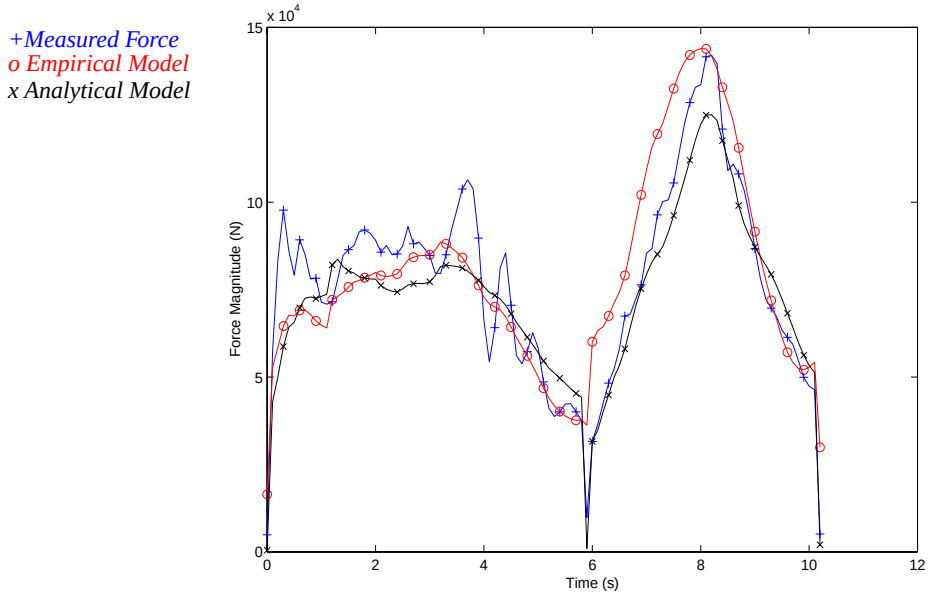

Figure 30: Comparison of the measured force magnitude to the predicted force magnitude for both the analytical and empirical models for two digs.

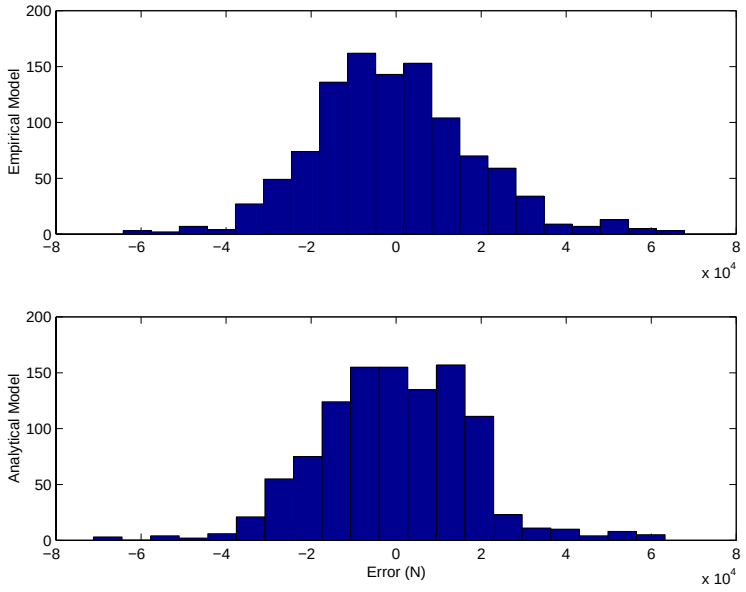

Figure 31: Histograms of the model errors for 23 digs. The mean absolute error in force magnitude for the analytical model is 13,980 N versus 14,790 N for the empirical model. Standard deviations are 10,790 N and 11,880 N respectively. The peak force observed in the test was approximately 142,000 N.

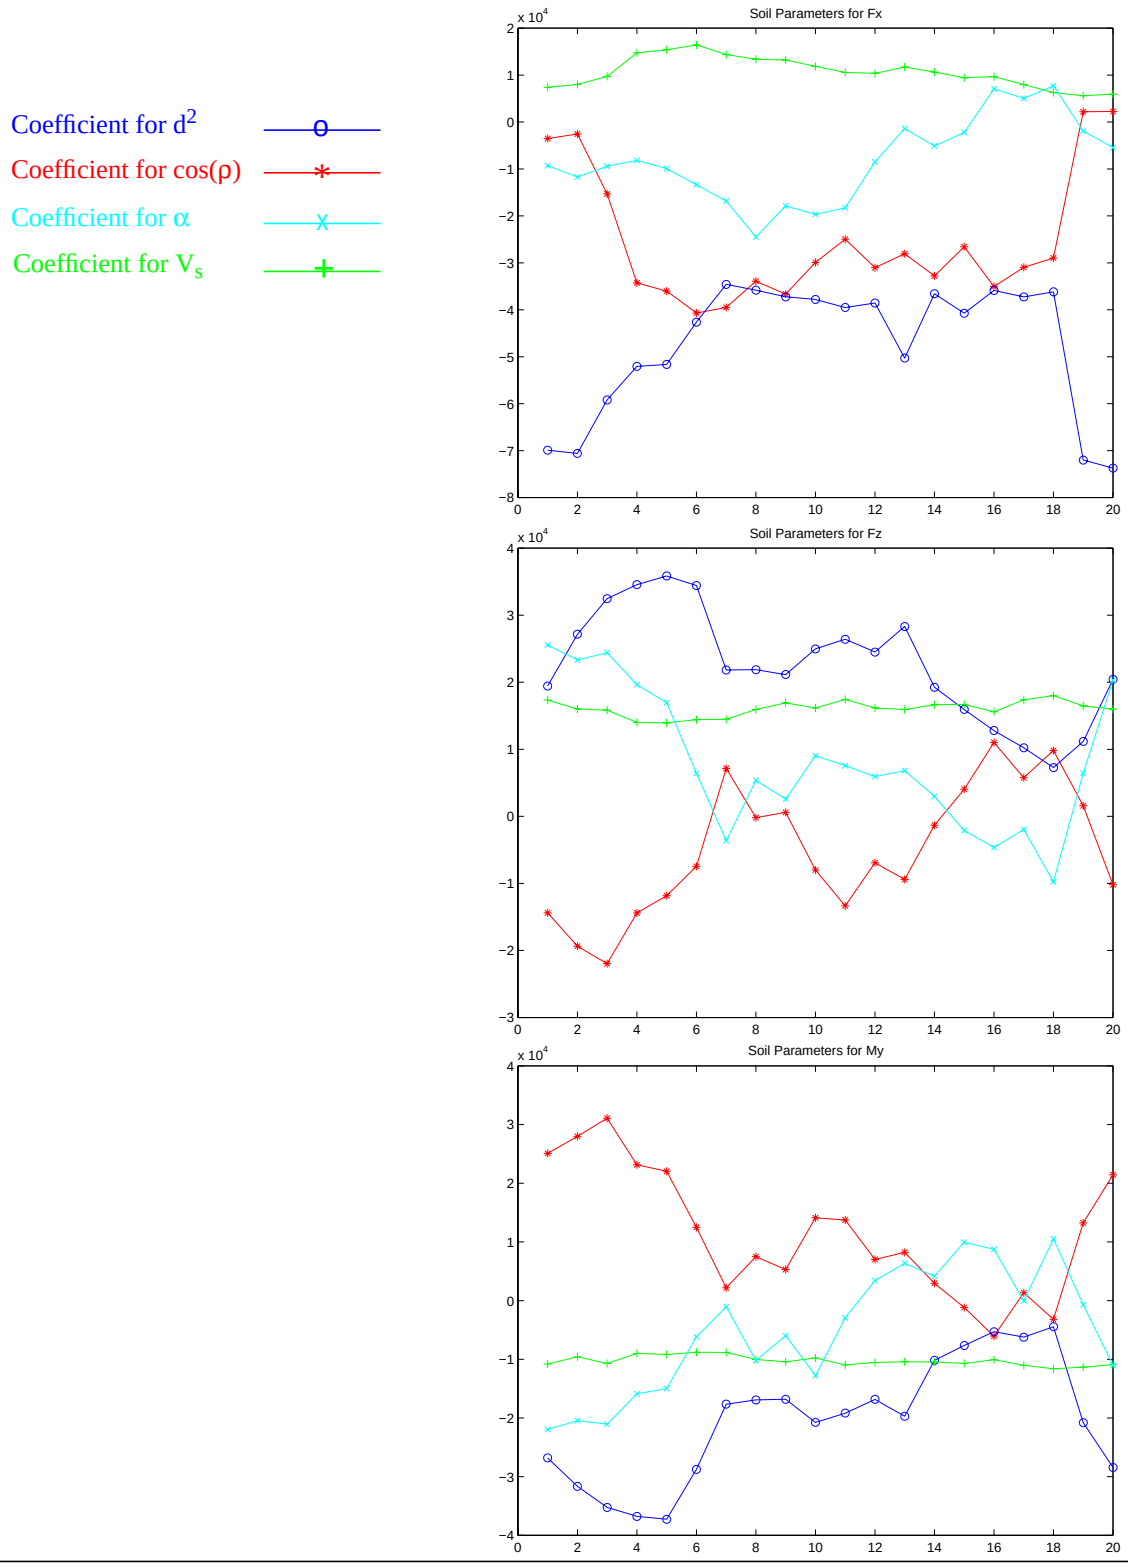

Figure 32: Extracted soil-tool properties for the empirical model for each dig. Each line represents one of the coefficients. Each component of the force requires a separate set of soil-tool properties. Note that the behavior of the properties is fairly erratic.

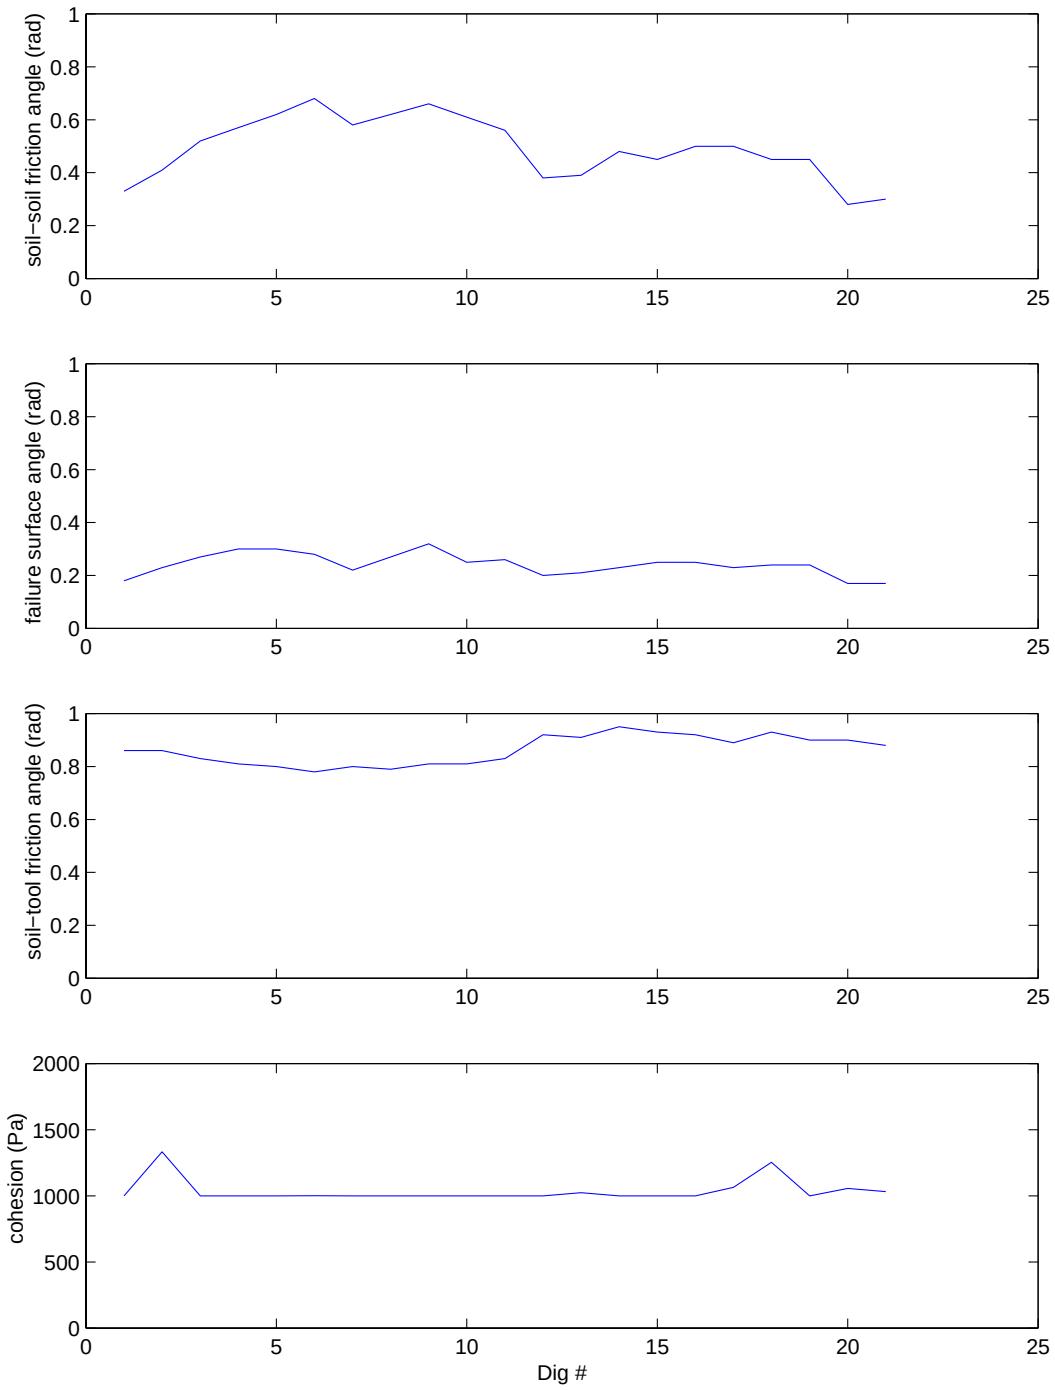

Figure 33: Extracted soil-tool properties for the analytical model for each dig. These parameters are used for all of the force components. The properties are relatively consistent.

4.4 Combined Dig Model Results

This section examines the results of the overall combined dig model. In general we have found that the model predicts the trajectory reasonably well most of the time. In the few cases where it does not, there is usually a good reason. For instance, the soil model assumes that the material is somewhat homogenous. However in our test site the bucket occasionally strikes inclusions such as boulders. Another problem we encountered is when the terrain is not entirely visible with the range sensor due to visual occlusions. To alleviate this problem we limit the distance of the machine from the dig face with the perception planning software. Probably the worst situations occurred however when encountering dust. The range sensors would detect the dust and our software would interpret this as soil. Not only would this throw off the prediction for the current dig, it would also corrupt the database of dig data for soil-tool property extraction.

Figures 34 through 37 show detailed results for a typical dig. In this comparison, the model has been started from the same exact position as the testbed. The model then predicts the trajectory that the bucket tip takes while executing Autodig given the intersection with the terrain map. Thus the model has to predict the commands that are executed by Autodig, the velocities that are achieved by the actuators, the forces that are encountered during digging, and the pressures in the cylinders that result from the forces. All of these internal predictions result in the overall prediction of the dig trajectory. Figure 34 shows how the predicted trajectory compares to the actual trajectory on the testbed.

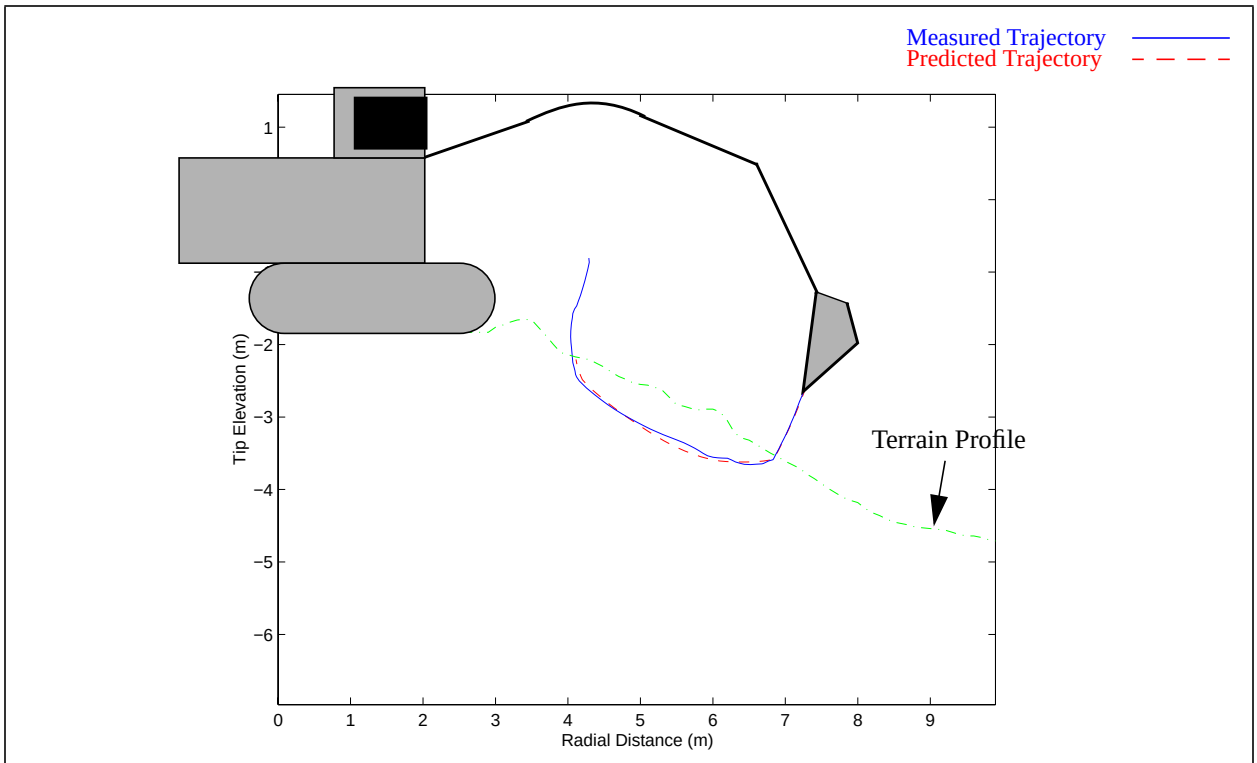

Figure 34: Bucket tip trajectory predicted by the model versus the actual dig trajectory experienced by the excavator testbed. The excavator is shown at the beginning of the dig cycle.

Figure 35 compares the commands generated by Autodig in the model versus the commands generated by Autodig on the testbed. Note that these commands are dependent on the pressures inside the actuators and the positions of the implements. Deviations between the commands generated in the testbed and the commands generated in the model are due to discrepancies in the modeled pressures and discrepancies in the motion of the implements.

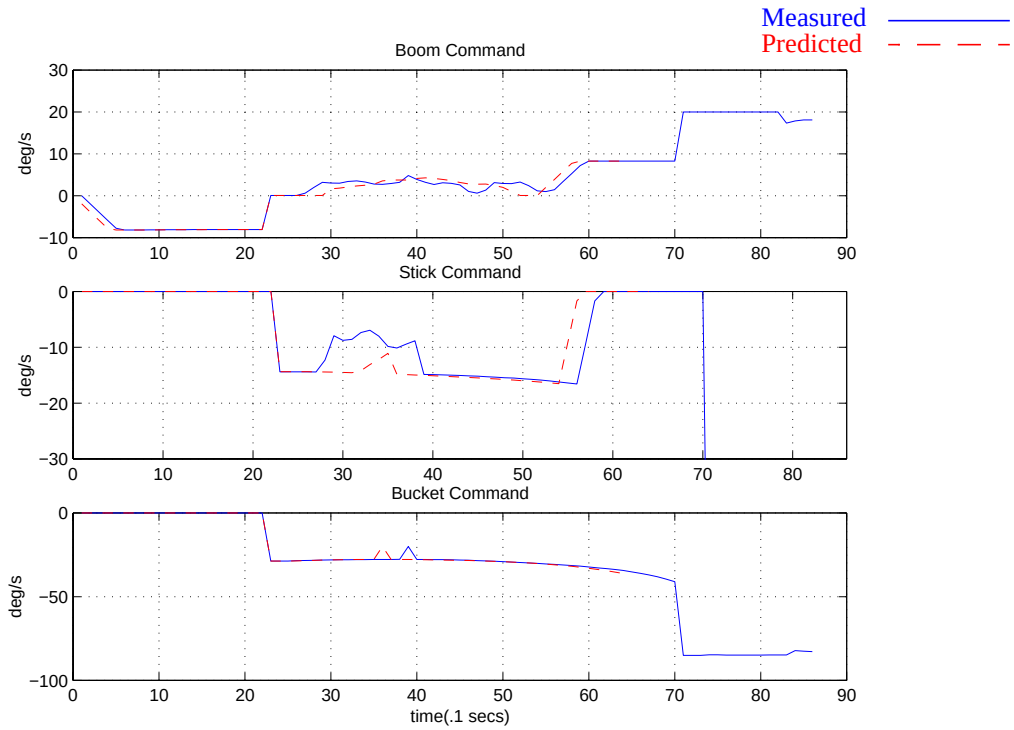

Figure 35: Autodig commands generated internally in the model versus autodig commands on the actual testbed. Autodig commands are based on the pressures in the cylinders and the position of the implements.

