Robotica (2001) volume 19, pp. 497–511. Printed in the United Kingdom © 2001 Cambridge University Press

An automated digging control for a wheel loader

Paul J. A. Lever*

Advanced Mining Systems Laboratory, Mining and Geological Engineering Department, The University of Arizona, Tucson, Arizona 85721 (USA)
Email: plever@mge.arizona.edu

(Received in Final Form: December 12, 2000)

SUMMARY
An Automated Digging Control System (ADCS) for a wheel loader is developed that utilizes a behavior-based control structure combined with fuzzy logic. This controller exhibits the real-time reactive responses necessary for executing digging tasks in an uncertain, unstructured and dynamic excavation environment. This paper presents field test results of a prototype ADCS that was developed and implemented on a Caterpillar 980G wheel loader. Test results show that the performance of the automated system is comparable to that of an expert human operator in a wide range of excavation situations.

KEYWORDS: Fuzzy behavior control; Automated digging; Robotic excavation

INTRODUCTION

Automating the dig component of the excavation cycle on earth moving machines such as wheel loaders, hydraulic shovels and mass-excavators, and cable shovels, has many potential benefits. Typically, when these machines are used in mining or construction applications they load large quantities of material (soil, rock, etc.) into a fleet of circulating trucks. Here, digging difficulty can vary dramatically and in these difficult digging situations effective loading performance is only achieved by expert operators. The dig time in these situations can double or triple which significantly reduces the output of the machine.

The use of an effective automated digging control system would give every machine operator the capabilities of an expert operator, and generate the following benefits. First, consistent operation over the duration of the shift, since the control system does not get tired or lose concentration. Second, improve machine availability because the controller will always operate the machine within design limits during digging. Third, reduced wheel slippage during digging. However, to achieve these benefits and also operate effectively in the harsh excavation environment, it is important that the design of an automated system meets the following criteria.

The sensors and actuators used should be limited to those currently available on a modern loading machine. For a wheel loader this includes electro-hydraulic actuation of bucket motions, bucket position sensors and measurement of a limited number of drive train parameters. Complex sensing and actuation systems may be prone to failure in the harsh environment. Next, the system should require no input from the operator related to characterizing digging difficulty. This would require operators to make a judgement concerning digging difficulty. In general, the subsurface characteristics of the material to be loaded and its potential interactions with the bucket during digging have the greatest effect on digging difficulty. Human operators cannot see below the surface. Thus, with no operator input the automated system must be able to adjust its digging trajectory by reacting to perceived changes in digging conditions.

Automatic digging control of loading machines is particularly difficult because they operate in dynamic and unstructured environments where conditions are unknown, extremely variable and difficult to detect. On the other hand, expert human operators can achieve sophisticated control of loading machines in these difficult environments. Repeated excavation experiences help the operator to learn machine operational skills and how to adapt their operating modes to the dynamic conditions. The complexities of the interactions between the excavation machine and its environment make it impractical or infeasible to develop mathematical models typically used in traditional control paradigms. Therefore, researchers at the University of Arizona have been developing an excavation control system that utilizes excavation knowledge gathered from skilled human operators. The Control Architecture for Robotic Excavation (CARE) is a hybrid architecture that employs a behavior-based control structure.[1–4] It has reactive control at the lowest level to generate primitive bucket actions, and task planning using finite state machines (FSM) that capture excavation knowledge required for behavior arbitration. Fuzzy logic combined with behavior-based control provide the excavation controller with the real-time reactive response necessary for digging task execution in an uncertain and dynamic environment.

Several years ago, the University of Arizona researchers started a project funded by Caterpillar Inc. to use CARE as the basis to develop, implement and test an Automated Digging Control System (ADCS) on a wheel loader. The implementation platform for the prototype ADCS was a Caterpillar 980G wheel loader (see Figure 1). This wheel loader weighs 29,497 kg, is 9.5 m long, 3.75 m high and has

* Now at: The University of Queensland Experimental Mine, Isles Road, Indooroopilly, Qld 4068 (Australia);
E-mail: p.lever@cmte.org.au

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

a 4.7 m^3 bucket. The criteria listed above were used for the designing ADCS.

In this paper, we show how the CARE approach has been used to develop the prototype Automated Digging Control System on the Caterpillar 980G. The ADCS utilizes only existing production sensors and actuators and has only modest computational needs. The first half of the paper details the control structure of the ADCS, while the remaining sections present data from field tests. These show that the performance of the automated system is comparable to that of an expert human operator over a wide range of excavation sites.

OVERVIEW OF RELATED AUTOMATED DIGGING CONTROL WORK

The many potential applications for automated earth moving systems has attracted a significant amount of research in this area. Typically, research has fallen into two major areas: digging process modeling and planning, and automated digging. A comprehensive summary of the current research in the field is given in Singh.[5] This section concentrates on work related to the automated digging direction.

In general, the simple trajectory planning and control approach is not effective, therefore several researchers measure forces during digging which are used to adjust the digging trajectory. Bullock and Huang[6–8] use these forces to initiate digging trajectory actions when fixed force thresholds are met. These techniques are not effective and often do not fill the bucket in a wide variety of excavation situations. Alternatively, other researchers have selected digging control actions using a set of control rules. The automated excavator (LUCIE) from the University of Lancaster[9–11] is an example of this approach. This excavator plans and attempts to follow an initial digging trajectory, then uses the rule set to react to excavation conditions.

A fuzzy logic controller has been developed by Sameshima *et al*.[12] which controls the actuation of each degree of freedom relative to bucket motion during the digging process. Thus the fuzzy rules are evaluated at each control cycle and joint velocity commands are the weighted output of the rules. The Autodig approach used by Rocke[13–15] uses the actual forces from hydraulic cylinder measurements. These forces are then related to forces inferred from bucket velocities. Commands for each degree of freedom for the bucket are generated from a lookup table based on how human operators control individual joints when digging in various soil conditions. Here, the soil condition must be provided to the system before digging and the material must remain relatively homogeneous for effective digging. Unexpected inclusions in the soil can be a problem for this system to deal with. Cannon[16] implements an augmented Autodig algorithm for dig execution in their Autonomous Loading System (ALS) which completely automates the task of loading trucks with a mass excavator.



Fig. 1. The Caterpillar 980G Wheel Loader Test Platform

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

Another Autodig approach by Shull[17] also uses actual forces measured from the hydraulic cylinders. These forces are used to determine a force vector passing through a point on the bucket that represents the resultant material forces resisting bucket motion. A target angle is also generated on the basis of accumulated energy and then bucket motion commands are generated in response to differences in the target angle and the force vector. This approach can cause the bucket to stall when high resistive forces are encountered in complex bucket rock interaction situations. Dasys *et al*.[18] describe an automated bucket loading system for a Load-Haul-Dump (LHD) machine that also uses feedback from the pressure sensors in the bucket's hydraulic cylinders. This system does not use a model of the muck pile to control bucket motions, but tilts or oscillates the bucket based on the "feel" of the material being loaded. The performance characteristics of this approach are not known.

A position-based impedance control approach for operator assistance during digging with a teleoperated mini-excavator was developed by Salcudean *et al*.[19] at the University of British Columbia. The system follows an operator specified digging path until material resistance impedes its progress, then the impedance controller tries to follow the path as closely as possible. Alternatively, Bernold[20] proposes an impedance controller where an optimal trajectory for the bucket is generated using a planning algorithm based on soil characteristics and soil-tool interaction.

Singh[21] proposes using a trajectory planner that uses a pure position based control system during the digging process. A prediction of forces that will be encountered during digging is used to reject trajectories that exceed the limitations of the hydraulic actuators. Predicting forces in soil with unknown inclusions or in blasted rock is extremely difficult if not impossible.

WHEEL LOADER AUTOMATED DIGGING CONTROL SYSTEM (ADCS)

The general excavation process can be viewed as a system where the excavation goals are the inputs and the output is material movement such as loading haul trucks. For example, at many mining sites the major excavation goal given to loading equipment is very simple. Dig blasted rock from a mine bench and load the rock into mine trucks as efficiently as possible. This requires the loading machine to dig into the bench face and fill the bucket, then move the loaded bucket to a position above the truck, dump, and finally return to a new bench face location to load the next bucket. For most practical purposes the mine bench can be considered an infinite source of rock where the loading machine progressively removes rock from the bench in a systematic manner. Usually, accurate placement and trajectory control of the bucket during digging is not very important when compared to many other types of production systems. For most loading machines the digging part of the operation has the largest variability in performance and requires the highest level of operator expertise. Since, numerous bucket loading cycles are required to complete an excavation goal, significant improvements in excavation productivity can be gained by improving the efficiency of the dig portion of the cycle.

Human operators use different manipulation strategies to fill the bucket, where the strategy selected corresponds to particular excavation conditions in each loading cycle that leads toward the completion of the goal. Here, the machine and excavation site interaction characteristics are critical in defining the bucket motions during bucket loading (digging) and the strategies selected by human operators. However, the operator cannot completely preplan the individual actions required to complete a particular bucket loading task since the operator has no way of knowing what subsurface bucket/material interactions are expected. An operator uses feedback (visual cues, machine vibrations, machine motion characteristics, sound, etc.) from the rock/bucket interaction to adjust machine actions that are most likely to fill the bucket. The operator skills required to recognize, understand and react to rock/bucket interaction feedback during digging can take up to 15 years to acquire. In addition, many operators never acquire the extensive skills required to dig effectively in extremely difficult digging conditions, which often results in over stressing the machine.

