Expert Systems for Planning Robotic Excavation

S.J. Fenves
N. Baker
J. Balash

1. INTRODUCTION

Shallow trench excavations have typically been performed by a simple backhoe. However, when this process is used to uncover leaking gas utilities, a high degree of danger exists to anyone close to the work area. An automated process to perform this excavation would reduce the element of danger to those making the repair.

REX, a Robotic EXcavator, is a tool being designed by Carnegie-Mellon University and the Dravo Corporation to excavate around buried pipes and other obstructions. REX will be truck-mounted and will consist of at least one robotic arm that allows controlled movement of scanning and excavating tools. REX is intended to be used primarily for performing excavations where an element of danger exists to workers. By using REX, workers can remain away from the excavation site and thus avoid danger. The need for this type of environment is particularly apparent in the case of gas line repair excavations.

Upon arriving at a given site location, REX will be oriented to the location, and started by an operator to excavate for a target object such as a leaking gas line. While excavating for the target object, other buried obstructions may have to be exposed without damaging them. It is expected that the operational version of REX will be able to excavate a volume with a surface area approximately three feet by four feet and to a depth of four feet within thirty minutes.

To accomplish its task, REX is composed of several tools. These are: scanning equipment to infer buried object sizes and locations; a primary excavator to remove soil in a rapid manner; and a secondary excavator to remove all soil adjacent to buried objects, in addition to excavating in areas where the primary excavator can not fit.

Scanning equipment may consist of several types of sensors. These instruments are able to detect different types of buried and exposed objects with varying degrees of accuracy. Examples of possible sensors include magnetometers, radar detectors, and sonar detectors.

Primary excavation equipment may consist of either a bucket ladder or an auger. In either case, the purpose of the primary excavator is to clear unobstructed volumes of soil in a rapid and coarse manner as well as providing a sump for the soil displaced during secondary excavation.

The secondary excavation equipment is responsible for fine excavation, and may consist of an "air knife", presently under development by the Dravo Corporation. The air knife essentially forces air at high velocity through an orifice to dislodge soil at the tool's tip. This type of excavator is termed non-destructive, as it can bump objects without damaging them while still performing the task of excavation.

The specific types of equipment required for REX are still under investigation, but the three activities of scanning, primary excavation, and secondary excavation will be maintained. There is clearly a need for monitoring and controlling these activities. These activities include planning the order of excavation operations, planning an operating location for each operation, and evaluating the results of each step of the excavation process. It is for this purpose that the expert systems described in this paper are being developed.

Expert systems are usually used for the codification of existing expertise, but this project has used the expert system framework to develop expertise in an uncharted new field. The flexible structure provided by the expert system framework allows easy representation of excavation strategies.

This paper will present the use of the expert system framework for planning robotic excavations, as discovered through the development of the REX project.

2. MONITOR/CONTROLLER SIMULATION

Since the development of REX is a first attempt at providing autonomous excavation, there is little past experience available to aid in the design of the system. Therefore, a simulation of the REX operations has been developed, so that characteristics of the excavation process can be learned. The capabilities and requirements of the prototype Monitor/Controller can then be determined using the information obtained from the simulation.

This chapter will present the scope of the Monitor/Controller simulation, and provide needed definitions for discussion. Concepts concerning the proposed Monitor/Controller are included and finally the implementation of the simulation is discussed.

2.1 Scope

The Monitor/Controller must initiate, coordinate, and monitor REX activities from the start of the excavation and progress until a predetermined target object is exposed.

The excavation will operate within a user-defined volume representing the site. This volume will contain the object or objects that are to be excavated. These objects do not have to be completely contained within the site, as it is desired to represent continuous pipe lines running through the site.

During the process, the Monitor/Controller must maintain a current picture of the excavation process. This requires storing information about the excavation site and those objects within the excavation.

The Monitor/Controller must also provide a user interface during the excavation process. This interface allows an operator to input information to the Monitor/Controller at the beginning of the excavation process and to view the excavation process as it occurs.

2.2 Definitions

The following description provides needed definitions for a detailed discussion of the excavation process. The definitions presented here are summarized in the Appendix.

All geometric objects represented in this simulation are cuboids, that is, volumes comprised by opposite parallel faces, with all interior angles of ninety degrees.

There are three types of cuboids used within the simulation. These are the site, domains, and objects (Figure 2-1).