Figure 36 shows a comparison of the implement velocities between the model and the testbed, and Figure 37 shows the resulting implement positions. Recall that the actuator velocities in the model are predicted by the neural network, and these velocities are dependent on the predicted pressures and the predicted commands. Additionally, the transformation from actuator velocity to the implement's angular velocity and angular position is dependent on the actuator position.

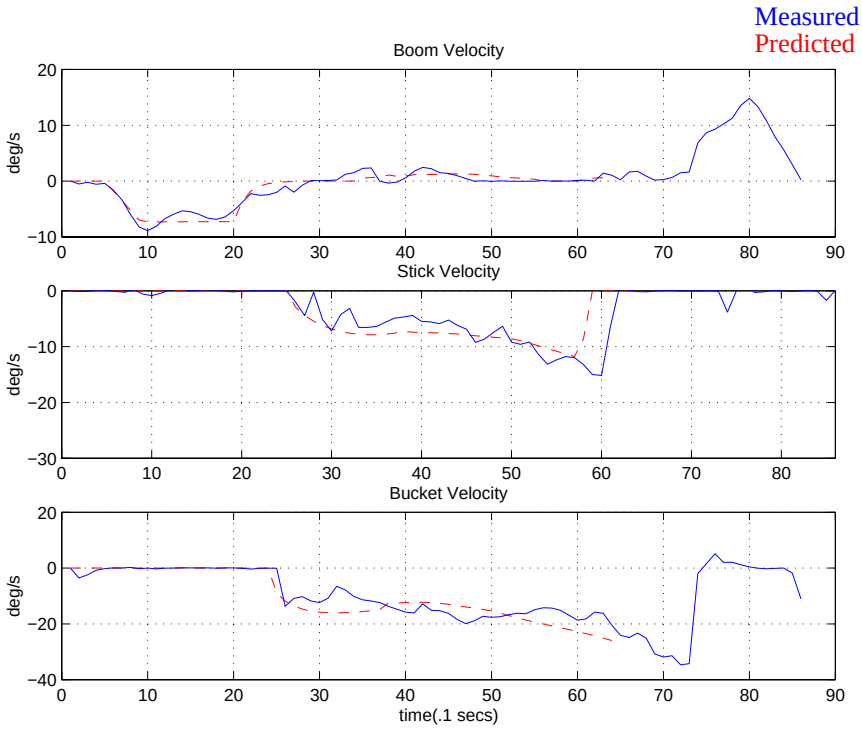

Figure 36: Model prediction of the implement angular velocities versus the implement velocities on the test-bed. The predicted velocities are generated by the vehicle model, and are dependent on the predicted autodig commands, the predicted pressures in the cylinders, and the predicted position of the cylinders.

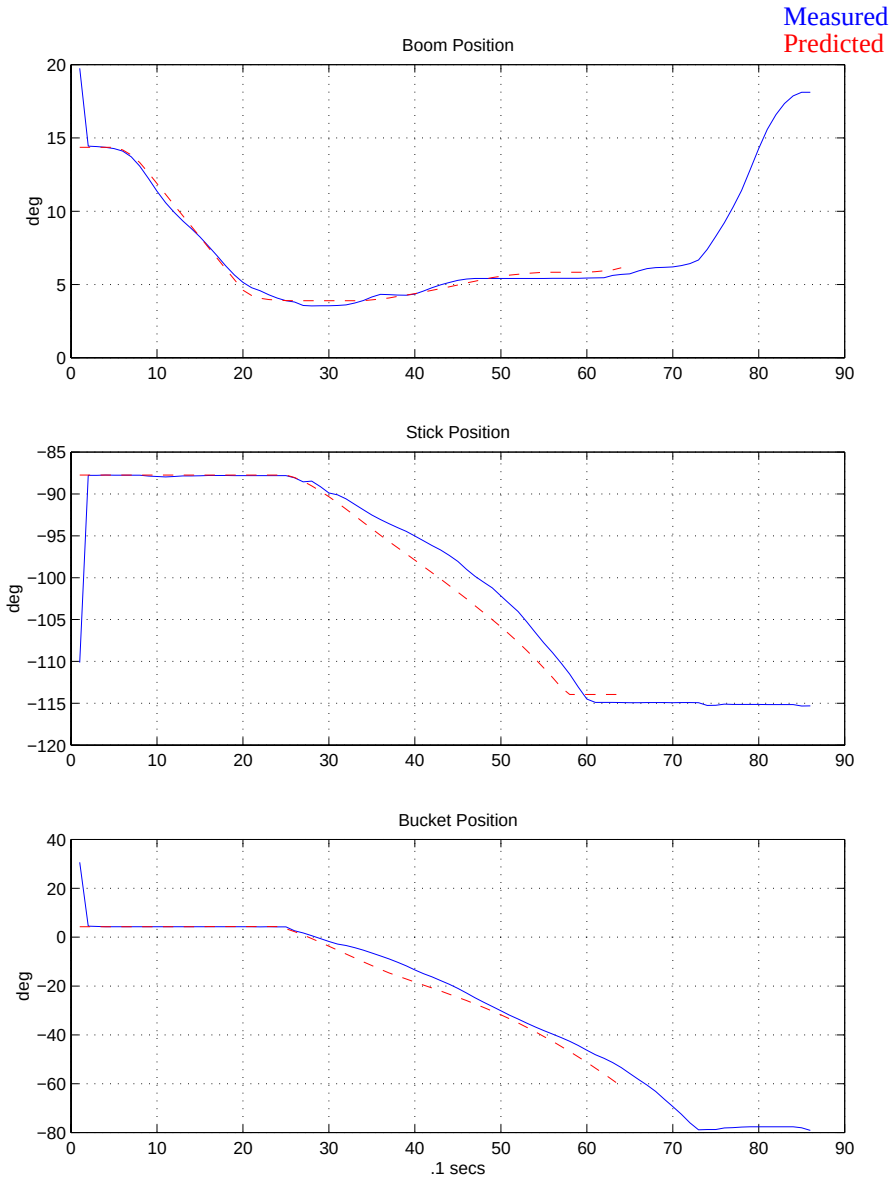

Figure 37: Model prediction of the implement angular position versus implement angular positions on the test-bed. The implement positions are a function of the implement velocities.

For a given time step, once the angular positions of the implements are predicted, then the pose of the bucket can be calculated using forward kinematics. The pose of the bucket will then dictate the intersection with the terrain map, and hence the soil-tool model is able to estimate the soil forces acting on the bucket. These forces are then transformed into pressures in the hydraulic actuators, which are then used to calculate the next Autodig command and the next implement position. Figure 38 shows the predicted pressures versus actual pressures.

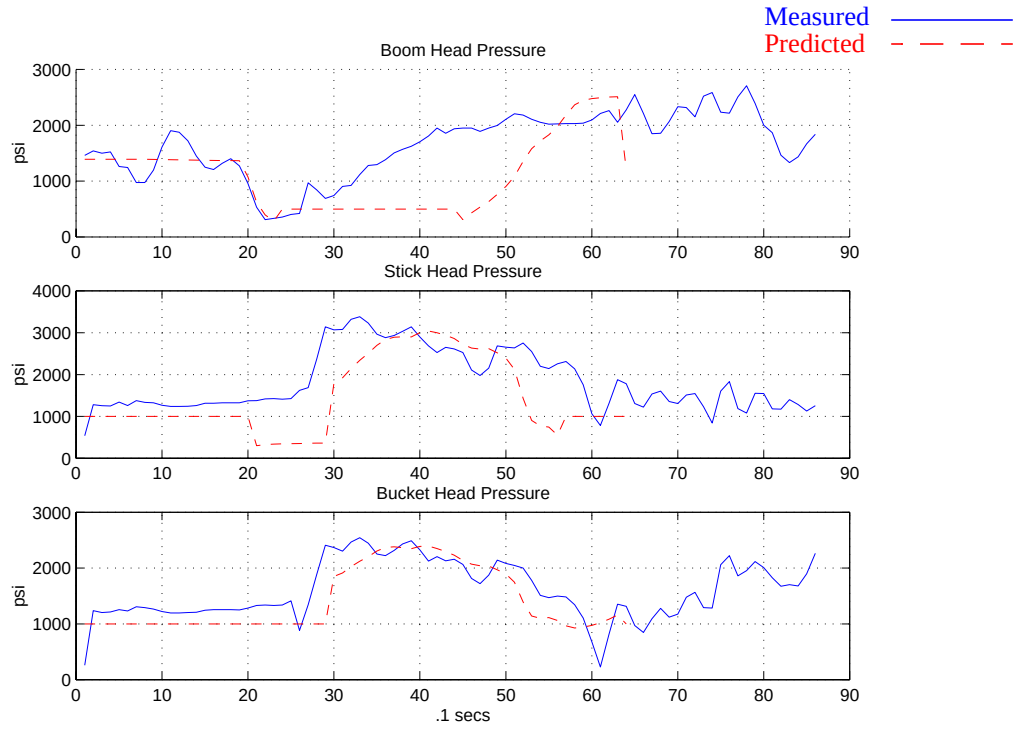

Figure 38: Modeled cylinder pressures versus actual cylinder pressures. The pressures are based on the forces acting on the bucket, which are calculated by knowing the position of the bucket and its intersection with the terrain.

From this discussion, it is clear that the model results depend largely on the ability to predict the digging forces. To illustrate the importance of the forces on the digging trajectory, the dig model was executed from the same position with three levels of soil hardness. The soil hardness was modified by multiplying the output of the analytical soil model by a factor of 2 and 3. Figure 39 shows the trajectories predicted by the model for each force level. Table 2 shows the resulting change in digging statistics. As would be expected, the time required to dig increased when changing the factor from 1 to 2. This is due to Autodig slowing down the bucket velocity as the force increased. This also caused the bucket to penetrate deeper increasing the swept volume. Increasing the force factor from 2 to 3 however decreased both the volume and the time. This is because the bucket was not able to penetrate as deeply due to the hardness of the ground, and less time was required to capture the material.

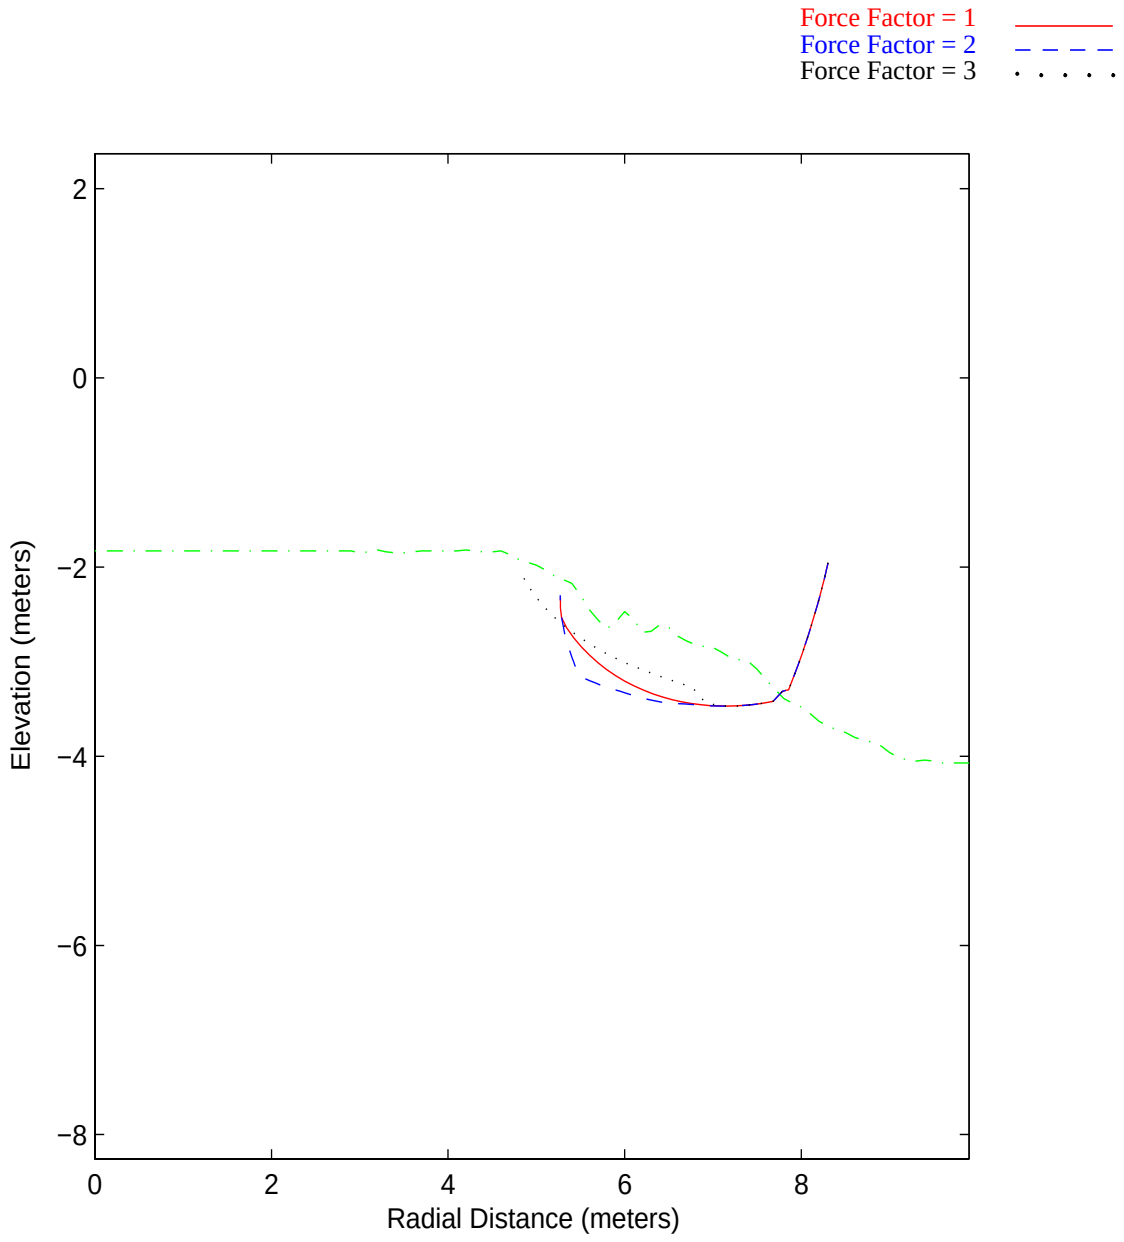

Figure 39: Predicted dig trajectories for three simulated force levels. Each simulated dig uses the same initial conditions and the same terrain map, however the output of the analytical soil model was multiplied by the force factor. This caused the dig model to produce three distinctly different trajectories.

Table 2: Comparison of simulated digging statistics for three force levels. For each simulated dig, the force output of the analytical model was multiplied by the force factor. This resulted in the following digging statistics.

Force Factor	Time (s)	Volume (m^3)	Energy (kJ)
1.0	5.5	2.2	113
2.0	8.1	2.5	234
3.0	7.8	2.1	253

Figure 40 shows a qualitative comparison between the predicted and measured trajectories for eight separate digs. As can be seen from the plots, the dig model appears to predict the trajectory reasonably well for a number of different terrain profiles and starting conditions. It should be noted here that we were unable to test the dig model in a large variety of soil conditions due to time constraints. Most of the material in these tests would be considered “easy digging”.

As previously indicated, the purpose of the model is to allow an evaluation of a number of candidate trajectories for selecting where to dig. As will be discussed in Chapter 5, the evaluation will be based on the time required to dig, the energy expended during digging, and the volume of material swept by the bucket. Figure 41 shows a comparison of these digging statistics compared to measured data. The model is able to match the measurements to within 12% to 15%. There appears to be a small systematic bias in that the volume is consistently underestimated. This bias is probably due to range sensor calibration errors, since the actual volume measurements seem excessive.

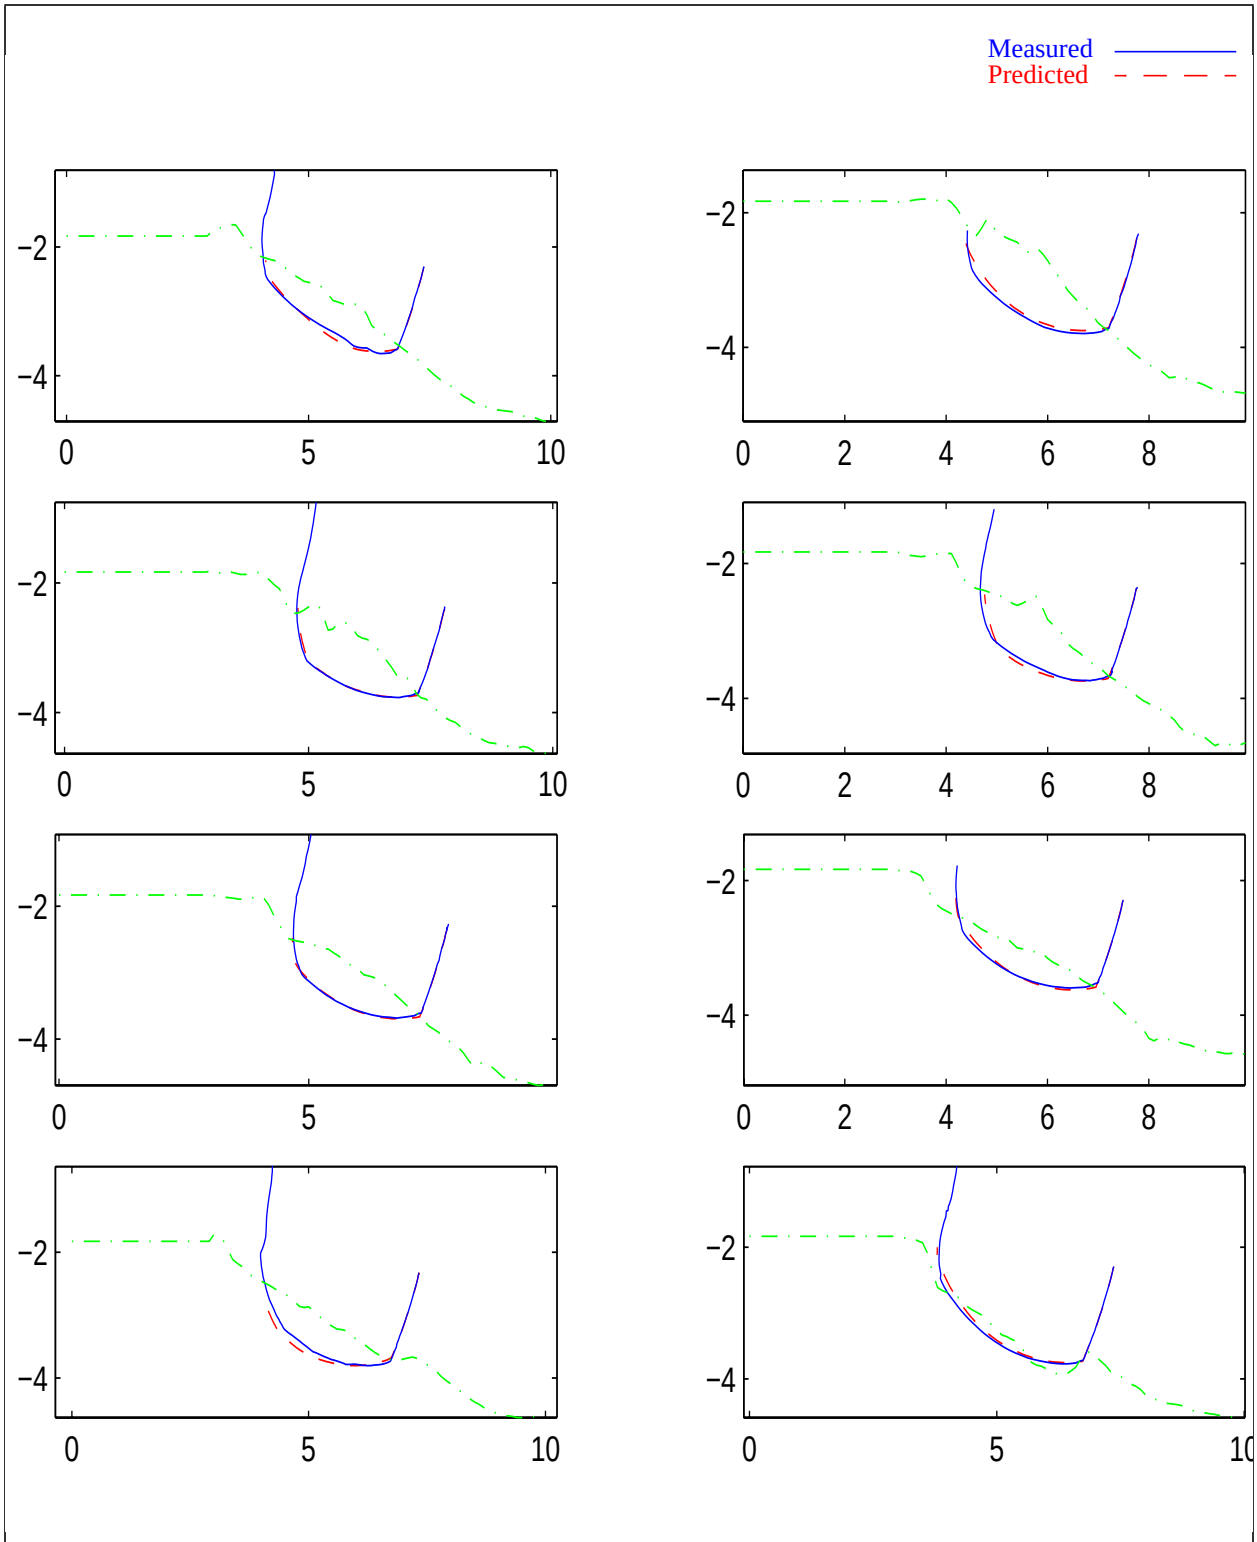

Figure 40: Comparison of modeled versus actual dig trajectories for 8 digs. The model predicts the trajectory reasonably well for an number of different terrain profiles and starting conditions.