With the goal of developing a robotic excavator the University of Arizona researchers have developed and tested a general Control Architecture for Robotic Excavation (CARE) using a simple laboratory-based excavator.[1–4] More recently a research project was completed with Caterpillar, Inc. where the objective was to take the concepts proposed in the CARE system and use them to develop, implement and test an automatic digging control system (ADCS) on a medium sized wheel loader (Caterpillar 980G Wheel Loader). This system would be capable of effciently filling the wheel loader bucket in the extremely diffcult and constantly varying digging conditions encountered in blasted rock mining excavation sites.

Several constraints were also specified for the proposed implementation. First, the system should essentially use only those sensors and actuators currently available on the production wheel loader. No additional sensors are to be used to provide the control system with information about the surface or subsurface of the excavation site. Second, the control system should be as simple as possible and be able to be implemented on computational platforms similar to those currently used on a production wheel loader. Third, the human operator will have control of the machine's throttle (Accelerator Pedal) during the automated digging process. This means that the ADCS cannot directly control horizontal bucket motion during digging. Also, the operator may retake control of the machine at any time during the automated digging process by manipulating the bucket joy sticks in the operator's cab.

Finally, the human operator interfaces with the ADCS in the following way. The loader operator lines the machine up perpendicular to the excavation face and positions the bucket bottom, parallel to and just above the excavation floor. With the machine in the first forward gear the operator initiates the ADCS systems and drives the machine towards the face. When the bucket contacts the face the control system starts executing bucket motions to efficiently fill the bucket. The ADCS determines when the bucket is full, lifts

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

it out of the excavation face and returns control to the operator. Note that during the dig cycle the ADCS must adjust bucket motions to prevent wheel slip, and also prevent undue stress on the machine.

1. Overview of ADCS control architecture

As the bucket moves during digging, the excavation and machine status changes continuously and sensed information is obsolete very quickly. Digging cycles can range from 4 to 20 seconds. This places a serious constraint on the time available for the control system to generate actuator commands. The control system has to respond rapidly to ensure that commands are appropriate to the current excavation situation. This was a major reason for selecting the behavior-based control technique as the basis for control design of the ADCS. Behavior control is a popular technique that has been used for control problems with similar characteristics to those required for rock excavation.[22–24] Complex tasks are decomposed into task specific behaviors where the controller requires no central model or plan, but is a set of concurrently executing behaviors which perform all functions from sensing to action. In addition, behavior control systems usually have some of the properties and components of Reactive or Reflexive Control.

The control architecture for the Automated Digging Control System (ADCS) contains three basic levels as shown in Figure 2. At the highest level is task planning where the digging task is decomposed into Fuzzy Behaviors using finite state machines (FSM). The middle level contains behaviors that are unlike a traditional behavior which executes a single action. The ADCS Behavior's consist of multiple Actions or Action sequences that are feasible to perform the particular task assigned to that Behavior. At the lowest level is the interaction between the control system and the machine/excavation environment.

Based on the inference of the current excavation situation, simple rules in each Behavior select one Action or Action sequence for the next bucket motion. Machine sensor data is input to fuzzy logic rules in the actions to evaluate the magnitudes of machine control variables for

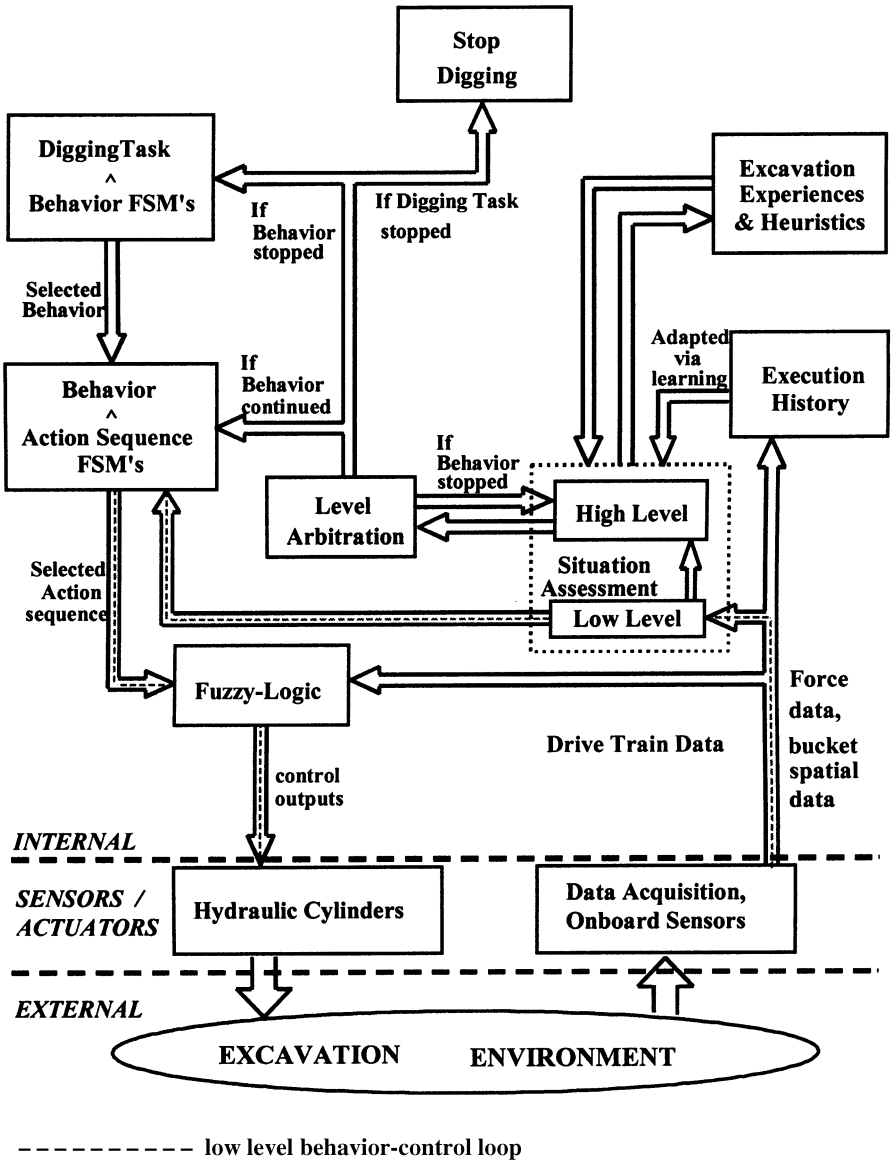


Fig. 2. The control architecture for the Automated Digging Control System (ADCS).

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

each fired action. After each control cycle, information from sensory data is fed to a Situation Assessment (SA) module in the currently executing Behavior. If the SA module determines that the current Behavior is still valid for the current excavation situation, then it is not necessary to interrupt the current behavior and an Action sequence is chosen using the active Behavior. However, if the excavation situation has changed significantly and the current Behavior is no longer correct, then the system jumps to a higher level SA module to select and fire another Behavior, or to generate a more comprehensive validation for continuing to use the current Behavior. If the higher level SA determines the digging task is complete then control is returned to the operator. Details of these control architecture components as they pertain to the ADCS are provided in the following sections.

2. Excavation situation assessment

As indicated earlier, our approach to developing the ADCS was to use only the sensors or actuators currently available on a production style Caterpillar 980G wheel loader. The current implementation of the ADCS meets this goal and therefore does not use any sensors such as a digital camera, laser range finders etc., to characterize the surface of the excavation site. However, during excavation the bucket is in contact with material below the surface whose characteristics are unknown. Data from these bucket/rock interactions provides the only source for information about subsurface material characteristics around the bucket. Moreover, the physical characteristics of the material are unknown until contact is made with the bucket which is further complicated since bucket motion changes these properties.

Excavation situation assessment (ESA) is the process that identifies the current excavation situation based on sensor data. Consider that the ADCS knows nothing about the surface or subsurface characteristics of the material it must excavate before bucket contact. The controller can only use data from the bucket/material interaction and from the associated responses of the machine to generate an appropriate digging bucket trajectory. Therefore, the process of inferring the current excavation situation is critical to effectively completing the digging task. However, the limited number of sensors and the dynamic environment means that information about subsurface material characteristics is incomplete and uncertain. Also, the mapping of sensor data sets to excavation situations is not unique and can be quite complex. For these reasons, situation assessment is not 100% accurate. The distributed ESA approach combined with Behaviors and Actions implemented with fuzzy logic attempts to minimize the effects of these inaccuracies.

Excavation situation assessment is distributed to individual Behaviors in the ACDS. That is, after each bucket motion, situation assessment is limited to examine only the validity of the current behavior. Performing a more complex global situation assessment would be computationally expensive. If after ESA the current behavior is valid, the system flows forward to the Behavior-to-Action FSM via Level Arbitration (LA). Otherwise, more extensive global ESA which uses a larger number of parameters is conducted to select a different Behavior or determine that the Task is complete.