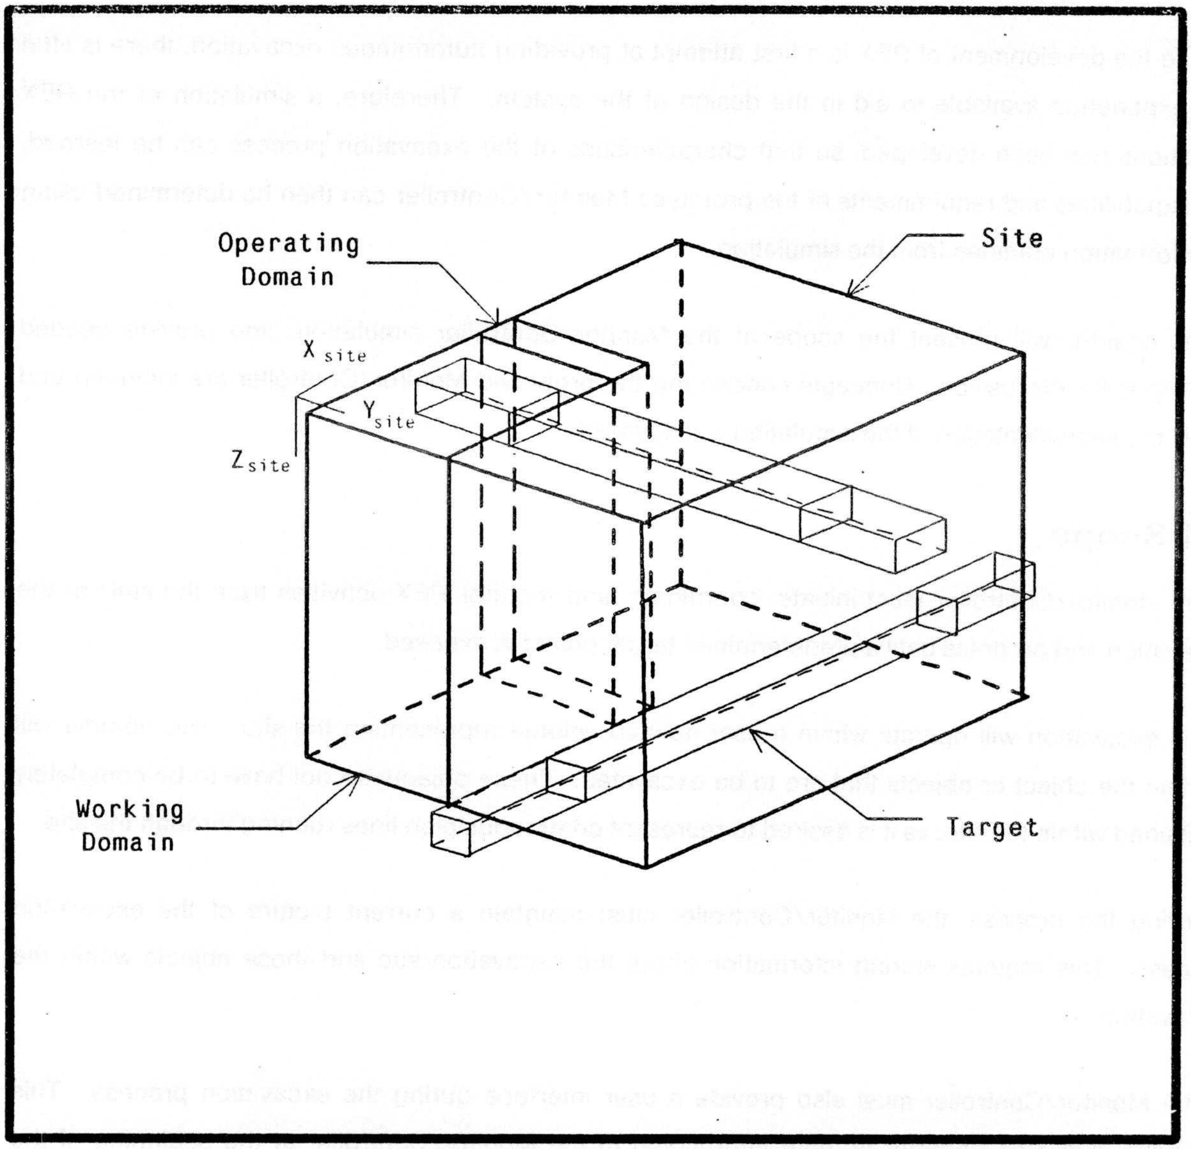

Figure 2-1: Simulation Definitions

The *site* is defined as the volume the excavation process is to be contained within. This volume is defined at the start of the simulation. It is represented by a Cartesian coordinate system, having the x and y axes on the surface, and the z axis extending into the soil.

Domains consist of two types of volumes: working domains, and operating domains. *Working domains* represent the space that has been excavated and therefore are available for tools to move and work within. *Operating domains* represent the volume that a tool is presently working within. Operating domains for excavation activities become working domains when completed.

The third type of cuboids are *objects*. These represent both the *target*, an object which the excavation is trying to uncover, and *obstructions*, those objects buried within the site that are not targets.

2.3 Proposed Monitor/Controller

As the name implies, the Monitor/Controller of REX must initiate, coordinate, and monitor REX activities during the excavation process. With these capabilities, the Monitor/Controller can guide REX to a safe, successful excavation of a target object. Inherent to these activities is the maintenance and storage of all information known by REX about the site and the buried objects within the site at any time during the excavation process. Using this information, the Monitor/Controller develops the strategy for a series of activities leading to completion of the excavation. It can also pass information to individual activities as required by their local strategies. Therefore, the Monitor/Controller acts as the process database, coordinates the principal activities of the excavation process, and monitors the activities' results.

To provide the capability of storing the process data, the Monitor/Controller requires a flexible "world model" (data structure) that can accommodate new information as it is found, and modify old information as it is updated by subsequent excavation processes. The current state of the excavation must also be represented so that the Monitor/Controller knows what it is doing, and what activities can be performed next.

2.3.1 Control

In the design of the Monitor/Controller presented here, it is assumed that the coordination of the principal activities of scan, primary excavation, and secondary excavation involves only the selection of the appropriate activities to perform. Each of these activities represents a major segment of the robotic excavation. *The Monitor/Controller does not know the strategies of any local processes associated with these major activities.* The Monitor/Controller need only communicate with and direct the principal activities shown in the State Transition Diagram (Figure 2-2). The three principal activities are each subdivided into three tasks: planning, operating, and evaluating. The planning task decides how and where to perform the activity. The operating task executes the plan, and the evaluating task determines the performance of the operation and selects the next activity.