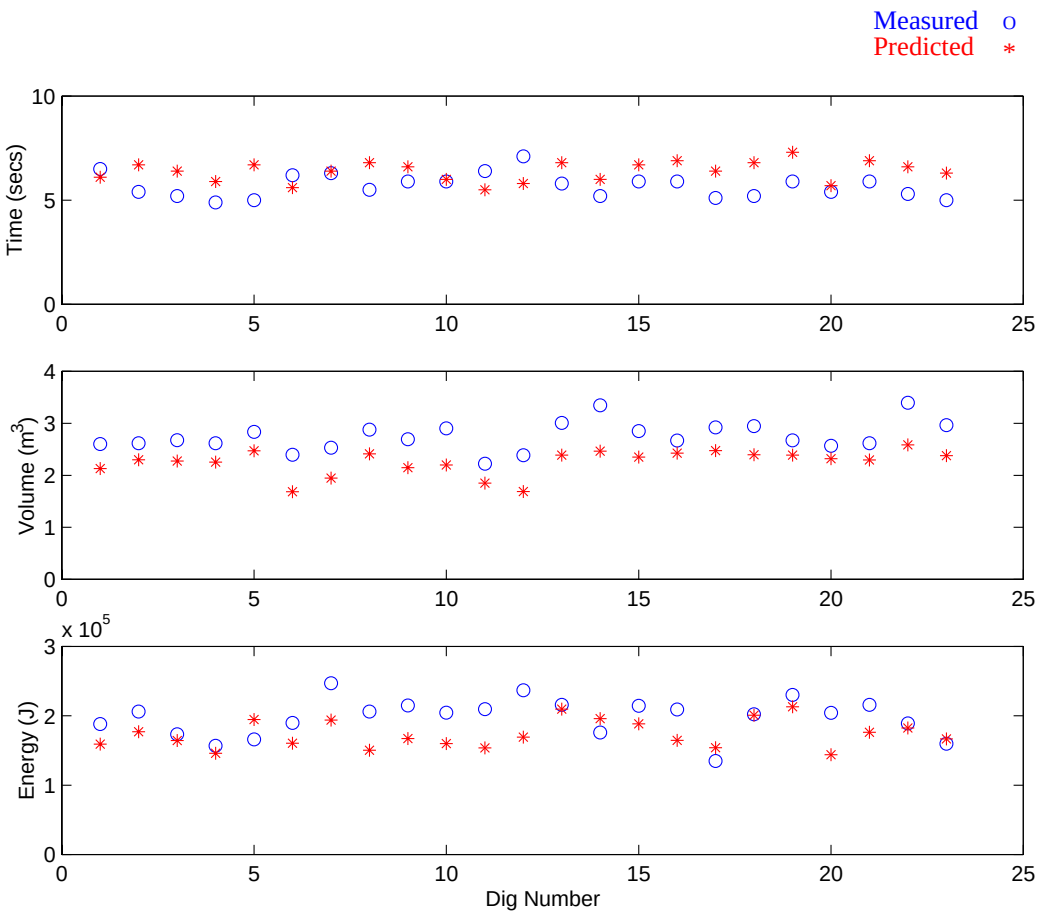

Figure 41: Comparison of digging statistics produced by the feed forward model to the statistics computed from actual test data. Mean errors in estimating time, volume and energy are 13.7%, 14.7% and 12.5%.

Chapter 5 Planning Digging Operations

The effectiveness of the truck loading process is dictated by many factors, such as the initial site preparation, the efficiency of the movements between the dig face and the truck, and the effectiveness of the dumping and the digging portions of the loading cycle. The focus of the work discussed in this thesis is to ensure that the digging portion of the cycle allows the machine to remain productive not only over a few cycles, but over an extended sequence of operations. In order to do this, a dig planning system is required to "manage" the excavation of the bench. That is, it needs to choose digging actions that fill the bucket quickly and efficiently. It must ensure that there is always enough material available in front of the machine so that a full bucket is obtainable. Yet it needs to remove the material completely so that piles of material are not left behind.

These requirements have led to the development of a dig planning system based on the following strategy:

- Choose dig locations that optimally fill the bucket.
- Clean the floor after the majority of the material has been removed.
- Track backwards when needed so that additional material may be excavated.

Section 5.1 of this chapter discusses the perception based planning system that was designed to fulfill this strategy.

In addition, the soil characteristics can have a major impact on how well the system performs. Chapter 3 discussed how the Autodig algorithm required the use of a soil hardness index that allowed it to adapt to various soil conditions. Section 5.2 discusses a means for automatically selecting this index based on the forces encountered during digging.

Finally, Section 5.3 discusses the software implementation of the entire excavation system.

5.1 Perception Based Dig Planning

In keeping with the strategy outlined above, the perception based dig planning algorithm was divided into three separate planners:

- Dig Location Planner - This planner is responsible for choosing the initial pose of the bucket prior to executing Autodig. The selected dig is based on being able to fill the bucket with the least amount of energy and in the smallest amount of time, without violating any boundary constraints such as digging below a desired floor elevation. This is how the majority of the material is removed from a given machine location.

- Cleanup Planner - The cleanup planner is responsible for dictating where the cleanup operation should begin based on the amount of material, and the distance of the material away from the machine. In Chapter 3 we discussed the use of a cleanup operation that slowly moved the bucket along a desired floor elevation to leave a relatively even and level floor. This planner selects where these cleanup actions should be executed.

- Tracking Planner - The tracking planner dictates how far the machine can be moved back during each planning cycle. The distance is based on being able to reach all of the material, and being able to see all of the material with the range sensors. The second condition is generally more confining because the angle of the repose of the bench can be quite steep. If the machine, and hence the range sensor is moved back too far, then the front of the bench can be occluded from view by the top of the bench.

All three plans are generated prior to each loading cycle. Thus a higher level process has the opportunity to use any one of the plans if it is privy to some information that should override the preference due to digging alone. For instance, if the truck has just been filled, and another truck is moving into place, it may be a good opportunity to track the machine backwards, or execute the more time intensive cleaning operation.

In general however, the dig planner signals its preference for which operation should be executed by providing two quantities. The first is the quality of the selected digging operation. This quantity will be described in more detail in Section 5.1.1. The second quantity is simply the amount the machine can be tracked backwards. Figure 42 illustrates a simple algorithm to decide which action should be taken.

```
if dig quality > minimum quality
    then execute a dig
else
    if tracking distance > minimum tracking distance
        then track the machine backwards
    else
        execute a cleanup
```

Figure 42: Simple algorithm used to select which action to execute based on the dig quality and minimum tracking distance provided by the dig planner.

The above algorithm shows a bias towards removing as much material as possible with the digging operation, and avoiding the time consuming cleanup operations. The machine attempts to dig if a good dig is possible. If not, it will try to back up so that more material is available. The motion of the machine backwards however is limited by the distance of the farthest material away from the machine. If the machine cannot be moved back some minimum distance, then a cleanup operation is executed. The cleanup operation removes the farthest material, potentially allowing the machine to track backwards on the next iteration.

The following sections discuss the three components of the perception based dig planner in more detail.

5.1.1 Dig Location Planner

The Dig Location Planner is used to decide where the next digging operation should take place, and give some indication as to the overall utility or quality of this digging operation. Its design is inspired by what we believe is the decision making process of a human when selecting a dig. A human operator generally has some overall strategy in mind for removing the material from the face, and then decides exactly where and how to dig based on the stiffness of the material, and the performance of the machine during the digging process.

Likewise, we have designed the Dig Location Planner as a multi-resolution planning system. A coarse planner is used to tessellate the dig face into smaller regions, and a region is selected based on some overall material removal strategy. Then the region is provided to a refined planner which searches within the region's limits for the locally optimal bucket pose with which to initiate Auto-dig. It does this by querying a dig evaluation routine for the quality of numerous candidate bucket poses. The dig evaluator bases the evaluation on the results of the feed forward dig model described in Chapter 4. Figure 43 illustrates the process for planning dig locations.

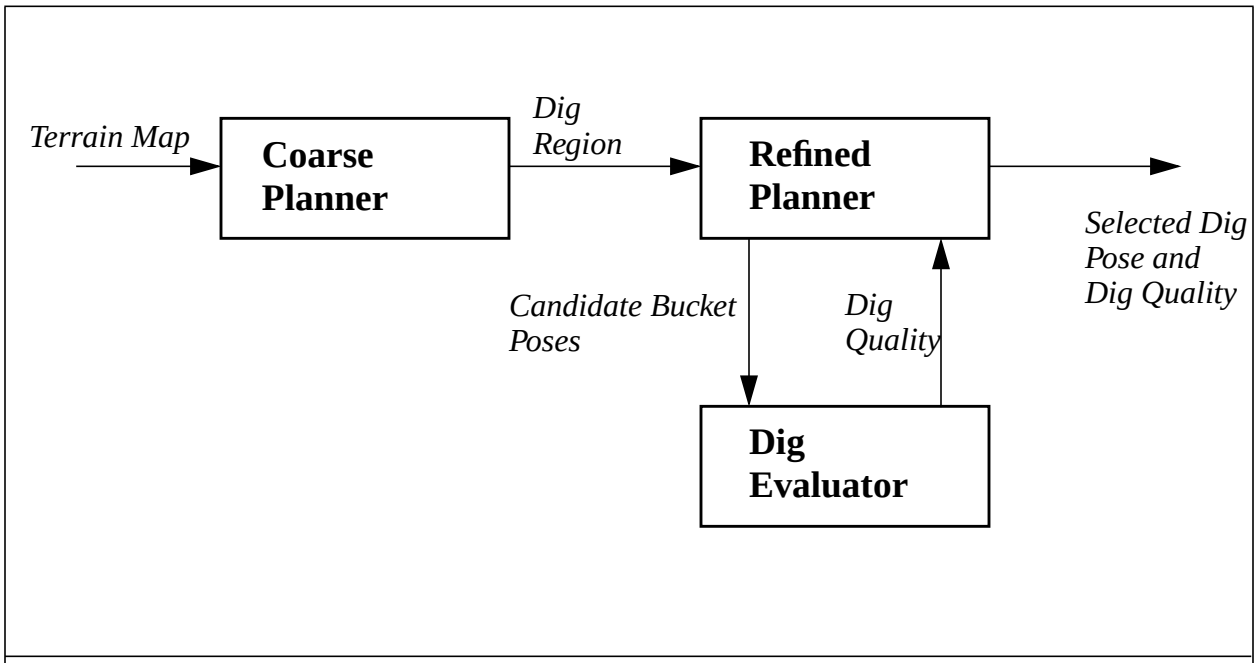

Figure 43: Coarse to refined dig location planning strategy. The coarse planner divides the workspace into dig regions. The refined planner searches the dig region for the best dig by evaluating each candidate. The dig evaluation is based on the forward model of the digging process.

5.1.1.1 Coarse Dig Location Planning

The purpose of the coarse planner is to divide the dig face into regions, and then for each loading cycle, select one of the regions in which to dig. In doing this, the coarse planner is generating an overall strategy in which to erode the bench, although it is not specifying exactly where the dig should take place. Figure 44 shows a coarse plan for a hypothetical bench configuration. The strategy shown here, and described below is based on the recommendations of an expert machine operator.

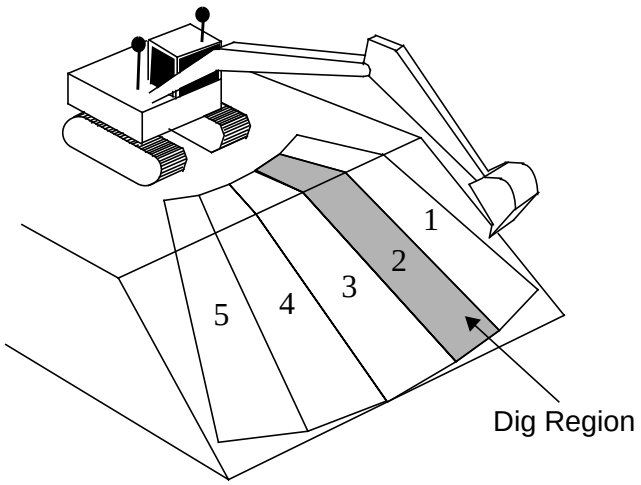

Figure 44: Coarse plan for typical bench configuration. The strategy encoded is based on recommendations from an expert machine operator.

First, the coarse planner needs to identify the boundaries of the workspace. The far radial boundary is defined by either the edge of the bench or the kinematic limit of the machine's reach, depending on which is closest to the machine's center. The edge of the bench is based on the terrain elevation exceeding some threshold above the desired floor level. If the edge of the bench is closer to the machine center than the kinematic limit, then the bench edge becomes the far radial boundary because digging should not penetrate the floor. Otherwise the kinematic limit determines the boundary because there is no reason to search outside of the machine's valid workspace.The angular extents of the boundaries are dictated by the stability of the machine. Typically the machine should dig within +/-30 degrees from parallel with the tracks in order to ensure that the machine is stable.

The workspace is now divided into angular regions, where the width of each region is approximately one bucket width wide, and are adjacent to each other at the far radial boundary while overlapping at the near radial boundary. The regions are ordered based on the location of the truck. If the truck is parked on the excavator's left, the digging is ordered from left to right. This is so the boom doesn't have to raise as high to clear the material when swinging to the truck. Digging the regions in order ensures that the bench is evenly eroded. This minimizes the pockets of material that are left behind which may not be removed efficiently.

We also experimented with dividing the regions radially into two and even three dig regions along one angular orientation. The reason behind this was that expert operators typically dig at the top of the bench first, and then reach towards the bottom. Digging is executed in this manner because it reduces the visual occlusions, and also reduces the force needed to lift the lower material. However we found that by forcing the dig to take place at an upper region first, we were artificially limiting the refined planner from obtaining a better dig that might have been in the lower region.

Instead, we found that it is best to allow the refined planner to search the entire radial distance for the best dig. Interestingly enough, the refined planner usually would select digs closer to the machine first, and then reach out farther as digging progressed.

5.1.1.2 Refined Dig Location Planner

The purpose of the refined planner is to search a selected dig region for the optimal bucket pose with which to begin digging. Again, only the start pose of the bucket need be specified since Autodig will guide the bucket through the rest of the digging process. The first task then is to discretize the region into a number of bucket poses. The initial pose of the bucket can be defined by two parameters as shown in Figure 45. The dimension *d* defines the distance that the bucket tip is located from the region's near radial boundary. The dimension α defines the angle of the bucket tip relative to the normal of the local terrain profile.

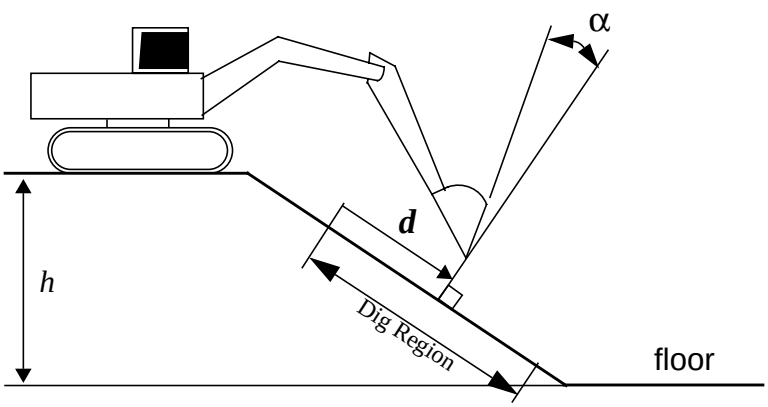

Figure 45: The dig pose is specified with two parameters *d* and α. The resolution of *d* that was used in our experiments was approximately 20 cm, and the resolution of α was 5 degrees.

An entire region is divided into a number of candidate digs, where each dig is specified by the candidate pair (d,α). Each candidate pair then is evaluated to determine the utility of the selected dig. This is a two part process. First the candidate pair is checked to ensure that it does not violate any geometric constraints. If not, the pair is evaluated in regards to an objective function. The refined planner selects the dig that results in the highest value of the objective function.

The first geometric constraint that the pose must satisfy is the *kinematic reachability constraint*. If the candidate pair is outside the machine's valid workspace, then it is rejected without need for further evaluation. Otherwise the bucket pose is sent to a dig evaluation routine as shown in Figure 43. The dig evaluation routine uses the feed forward dig model described in Chapter 4 to sim-

ulate the effect of selecting the candidate bucket pose. The dig model is initiated at the candidate bucket pose, and run forward in time until it completes the simulated dig. If the simulated dig ever penetrates the desired floor elevation, then it has violated the *geometric shape constraint* and is rejected without need for further evaluation.

During the simulation, the dig model calculates the digging statistics of time required to dig, energy expended during digging, and the volume of material swept into the bucket. These quantities are used to determine the value of the objective function which we term the *dig quality.*

The dig quality is a measure of the desirability of executing a particular dig. The dig quality is calculated by combining the three functions ξ, ψ, and ζ shown in Figure 46. These functions are dependent on swept volume, the time, and the energy required for digging respectively. When the volume is less than V_{min}, then ξ=0. As the volume increases, ξ increases linearly until the volume reaches V_{max} which corresponds to the capacity of the bucket. This function essentially gives higher values to increased bucket volumes up to the maximum bucket capacity. Similarly, we would like to execute digs that use the least amount of time and energy. Therefore ψ and ζ are both maximum at zero, and decrease linearly as the time and energy increase.

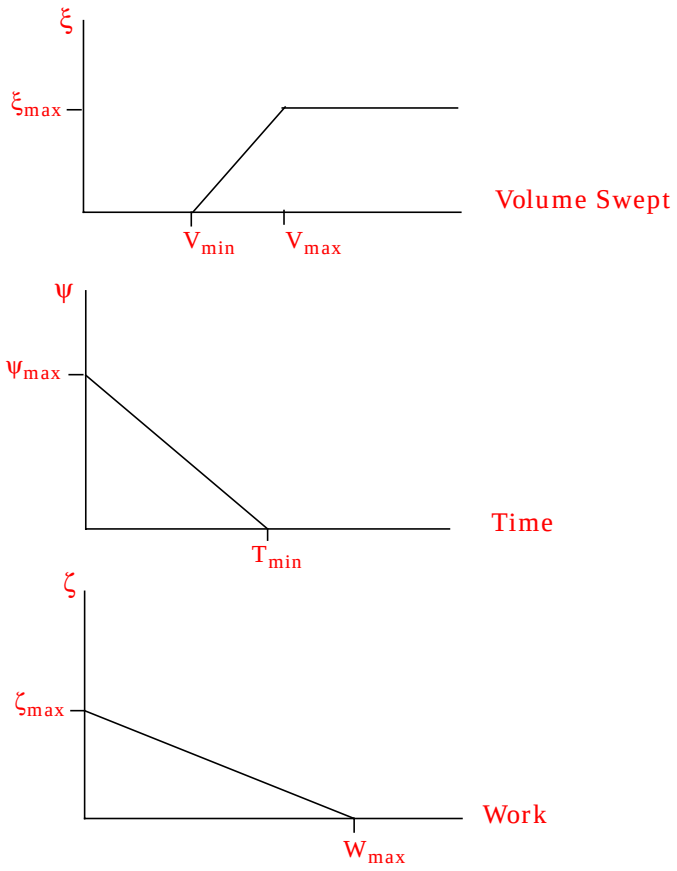

Figure 46: Evaluation functions for selecting a dig. V_{max} corresponds to the bucket capacity. The values of V_{min}, T_{min}, and T_{max} are based on the range of acceptable criteria for a digging operation. ξ_{max}, ψ_{max}, and ζ_{max} are tuned to get a desirable result. The overall dig quality is a linear combination of these functions.