The application of ESA is dependent on specific sensing strategies and the data provided by these sensor systems. The ACDS uses output from several situation assessment sub-modules. Output from these submodules are input to the ESA module which is rule based and derived from excavation heuristics and experiences. The sub-modules include the following:

- Material Force Characterization (MFC) sub-module: Uses bucket/material interaction forces experienced by the machine to define the material characteristics around the bucket.
- Bucket Spatial Analysis (BSA) sub-module: Uses bucket spatial position information as a function of the machine's digging capabilities at different cycle points to define the machine's excavation situation.
- Wheel Slip Analysis (WSA) sub-module: Defines the excavation situation relative to the possibility of wheel slip.
- Bucket Fill Rate (BFR) sub-module: Uses bucket spatial displacement and forces experienced during excavation to define how much material is in the bucket and the rate of material acquisition.

3. Excavation tasks and behaviors

The task decomposition strategy in the original Control Architecture for Robotic Excavation (CARE) is to resolve Goal complexity and narrow task specification to multiple sub-elements by using finite state machines (FSM). At the highest level the system is given an excavation Goal, for example, *excavate material from a bench and load it into trucks*. This Goal is then decomposed into an excavation strategy or strategies that consist of a set of suitable excavation Tasks and Task sequences. As described earlier, the required ADCS implementation for the Caterpillar 980G wheel loader has only one excavation task. The task is "Dig into a mine bench or rock pile and fill the bucket as efficiently as possible". Thus the ADCS requires only a subset of the CARE task decomposition structure which is shown in Figure 3. Here, each excavation Task is completed by a sequence of excavation Behaviors. Depending on sensory feedback, many alternative Behavior sequences can be selected for Task completion, and the proper sequence

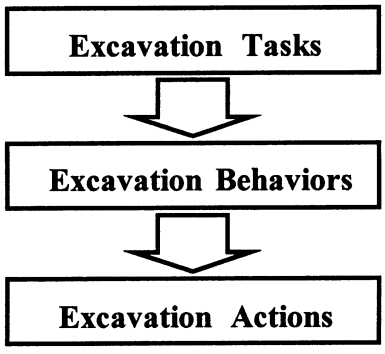


Fig. 3. Control Decomposition from Excavation Task to Excavation Actions.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

can only be chosen on-line in real-time based on situation assessment and Behavior arbitration. The use of Behavior sequences provides a mechanism to prevent ineffective and inefficient task executions caused by factors such as endless looping, or the selection of completely inappropriate Behaviors for the current digging context.

It should be noted that the Behaviors here are the basic elements for excavation control. They represent the ability of the excavator to perform various desired tasks. Existence of the Behaviors is dependent on the particular structure of the machines as well as the control approach, but are independent of the tasks. Once an excavator and its control algorithm are given, these Behaviors exist no matter what digging tasks will be assigned to the machine.

4. *Task formulation*

As indicated above the automated digging system requires only a single task $T_{Autodig}$. Thus if S_T represents the set of all the Tasks required to implemented the automated digging system, then $S_T=\{T_{Autodig}\}$. Each Task is accomplished by a set of excavation Behaviors. To specify all the feasible sequences of Behaviors for Task completion, an FSM is used,

$$T_{Autodig}=(s_i^0, S_i, S_i^f, H_i, \Delta_i), \tag{1}$$

where S_i is the set of states of $T_{Autodig}$, $s_i^0 \in S_i$ is the initial state, $S_i^f \subset S_i$ is the set of terminal states indicating the Task completion, $H_i=\{B_i^1, \ldots, B_i^{m_i}\}$ is the set of Behaviors applicable to $T_{Autodig}$, and $\delta_i: S_i \times H_i \rightarrow S_i$ is the state transition function, i.e., $\delta_i(s_j, B)=s_k$ indicates that if Behavior $B \in H_i$ is executed under the current state s_j, then the next state will be s_k where $s_j, s_k \in S_i$.

Figure 4 shows the FSM for the Task "Remove material covering the surface of a particle", where the Behavior set $H=\{B_1$: *entry-point-mover*, B_2: *under-particle-follower*, B_3: *over-particle-follower*, B_4: *horizontal-digger*, B_5: *bucket-lifter*, B_6: *bucket-extractor*$\}$.

In many cases multiple Behaviors may be applicable at a given state. Let $H_i(s)$ define all feasible Behaviors under state s, i.e.,

$$H_i(s)=\{B \mid B \in H_i \text{ and } \delta(s, B) \text{ is defined}\} \tag{2}$$

Feedback from sensors, and the excavation needs to prevent wheel slip, maximize bucket fill rate and minimize machine stress are the basis for choosing the most appropriate Behavior among $H_j(s)$.

5. *Behaviour formulation*

The behaviors in the Automated Digging System are carried out using a set of primitive machine actions. Let $S_B=\{B_1, \ldots, B_M\}$ represent the set of all the automated digging system Behaviors. A FSM is also used to specify all the feasible sequences of Actions executed by the Behavior,

$$B_i=(s_i^0, S_i, S_i^f, A_i, \delta_i) \tag{3}$$

Where s_i^0 and S_i are the initial state and the set of states of B_i, respectively; S_i^f is the set of terminal states indicating the Behavior completion; $A_i=\{a_i^1, \ldots, a_i^{p_i}\}$ is the set of Actions applicable to B_i, and $\delta_i: S_i \times A_i \rightarrow S_i$ is the state transition function, i.e., $\delta_i(s_j, a)=s_k$ indicates that if Action $a \in A_i$ is executed under the current state s_j, then the next state will be s_k.

When the excavation strategies of expert human operators are represented by FSMs, the event/state selection decisions in the task decomposition can be indeterminate. This is because the decisions human operators make during the execution of excavation tasks are not based on precise thresholds but are often on degrees of eligibility. Thus, multiple events may be eligible simultaneously at a state (i.e. Behavior/Task conflict). In addition, execution of the same event at a state can lead to different states, depending on the execution effect. Clearly, the task-decomposition FSMs requires a flexible decision making approach that involves arbitration of both (1) events and (2) states. This is accomplished by the situation assessment process described previously.

One approach for decomposing tasks into behaviors is to use a large number of small and very focused behaviors. The advantage of this approach is that behaviors execute rapidly and have well defined outcomes. The disadvantage is that behavior arbitration becomes more complex and time consuming. An alternative approach is to use a fewer

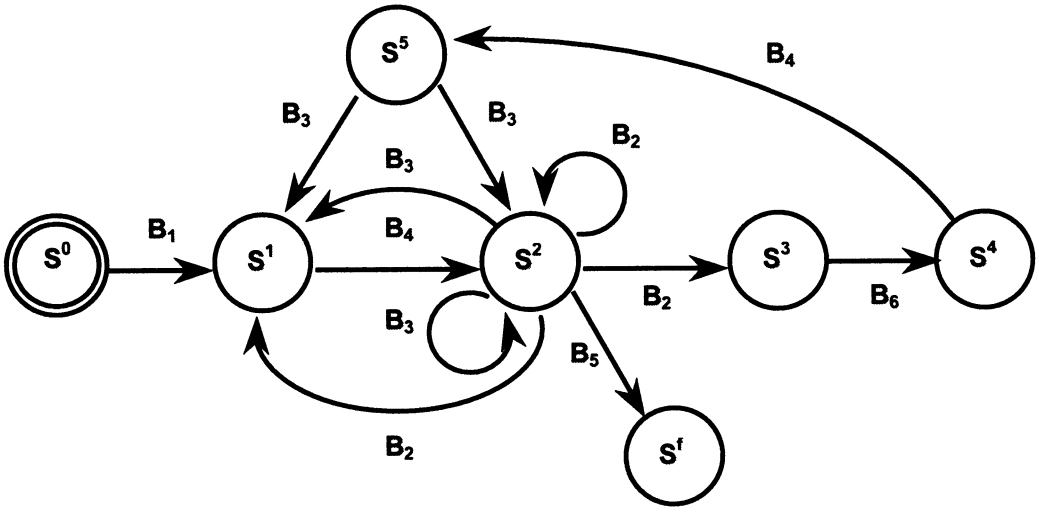


Fig. 4. The finite state machine for the Behavior – *Remove material covering the surface of a particle.*

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

number of behaviors with more extensive capabilities. The ADCS uses the latter construct where the excavation Behaviors have the following features: (i) the functional scope of each Behavior is broad and cannot be achieved using a single machine action; (ii) many terms in the function definitions are ambiguous. Thus, excavation behaviors use multiple Actions or Action sequences in each Behavior to extend their functional scope. Also, they evaluate each excavation Action and assess excavation situations by applying fuzzy logic in order to operate with incomplete and uncertain information. As a result, these Behaviors are very flexible and able to adapt to this dynamic and unpredictable excavation environment. This concept has been key to the successful completion of complicated excavation tasks. It also provides a solution that requires limited computational resources. For example, in the case of the ADCS implemented on the CAT 980G test platform, only a small number of behaviors was required (approximately 5). In addition, each behavior uses a limited number of actions (approximately 5–7) which on average are implemented with between 4 to 7 fuzzy rules.