Figure 2-2 also shows a fourth activity not yet discussed, preplan. Each of the four activities are now addressed.

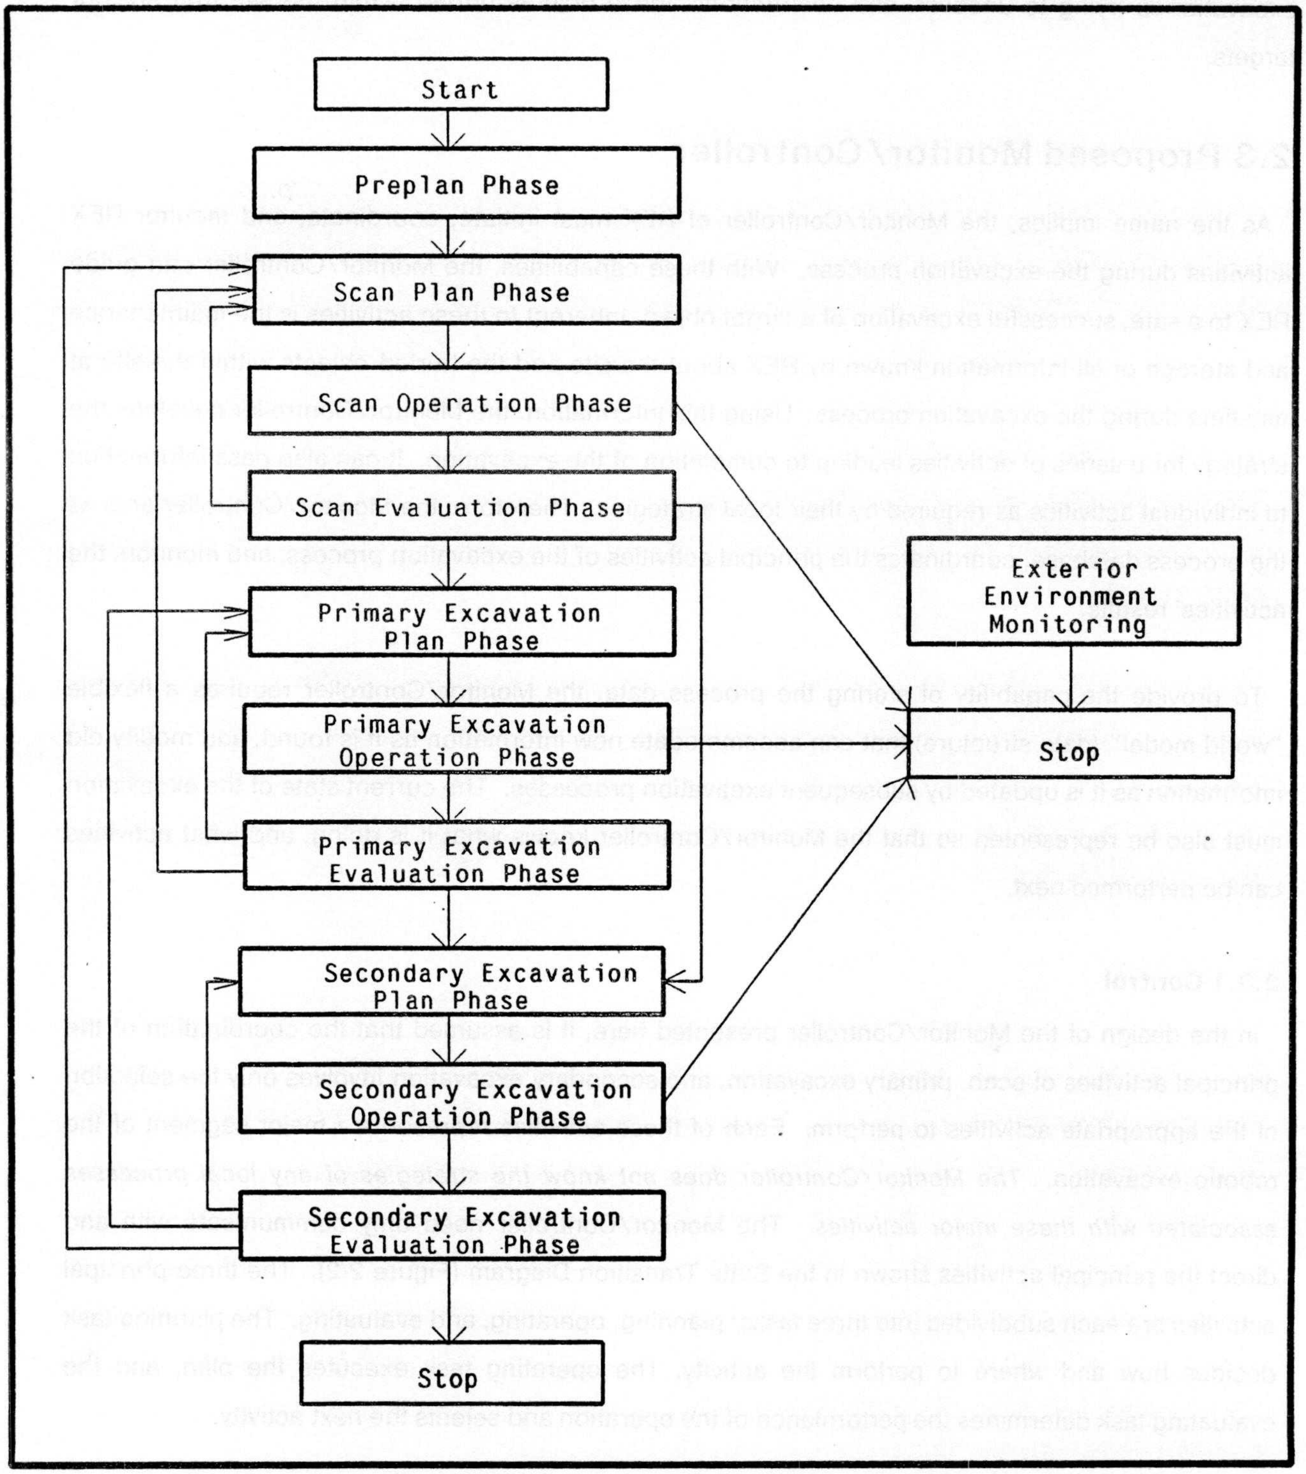

Figure 2-2: State Transition Diagram

Preplan. When arriving at a given site location, there may already be certain information known from previous excavations, and/or from utility maps. The preplan phase allows this information to be entered into the "world model" so as to make the excavation process more efficient. Two types of information are made available to the preplan. The first is information about the site that is known prior to excavation. This information may be approximate, but the ability to provide the Monitor/Controller information about suspected objects within the site assists the planning process. Additionally, a target object must be identified, the one which the excavation will attempt to completely uncover. All of this information may be determined offsite. The second type of information orients REX to the site and must be determined at the site. With this information, a three dimensional representation of the site and of the possible buried obstructions can be formulated. Using this information, the Monitor/Controller controls the excavation process in order to excavate the target, while verifying the premapped knowledge.

Scan. The scan activity is represented in such a fashion that the number and kinds of sensors employed is not a factor in the functioning of the Monitor/Controller. The scan activity is used to provide the Monitor/Controller with information about potential objects buried in the site. With this information, the other activities can be selected to avoid potential collisions of the excavator(s) with buried objects. The sensor data contains processed information from the individual sensors and represents those objects detected by the sensors. This data is evaluated to update the "world model" maintained by the Monitor/Controller.

Primary and Secondary Excavation. The primary and secondary excavation activities are performed in a similar manner. The difference between the two operations is the tool used for the activity. The primary excavator will remove soil more rapidly than the secondary excavator. Again, the type of excavating equipment is not a factor in the functioning of the Monitor/Controller, as long as it has information about their gross operating characteristics and dimensions. The plan phase of the excavation activity uses the current "world model" to determine the volume to excavate. The volume is then excavated and the evaluation task determines the success of the operation. In addition, the evaluation task updates the "world model" with current locations of previously known objects, and adds any new objects found during excavation.

2.3.2 Monitoring

Monitoring of REX activities involves several different concepts ranging from software/hardware monitoring of process activities, to potential human monitoring of the excavation process.

Process Monitoring. First, the Monitor/Controller must monitor the results of each activity when completed or interrupted, to decide when the excavation is complete, or to decide the next activity to invoke.