The dig quality is then calculated by the equation:

$$Quality = w_v\xi + w_t\psi + w_e\zeta$$

where w_v, w_t, and w_e, are weightings based on the relative importance of each of the factors. For our experiments we specified w_v=3, w_t=2, and w_e=1. Additionally, in order to allow any of the factors to veto a particular dig, the quality is set to zero if any of the factors are zero.[1]

The refined planner conducts an exhaustive search to determine which candidate pair has the best overall quality. Figures 47 and 48 illustrate the evaluation process for two separate terrain topologies. The terrain in Figure 47 is fairly regular, whereas the terrain in Figure 48 is undulating due to previous digs.

The plots at the top left of the figures show the terrain shape, the floor elevation, the boundaries of the coarse plan, and the trajectory of the dig that was selected. The plot at the top right shows the combinations of d and α with regard to the geometric constraints. Combinations that are colored in blue correspond to digs that are reachable but would violate the floor boundary. Configurations that are reachable and that do not penetrate the floor are colored in red. The valid configurations are evaluated in terms of the dig quality which is illustrated in the remaining plots.

In Figure 47, there are no shape violations for the terrain since the coarse plan prevents the machine from penetrating the floor. In general, as the value of d increases, the volume, time, and energy also increase. Clearly the dominant term in the dig quality is the volume. Note that once the swept volume surpasses V_{max}= 2.0 m^3, the evaluation function levels off. The selected dig then is dictated mainly by the time required to dig. This is further illustrated by the fact that the maximum swept volume for all of the candidates is 2.7 m^3, however the volume of the selected dig only has 2.1 m^3. In Figure 48, there are several shape violations for large values of d. Also note how the evaluation function has become more complicated because of the undulating terrain.

1. In [Singh 98] we describe using a quality function which is a multiple of the three factors. We found however that it was too difficult to adjust the factors to obtain a desired result.

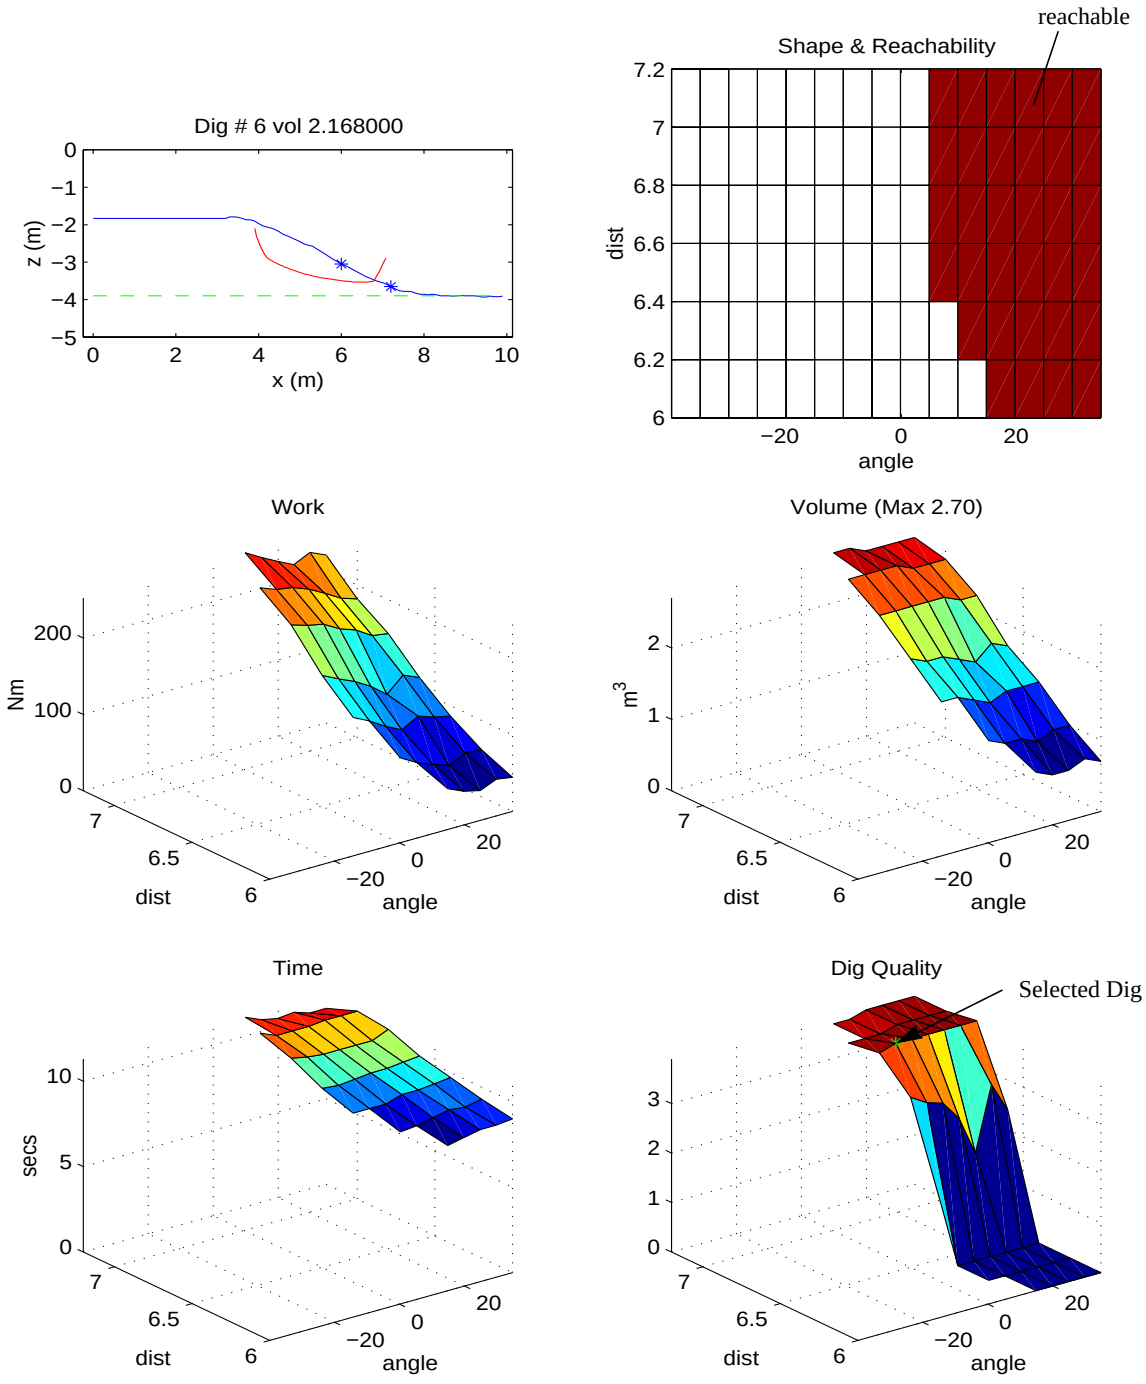

Figure 47: Plots showing the evaluation of a relatively even terrain. The plot at the top left shows the terrain, the floor elevation, boundaries of the coarse plan indicated by *'s, and the selected dig. The plot at the top right shows the combinations of (d,α) with regard to the geometric constraints. Elements in blue are reachable but violate the shape constraint. There are no shape violations for this terrain. Elements in red do not violate either constraint. The rest of the plots show the energy, volume, time, and the combination of these values into the overall evaluation (dig quality).

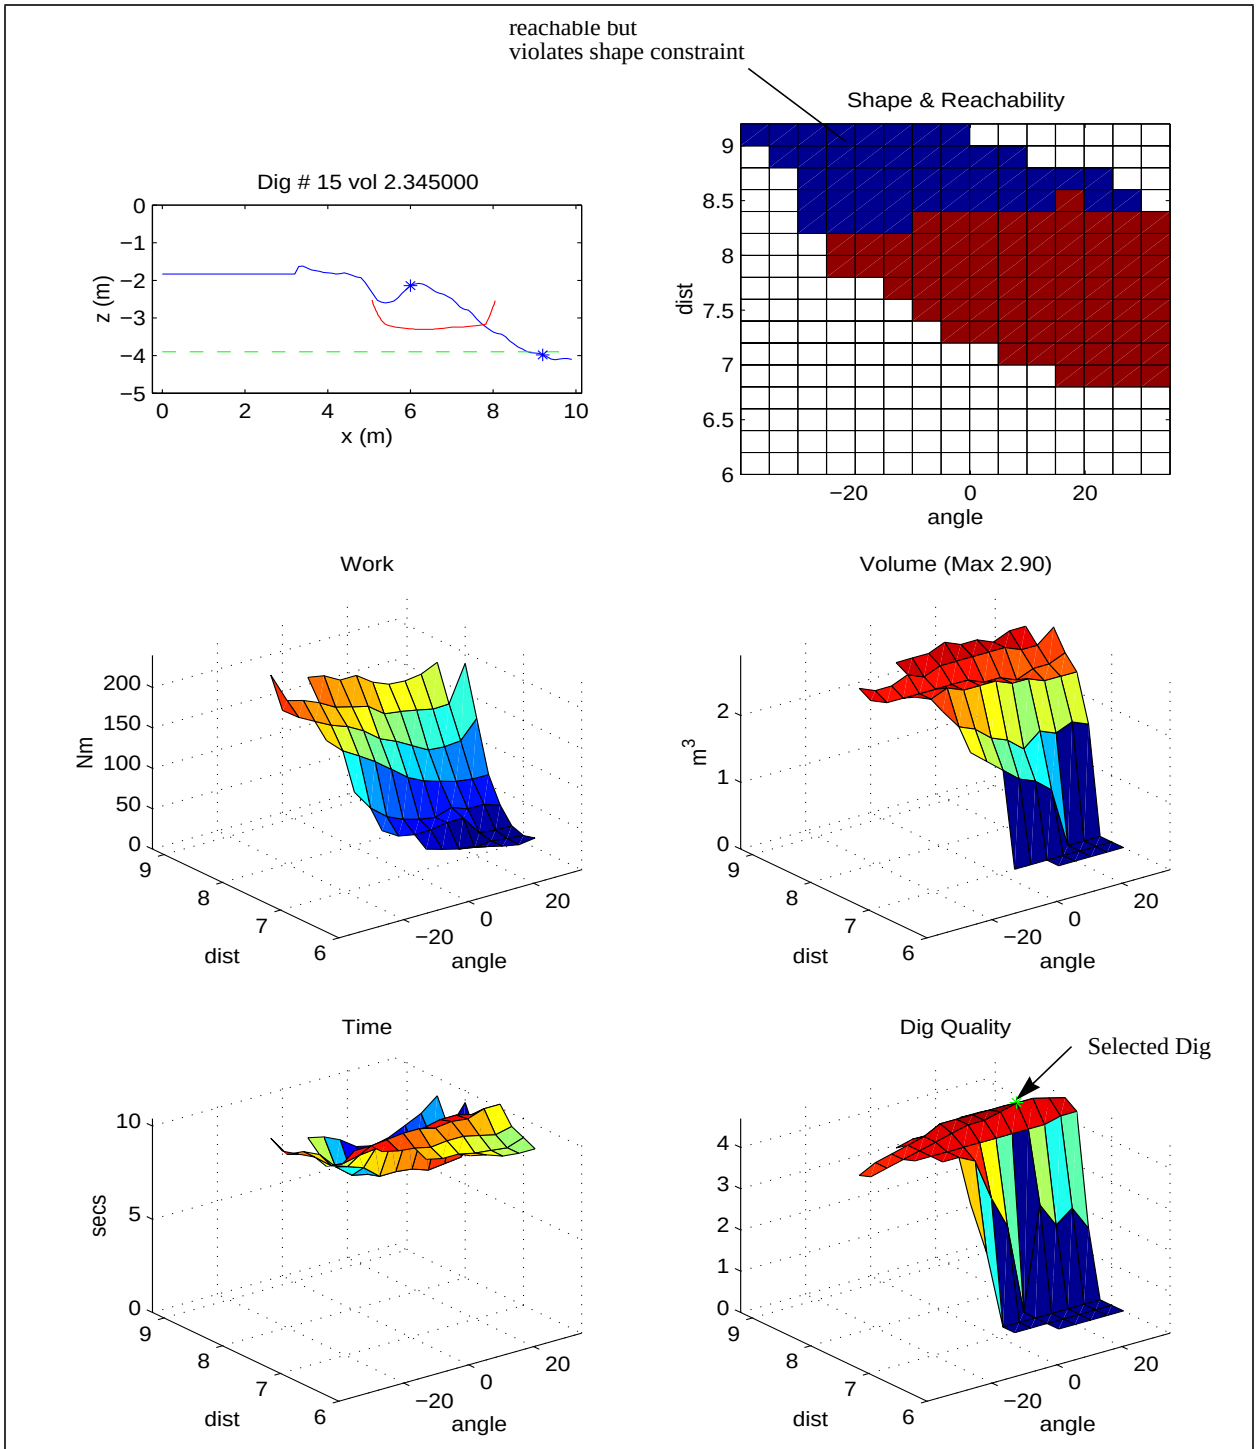

Figure 48: Plots showing the evaluation of an undulating terrain. The blue elements in the top right plot show which combinations of *d* and α are reachable but violate the shape constraint. Only the red elements which do not violate either shape or reachability constraints are evaluated further. The evaluation function is more complicated for this topology.

5.1.2 Cleaning and Tracking Planners

The purpose of the cleanup planner is to decide where it is most desirable to execute a cleanup operation, and the purpose of the tracking planner is to determine how far the machine can be moved backwards. Both of these functions are dependent on the distance to the farthest material. To reduce the amount of time needed to clean the floor, the cleanup planner will extend just beyond this material. The tracking planner determines how far the machine can be moved back and still be able to reach the material and see the material with the sensors.

The first step in generating either plan is to examine the terrain to look for piles of material that need to be cleaned up. To do this, the terrain elevation is examined along a number of swing angles to look for material that exceeds a threshold in elevation. Starting from the far reach of the machine and moving radially inward, when the average terrain elevation exceeds the threshold elevation, then this corresponds to the maximum distance for a given swing angle. These maximum distances are the only terrain dependent features needed to specify both plans. Increasing the terrain elevation threshold above the floor causes the machine to spend less time cleaning. This improves productivity, however it causes the resulting floor to be rougher. This value needs to be tuned based on the overriding objectives.

The operation of the planners is illustrated in Figure 49. The cleanup plan along each swing angle extends just beyond the material boundary as dictated by the elevation threshold. Some cleanup plans are eliminated from consideration. First, if the length of a cleanup plan along a particular swing angle is too short to be worth executing, then this cleanup plan is eliminated. Second, when a cleanup has been executed along one swing angle numerous times, and still it’s limiting the machine’s backwards movement, then this region is considered “not-cleanable”, and is eliminated. Finally when excavation creates a wall as shown in Figure 49, the region adjacent to the wall is not cleaned because the wall must be allowed to settle and form a natural angle of repose. The cleanup planner then selects the region that has the longest length.

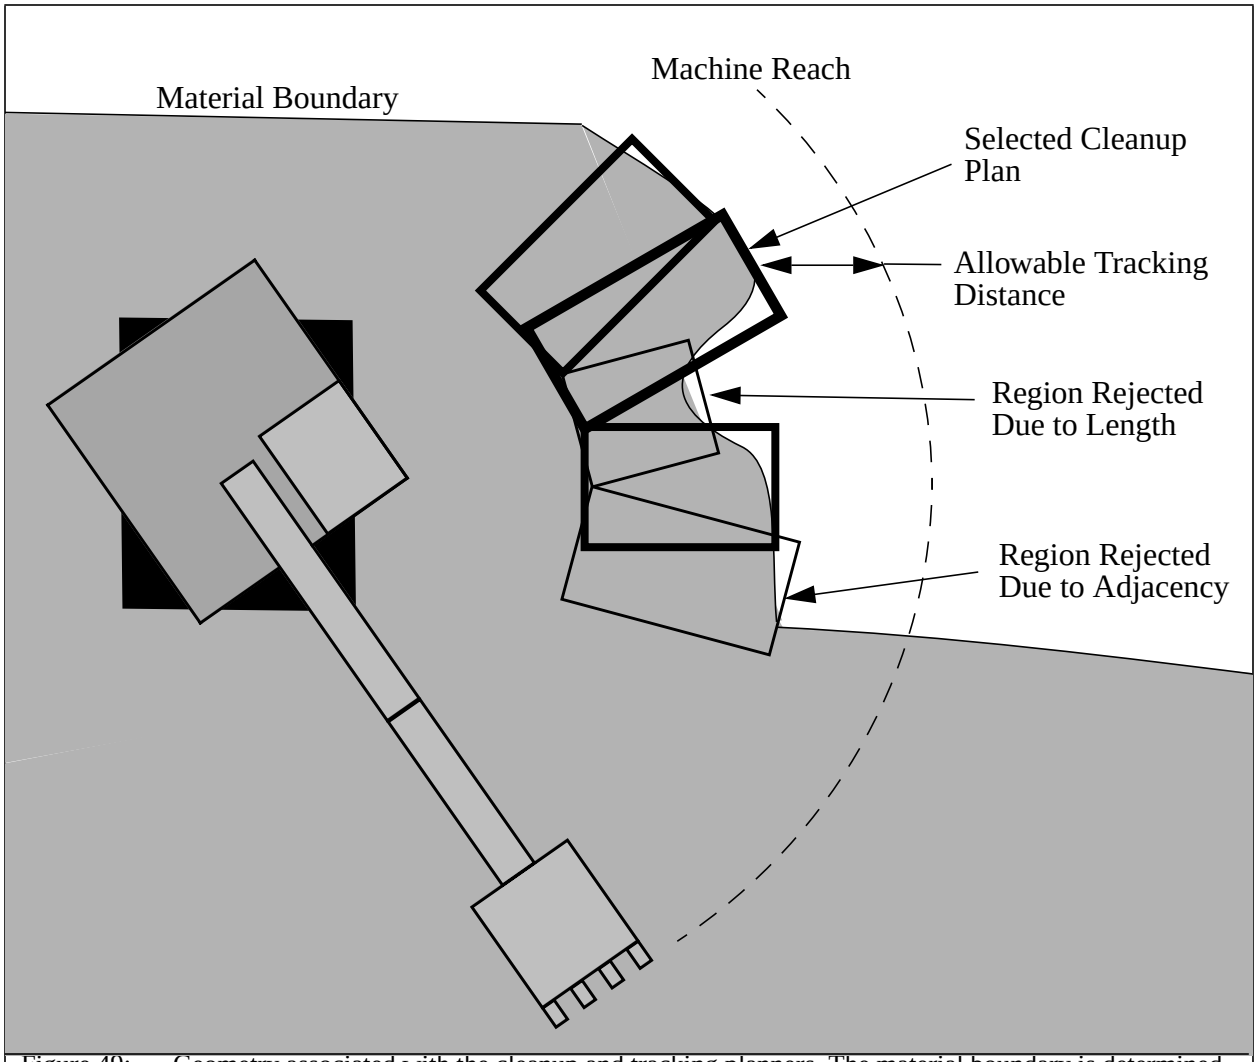

Figure 49: Geometry associated with the cleanup and tracking planners. The material boundary is determined by searching for terrain elevations that exceed a threshold. The cleanup plans extend just beyond the material boundary. The tracking distance is limited so that all of the cleanup plans can be reached. Some cleanup plans are eliminated and are not used in calculating the tracking distance.

The allowable tracking distance is the shortest horizontal distance between the machine's reach and the farthest cleanup plan as illustrated in Figure 49. This is the maximum distance that the machine can be moved back and still be able to reach all of the material. In addition, a further constraint is added that reduces the occlusion of the dig face from the view of the sensor. Figure 50 illustrates how the top of the bench can occlude the rest of the bench from view if the machine is moved back too far. To reduce this problem, a safe tracking distance is calculated based on the material laying at some maximum angle of repose. The repose angle is assumed to start at the far edge of the material, and is extended to the height of the sensor. The maximum tracking distance is the horizontal distance between the sensor and this projected line. During our experiments we found that the use of a maximum repose angle of 45 degrees generally ensured that most of the

material would remain in view. Of the two constraints, machine reach and the material remaining in view, the material remaining in view generally was the more restrictive and limited the distance that the machine could be moved back. Figure 51 shows the trigonometry that is used to calculate the backwards tracking distance for both constraints.