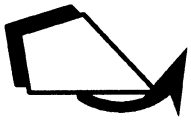
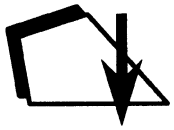
6. Bucket actions

Wheel loaders use hydraulic cylinders to achieve vertical and rotational bucket motions and the vehicles forward and reverse drive for the horizontal bucket motion. Usually, one set of cylinders controls bucket translation in a vertical plane (lift), and a second cylinder set rotates the bucket in the vertical plane (tilt) for a particular orientation as shown in Figure 5. During digging, bucket motions are within a vertical plane parallel to the major axis of the machine. Typically, translation of the bucket in the horizontal and vertical directions are called crowd and hoist respectively, while the rotation is curl. It should be noted that the implement linkage configurations of most wheel loaders do not generate purely vertical or rotation bucket motions when the lift and tilt cylinders are individually activated.

Obviously, digging tasks are formed by a combination of different crowd, hoist and curl motions. Table I lists the fundamental bucket motions used during digging control. Here bucket translations generated by the control algorithm are linear movements of the bucket blade, while rotations are around the blade edge as shown in Figure 5. Also the motion directions are with respect to the local bucket frame. Thus, the same action pattern executed with different bucket orientations may cause different loading effects.

The basic bucket actions listed in Table I specify only the general form of the action trajectory and not the specifics of motion. Thus for these actions to be machine executable the ADCS must specify an action's motion as a function of translational and rotational increments (Δ_y, Δ_z, Δ_r) as well as bucket velocity (BV). However, on the actual wheel loader the hydraulic systems that perform bucket motions are described by very complex dynamic non-linear models that require excessive computation time which makes them impossible to use for practical real-time control. This configuration makes accurate position control of the bucket very difficult and the specification of bucket actions in the form described above cannot be used. For the wheel loader ADCS it is more effective to specify bucket actions as a function of the wheel loaders actuator parameters.

Table I. Basic Bucket Action Motions used by the ADCS.

 A1	Move forward parallel to bucket bottom
 A2	Move forward and upwards, no rotation
 A3	Move forward (up) and rotate upwards
 A4	Move upwards perpendicular to bucket bottom
 A5	Rotate up
 A6	Move downwards perpendicular to bucket bottom
 A7	Move forward and rotate down
 A8	Rotate down

On the 980G wheel loader, separate lift and tilt electro-hydraulic valves control the motion of the lift and tilt hydraulic cylinders. Actuation of the lift and tilt hydraulic cylinders is achieved by specifying a valve position for lift and tilt respectively. Thus, a value of 100% would indicate fully open and 0% would be closed. The onboard electro-hydraulic control system then drives the valve to the required position. As mentioned earlier, the motion of the lift or tilt cylinder resulting from a valve setting request can vary significantly since it is affected by many different factors. These include, engine speed driving the hydraulic pump, concurrent hydraulic fluid draw from other onboard systems, reactive forces on the bucket, and the temperature and pressure of the hydraulic fluid.

Data from two optical encoders on the implement (bucket) assembly are used to resolve the position and orientation of the bucket (relative to the machine frame) and therefore also the extension lengths of the lift and tilt hydraulic cylinders. Head and rod hydraulic cylinder end pressures are measured to determine representative lift and tilt cylinder forces as follows:

$$F_{cyl} = P_h A_h - P_r A_r \tag{4}$$

where, F_{cyl} is the force exerted by the cylinder, P_h, A_h and P_r, A_r are the pressure and area at the cylinder head and rod end

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

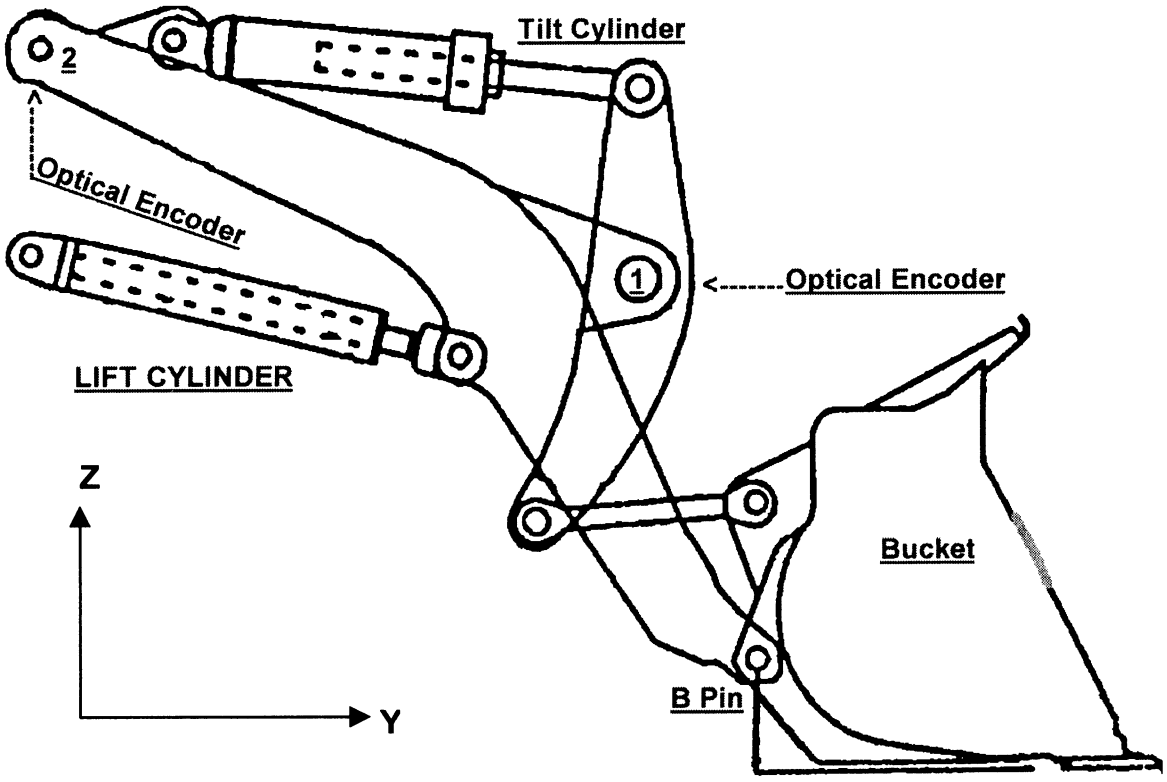


Fig. 5. Bucket manipulation linkage for the CAT 980G Wheel Loader.

respectively. The lift and tilt cylinder forces are then resolved into bucket forces. Drive train speeds are measured and translated into vehicle speed which also generates forward or reverse horizontal displacement and machine position. However, using only drive train measurements the occurrence of wheel slip will give errors in machine position and thus bucket motion.

Therefore, the inputs to the ACDS are:

- Lift and tilt hydraulic cylinder rod and head end pressures (Bucket forces).
- Two joint angles from the implement assembly (Bucket position and orientation relative to frame).
- Drive train angular velocity (Bucket horizontal position) and the ACDS outputs are:
- Electro-hydraulic valve position of the lift and tilt hydraulic cylinders (valve openness).

Note that the human operator controls the engine throttle. Therefore the ACDS does not control loader forward motion and the magnitude of the crowd force during digging.

The ACDS uses fuzzy logic to determine the magnitudes of basic bucket actions. That is, since the ACDS is not capable of direct bucket position control, fuzzy logic is used to generate crisp values for the lift and tilt electro-hydraulic valve positions. This approach allows the system to function with uncertain, inaccurate and dynamic information about both the excavation site characteristics and bucket/rock interaction situation Also, fuzzy logic is an ideal tool to generate bucket Actions where ambiguous and incomplete linguistic information describes the knowledge and strategies of human experience and heuristics through fuzzy rules.

In general, an Action A_i can be implemented as a set of fuzzy decision rules as,

Rule 1 – IF B is S_{i1} and F is Q_{i1} and ΔF is ΔQ_{i1}, THEN M is U_{i1},

•
•
•

Rule n_i – *IF* B is S_{in_i} and F is Q_{in_i} and ΔF is ΔQ_{in_i}, THEN M is U_{in_i}

where n_i is the number of rules for A_{i}, and S_{ij}, Q_{ij}, ΔQ_{ij}, and U_{ij} are linguistic terms for bucket position/orientation, force data and their variation, and electro-hydraulic valve setting commands respectively. B is a bucket orientation vector; F, ΔF and M are the force vector, the force variation vector, and the valve setting vector, respectively:

$$F=(F_{\mathrm{x}}, F_{\mathrm{y}})$$
$$\Delta F=(\Delta F_{\mathrm{x}}, \Delta F_{\mathrm{y}})$$
$$M=(\Theta_L, \Delta_T)$$
$$B=(\Theta_B, Z_B, Y_B).$$

In M, Θ_L and Θ_T are the desired lift and tilt valve positions, respectively. While in B, Θ_B is the angular deviation of the bottom of the bucket from the horizontal and Z_B and Y_B are the respective vertical and horizontal position on the lower pivot point of the bucket in the loader's coordinate system. The variations in force readings are utilized to detect collisions between the bucket and rock particles, since it is difficult to differentiate between a contact with an oversized particle and general resistance to bucket motion by using force values alone.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

As an example, the basic bucket excavation Action "rack the bucket" is defined by the following set of fuzzy-logic rules:

Rule 1 – IF F_y is *L*, THEN Θ_T is *PL*.
Rule 2 – IF F_y is *M*, THEN Θ_T iS *PM*.
Rule 3 – IF F_y is *S*, THEN Θ_T is *PS*.
Rule 4 – IF F_y is *ZR*, THEN Θ_T is *PS*.
Rule 5 – IF ΔF_T is *NL*, THEN Θ_L is ZR.
Rule 6 – IF ΔF_L is *NM*, THEN Θ_L is *PS*.

where (PL, PM, PS, ZR, NL, NM, NS) is a set of linguistic terms (Positive Large, Positive Medium, Positive Small, Zero, Negative Small, Negative Medium, Negative Large).