Operator Interface. The Monitor/Controller needs to provide to an operator the capability to graphically view the database that REX is using. The operator should be able to view the data from any angle and distance, thus displaying the current state of the excavation as the Monitor/Controller knows it. This information could be compared to a remote television camera to verify that the data contained in the Monitor/Controller is consistent with the present state of excavation. Additionally, the monitor should give the operator the ability to stop any activity at any time, or if the operator detects something is wrong.

Status Monitoring. The monitoring must also include checking the self-health of the equipment, so that if some component is malfunctioning, the activity can be modified and/or stopped as needed. The monitor will also need to include some external environmental monitoring, such as a gas detector to detect major leaks if the process breaks a gas line and a perimeter sensor to detect if a person is within an operational danger zone.

2.4 Simulation of Monitor/Controller

2.4.1 Need for Simulation

Given the complex requirements for the Monitor/Controller, it was decided to develop a simulation of the process so as to better define the requirements of the eventual prototype Monitor/Controller and to determine how the different parts of the project interact. The simulation would allow the development of the Monitor/Controller prototype to progress without building expensive equipment that ultimately may not be suited for the job. Having a simulation of the process would also provide a test bed to try new ideas and concepts and determine their effects on the excavation process. Finally, a simulation would provide many of the components needed to produce a graphic user interface for the excavation process from which an operator's station could be developed.

In conjunction with the above needs for the simulation, there was the overlying objective to use the simulation to determine where the control of each activity should reside. At some point, control should be separated between global level issues and local issues of the individual tool. By collecting statistical information about the excavation process using various models of control, advantages and disadvantages of local vs. global control could be evaluated.

The possible separation of control between local and global levels occurs in the planning and evaluation tasks of each activity. At all levels of control, the Monitor/Controller determines the order of the activities. The distinction at the planning level occurs in selecting the operating domain for a given activity. During the evaluation task, the issue of local vs. global control arises again when it must be decided how much evaluation should take place at each of the levels.

The simulation was designed to simulate three levels of control: volume, sequence, and element, as sketched in (Figure 2-3). The figure shows that the control space of the Monitor/Controller is larger at the element strategy than for the volume strategy. These strategies correspond to the local and global strategies respectively, as discussed earlier. The figure also shows that certain knowledge is always needed by the Monitor/Controller regardless of the strategy chosen.