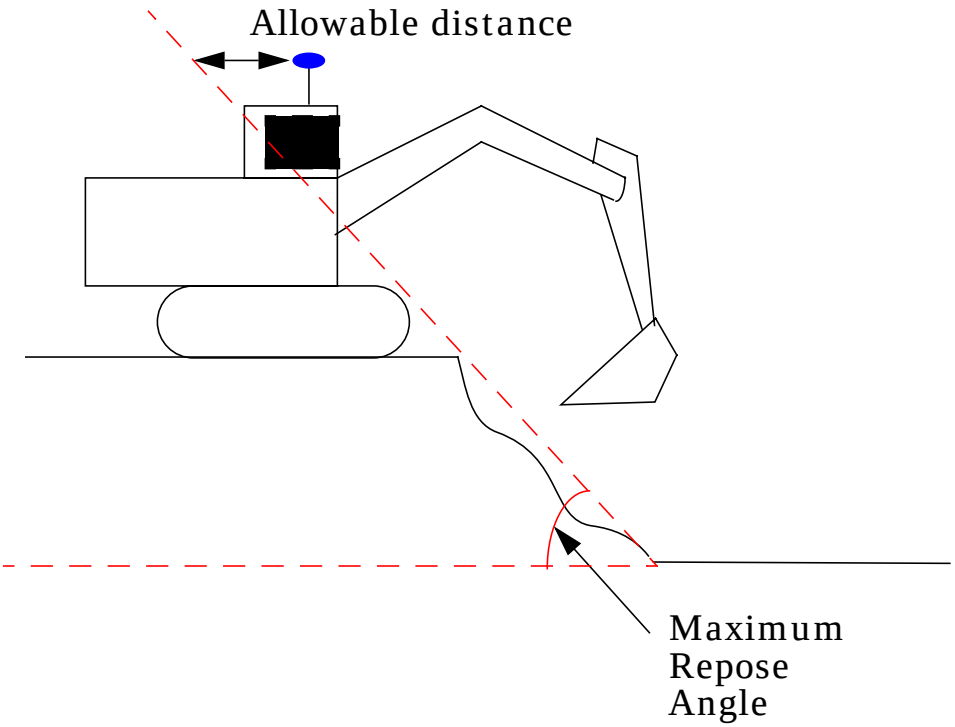

Figure 50: The maximum tracking distance is further constrained based on being able to view the material with the sensor assuming a maximum repose angle for the material (45 degrees). Tracking back too far would cause the dig face to be occluded by the top of the bench.

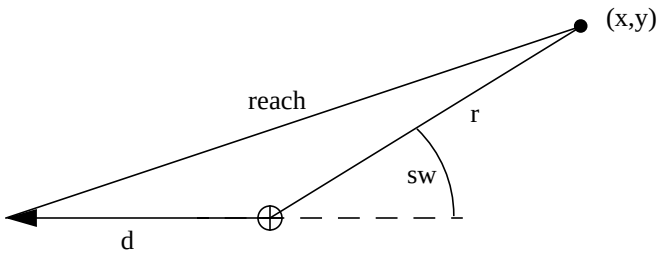

Top View: Machine Reach Constraint

$$r = \sqrt{x^2 + y^2}$$

$$d = r\cos(\pi - sw) - \sqrt{r^2\cos^2(\pi - sw) - (r^2 - reach^2)}$$

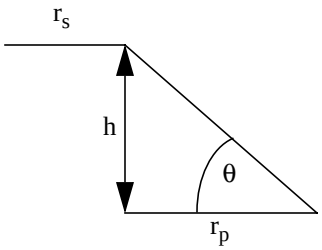

Side View: Sensor View Constraint

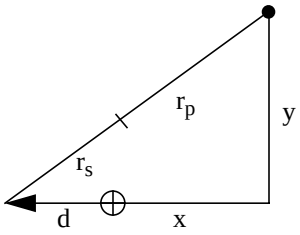

Top View: Sensor View Constraint

$$r_p = \frac{h}{\tan\theta}$$

$$d = \sqrt{(r_p + r_s)^2 - y^2} - x$$

Figure 51: The trigonometry involved in calculating the maximum travel distance based on machine reach and the ability to view all of the material. The circled plus corresponds to the machine's current center, *(x,y)* is the position of the material's edge, *h* is the height of the sensor from the floor, *sw* is the swing angle, θ is the maximum repose angle, r_s is the radial distance to the sensor center from the machine's center, and *d* is the distance the machine can be moved back.

5.2 Modifying the Soil Hardness Index

Recall that the soil hardness index in Autodig controls how much force is applied to the ground during the digging process. The soil hardness index is another control parameter effecting performance, and the optimal index depends on the characteristics of the material in which the machine is digging. The soil hardness index could be another search parameter in the planning process, however adding another degree of freedom would be too costly in terms of computation. Alternatively, we would like to find a relationship between the soil-tool properties that are already being estimated, and the optimal index. Once a relationship is established, then the soil hardness index may be selected on-line by estimating the soil-tool properties.

To illustrate the effect of the soil hardness index, the model of the digging process was used to generate a trajectory of the bucket through the ground for two separate settings. The top plot in Figure 52 shows the effect of the index in a soft soil, and the bottom plot shows the effect in a much harder soil. In the soft soil, both digs obtained a full bucket of dirt. However, the higher index caused the bucket to go too deep, and increased the time required to dig by 60% while increasing the energy required to dig by 75%. In the harder soil however, the lower index was only able to fill the bucket by 50%, while the higher index was able to obtain a full bucket.

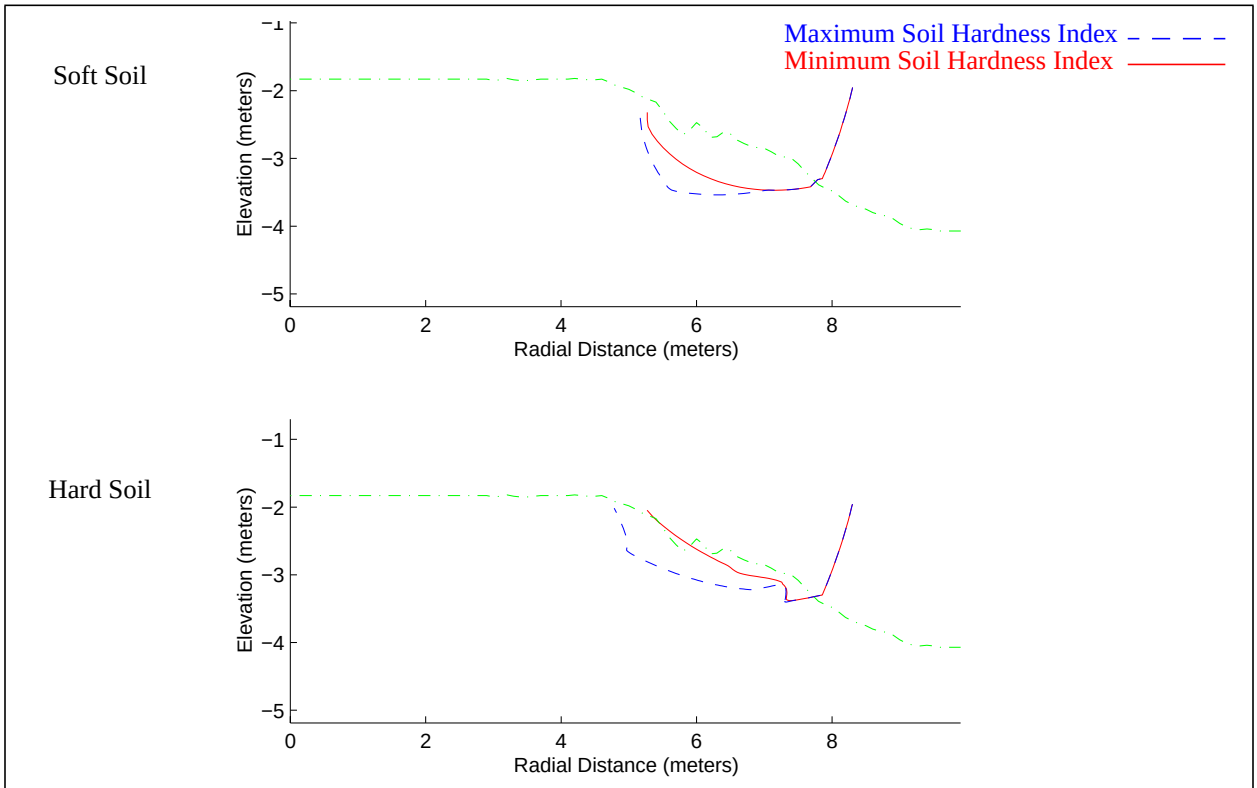

Figure 52: Simulation results showing the effect of Autodig's soil hardness setting in different soil conditions with the same terrain profile. The plots show the trajectory of the bucket tip through the ground for two different soil hardness settings. The top plot illustrates the effect in a soft soil, while the bottom shows the effect in a harder soil. The higher index is required in the hard soil to get a full bucket of dirt. However it wastes time and energy in the soft soil.

Therefore it is expected that a function that relates the optimal soil hardness index to the soil-tool properties should increase as the stiffness of the soil increases. The first step in determining this function is to find a suitable basis vector that relates the soil-tool properties to the optimal index. One simple alternative is to use the force that would be exerted on the bucket with the soil characteristics given an average intersection geometry. That is:

$$\text{Average Soil Force} = \Psi\Gamma_{typical} \quad \textbf{(25)}$$

where $\Gamma_{typical}$ ={depth=1 meter, rake angle=-45 degrees, terrain angle=30 degrees, and swept volume=.75 m^3}.

We used simulation to test this idea since a large variety of soils were not available at our test site. The analytical soil model was used to simulate the forces during digging, and an exhaustive search was conducted in each simulation for the optimal soil hardness index as defined by the dig quality function. The soil forces were modified by adjusting the coefficients in the analytical soil model, and the tests were conducted over a large number of terrain topologies and soil coefficients. For each test, the soil-tool properties were estimated for the empirical model, and these properties were used to calculate the average soil forces as given in Equation 25.

Figure 53 shows the results of the simulations. The optimal soil hardness index for each soil condition and terrain topology is plotted versus the average soil force. The dashed line shows the results of a linear regression between the data points. There does appear to be a fairly strong relationship between average soil force and the optimal index.

There are numerous alternatives that could be used for the independent variables for this procedure, and many different function approximation methods. For instance, a multi-dimensional regression for each soil parameter may provide a stronger relationship. Memory based learning or neural networks could be used for the function approximation.

The results of these simulations show that the use of the soil-tool properties for setting the soil hardness index appears to be a viable alternative. Due to lack of time, we could not pursue this subject any further. Additional experimentation would be needed with real test data and a statistical analysis conducted to determine the validity of this approach.

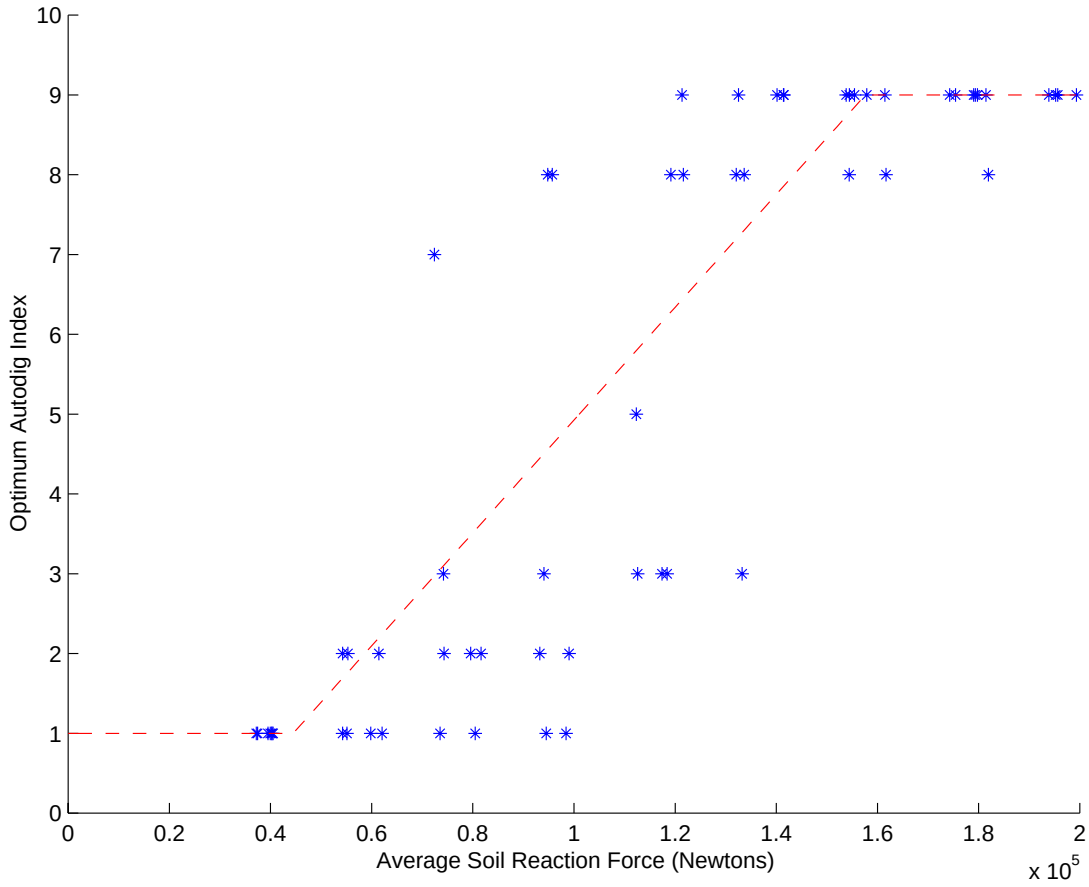

Figure 53: The optimum Autodig soil hardness index for a variety of simulated soil conditions and terrain topologies. The dashed line shows a linear regression between the average soil force and the optimum autodig index. Note that as the average soil force increases, the optimum autodig index increases.

5.3 System Implementation

In Chapter 1, the overall software architecture for the system was introduced. This section examines in more detail the implementation of the excavation system alone. A detailed view of the software architecture is shown in Figure 54. The system was designed so that each module could be developed and tested separately, and so that they could be run as separate tasks or on separate processors as needed computationally. The software system consists of a motion planning module, a dig planning module, a soil estimation module, a control interface, and a sensor interface. The control and sensor interfaces are responsible for communicating with the actual hardware on the machine, while the rest of the modules are higher level processes that determine the actions the machine should take based on the state of the machine and its environment.

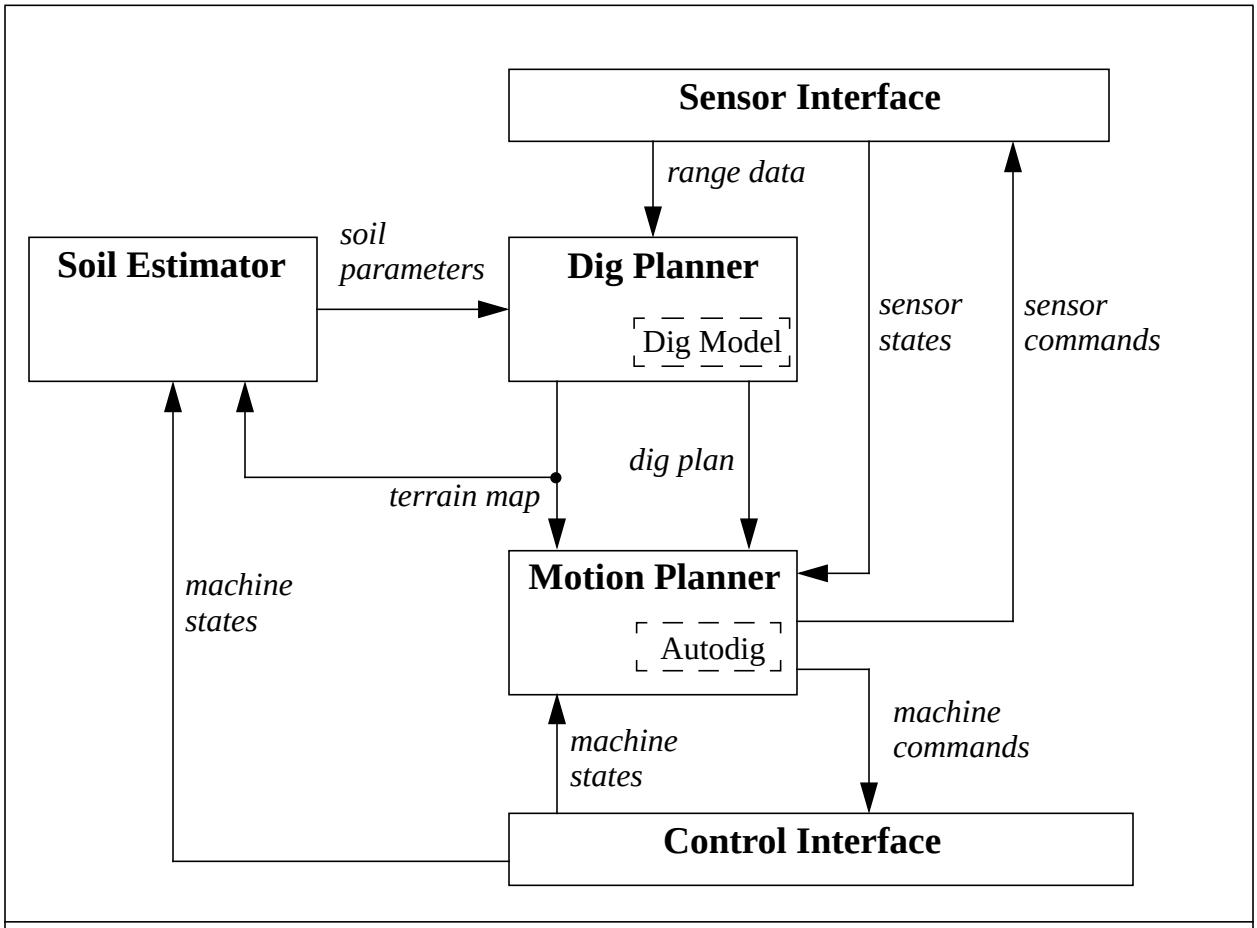

Figure 54: Software Architecture. The Motion Planner controls all the movements of the machine including the execution of Autodig. The Soil Estimator monitors the digging process, and estimates the soil characteristics. These characteristics are used by the Dig Planner to generate the dig plan. The dig plan consists of an optimal dig location, a cleanup plan, a tracking distance, and the dig quality.

The Dig Planner module is responsible for generating a plan for excavation based on the shape of the terrain obtained from the range sensors. The Dig planner receives the range data from the Sensor Interface, and packages it into the terrain map structure that was previously described. The terrain map is then sent on to the Motion Planner and the Soil Estimator. The Dig Planner receives an indication of the soil characteristics from the Soil Estimator, generates the next plan of action for eroding the bench, and sends this on to the Motion Planner. The dig plan consists of an optimal dig location and associated dig quality, a plan for cleaning the floor, and a maximum allowable distance for tracking the machine backwards.

The Motion Planner controls the movements of the machine and the sensors by communicating with the hardware interface modules. It receives the terrain map and plan from the Dig Planner, and then maneuvers the machine to execute the plan. It decides on which plan to execute based on the algorithm shown in Figure 41. If it decides to execute a digging action or cleanup action, then it moves the machine to the specified location. It then uses the terrain map and Autodig to execute the dig. If instead the machine should be tracked backwards, it would execute this action by the specified amount.

The Soil Estimator module is responsible for monitoring the digging process to determine the soil characteristics. It does this by monitoring the machine states during the digging process through the control interface, transforming the machine states to bucket poses, intersecting the bucket poses with the terrain map, and finally extracting the soil parameters as described in the previous chapter. This module would also be responsible for setting the Autodig soil hardness index, however it was never implemented.

Chapter 6 will describe the test results obtained with this system.

Chapter 6 Dig Planning Results

During the last year of the project we executed over 1700 dig cycles in developing the dig planning system. This is equivalent to loading over 280 trucks, or moving approximately 4500 tons of material. The purpose behind this experimentation was to extend the amount of time that the machine could operate autonomously without any human intervention, and improving the productivity and efficiency of the excavation process.

Great strides were made in both of these areas during this time. This was due to the culmination of a significant effort into improving the reliability of the hardware systems, improvements to the robustness of the software, and finally improvements in the dig planner's ability to manage the erosion of the bench.

This chapter describes the achievements that were made in the area of extended operations. It will also describe several experiments which compared the digging performance of the automated machine with the performance of an expert human operator.

6.1 Extended Operation Results

Extended operation experiments were conducted with the excavator testbed described in the first chapter. The objective of these experiments was to extend the ability of the machine to dig continuously without any assistance external to the testbed itself. Once the computers on board the machine were initialized, the machine required no further assistance. The computers handled all of the strategic decision making, and executed all of the machine's actions such as digging, dumping, and tracking.