6. Excavation behaviors using fuzzy actions

In the ADCS excavation Behaviors use FSMs to organize fuzzy Actions using the same approach used to formulate task FSMs. As an example, the Behavior under-particle-follower is built using the FSM in Figure 6.

Our observation indicates that several basic bucket digging actions, or combinations thereof, are executed to cover most of the different conditions encountered during excavation. Based on their perception of the excavation situation, human operators use sets of action sequences compiled from the basic actions presented in Table I. Table II presents typical action sequences selected by several different operators. Here, we can see that the action sequence for the *easy-to-dig* rock is actually a simple variation of the smooth scooping action executed when digging in unconsolidated soil or sand. If the material is difficult to load, the operator often tries an alternate attack point by lifting the bucket (*move_vertical* action) to penetrate further into the rock pile. However, when the muck is very difficult to dig, the bucket gains only a small increase in load with each forward action when using action sequence 2. This results in a small bucket load when the hoist limit or the top of the pile is reached. Thus, action sequence 3 is applied which consists of *moveforward_down* and *moveforward_up* actions to penetrate the rock pile by wiggling in between the rock particles.

The ADCS uses simple rules for on-line arbitration between particular Actions/Action-sequences based on the functional definition of the Behavior, current force data, bucket spatial status, wheel slip status and bucket fill rate. A particular Action sequence is activated each time a Behavior

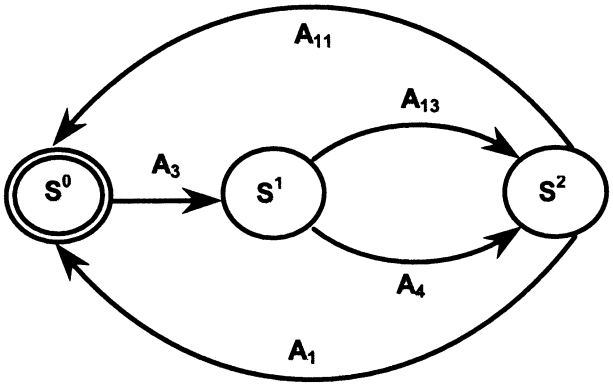
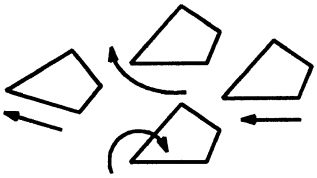


Fig. 6. FSM for the Behavior *under-particle-follower.*

Table II. Bucket excavation actions relative to excavation situation.

Excavation	Action
easy to dig	 (1)
difficult to dig	 (2)
extremely difficult to	 (3)

is executed. It is noted that, sometimes it is difficult to arbitrate definitely that a particular Action sequence will be better than another because of the ambiguity in the inferred excavation status. Thus, it would be more appropriate to fuse all suitable Actions to find the resultant of the motion control variables rather than simply select a single Action that has the highest "suitable rate". However, the Actions to be fused cannot be contradictory in function. For instance, if they have two totally opposite movements then they cannot be fused together since the resultant action may be neutral with no step towards the goal.

A fusion procedure can be executed in two steps. First, the activation level of each admissible Action is calculated according to the firing strengths of its rules, and a threshold is used to remove any Action with an activation level below the threshold from its participation in the action fusion. Second, a single numerical motion Action is obtained by fusing individual decisions made by those Actions whose activation levels are higher than the threshold. The resulting motion action is actually the sum of individual Actions weighted by their activation levels.

The activation level ω_i of an excavation Action A_i at a particular state is obtained using the firing strength of its decision rules. That is,

$$\omega_\iota = \max\{w_{\iota 1}, \ldots, w_{\iota n_\iota}\} \tag{5}$$

where w_{ij}'s are the strengths of the individual rules of Action A_i:

$$\begin{aligned} \omega_{\iota 1} &= \min\{\mu_{S_{\iota 1}}(F), \ldots, \mu_{\Delta S_{\iota 1}}(\Delta F)\}, \\ &\ldots \\ \omega_{\iota n_\iota} &= \min\{\mu_{S_{\iota n_\iota}}(F), \ldots, \mu_{\Delta S_{\iota n_\iota}}(\Delta F)\}, \end{aligned} \tag{6}$$

where *m* represents the membership function of a fuzzy set.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

If the activation level of an Action is less than a threshold ω_c, then this Action will not be considered for action fusion. In other words, if all the preconditions of its rules cannot be matched strongly by the current state, this Action is not suitable under the current situation, and should not be included in the action fusion. All Actions with activation levels above the threshold ω_c will participate in action fusion. A simple action-fusion strategy is to use the sum of all individual decisions weighted by the corresponding activation level. In addition, preference values can also be introduced to discriminate between Actions at each state. If the center of area (COA) defuzzification algorithm is used for each fuzzy Action, then the individual motion command M_i^* from a qualified Action $A_i(\omega_i > \omega_c)$ is calculated by

$$M_i^* = \frac{\int m\mu_{U_i}(m)dm}{\int \mu_{U_i}(m)dm} \tag{7}$$

$$\mu_{U_{ik}^e}(m) = \omega_{ik}\mu_{U_{ik}}(m),\ \mu_{U_i}(m) = \max\{\mu_{U_{i2}^e}(m),\ \mu_{U_{i2}^e}(m),\ \ldots,\ \mu_{U_{in_i}^e}(m)\} \tag{8}$$

where the integration is carried out over the space of all admissible motion commands. Finally, the fused motion command M^* is evaluated by

$$M^* = \frac{\sum \omega_i M_i^*}{\sum \omega_i}, \quad \omega_i > \omega_\chi \tag{9}$$

where the summation is over all fuzzy excavation Actions with the activation level higher than the given threshold.

EXPERIMENTAL RESULTS

The proposed automated digging control approach has been implemented as a prototype on a Caterpillar 980G wheel loader at the University of Arizona San Xavier Mining Laboratory (SXML). The performance of the digging control system was evaluated by comparing it to the performance of expert human operators for a wide variety of excavation environments. This was accomplished by selecting digging situations that represent typical operating conditions for a wheel loader at mining, quarrying and construction sites. To meet the wide range of digging conditions needed to test the performance of the automated digging system, several mine and test locations close to the UA SXML were chosen. These locations include the ASARCO Mission Mine Complex, the Cyprus-Amax Sierrita Mine, and Caterpillar's Tucson Proving Grounds (TPG). Detailed testing was completed in material piles and mine bench faces with the following material characteristics, (1) loose top soil, (2) compacted clean soil, (3) 2″ crushed rock; damp, with finer mixtures, (4) 8″ clean rock. (5) compacted damp mixture of rock & soil, (6) compacted soil with rock, (7) rock bank with large, medium & small sized rock, (8) oversized boulders in finer soil mixture, (9) piled blasted-rock, (10) mine bench with blasted rock. Results from only a select number of test sites are presented here.

1. Test procedure

Test sites were carefully selected to ensure that the digging conditions at each test site (or difficulty) were similar during automated and expert human operator digging cycles. This can be a problem at mining sites where the wheel loader is digging into a near vertical face of blasted material called a mining bench. Here, digging difficulty can be extreme when digging into a standing bench face; however if the bench face has subsided then easy to dig material collects at the toe of the bench. This means that the loader can encounter loose easy to dig material during one digging cycle and then tightly coupled rock particles that are diffcult to dig on the very next cycle.

At each test site 30 digging cycles were completed under automated digging control and then a further 30 digging cycles with the expert human operator. After, each dig cycle the material in the bucket was dumped away from the face to ensure that each dig cycle was in undisturbed material. Every attempt was made to ensure that the automated system and the human operator encountered the same distribution of digging conditions over their 30 dig cycles. To access machine performance a wide array of data was collected from onboard sensors during each dig cycle. These data included wheel loader cycle times, kinematics, hydraulic cylinder and system pressures, drive train speeds and torques, and engine parameters. The onboard payload monitoring system determined the bucket payload from each dig cycle and as a check every fifth bucket payload was checked using external vehicle scales. The sampling rate for all data was 10 Hz. It should be noted, that the exact same code for the CAT 980G automated digging system prototype was used for all the automated digging tests across all material types.