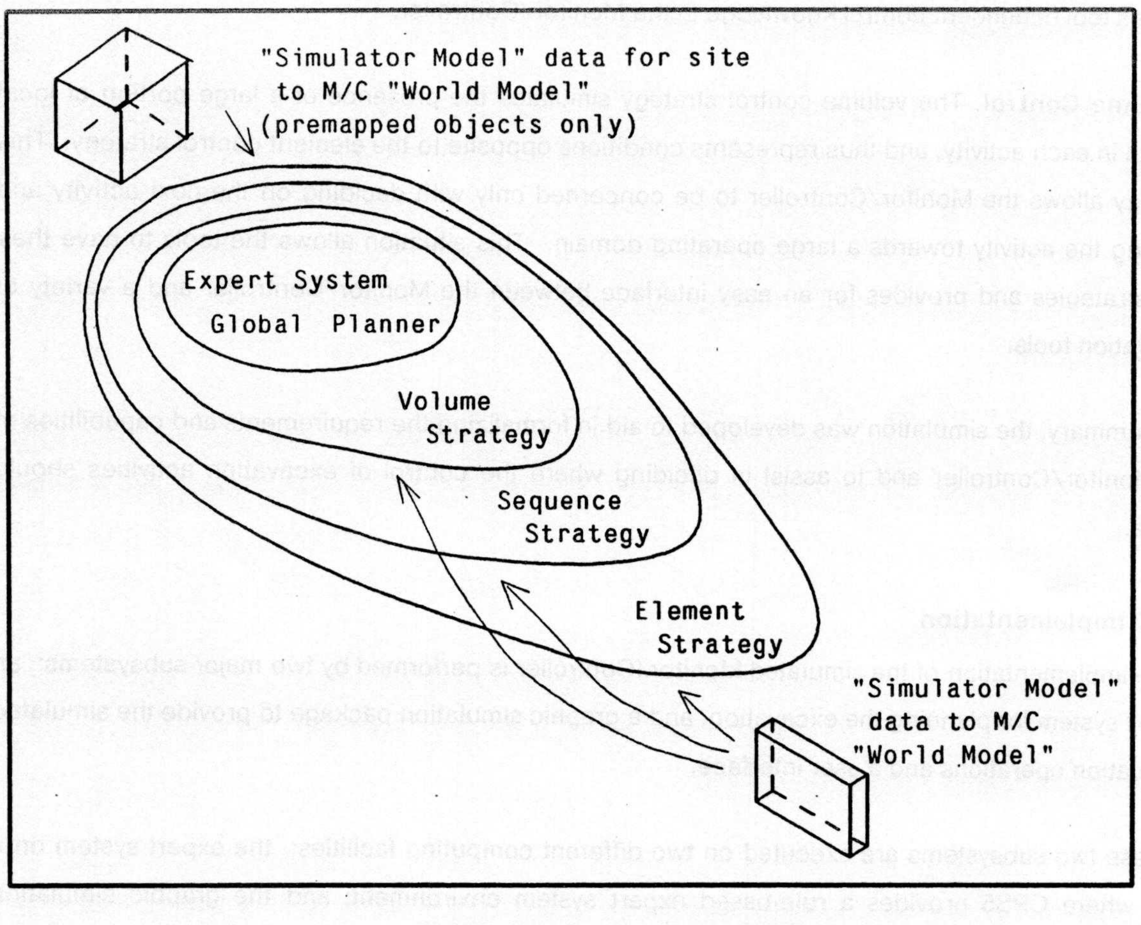

Figure 2-3: Strategy Control Options

Element Control. The element control strategy models the presence of control knowledge residing primarily within the Monitor/Controller. This method assumes that a given activity has very little control knowledge and that the Monitor/Controller contains this information. The volume that the

Monitor/Controller directs the activity to work on is composed of the tool's reach, as defined by the gross operating sizes input by the user during creation of the simulation. These dimensions represent a maximum surface area the tool can "reach" without movement of the robotic arm behind the tool. A depth is also associated with the tool, indicating the volume this tool can "reach" without movement of the arm. The Monitor/Controller decides how to move the tool to operate within the complete operating domain chosen for a given activity.

Sequence Control. Sequence control strategy is implemented to exhibit less local control within the Monitor/Controller than the element control strategy. The volume that the Monitor/Controller selects for the tool to operate within is still based on some of the tool's properties, but at least one of the constraints has been removed. While some control has been shifted to the local level, there still remains tool dependent control knowledge in the Monitor/Controller.

Volume Control. The volume control strategy simulates the presence of a large portion of local control in each activity, and thus represents conditions opposite to the element control strategy. This strategy allows the Monitor/Controller to be concerned only with deciding on the next activity and pointing the activity towards a large operating domain. This situation allows the tools to have their own strategies and provides for an easy interface between the Monitor/Controller and a variety of excavation tools.

In summary, the simulation was developed to aid in formalizing the requirements and capabilities of the Monitor/Controller and to assist in deciding where the control of excavation activities should reside.

2.4.2 Implementation

The implementation of the simulated Monitor/Controller is performed by two major subsystems: an expert system for planning the excavation, and a graphic simulation package to provide the simulated excavation operations and a user interface.