Most of the development process was conducted on a bench in a contrived worksite that allowed repeated experiments in a relatively non-variant environment. Some portion of the tracks always rested on a concrete foundation to ensure the vehicle's stability, and the material was homogeneous and easy to dig. However during the last two months of the project we moved the excavator from our development bench to a remote worksite where the environment was much less structured. The material included boulders, old tires, and other inclusions which increased the level of difficulty for digging. The purpose of moving to the remote site was to test our algorithms in a more realistic environment, and to provide us with larger amounts of material so that the machine could dig for longer periods of time without being interrupted.

Several modifications were required to the remote worksite in order for our system to function properly. First we had to level the bench to improve the stability of the excavator and to minimize the amount the machine tilted. Since the excavator was not equipped with a positioning system or inclinometers, the contour of the floor that it digs is relative to the surface upon which the machine sits. Large variations in the contour of the surface would cause large variations in the floor. Additionally all of the vegetation such as tall weeds and bushes had to be removed initially from the bench. Otherwise the vegetation would show up in the laser scans, and would not be distinguishable from the soil. This could cause large errors in modeling the digging process and estimation of the soil-tool properties.

We also found that the material in the remote site was much dryer than the development site resulting in large amounts of dust. This was a major problem because the laser sensors would image the dust, and the planning system was not able to distinguish the cloud from the soil. Special lasers which were better at seeing through the dust were used to help alleviate this problem although it was not eliminated totally. So when the dust conditions were particularly bad, a pause was added to the cycle after digging which allowed the dust clouds to settle prior to sensing the terrain with the lasers. This would not be suitable in a product implementation and would need to be rectified. Other sensing technologies such as high resolution radar or further advancements in dust penetrating lasers may be the ultimate solution.

Since the focus of the work was on digging, we did not place any effort on optimizing the motions between the dig face and the dump location. This technology was being developed concurrently on the project and was not a consideration in our work. In fact to simplify the logistics of moving the material, we did not load the material into trucks. The material was cast into spoil piles to the side of the digging area.

In order to analyze the performance of the dig planning system during these extended runs, we measured several quantities that are particularly important in the digging portion of the cycle. Measured parameters included the amount of time that the bucket was in contact with the ground, the weight of the material in the bucket after a dig, and the amount of energy that was expended during digging. Note that these quantities are related to the values used in the dig planner for determining the overall dig quality.

To measure the time that the bucket was in contact with the ground, we used a terrain map that was generated prior to each dig. While digging, the position of the bucket tip was calculated using the sensed resolver joint angles and forward kinematics. If the bucket tip position was below the terrain elevation, then time was added accordingly. To calculate the energy of the dig, the displacement of each cylinder multiplied by the cylinder pressure force was integrated over time. The weight of the material was obtained by stopping the machine immediately after the digging process, waiting for the pressure fluctuations to die down, and finally conducting a static analysis on the linkage using cylinder pressures to determine weight.

Figure 55 shows the excavator working in the remote worksite during an extended run. The first graphic shows the bench after 12 dig cycles and the second graphic shows the bench after 96 cycles. It can be observed that the machine has moved backwards several times and eroded a large tract of material parallel to the tracks. The ground left behind is fairly level but would not be suitably clean for trucks to traverse. A subsequent levelling operation would be necessary.

Figure 55: The excavator during an extended digging test at the remote worksite. The top image was taken after 12 digs, and the bottom image was taken after 96 digs. Note how the machine has moved farther back on the bench. The spoil piles are shown on the right.

Table 3 shows a summary of three of the longest duration runs. Test #1 corresponded to a bench that was approximately one meter higher than the others. This was the longest run with 114 dig cycles and 3 backups. The best runs however corresponded to tests #2 and #3. The improved results were due to reduction in dust levels, and improvements that were made in the planning software. Tests #1 and #2 were halted due to malfunctioning hardware and software. Test #3 however had to be halted because the excavator completely eroded all of the removable material along its path.

Table 3: Summary of the three longest extended runs. The number of digs and the number of times the machine was tracked backwards is shown along with the average weight, time, and energy per dig cycle. Test 3 was halted because the bench was completely removed.

Test Number	Digs	Backups	Weight (N)	Time (s)	Energy (kJ)
1	114	3	18917	10.2	225
2	109	5	22654	6.7	213
3	109	6	21300	7.6	252

Figure 56 shows the digging statistics corresponding to test #3. The plots depict the weight, time, and energy recorded for the digs in the order in which they occurred. The time plot exhibits six groupings of time larger than 10 seconds. These groupings correspond to the cleanup actions just prior to tracking backwards. Note that the higher values of energy generally correlate to the cleanup actions. Given that 25,000 N is approximately a full bucket of material, then 6% of the buckets were less than half full, and 2% of the buckets were less than one third full.

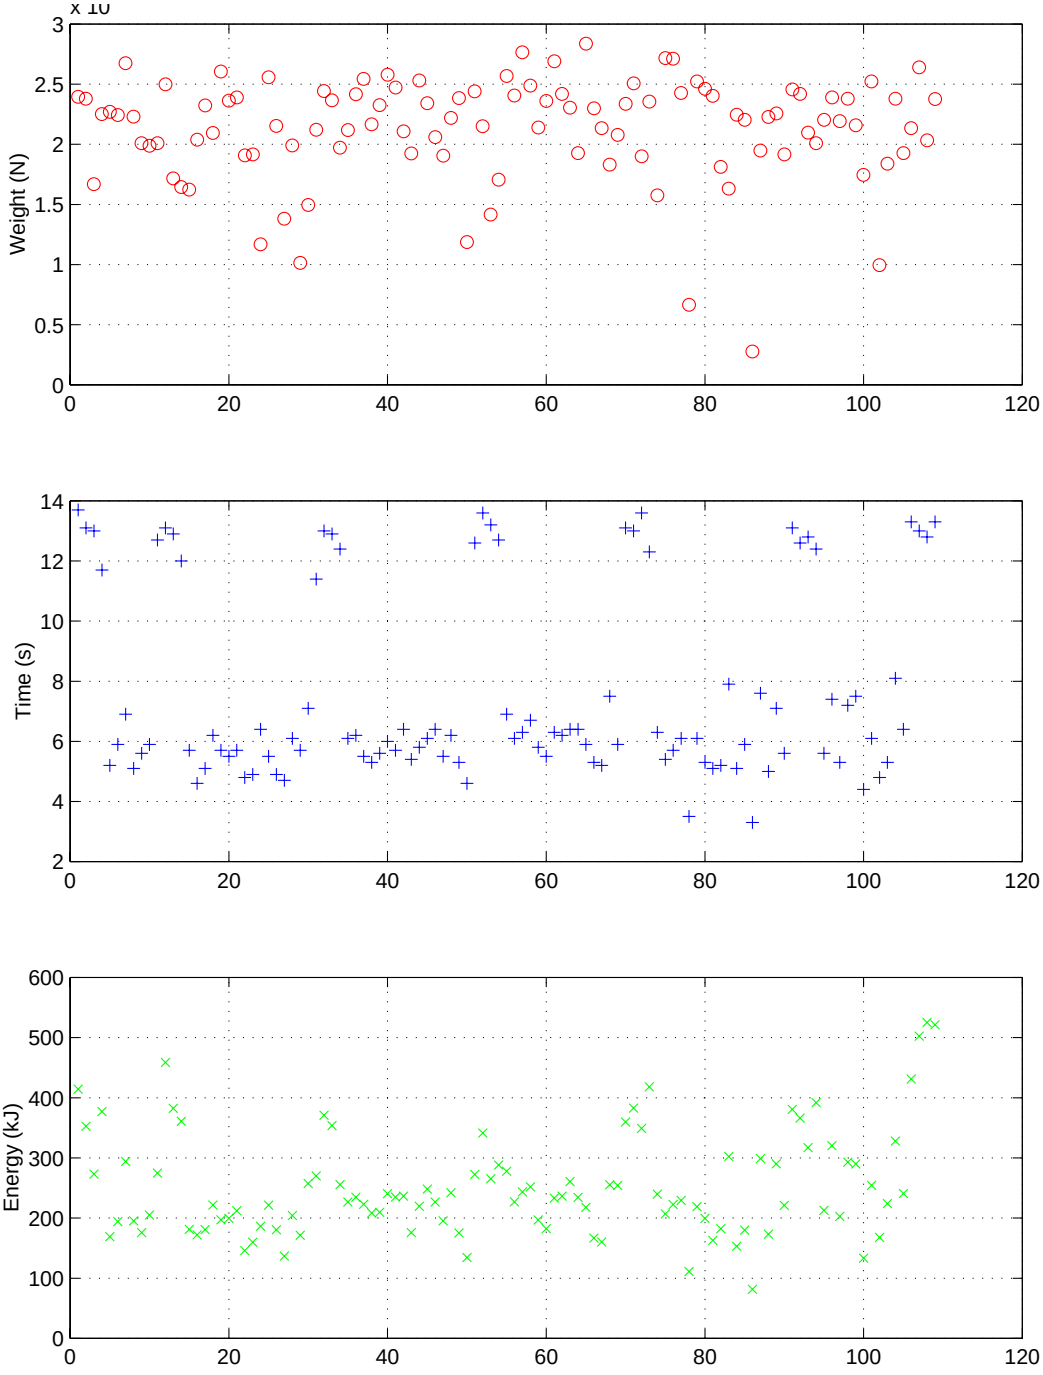

Figure 56: Digging statistics corresponding to test #3. The digs with cycle times greater than 10 seconds correspond to cleanup actions. 94% of the bucket payloads were over half full.

Another way to examine this data is to compare the statistics for the cleanup actions and the dig actions. The comparison is shown in Table 4. This illustrates that the cleanup actions are extremely inefficient in terms of time and energy. Further enhancements in the cleanup execution algorithm could yield dramatic improvements in the autonomous system.

Table 4: Comparison of statistics for digging and cleanup actions for test #3. The average weight time and energy is summarized for each action. The cleanup actions are much less effective.

	Weight (N)	Time (s)	Energy (kJ)
Dig Actions	21201	5.8	212
Cleanup Actions	21590	12.8	366

6.2 Comparison to Expert Human Operator

Given that we can successfully run the machine autonomously for extended periods of time, we must now examine how well the system performed. Our analysis will be based on the digging performance of a human operator who is recognized as an expert in the earth moving industry. Additionally, it would be quite useful to understand the benefits of using perception based dig planning in autonomous operation. For comparison purposes, we created a predefined excavation sequence that could be used instead of planning the digging operations. In this method we predefined where the machine should dig, always cleaned up the floor with a fixed pattern, and always tracked back a given distance. In this section we will compare the digging performance of the expert human operator to the automated system with the perception based dig planner, and to the automated system with the predefined excavation sequence.

Again we will use the quantities of weight, time, and energy to make our comparisons. In order to monitor the performance of the human expert, we took a laser scan before each dig cycle. This allowed us to determine when the bucket entered and left the ground. We also paused after each dig so that we could use the pressures in the cylinders to weigh the material in the bucket. During the actual digging operation, the pressures and displacements of the cylinders were monitored so that the energy could be calculated.

For autonomous operation without the use of a perception based dig planner, we needed to predetermine the excavation process. The dig locations were specified such that one set of implement joint angles were used for all of the digs, and the swing angle was incremented from left to right similar to the method used by the coarse dig location planner. After digging once in all of the regions, a cleanup action was then executed in each region. The cleanup was initiated at the far reach of the machine because without perception there is no means of knowing the distance to the material. Finally the cleanup actions were followed by tracking the machine backwards a typical distance selected by the perception based dig planner (1.3 meters), upon which the process started over.

We experimented with two methods for predefining the implement joint angles for initiating the digs. The first method was to have a human operator specify the joint angles by manually placing the bucket in a configuration that he thought would be generally favorable for all of the digs. The results of this were so poor that the test was discontinued. The problem with this method could be attributed to differences in how the human operator might dig and how Autodig behaves. The human operator specified the location based on how he would execute the dig, even though perhaps this was an unfavorable position for Autodig. The second method was to specify a set of joint angles that was close to the median of the joint angles selected by the perception based dig planner. This is akin to having a 'virtual expert' specifying the predefined start positions. This method seemed to work better, and will be used in the rest of this discussion.

Surely it would have been possible to modify these predetermined values allowing some improvement in the performance of the machine. However once the terrain situation changed such as increasing the height of the bench, then a new set of values would need to be found. Rather than spend a lot of effort looking for the best combination for this particular terrain situation, the test was intended to illustrate our best guess at reasonable values.

Table 5 shows a comparison of the digging statistics for the expert operator, the autonomous machine with perception based dig planning, and the autonomous machine with a predefined excavation sequence. The comparison was made over 41 dig cycles such that it was necessary to track the machine backwards several times.

Table 5: Digging statistics comparing expert human operator to automated systems. The statistics correspond to the average per cycle for 41 dig cycles.

	Weight (N)	Time (s)	Energy (kJ)
Expert Operator	27412	7.9	350
Perception Based Planner	23992	7.5	272
Predefined Excavation Sequence	19397	6.5	147

The first thing to notice is that the performance of the perception based dig planning system appears to be better than the extended runs given in Table 3. This is due to a minor change in the floor cleanup operation which was never tested in a longer sequence due to lack of time. The change was to detect the floor during the first cleanup execution, and then to use this elevation for all subsequent cleanups. Prior to this, the floor was detected on every cleanup operation. This change caused the cleanup operations to attempt to dig deeper into the ground in some instances, resulting in more material in the bucket. This change was not implemented for the predefined excavation sequence.

From the table it is clear that the expert operator managed to capture more material in the bucket on average than the autonomous systems. Thus the productivity of the expert is superior. If we

assume that it would take 12 seconds to move between the dig and dump locations, then the productivity can be estimated by:

$$\text{productivity} \cong \frac{\text{weight}}{12.0 + \text{time}} \tag{26}$$

The productivity of the expert is approximately 12% higher than the perception based dig planner, and 31% higher than the predefined excavation sequence. Since the perception based dig planner used a different cleanup method than the predefined excavation sequence, it is only fair to compare these two during the normal digging operations. During digging only, the average payload for the perception based dig planner was 42% higher than the predefined excavation sequence.

Neglecting the energy required to move between the dig and the dump locations, the energy efficiency can be calculated by:

$$\text{energy efficiency} \cong \frac{\text{energy}}{\text{weight}} \tag{27}$$

Thus the perception based dig planner is 12% more efficient than the expert while the predefined excavation sequence is 32% more efficient than the expert. This is an interesting result. There appears to be a direct trade-off between productivity and energy efficiency.

Histograms of the measurements are shown in Figures 57 through 59. The distribution of weights shown in Figure 57 illustrates that the variance of the bucket sizes for the predefined excavation sequence is significantly larger than for either the expert operator or perception based dig planner. In fact, when using the predefined sequence, we noted that initially the buckets were adequately filled, but the bucket fill quickly deteriorated as testing progressed. Figure 60 shows the sequence of weights obtained for the predefined excavation sequence and the perception based planner. The bucket fill factor for the perception based dig planner is fairly constant whereas the weights for the predefined sequence is rapidly decreasing.

One possible reason for this distribution is due to the changing shape of the terrain. Prior to beginning the test, the bench was smoothed with the excavator bucket. Hence a full bucket was easily obtainable. However as digging progressed the terrain undulations increased, making the choice of the dig locations more critical. The perception based dig planner is able to consider the undulations in the terrain when deciding where to dig and is thus able to maintain a high bucket fill factor throughout the test.

The distribution of dig execution times shown in Figure 58 illustrates the impact of using a separate cleanup operation. Note that the histograms for the perception based dig planner and the predefined sequence both have basically two distributions. The distributions with the lower execution times correspond to normal digging operations, whereas the higher execution times are due to the cleanup operations. With the expert operator, these two functions that we have treated separately are blended together. It is difficult to tell when the expert is cleaning the floor and when he is digging normally.

Figure 59 shows the histograms of energy expended during digging. Again, the higher distribution of energies for the perception based dig planner coincide with the cleanup operations. This result is not as dramatic in the predefined excavation sequence because the cleanup operations were not as effective. Again, the reason for this difference is the method used for cleaning the floor.

From this testing we can conclude that the perception based dig planning methodology has added significant value in terms of productivity as compared to simply using “canned” data points in a predefined excavation sequence. We have also shown that the perception based planning method has enabled the automated system to approach the productivity of an expert human operator during digging while exceeding energy efficiency levels. Further improvements could be gained in the automated system by improving the cleanup operations.

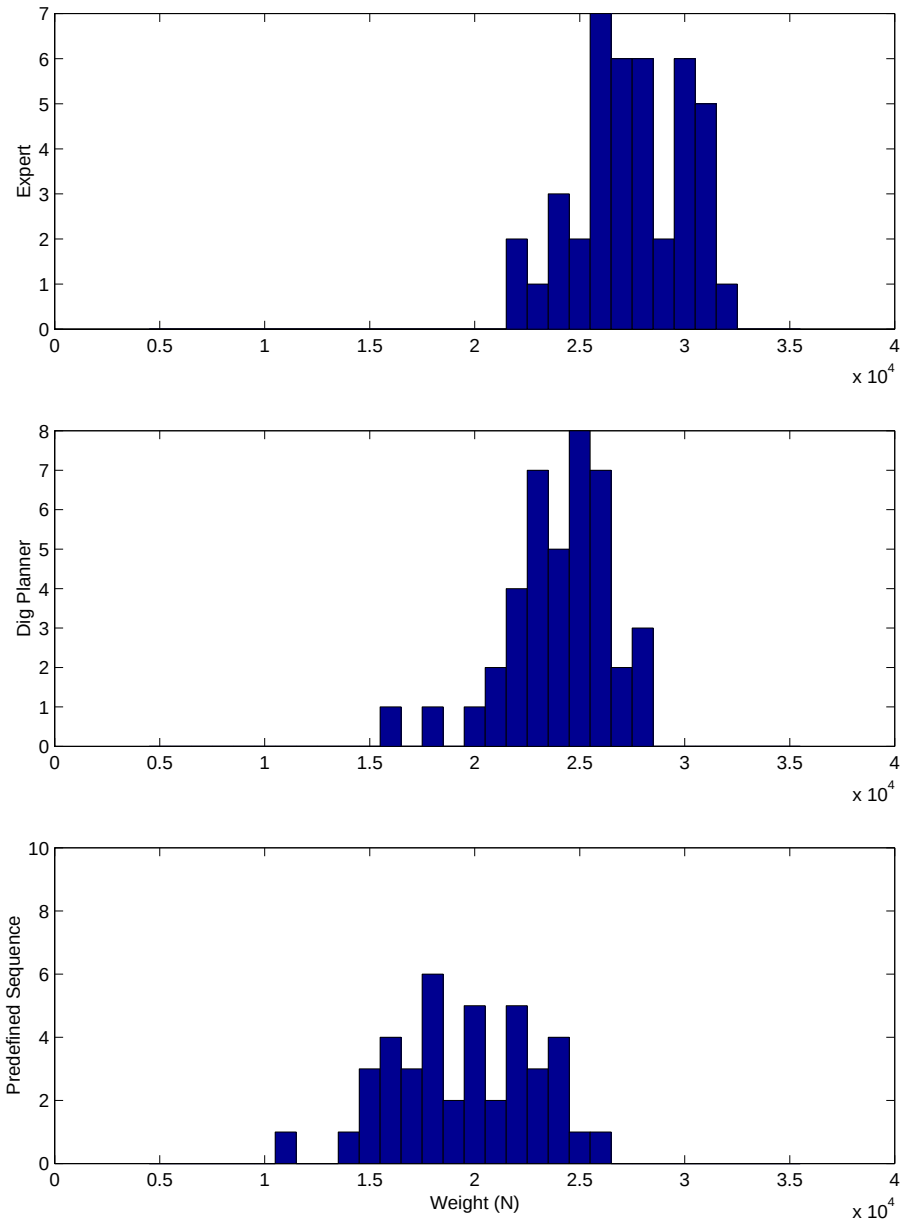

Figure 57: Histograms comparing the weight of material captured. The variance of the weights is larger for the predefined excavation sequence than for either the dig location planner or the expert operator.