2. Force and trajectory analysis results

The automated digging system has no prior information about the material or the configurption of the digging site before the dig cycle begins. Machine data acquired from bucket/material interaction during digging is the only source for this information. However, at many mining sites the digging characteristics can change dramatically from one dig cycle to the next, and the lab studies have shown that the location of one medium size particle can significantly effect the digging difficulty. Therefore, specific information gathered from previous digging cycles cannot be used to plan future digging trajectories. As discussed previously the solution is a reactive system that combines data collected at the bucket/material interface with specific digging knowledge to adjust the digging trajectory. Thus, the prototype automated digging system has no preplanned digging trajectory before starting the dig cycle and the actual dig trajectory executed is a representation of that particular cycle's digging difficulty. Figure 7 below demonstrates the variation in digging trajectories executed by the automated digging system in consecutive dig cycles at two different sites (Blasted rock production bench face and 8 inch rock pile). Figures 7 and 8 also show the wide variation in bucket trajectories in the same material for consecutive dig cycles.

As described earlier, situation assessment is the process that relates forces experienced at the bucket/rock interface, to excavation situations that drive the digging trajectory. When digging in soft sand the forces experienced by the bucket during the dig cycle are quite predictable. Therefore,

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

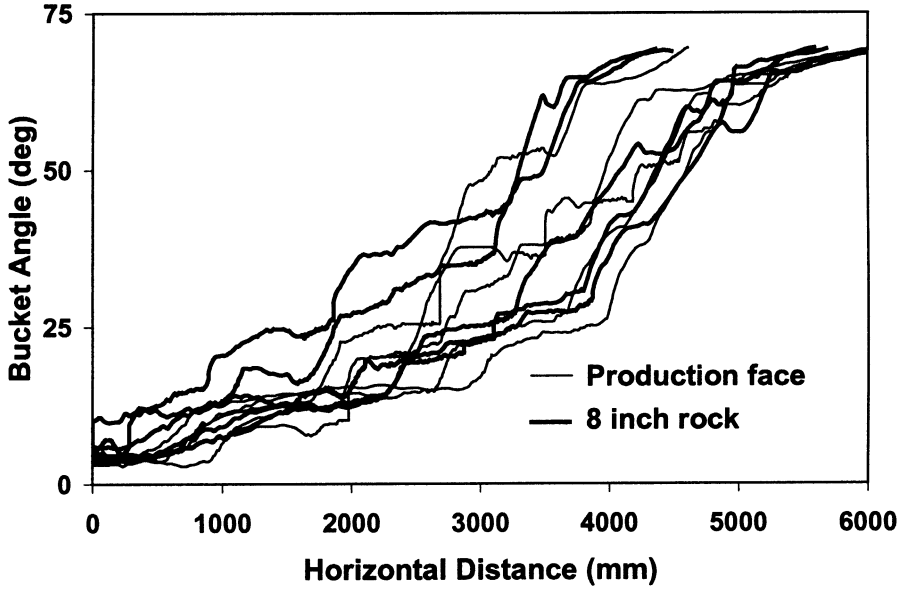


Fig. 7. Change in bucket angle orientation for 5 consecutive automated digging tests in 8 inch rock pile and mine production bench face.

the excavation situations determined by the automated system during digging are consistent and as expected the resulting dig trajectories are similar. However, when digging in rock piles with a larger homogenous particle size such as the 8 inch rock pile test site, the location of one, or a small set of interacting particles can dramatically affect the perceived excavation situations. Since real-time sensor systems that can accurately map the subsurface of the rock pile do not exist, it is impossible to predict when these excavation situations will occur. The automated system must be able to detect these excavation situations and adjust the digging trajectory to effectively fill the bucket. It should be noted that the bucket width on the 980G is 137 inches and an 8 inch wide particle is only $1/15^{th}$ of the bucket width. Figure 9 demonstrates how the situation assessment system is able to detect a change in the excavation situation and adjust the bucket trajectory. At the start of the cycle the horizontal forces experienced by the bucket is very similar until point A where test 105 experiences a dramatic force increase. The automated system reacts to adjust the dig trajectory which begins to occur at point B. Similar trajectory changes also occur later in both dig cycles.

Alternatively Figure 10 demonstrates how the vertical forces experienced by the bucket for two different dig cycles in the same material (8 inch rock pile) are similar and the bucket trajectories are very close.

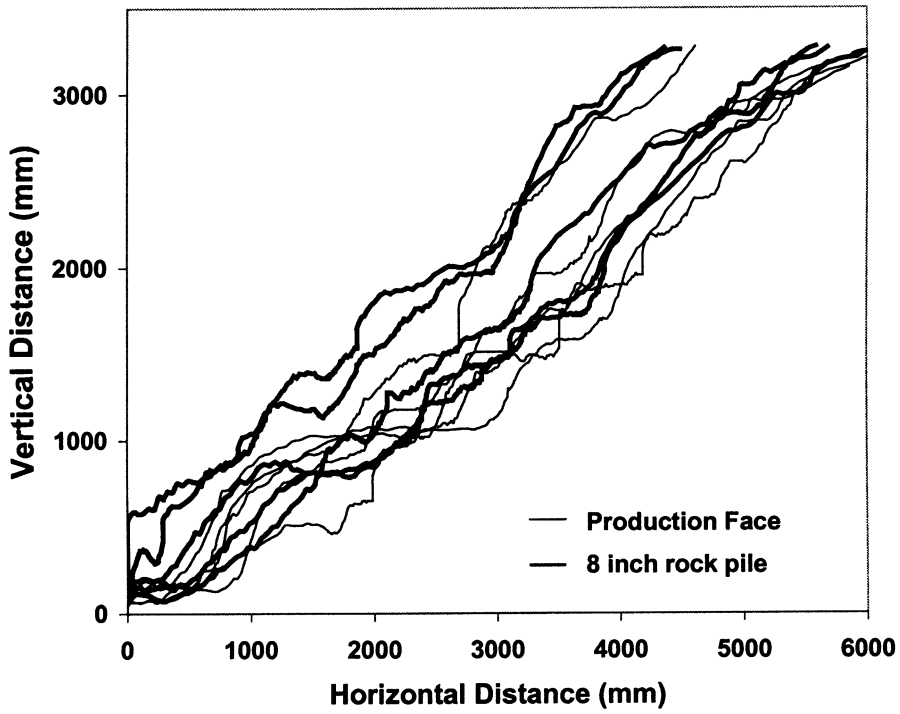


Fig. 8. Bucket trajectory for 5 consecutive automated digging tests in 8 inch rock pile and mine production bench face.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

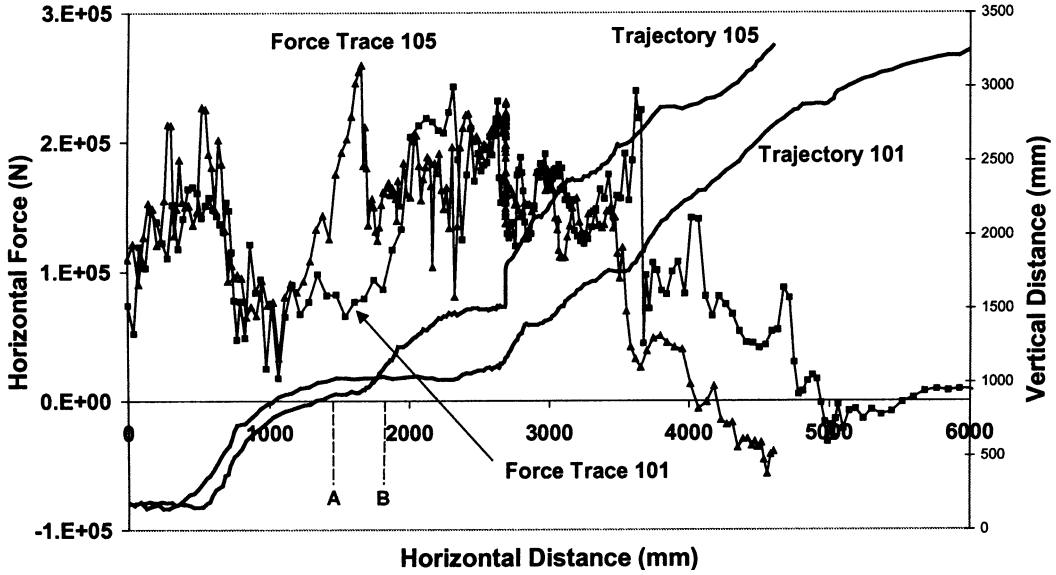


Fig. 9. Bucket trajectory and horizontal bucket forces for two dig cycles (101 and 105) at a mine production bench face.

Finally, Figure 11 shows the vertical bucket forces and bucket trajectories from two automated dig cycles in the 8 inch rock pile. In dig cycle 102 the vertical bucket forces experienced at the start of the cycle are low, enabling a trajectory where the bucket penetrates horizontally into the rock pile quickly. Conversely, in dig cycle 105, high initial vertical bucket forces indicate difficulty penetrating the rock pile and the bucket starts searching for a location to penetrate more deeply into the pile (See point C). When this is not successful the bucket then follows a shallower trajectory up the surface of the rock pile.