These two subsystems are executed on two different computing facilities: the expert system on a VAX where OPS5 provides a rule-based expert system environment; and the graphic simulation package on an IRIS Workstation utilizing its capabilities of excellent graphic resolution and speed of display. A communication package allows the two subsystems to communicate by transferring a file containing information about the objects and the domains (Figure 2-4). Through the communication package, the progression of the simulation flows between the two machines. The expert system plans an activity, the graphic simulation "executes" the plan, and the expert system evaluates the operation and plans the next activity.

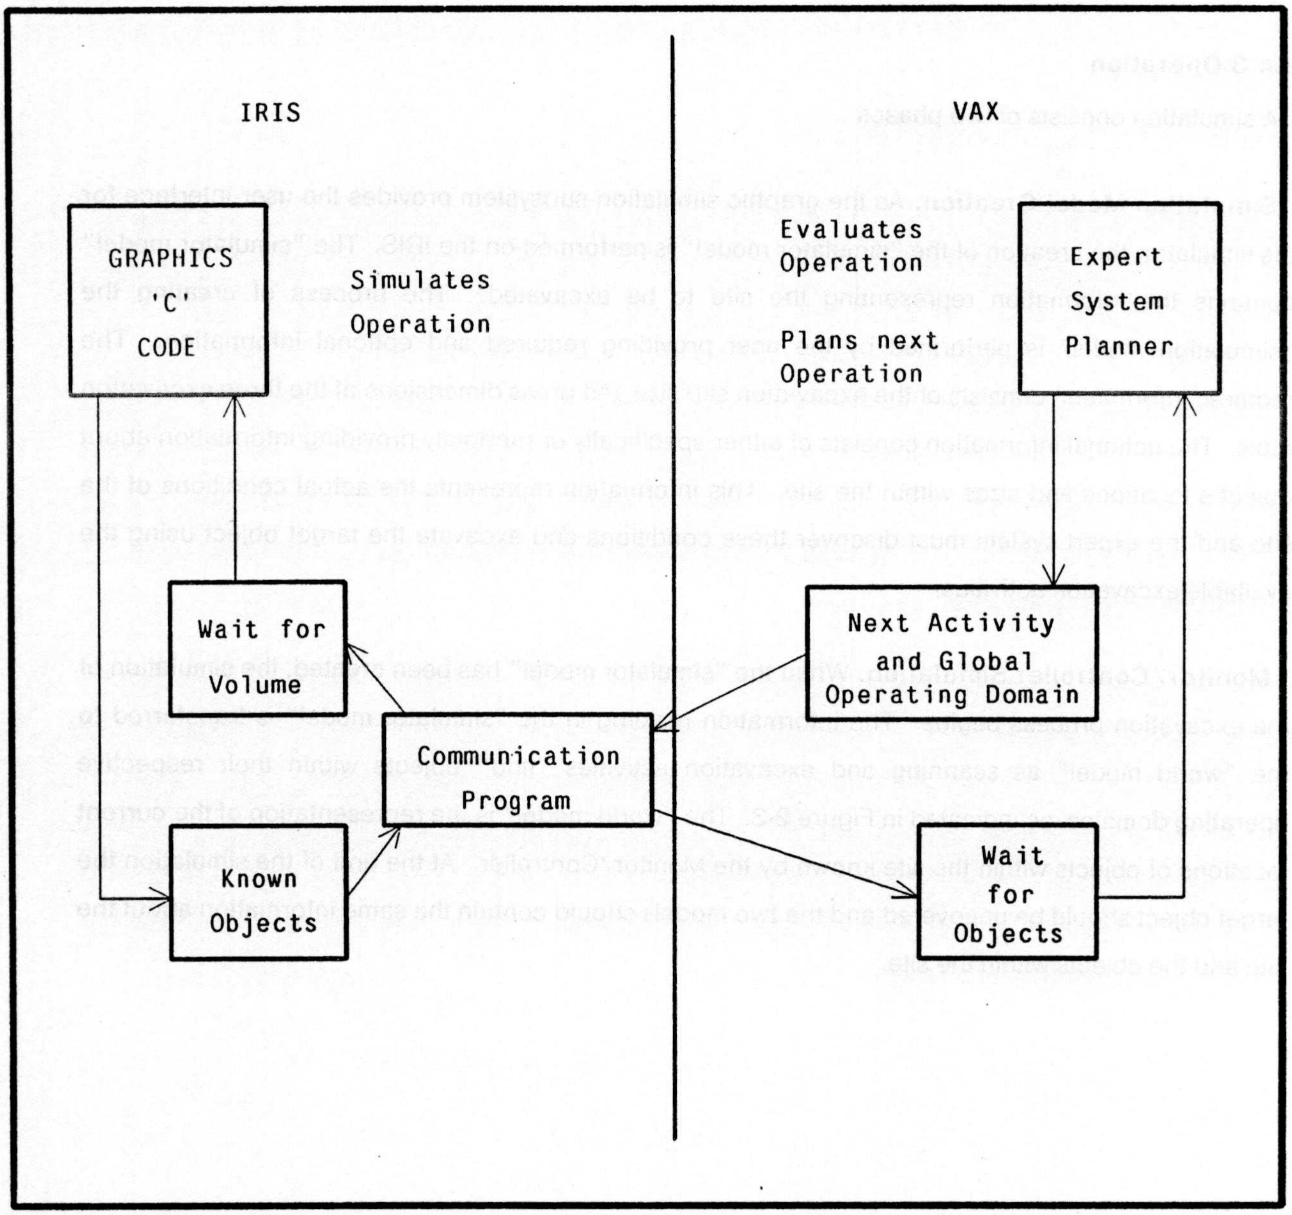

Figure 2-4: Simulation Implementation

The file written by the expert system during the plan activity describes the current state, the operating domain to work within, and an orientation for the tool to progress. The graphic simulation program reads this file and simulates the operation phase. This results in the "finding" and/or "redefinition" of the objects within the site. The objects and their current known geometric attributes are then placed into the file to be transferred to the expert system allowing the activity to be evaluated. The cycle continues until the expert system decides that the excavation is complete, or the target is unreachable.