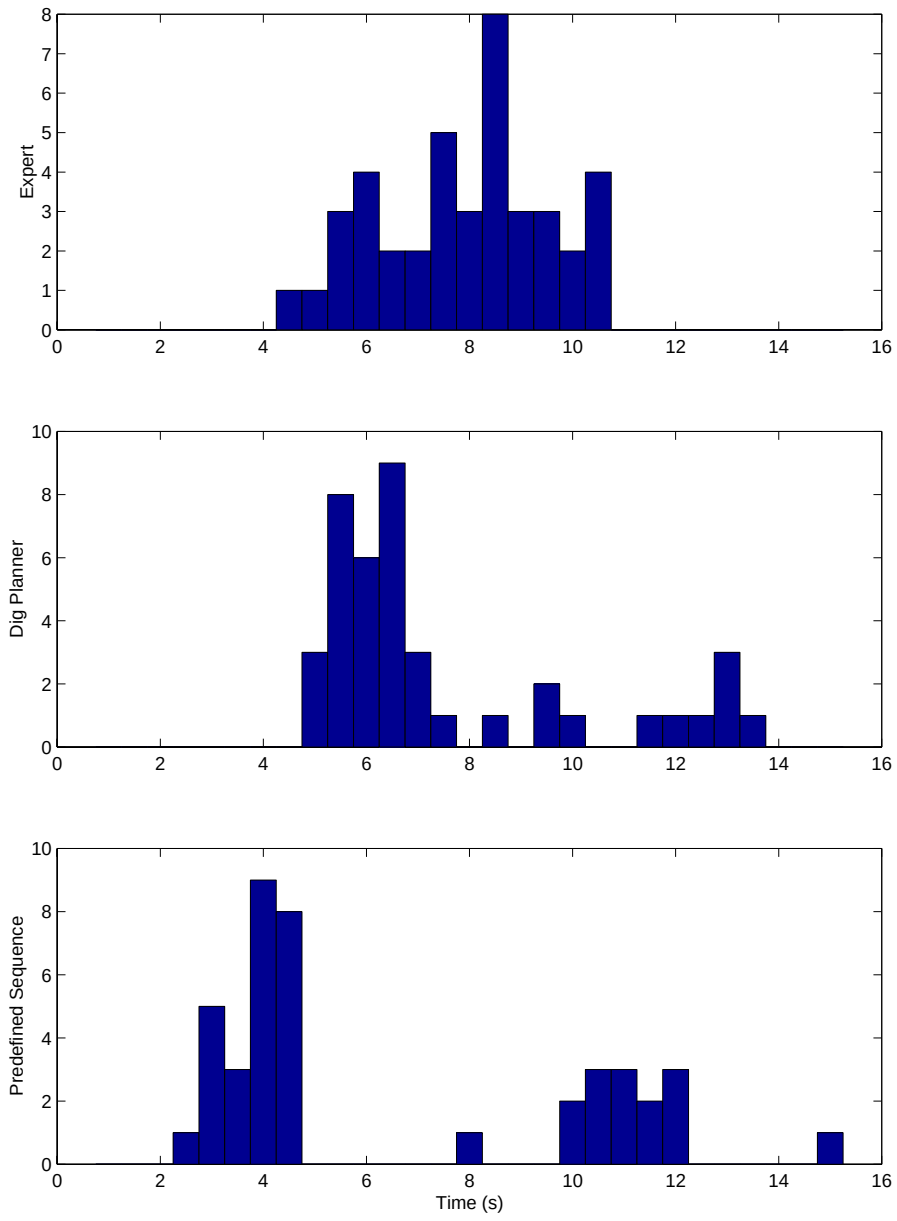

Figure 58: Histograms comparing the time while in contact with the ground. The predefined excavation sequence and the dig location planner both have two distributions. The lower distribution is due to digging normally, while the higher distribution is due to the cleanup operations.

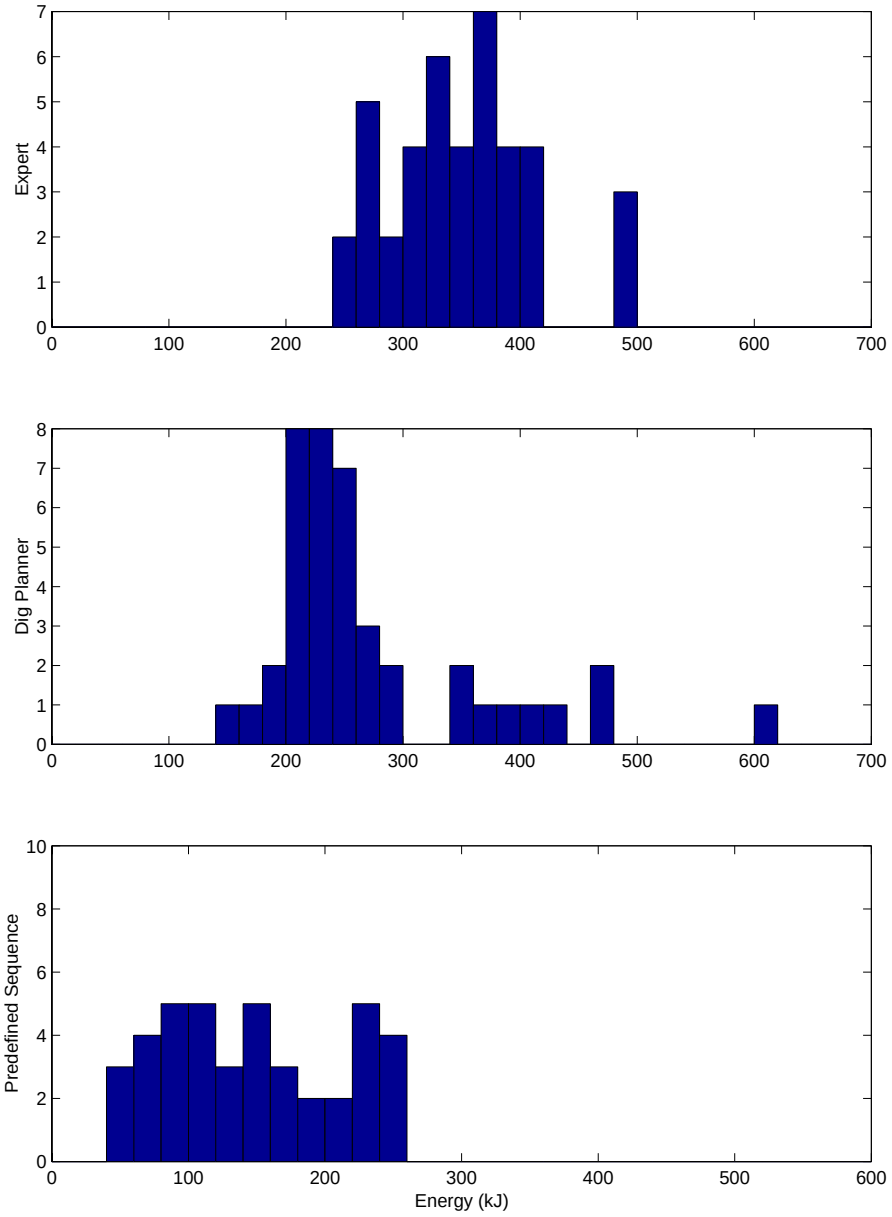

Figure 59: Histograms comparing the energy expended during digging. The cleanup operations for the perception based dig planner result in significantly higher energies.

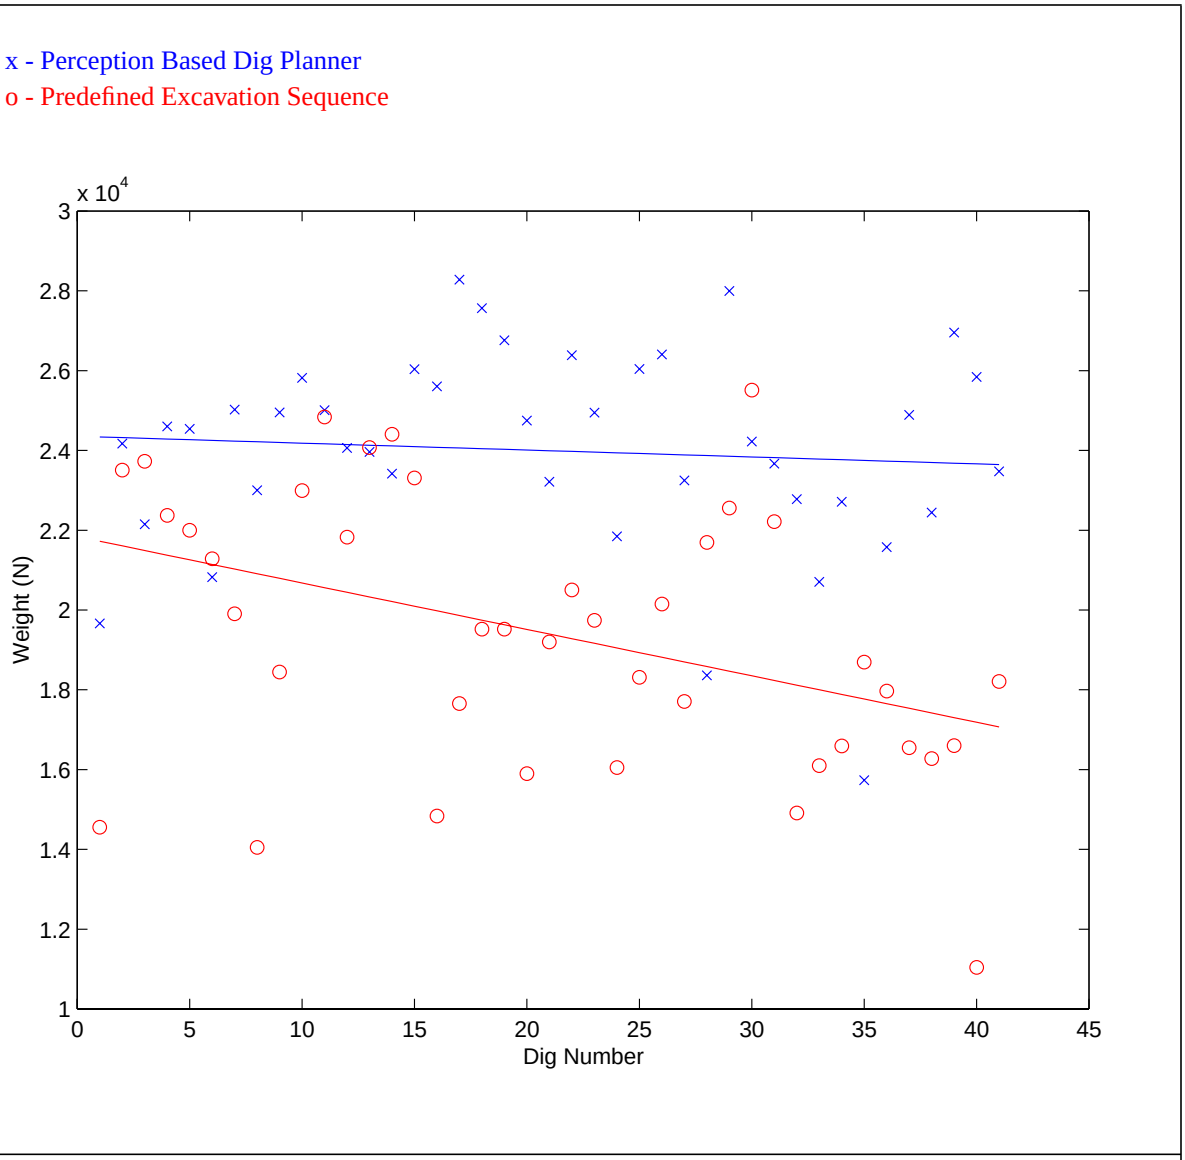

Figure 60: Comparison of the sequence of bucket weights obtained with the perception based dig planner and using a predefined excavation sequence. Note that the perception based dig planner maintained a consistent bucket fill factor whereas the bucket fill factor for the predefined sequence is rapidly decreasing.

Chapter 7 Conclusions

There have been several advances made in this research related to improving the efficiency of digging, development of computationally tractable models for the excavation process, planning locally optimal dig sequences subject to constraints, and combining planning methods to manage the erosion of a bench for the sake of extended operations. This chapter summarizes the research described in this thesis, describes some directions for future work, and lists the major accomplishments.

7.1 Summary

Chapter 1 introduced the Autonomous Loading System and described the truck loading process in a mass excavation scenario. The overall operation of the automated system was described, including the role of the perception based dig planner. Chapter 2 outlined previous research related to automated dig execution, excavation modeling, and planning earthmoving operations.

Chapter 3 contained a discussion on the operation of the closed loop control law called Autodig, which was utilized to control the machine during the excavation process. This control law allowed the machine to quickly obtain a full bucket of material as long as Autodig was initiated where sufficient material was available, and the control parameters were properly tuned. Prior to this research, a human operator was required to adjust the control parameters. We identified that one of the most important parameters that influenced the digging process was the stick angle at which to end the dig. We implemented a perception augmentation to Autodig which automatically ended the dig cycle when a desired amount of volume was swept by the bucket. In Chapter 5 we proposed a means by which the hardness of the soil could be estimated, and used to automatically tune the soil hardness parameters inside Autodig. Finally we added a position based enhancement to Autodig which caused the bucket to track a level floor for cleaning operations.

Chapter 4 discussed the development of a computationally tractable forward model of the excavation process. The purpose of this model was to provide a means for estimating the utility of a candidate machine configuration for initiating Autodig. This model consisted of a hydraulic actuator model, a model of the soil reaction forces, and the actual Autodig algorithm. Since a large number of candidate configurations are investigated prior to every dig, it was necessary to find a suitable trade off between model accuracy and computational complexity. Thus for the hydraulic actuator model, a combination of neural networks were used to predict the velocities of the cylinders given the commands issued from Autodig and the forces acting on the machine.

The motion of the hydraulic cylinders dictates the position of the bucket in space, and thus the bucket can be intersected with the terrain profile in order to estimate the resistive forces of the soil. We investigated two different means for modeling the resistive forces. The first means was an analytical model based on the Fundamental Earthmoving Equation in soil mechanics which was modified to more adequately reflect the bench loading application. The second means was an empirical model based on a linear combination of terms found in the analytical model. Since the soil conditions can be quite diverse, we implemented a means for estimating the soil conditions for both model types. Both models were capable of estimating the resistive forces with good accuracy. However extraction of the soil parameters for the empirical model required much less time.

Once the resistive forces on the bucket are predicted, the pressures in the hydraulic cylinders can be estimated. These pressures are then used both as an input to the vehicle actuator model, and as an input to Autodig. Autodig uses the pressures and the predicted position of the implements to calculate the actuator commands. Model operation therefore consisted of initiating the model in the candidate machine configuration, and then integrating the predictions forward in time to calculate the trajectory. During this prediction several key statistics are compiled: the energy

required to dig, the time expended during digging, and the volume of material swept into the bucket.

Chapter 5 discussed the implementation of the planning system for selecting where to dig. The planning system consisted of three separate planning methods. This included an optimal dig location planner, a cleanup planner, and a tracking planner.

The purpose of the optimal dig location planner was to find the optimal dig location relative to a cost function and subject to kinematic and material shape constraints. This method consisted of a two part planning process: a coarse and refined planner. The coarse planner divided the dig face into separate regions, and the refined planner then searched for the optimal dig location within a region. The use of a two part planning strategy reduced the search space for the refined planner, and caused the material to be removed in an orderly fashion based on some overall material removal strategy. The refined planner used the forward model of the excavation process to determine if a candidate dig would violate any material shape constraints (such as a level floor), and for calculating a cost function based on time, energy, and volume. The refined planner selected the dig which optimized this function.

A cleanup planner and a tracking planner were also implemented. The cleanup planner determined where a cleanup action should be initiated based on how far significant amounts of material were located relative to the machine's position. The tracking planner determined how far the machine could be tracked backwards, and still be able to both reach the material and perceive the material with the range sensors.

The planning and execution system was implemented on a 25 ton commercial excavator testbed in a relatively unstructured work environment. Chapter 6 discussed how the system's performance was comparable to an expert human operator, and significantly better than a simple heuristic based approach. It was also demonstrated that the machine could be operated for several hours at a time without human assistance.

7.2 Future Work

Several opportunities exist for extending this work beyond what has been presented in this thesis. First, the algorithms could be generalized to include other earthmoving machines. In [Singh 98] we show how the coarse to refined planning process could be extended to a wheel loader. The algorithm however used a heuristic approach for estimating the quality of a selected dig point instead of modeling the digging process. Part of the reason for this was due to the lack of a testbed for assessing the digging forces. Actual implementation on a machine would allow a comparison between the heuristic and model based approaches. Furthermore it may be interesting to see how this planning methodology could apply to a dramatically different machine such as a dozer.

It would also be beneficial to experiment with this system in a diverse set of soil conditions. Our experiments were conducted in material that would be considered fairly easy to dig. Harder and more heterogeneous materials such as blasted rock may prove to be difficult to model. Experi-

ments in varied materials would also allow further exploration of automatically setting the Autodig soil hardness index based on estimated soil-tool properties.

Finally, in Chapter 6 we showed that the cleanup operation severely degraded the performance and efficiency of the automated machine. Further work could be conducted on improving this process, or perhaps eliminating the process entirely through the use of an impedance based control. The use of impedance control would allow the digging process to conform to a desired shape so that the floor elevation could be maintained while removing the bulk of material. This appears to be closer to the methodology used by an expert operator since it is difficult to distinguish when he is digging versus cleaning the floor.

7.3 Major Accomplishments

There are four main accomplishments achieved through this research. First, we have augmented the Autodig control algorithm with perception based enhancements for improving the consistency of the digging process. Second, we have developed an excavation model that captures machine dynamics, soil-tool interaction forces, and control system dynamics. The model has been proven to be both reasonably accurate and computationally fast, and is capable of adapting to the characteristics of the soil encountered at the work site. Third, we have developed a perception based planning system which uses the model for managing the excavation of the bench. The planner consists of a two part planning process for selecting optimal dig locations, and plans the cleanup and tracking operations. Finally, we have implemented the automated excavation system on a large commercial hydraulic excavator, and demonstrated that the system approaches the performance of an expert human operator.

In a broader sense, we have proven that an automated earth moving machine can maintain high levels of productivity and efficiency for extended periods of operation. Even though an expert human operator may be able to exceed the capabilities of the automated machine for a brief period of time, the automated machine can sustain these levels with limited interruption. Certainly there exist many challenges ahead in making the system robust for commercial viability. However we have shown that the automation of extended earthmoving tasks is technically achievable.

Appendix Actuator Forces and Bucket Forces

This appendix describes the mathematical relationships between the soil forces acting on the tip of the bucket and the forces acting on the hydraulic actuators. Section A.1 covers the transformation of bucket forces to actuator forces. Section A.2 covers the transformation of actuator forces to bucket forces.

A.1 Bucket Forces to Actuator Forces

Chapter 4 discussed how the soil-tool model predicts the forces acting on the bucket tip. These forces must subsequently be converted into forces acting on the hydraulic actuators for use in the actuator model. The following equation represents the force relationships in joint space [Craig 89].

$$M(q)\ddot{q} + V(q,\dot{q})\dot{q} + G + J^T f = \tau \quad \text{(A1)}$$

M is the inertia matrix, V is a matrix containing the centripetal and Coriolis terms, G is a vector containing the torques on the joints due to gravity, f is a vector containing the forces and moments on the bucket, τ is a vector representing the torque on each joint, J is the Jacobian of the mechanism, and q represents the joint angles. Assuming that the acceleration and velocity terms are negligible during digging, the equation reduces to:

$$\tau = (G + J^T f) \quad \text{(A2)}$$

This equation shows that the torques on the joints can be calculated by knowing the torques on the joints due to gravity, the Jacobian which is a function of the joint angles, and the forces. Since we are only concerned about planar forces during digging, the force vector corresponds to (F_x,F_z,M_y)', and the Jacobian corresponds to:

$$\begin{bmatrix} (-l_2S_2 - l_3S_{23} - l_4S_{234}) & (-l_3S_{23} - l_4S_{234}) & (-l_4S_{234}) \\ (-l_2C_2 - l_3C_{23} - l_4C_{234}) & (-l_3C_{23} - l_4C_{234}) & (-l_4C_{234}) \\ 1 & 1 & 1 \end{bmatrix} \quad \textbf{(A3)}$$

where l corresponds to joint length, S and C correspond to the sine and cosine of the summation of the joint angles indicated by the indices. The indices 2,3,and 4 correspond to the boom, stick, and bucket respectively. The torques due to gravitation can be calculated by:

$$G_i = g \sum_{n=i}^{4} l_{in} m_n \cos\psi_{in} \quad \textbf{(A4)}$$

where l_{in} corresponds to the distance from joint i to the center of gravity of link n, m_n corresponds to the mass of link n, and ψ_{in} is the angle between link i and the center of gravity for link n. See [Singh 95] for these derivations.

Given these equations, the torques on the boom, stick, and bucket joints can be calculated. These torques must be subsequently converted into forces acting on the actuators. Figure A1 shows a simplified drawing representing the boom and stick links.