This type of significant differences in digging trajectories followed by the ADCS can be seen in Figures 12 and 13 below. Here, both trajectories are from dig cycles at the same location within a few cycles from each other. Drastic changes in bucket trajectories within close digging proximity of each other are typical of conditions found in mines with blasted rock. The bucket trajectory in Figure 12 shows the typical smooth digging arc that buckets follow when resistance to motion is within the machine's force limits. On the other hand, Figure 13 shows a more complex trajectory where the bucket encounters significant resistance to motion at different points in the cycle. At point A the bucket encounters a hard toe at the start of the dig cycle and the bucket is lifted immediately. The bucket then starts to penetrate the pile at point B and in fact digs downwards to increase the bucket fill rate. Another resistance point is encountered at point C with an appropriate change in the bucket's trajectory. Additional high dig resistance situations are also encountered during the completion of the cycle.

3. Productivity comparisons with expert human operators

The performance of the ADCS on the Caterpillar 980G Wheel Loader was evaluated relative to the performance of an expert operator in a wide range of excavation site

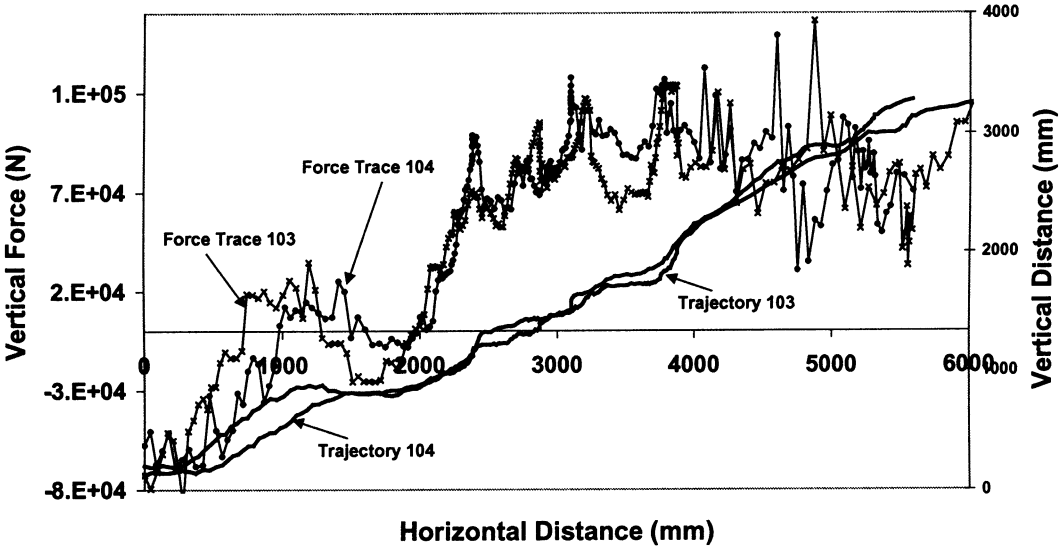


Fig. 10. Bucket trajectory and vertical bucket forces for two consecutive dig cycles (103 and 104) in a rock pile with 8 inch rock particles.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

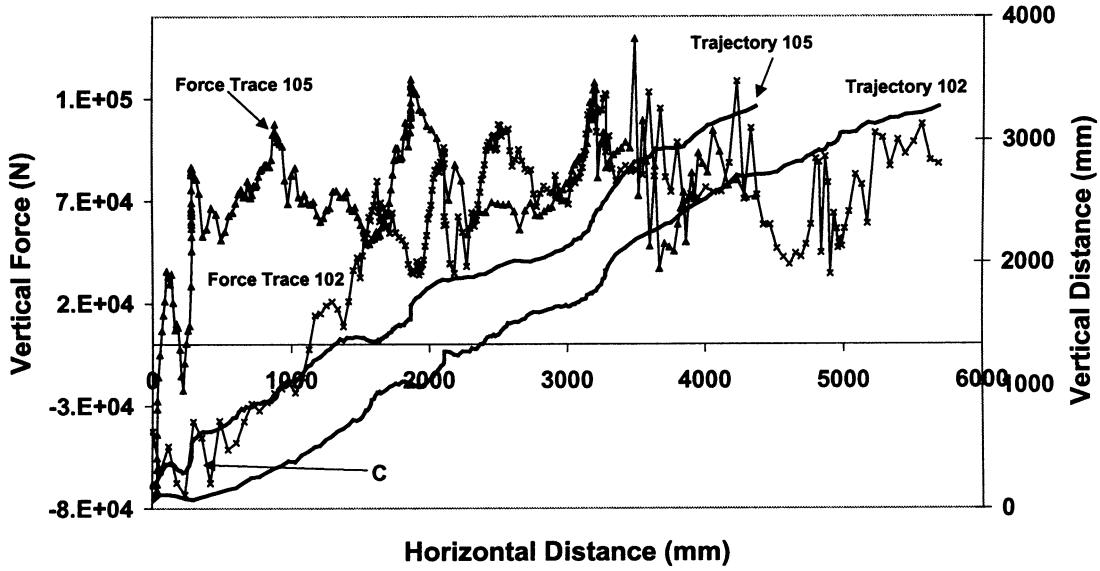


Fig. 11. Bucket trajectory and vertical bucket forces for two consecutive dig cycles (102 and 105) in a rock pile with 8 inch rock particles.

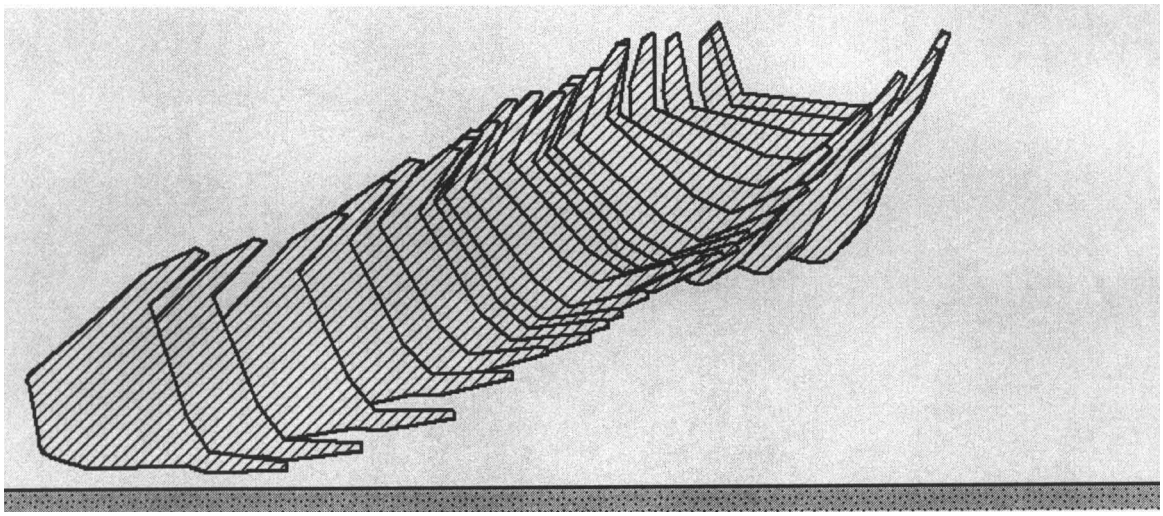


Fig. 12. A simple dig trajectory in a production mine bench with blasted rock.

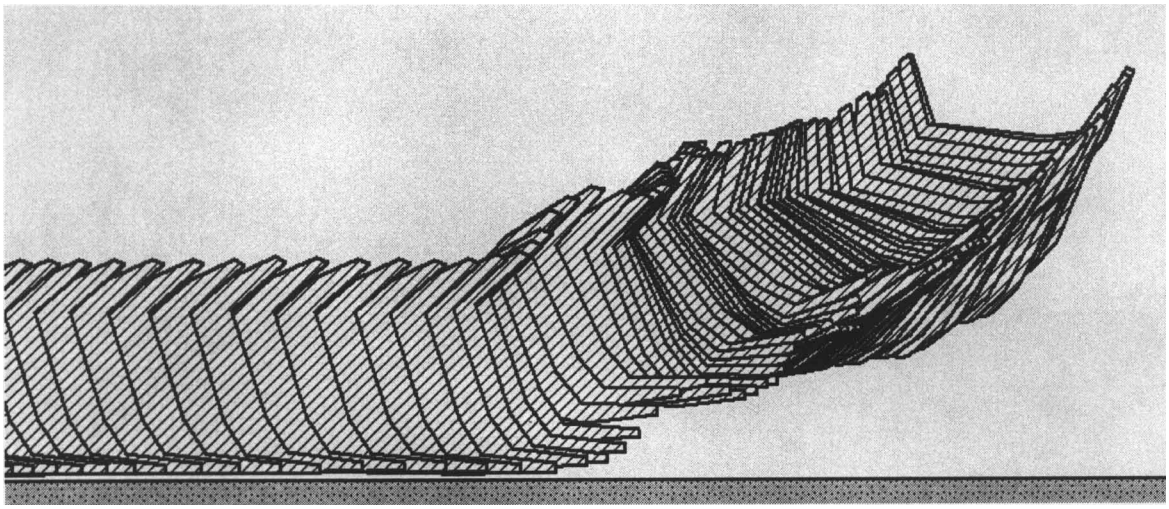


Fig. 13. A complex dig trajectory in a production mine bench with blasted rock.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

Table III. ADCS vs Expert Operator Caterpillar 980G Wheel Loader Productivity Test Data.