2.4.3 Operation

A simulation consists of two phases.

Simulation Model Creation. As the graphic simulation subsystem provides the user interface for the simulator, the creation of the "simulator model" is performed on the IRIS. The "simulator model" contains the information representing the site to be excavated. The process of creating the "simulation model" is performed by the user providing required and optional information. The required information consists of the excavation site size and gross dimensions of the three excavation tools. The optional information consists of either specifically or randomly providing information about object's locations and sizes within the site. This information represents the actual conditions of the site and the expert system must discover these conditions and excavate the target object using the available excavation activities.

Monitor/Controller Simulation. When the "simulator model" has been created, the simulation of the excavation process begins. The information residing in the "simulator model" is transferred to the "world model" as scanning and excavation activities "find" objects within their respective operating domains, as indicated in Figure 2-3. The "world model" is the representation of the current locations of objects within the site known by the Monitor/Controller. At the end of the simulation the target object should be uncovered and the two models should contain the same information about the site and the objects within the site.

3. EXPERT SYSTEM GLOBAL PLANNER

One of the major subsystems in the simulation is the expert system. It provides the knowledge to guide REX through the excavation process and the knowledge to perform the planning and evaluation of activities. The methods used to develop the expert system are addressed below. For a more thorough discussion of the global planner see reference [Baker 85].

3.1 Purpose

The primary purpose of the expert system global planner is the sequencing of excavation activities, as shown in Figure 2-2. The sequencing is achieved by providing strategy planning and process evaluation for each activity.

As the amount of data and the specific knowledge needed were unknown prior to the development of the simulation, the expert system framework provided a modularized and an easy to understand knowledge base that was easy to generate and is very flexible to change. These advantages allowed the development of the simulation to proceed while additional knowledge was being gathered. The simulated process is easily visualized through the graphic subsystem which helps verify the knowledge. For these reasons, the expert system environment was chosen to provide the strategy planning and process evaluation for the excavation simulation.

3.2 Implementation

A complete expert system consists of six major components: an user interface, an explanation facility, a knowledge acquisition facility, a context, an inference mechanism, and a knowledge base. The expert system contains portions of these six components.

The user interface is provided by a graphic simulation package and is not discussed in this report. The explanation facility in this system is very limited. It essentially involves output during execution of the expert system informing the user as to the current state of operation and in some cases why a state was chosen. Technically, a knowledge acquisition facility does not exist in this system. However, the knowledge represented is easy to read and understand, therefore adding new knowledge is fairly straightforward.

The three other components of an expert system: the context, inference mechanism, and knowledge base are now discussed in detail as they constitute the essential ingredients of the system.

3.2.1 Context

The context contains information describing the current state of the excavation. The site size and tool sizes are each represented by their own working memory elements containing the dimensions in the three primary directions referenced to site coordinates. These sizes are obtained from the graphic simulation, where the user is requested to input them.

Buried objects within the site are represented by two methods: two point, and face representation. Two point representation requires storing the coordinates of two points along an objects axis and the storage of the cross sectional dimensions. Face representation requires storing the normal coordinate value for each of the six faces of a given object. This method is possible due to an original assumption that all objects are parallel to the site axes, and are at constant depth. The two point representation is used in the communications subgoal to allow direct communication to the graphic simulation subsystem which uses the two point representation. This subgoal performs the translation into the face representation as it is used within the rest of the expert system for it allows calculations of free space, i.e., volume between objects, to be performed more easily and removes any distinction based on an object's orientation required with the two point representation.

Context information for buried objects also contains information indicating whether the object is the target or an obstruction. Additionally, the object's status value used by the simulation is included so that all required information about an object is represented identically within both subsystems.

The working and operating domains are represented in the context. These volumes represent the space where excavation has already occurred, and the space where a given activity is presently operating, respectively. Each domain is represented in a manner similar to the objects using the face representation. Also, each domain is given a name to distinguish it from the others, as an aid in planning, and a status value to indicate whether the domain is active or has been completed.

Lastly, the context contains a representation of all the possible states of excavation as outlined in the State Transition Diagram (Figure 2-2). A further description of this mechanism is contained in the discussion of the control structure.

3.2.2 Control Structure

As OPS5 provides a general inference mechanism internally, this discussion will entail the specific control structure used by the expert system, so that the knowledge represented is executed as required. The control structure is discussed on two levels: a outer level representing the transitions between activities, and an inner level describing the control structure within each activity.

Outer Level. The outer level control structure is multiply sequential and is an internal representation of the State Transition Diagram (Figure 2-2). From a given state, shown by the boxes within Figure 2-2, there exist possible or legal transitions to other states. Some states have only one transition, while others have many and can even call the same activity again. The working memory of the expert system is set up to reflect this diagram. Therefore, working memory elements represent the major phases of each of the four major excavation activities. Working memory also represents the legal paths connecting these states. Each path has a status attribute indicating whether the path is enabled (transition currently allowed) or not enabled (transition not allowed at that time).