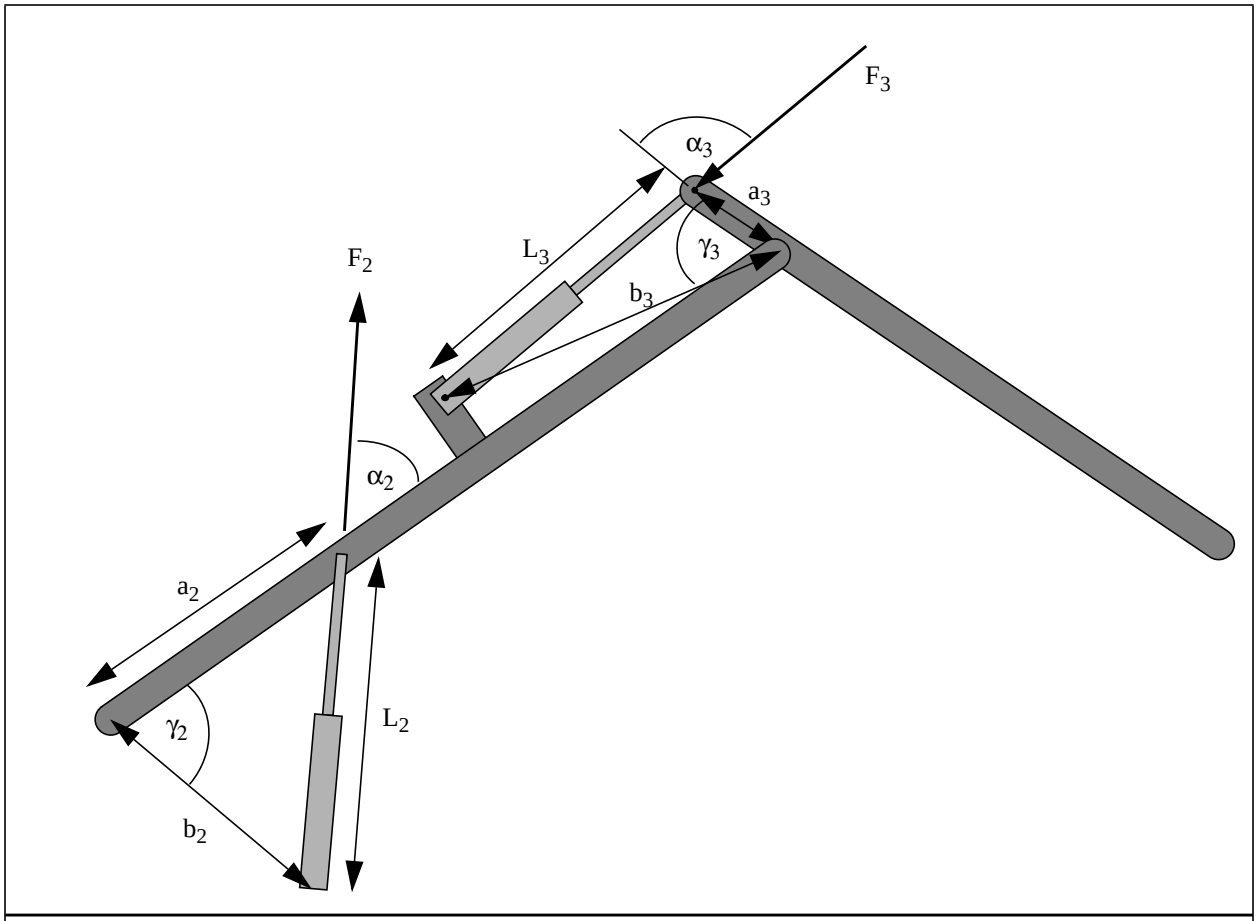

Figure A1: Simplified drawing representing the boom and stick links. *L* corresponds to the actuator length, *b* is the distance from the rotational joint to the base of the actuator, *a* is the distance from the joint to the end of the actuator, γ is the included angle between *a* and *b*, α is the angle between the link and the actuator, and *F* is the force on the actuator. Note that the forces are oriented to cause a positive torque.

Distances a_i and b_i are known from the linkage geometry, and the angles γ_i can be calculated by knowing joint angles. Distances L_i may be calculated using the law of cosines:

$$L_i = \sqrt{\left(a_i^2 + b_i^2 - 2a_i b_i \cos\gamma_i\right)} \quad \textbf{(A5)}$$

Again using the law of cosines and a trigonometric identity:

$$\sin\alpha_i = \sqrt{1 - \left(\frac{b_i^2 - L_i^2 - a_i^2}{2L_i a_i}\right)^2} \quad \textbf{(A6)}$$

The force due to the torque is the torque divided by the moment arm. Therefore the actuator force equation for the boom and stick is:

$$F_i = \frac{\tau_i}{a_i \sqrt{1 - \left(\frac{b_i^2 - L_i^2 - a_i^2}{2 L_i a_i} \right)^2}} \qquad \textbf{(A7)}$$

The equation for calculating the bucket actuator force is different because the actuator transmits the force to the bucket through a four bar linkage. Figure A2 shows a simplified drawing of the linkage.

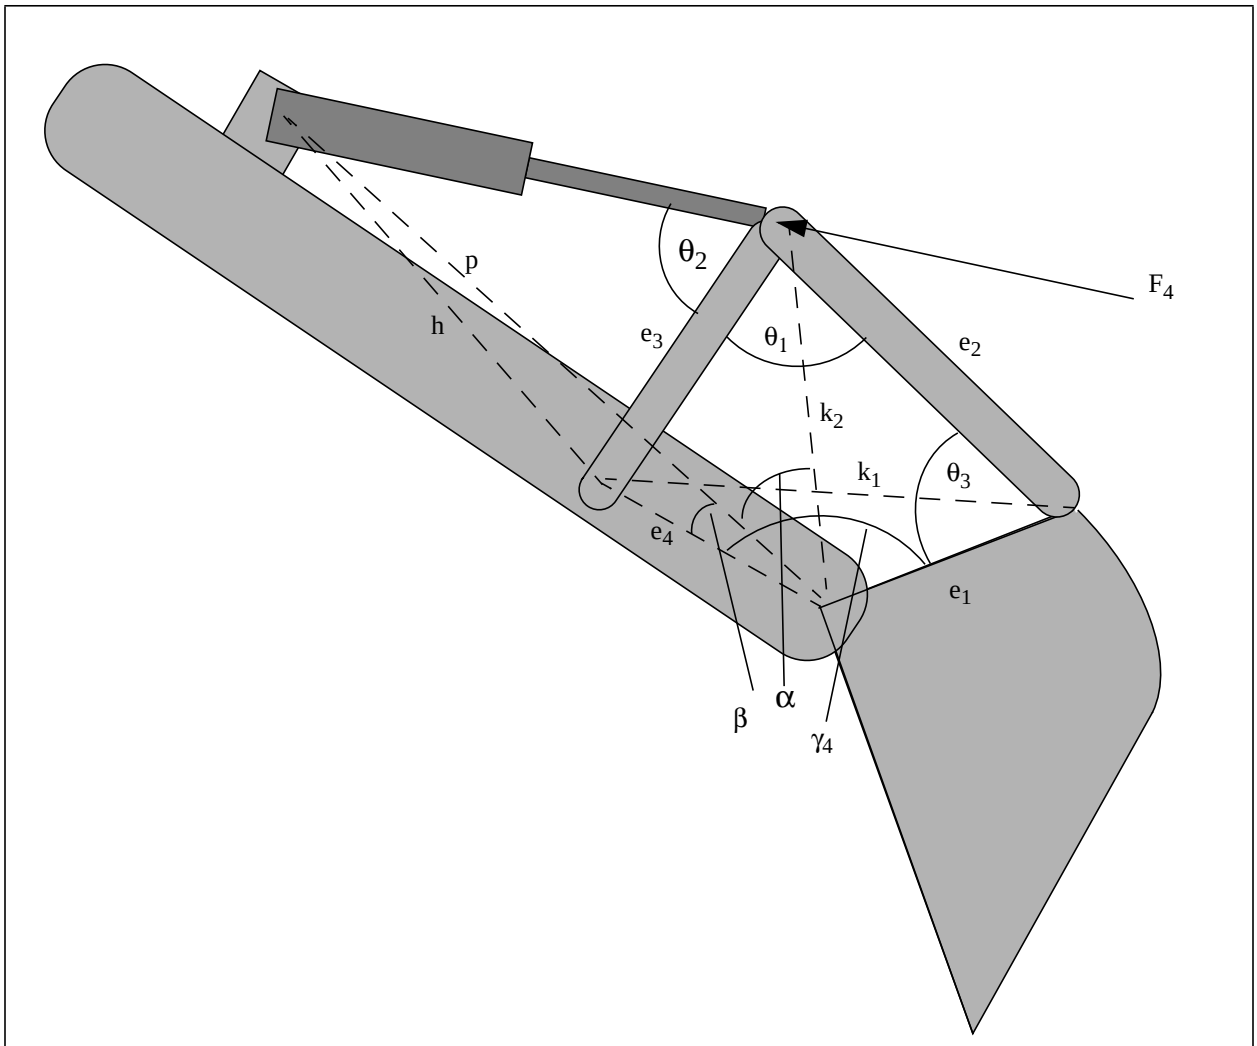

Figure A2: Simplified drawing representing the bucket linkage. The values of e_i correspond to the lengths of the members in the four bar mechanism, γ_4 is known by the bucket joint angle, and F_4 is the actuator force.

k_1, θ_1, θ_3, k_2, α, L_4, and θ_2 can be found through successive use of the law of cosines:

$$k_1 = \sqrt{(e_4)^2 + (e_1)^2 - 2e_4 e_1 \cos\gamma_4} \quad \textbf{(A8)}$$

$$\theta_1 = \operatorname{acos}\left(\frac{(e_3)^2 + (e_2)^2 - (k_1)^2}{2e_3 e_2}\right) \quad \textbf{(A9)}$$

$$\theta_3 = \operatorname{acos}\left(\frac{(e_2)^2 + (k_1)^2 - (e_3)^2}{2e_2 k_1}\right) + \operatorname{acos}\left(\frac{(e_1)^2 + (k_1)^2 - (e_4)^2}{2e_1 k_1}\right) \quad \textbf{(A10)}$$

$$k_2 = \sqrt{(e_1)^2 + (e_2)^2 - 2e_1 e_2 \cos\theta_3} \quad \textbf{(A11)}$$

$$\alpha = \operatorname{acos}\left(\frac{(e_4)^2 + (k_2)^2 - (e_3)^2}{2e_4 k_2}\right) - \beta \quad \textbf{(A12)}$$

$$L_4 = \sqrt{(k_2)^2 + p^2 - 2k_2 p \cos\alpha} \quad \textbf{(A13)}$$

$$\theta_2 = \operatorname{acos}\left(\frac{(L_4)^2 + (e_3)^2 - h^2}{2e_3 L_4}\right) \quad \textbf{(A14)}$$

Now that the geometry is fixed, we are ready to calculate the forces. Note that there are two members attached to the end of the actuator, and these are the members that transmit the force to the bucket. Summing the forces at the end of the actuator in the direction of F_4:

$$F_4 + -F_{e_3}\cos\theta_2 + -F_{e_2}\cos(\theta_1 + \theta_2) = 0 \quad \textbf{(A15)}$$

Summing the forces on the actuator in the direction perpendicular to F_4:

$$F_{e_3}\sin\theta_2 + F_{e_2}\sin(\theta_1 + \theta_2) = 0 \quad \textbf{(A16)}$$

Combining the equations by eliminating F_{e3}:

$$F_4 = F_{e_2}(\cos(\theta_1 + \theta_2) - \sin(\theta_1 + \theta_2)\cot\theta_2) \quad \textbf{(A17)}$$

Finally recognizing that F_{e2} must resist the torque:

$$F_{e_2} = \frac{\tau}{e_1 \sin\theta_3} \quad \textbf{(A18)}$$

The equation for the torque on the bucket actuator becomes:

$$F_4 = \frac{\tau_4(\cos(\theta_1 + \theta_2) - \sin(\theta_1 + \theta_2)\cot\theta_1)}{e_1 \sin\theta_3} \quad \textbf{(A19)}$$

A.2 Actuator Forces to Bucket Forces

In order to extract the soil-tool properties from actuator pressure data, the actuator pressures must be converted into soil forces at the bucket tip. This can be accomplished by inverting some of the equations in A.1. First, the pressures must be converted into an actuator force by multiplying by the cross sectional areas.

$$F_i = P_{i1}A_{i1} - P_{i2}A_{i2} \tag{A20}$$

where P_{ik} is the pressure acting on end k in actuator i, A_{ik} is cross sectional area corresponding to end k of actuator i, and F_i is the actuator force.

These forces can then be converted to torques by inverting equations A7 and A19. For the boom and bucket:

$$\tau_i = {}^{F}{}_i a_i \sqrt{1 - \left(\frac{b_i^2 - L_i^2 - a_i^2}{2L_i a_i} \right)^2} \tag{A21}$$

where the variables are described in Figure A1. And for the bucket:

$$\tau_4 = \frac{F_4(e_1 \sin\theta_3)}{(\cos(\theta_1 + \theta_2) - \sin(\theta_1 + \theta_2)\cot\theta_1)} \tag{A22}$$

where these variables are described in Figure A2, and the values calculated by equations A8 through A14.

Now that the torques are known, the forces can be calculated by inverting equation A2. Note that the Jacobian matrix is full rank and guaranteed to be invertible because the stick angle limit ensures that there are no singularities.

$$f = J^{-T}(\tau - G) \tag{A23}$$

References

[Bernold 93] Bernold, L., "Motion and Path Control for Robotic Excavation", *Journal of Aerospace Engineering*, Vol. 6, No. 1, 1993.

[Bisse 94] Bisse, E., Hemami, A., and Boukas, E., "Optimal Excavation Path Planning for Scooping by a Bucket", in *Proc. of the 6th Canadian Symposium on Mining Automation*, Montreal, QC Canada, October 1994.

[Bradley 93] Bradley, D., Seward, D., Mann, J., and Goodwin, M., "Artificial Intelligence in the Control and Operation of Construction Plant - the Autonomous Excavator", *Automation and Construction*, Vol. 2, No. 3, 1993.

[Bullock 89] Bullock, D. and Oppenheim, I., "A Laboratory Study of Force-Cognitive Excavation", in Proc. *Sixth International Symposium on Automation and Robotics in Construction*, June 1989.

[Bullock 92] Bullock, D. and Oppenheim, I., "Object Oriented Programming in Robotics Research for Excavation", *Journal of Computing in Civil Engineering*, Vol. 6, No. 3, July 1992.

[Cannon 99] Cannon, H. and Singh, S., "Models for Automated Earthmoving", in *Proc. of the 5th International Symposium on Experimental Robotics*, 1999.

[Gill 68] Gill, W. and Vanden Berg, G., *Agricultural Handbook*, No. 316: *Soil Dynamics in Tillage and Traction*, US Department of Agriculture, 1968.

[Hemami 92] Hemami, H., "Study of Bucket Trajectory in Automatic Scooping with LHD Loaders", in *Transactions Institution Mining, Metallurgy*, section A, 102, 1992.

[Hemami 94] Hemami, H., "Modeling, Analysis, and Preliminary Studies for Automatic Scooping", *Advanced Robotics*, pages 1-19, 1994.

[Hettiaratchi 67] Hettiaratchi, D., and Reece, A., "Symmetrical Three Dimensional Soil Failure", *Geotechnique*, Vol. 4, No. 3, 1967.

[Huang 93] Huang, X. and Bernold L., "Robotic Rock Handling during Backhoe Excavation", *Automation and Robotics in Construction*, 1993.

[Koivo 92] Koivo, A., "Controlling an Intelligent Excavator for Autonomous Digging in Difficult Ground", in *Proc. of the 9th International Symposium on Automation and Construction*, Tokyo, June 1992.

[Krishna 99] Krishna, M. and Bares, J., "Constructing Hydraulic Robot Models using Memory-Based Learning", to appear in *ASCE Journal of Aerospace Engineering*, 1999.

[Lawrence 95] Lawrence P., Salcudean, S., Sepehri, N., Chan, D., Bachmann, S., Parker, N., Zhu, M. and Frenette, R., "Coordinated and Force-Feedback Control of Hydraulic Excavators, in *Proc. of the International Symposium on Experimental Robotics*, July 1995.

[Lever 94] Lever, P., Wang, F., and Chen D., "Intelligent Excavator Control for a Lunar Mining System", in *Proc. ASCE Conference on Robotics for Challenging Environments*, Albuquerque, NM, February 1994.

[Lever 95] Lever, P., Wang, F., and Chen, D., "A Fuzzy Control System for an Automated Mining Excavator", in *Proc. International Conference on Robotics and Automation*, 1995.

[Luengo 98] Luengo, O., Singh, S., and Cannon, H., "Modeling and Identification of Soil-Tool Interaction in Automated Excavation", in *Proc. of the International Conference on Intelligent Robots and Systems*, Victoria, B.C., Canada, October 1998.

[Luth 65] Luth, H. and Wismer, R., "Performance of Plain Soil Cutting Blades in Soil", *Transactions of the American Society of Agricultural Engineers*, 1965.

[McKyes 85] McKyes, E., *Soil Cutting and Tillage*, Elsevier, 1985.

[Merritt 67] Merritt, H., *Hydraulic Control Systems*, John Wiley & Sons, 1967.

[Mitchell 97] Mitchell, T., *Machine Learning*, pages 81 - 127, McGraw-Hill, 1997.

[Press 88] Press, H., Vetterling, W., Teukolsky, S., and Flannery, B., *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 1988.

[Reece 64] Reece, A., "The Fundamental Equation of Earth Moving Mechanics", in *Proc. of the Institution of Mechanical Engineers*, 1964.

[Rocke 94] Rocke, D. "Control System for Automatically Controlling a Work Implement of an Earthmoving Machine to Capture Material", US Patent 5528843, 1994.

[Rocke 95] Rocke, D. "Automatic Excavation Control System and Method", US Patent 5446980, 1995.

[Rowe 99] Rowe, P., "Adaptive Motion Planning for Autonomous Mass Excavation", Ph.D Thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, 1999.

[Salcudean 97] Salcudean, S., 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*, Monterey, CA, July 1997.

[Sameshima 92] Sameshima, M. and Tozawa, S., "Development of Auto Digging Controller for Construction Machine by Fuzzy Logic Control", in *Proc. of the Conference of Japanese Society of Mechanical Engineers*, 1992.

[Sarata 93] Sarata, S., "Concept of an Autonomous System for Piled Ore Shoveling", in *Proc. of the International Symposium on Mine Mechanization and Automation*, Lulea, Sweden, June 1993.

[Seward 88] Seward, D., Bradley D., and Brasserie, R., "The Development of Research Models for Automatic Excavation", in *Proc. of the 5th International Symposium on Automation and Robotics in Construction*, 1988.

[Seward 92] Seward, D., Bradley, D., Mann, J., and Goodwin, M., "Controlling an Intelligent Excavator for Autonomous Digging in Difficult Ground",in *Proc. of the 9th International Symposium on Automation and Construction*, Tokyo, June 1992.

[Shi 95] Shi, X., Wang, F., and Lever, P., "Task and Behavior Formulations for Robotic Rock Excavation", in *Proc. of the International Symposium on Intelligent Control*, 1995.

[Shi 96] Shi, X., Lever, P., and Wang, F., "Experimental Robotic Excavation with Fuzzy Logic and Neural Networks", in *Proc. of the International Conference on Robotics and Automation*, 1996.

[Siemens 65] Siemens J., Weber, J., and Thornburgh, T., "Mechanics of Soil as Influenced by Model Tillage Tools", *Transactions of the American Society of Agricultural Engineers*, 1965.

[Singh 95] Singh, S., "Synthesis of Tactical Plans for Robotic Excavation", Ph.D Thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, January 1995

[Singh 97] Singh, S., "The State of the Art in Automation of Earthmoving", *ASCE Journal of Aerospace Engineering*, Vol. 10, No. 4, October 1997.

[Singh 98] Singh, S. and Cannon, H., "Multi-Resolution Planning for Earthmoving", in *Proc. of International Conference on Robotics and Automation*, May 1998.

[Stentz 98] Stentz, A., Bares, J., Singh, S., and Rowe, P., "A Robotic Excavator for Autonomous Truck Loading", in *Proc. of the International Conference on Intelligent Robots and Systems*, 1998.

[Vaha 91] Vaha, P., Skibniewski, M., and Koivo, A., Excavator Dynamics and Effect of Soil on Digging, in *Proc. of the International Symposium on Automation and Robotics in Construction*, Stuttgart, Germany, 1991.

[Vaha 93] Vaha, P. and Skibniewski, M., "Dynamic Model of an Excavator", *ASCE Journal of Aerospace Engineering*, Vol. 6, No. 2, April 1993.

[Wadhwa 80] Wadhwa, D., "Force Prediction Equation Using Distorted Model as an Analog Device", *Journal of Terramechanics*, Great Britain, 1980.

[Yong 77] Yong, R. and Hanna, A., "Finite Element Analysis of Plane Soil Cutting", *Journal of Terramechanics*, 1977.