	8 inch Rock Pile		Mine Production Bench		Shot Rock	
	ADCS	**Expert Operator**	**ADCS**	**Expert Operator**	**ADCS**	**Expert Operator**
Ave. Payload (lbs*1000)	15.9	17.1	14.8	16.8	13.4	13.9
Maximum	18.0	19.2	19.2	17.4	16.3	15.4
Minimum	10.2	15.7	8.8	15.8	9.8	12.1
Standard Deviation	1.45	0.8	2.6	0.9	1.6	1.2
Energy/Payload(ft.lbs/lbs)	29.3	21.9	29.6	29.76	36.63	37.12
Ave. Dig time (seconds)	10.18	6.2	9.35	11.8	11.4	10.4
Maximum	14.25	12.8	15.5	14	16	11.7
Minimum	4.5	3.5	5	9.5	8.5	8.5
Standard Deviation	2.35	2.2	2.16	2.2	2.2	1.2
Est. Prod./hr (tons/hr)	1,017	1,272	971.2	1,015	815	881

conditions. To test the versatility of the ADCS, the prototype code was not changed across all the tests, and no external input information was provided to the system other than sensor data collected during automated digging tasks. Thus, no input information relating to the characteristics of each excavation site was given to the system.

Table III presents productivity data for the Caterpillar 980G wheel loader from three test excavation sites with a wide range of digging difficulty. The 8 inch rock pile is generally the easiest to dig, with the production mine bench rated as difficult digging, and the blasted rock site as extremely difficult digging. Each test consists of around 30 dig cycles (30 for the ADCS and 30 for the Expert Operator). Note, to generate production output in tons/hr, constant times were used for the other truck loading components of the cycle.

The data in Table III demonstrates that the ADCS performed very well relative to the expert human operator. As the digging conditions become more diffcult the ADCS tended to match the operator's performance more closely. In addition, the expert human operator understands he is being evaluated and therefore attempts to perform at the top of his capabilities for the 30 test cycles. In reality, a human operator's performance is expected to vary considerably over the duration of an 8 hour shift as they tire and lose concentration. However, the ADCS does not get tired and thus will maintain its performance level over the shift duration.

The difference in performance seen in the 8 inch rock pile test can be attributed to controller communication bandwidth issues related to the hardware used for the prototype ADCS. Currently, the prototype can only acquire sensor data and execute an actuator command every 50 ms. Our tests indicate that this bandwidth problem with the hardware implementation of the ADCS prototype means the system cannot react fast enough to execute the required short cycle times.

Further examination of the data in Table III indicates that although the ADCS and expert operator may generate similar production outputs in tons per hour, they attain those outputs in different ways. For example, in the mine production bench test the energies per payload are identical however the ADCS has shorter on average cycle times and smaller payloads. In principle, the ADCS should perform better than an expert operator because the operator is not capable of sensing the bucket forces experienced by the loader in a quantifiable manner. However, the expert operator is using visual information about the surface characteristics of the excavation site. Our analysis of the production data in Table III combined with the force feedback and bucket trajectory data, indicates that the addition of limited excavation site surface data to ADCS could yield significant improvements in performance.

CONCLUSIONS

This paper presents some initial results for a behavior-based control architecture that uses fuzzy logic in the formulation of behaviors and actions to execute automated digging tasks. Excavation tasks are implemented as finite state machines that define feasible sequences of behaviors for their completion, and feasible fuzzy action sequences to complete the behaviors. This approach is goal-directed and utilizes effective behavior fusion through the inherent state memory characteristics of FSM to place behaviors and actions in the appropriate task context. This control architecture was used to implement an automatic digging control system (ADCS) prototype on a Caterpillar 980G wheel loader. Test results in a wide variety of excavation conditions show that the ADCS is able to react to changes in excavation situations, adjust its digging trajectory and fill the bucket efficiently. Without any prior input about excavation site conditions the system generates effective dig trajectories that account for variations in digging difficulty at each site and when moving from site to site.

Field tests show that the ADCS can be as productive as an expert human operator performing at their highest level. From this it can be inferred that the automated system will be more productive than the average operator found at most excavation sites. The ADCS has additional benefits which include expert operator digging performance for all operators, consistent operation over the duration of a shift,

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press

increased machine availability since ADCS maintains operating modes within design limits, and reduced wheel slip during digging.

REFERENCES

1. X. Shi, P. J. A. Lever and F. Y. Wang, *Autonomous Rock Excavation – Intelligent Control Techniques and Experimentation*, Series in Intelligent Control and Intelligent Automation, **Vol. 10** (World Scientific Press, Singapore, 1998).
2. P. J. A. Lever, F. Y. Wang, D. Chen and X. Shi, "Autonomous Robotic Mining Excavation using Fuzzy Logic and Neural Networks", *J of Intelligent and Fuzzy Systems* **3**, 31–42 (1995).
3. X. Shi, P. J. A. Lever and F.Y. Wang, "Experimental Robotic Excavation with Fuzzy Logic and Neural Networks", *Proc. of IEEE Intl. Conf. on Robotics & Automat.*, Minneapolis, Minn. (1996) **Vol. 1**, pp. 957–962.
4. X. Shi, P. J. A. Lever and F. Y. Lever, "Fuzzy Behavior Integration and Action Fusion for Robotic Excavation", *IEEE Trans. on Industrial Electronics* **43**, No. 3, 395–402 (1996).
5. S. Singh, "The State of the Art in Automation in Earthmoving", *ASCE Journal of Aerospace Engineering* **10**, No. 4, 179–188 (October, 1997).
6. D. Bullock and I. Oppenheim, "A Laboratory study of Force-Cognitive Excavation", *Proc. Sixth International Symposium on Automation and Robotics in Construction* (June, 1989).
7. D. Bullock and I. Oppenheim, "Object Oriented Programming in Robotics Research for Excavation", *Journal of Computing in Civil Engineering* **6**, No. 3, 370–385 (July, 1992).
8. X. Huang and L. Bernold, "Control Model for Robotic Excavation and Obstacle Handling", *Robotics for Challenging Environments, Proc of ASCE Specialty Conf.* (1994) pp. 123–131.
9. D. Seward, D. Bradley and R. Brasserie, "The Development of Research Models for Automatic Excavation", *Proc. of the 5th International Symposium on Automation and Robotics in Construction* (1988).
10. D. Seward, D. Bradley, J. Mann and M. Goodwin, "Controlling an Intelligent Excavator for Autonomous Digging in Difficult Ground", *Proc. of the 9th International Symposium on Automation and Robotics in Construction* (1992).
11. D. Bradley, D. Seward, J. Mann and M. Goodwin, "Artificial Intelligence in the Control and Operation of Construction Plant – the Autonomous Excavator", *Automation and Construction* **2**, No. 3, 217–228 (1993).
12. M. Sameshima and S. Tozawa, "Development of Auto Digging Controller for Construction Machine by Fuzzy Logic Control", *Proc. of the Conference of Japanese Society of Mechanical Engineers* (1992) pp. 271–287.
13. D. Rocke, "Control System for Automatically Controlling a Work Implement of an Earthmoving Machine to Capture Material", *US Patent 5528843* (1994).
14. D. Rocke, "Automatic Excavation Control System and Method", *US Patent 5446980* (1995).
15. D. Rocke, "System and Method for Automatic Bucket Loading using Crowd Factors", *US Patent 5968103* (1999),
16. H. Cannon, "Extended Earthmoving with an Autonomous Excavator", *M.S. Thesis* (Robotics Institute, Carnegie Mellon University, l999).
17. A. G. Shull., "System and Method for Automatic Bucket Loading using Force Vectors", *US Patent 5974352* (1999).
18. A. Dasys., L. Geoffroy, and A. Drouin, "Sensor Feedback Control for Automated Bucket Loading", *US Patent 5941921* (1999).
19. S. Salcudean, S. Tafazoli, P. Lawrence and I. Chau, "Impedance Control of a Teleoperated Mini Excavator", *Proc. of the 8th IEEE International Conference on Advanced Robotics* (1997).
20. L. Bernold., "Motion and Path Control for Robotic Excavation", *Journal of Aerospace Engineering* **6**, No. 1, 1–18 (1993).
21. S. Singh., "Synthesis of Tactical Plans for Robotic Excavation", *Ph.D. Thesis* (Robotics Institute, Carnegie Mellon University, 1995).
22. R. A. Brooks, "A Robust Layered Control System For A Mobile Robot", *IEEE Journal of Robotics and Automation* **R-A-2**, No. 1, 14–23 (1986).
23. E. Gat, R. Desai, R. Ivlev, J. Loch and D. P. Miller, "Behavior Control for Robotic Exploration of Planetary Surfaces", *IEEE Transactions on Robotics and Automation* **10**, No. 4, 490–503 (1994).
24. M. Watanabe, K. Onoguchi, I. Kweon and Y. Kuno, "Architecture of Behavior-based Mobile Robot in Dynamic Environment", *Proceedings of IEEE Conference on Robotics and Automation* (1992) pp. 2711–2718.

https://doi.org/10.1017/S0263574701003435 Published online by Cambridge University Press