Each state within working memory has a status attribute. This status attribute indicates whether the given state is active (execution currently resides there), is inactive (execution currently not residing there), or completed (execution has just finished in this state). When any state achieves the status of completed, a production fires to look for the enabled path from the completed state to the next state. This production makes the next state active and the old state inactive. Therefore, those states which have multiple exit paths require process evaluation to decide which path to enable, so that the process continues.

Inner Level. The inner level control structures for the four primary activities will now be described. With the exception of the preplan activity, each activity has three states: plan, operate, and evaluate. The control mechanism for the preplan activity, and then for plan, operation, and evaluation states is discussed.

The control structure for the preplan activity is sequential and is accomplished through subgoals. The first step is to perform the communication subgoal.

The communication subgoal first waits for a file from the graphic simulation package. Upon detection of this file, it reads the file, and makes or updates objects using information in the file. The file for the preplan also contains the information describing the site size and tool sizes. The communication subgoal then determines the locations of the potential operating domains, deletes the file, and passes control back to the caller.

After the communication subgoal has executed, the preplan activity state becomes completed and always chooses the scan state to become active. This is done to gather more information about the site prior to commencing excavation.

The plan states for each of the activities, scan, primary excavation, and secondary excavation, perform in a similar manner. Each plan state has the responsibility of selecting the next operating domain, deciding on a tool direction to perform the operation, and indicating to the user the selected operating domain. The control mechanisms used to accomplish these responsibilities include sequential and data driven control components. The sequential control component insures that the responsibilities are performed in the order listed above. The data driven control components then decide on the next operating domain and tool direction. The knowledge required to perform this is discussed under the domain knowledge.

The operation states for the activities are sequentially driven. Each operation state first writes a file containing the activity's name, the current operating domain chosen by the plan state, the location and size of the operating domain, and the tool orientation to use in performing the activity. Then the operation state calls the communication subgoal to wait for the file returned by the graphic simulation. The communication subgoal also reads and interprets the contents of the returned file. The operation state then removes the operating domain from working memory, and passes control on to the evaluation state of the activity.

Evaluation states for the activities are controlled in a manner similar to the plan states. Each evaluation state must initialize all possible paths from the state to be not enabled, evaluate the results from the operation, determine the next activity in the excavation process, enable the path to it, and notify the user of this choice. The control structure for these states again involves sequential and data driven control components. The sequential control component insures that the above requirements are performed in the listed order, and the data driven control components perform the evaluation and determination of the next activity. The knowledge required to perform this decision is discussed in the domain knowledge.

3.2.3 Domain Knowledge.

The domain knowledge present in this expert system performs the two tasks of planning and evaluating activities. These two areas of domain knowledge will be discussed and an example of the domain knowledge presented.

Planning Knowledge. The knowledge contained within the planning states provides the ability for the expert system to select the next operating domain and to choose a tool direction for operating on the domain. The rules comprising the knowledge are dependent on the presence or absence of a target object, the orientation and location of the target object, the present activity, and on the number and location of working domains completed.

As stated before, preplan causes a scan to be executed. If a target object does not exist after the first scan, the system requires that a second scan be accomplished. Without knowing the suspected location of the target, the expert system does not know what the excavation is trying to achieve. If the target is still not found on the second scan, the current knowledge will perform a top slice excavation to a tolerance above the closest known object. Another scan is performed after this excavation. If the target object has not been identified after this operation, the current expert system stops the excavation process.

Given that a target has been identified and that its location and orientation are known, the planning rules will select the next operating domain based on the active activity and the working domains previously completed. The current system first looks to perform operations on the top slice. If the top slice is too shallow (less than a depth of eighteen inches), it becomes inactive and is not used. If the top slice can be acted upon, the activity is selected within this region.

Once the top slice has been acted upon, it achieves a status of completed and further planning rules attempt to operate on one of the four site corners. The origin corner, as defined by the site coordinate system is looked at first. If there is adequate room in this corner for the activity to operate, it is chosen. If this corner is too small for the current activity's tool to operate within, the corner along the x or y axis is looked at next, depending on the orientation of the target object. Since excavating the target is the overall goal, the major thrust of the excavations will occur in a direction parallel to the target, rather than excavating everything possible. This allows for a more efficient excavation where only necessary material is removed. If the corners have been investigated, the system determines if there is a volume of material parallel to the target that has not been excavated. These potential volumes are represented by middle regions. If the faces of the corner locations meet, or go to the edge of the site, the middle regions are not used.

Once all the primary excavations have been performed based on the above criteria for looking at domain regions, the secondary excavator is called to remove the soil around the target starting from each of the completed working domains. Each working domain has a face parallel to the target's direction and the secondary excavator advances this face beyond the target. When this has been completed, the expert system decides that the excavation is complete and stops.

An example of a planning rule for performing primary excavation on the top slice is shown below. This rule shows how the knowledge is dependent on the current active state, the working domains completed, or not completed in this case, and on the geometry of the domains.

```
IF
state is primary excavator
and subgoal is plan
and closest object is at depth > 18"
and top slice has not been excavated

THEN
make operating domain top slice active
          length = site length,
          width = site width,
          depth = tolerance above closest object.
and make tool heading = +z
          width = tool x dimension
          length = tool y dimension
          depth = tool z dimension
and tell user operating domain is top slice
```

The actual rule associated with this translation is now given. Although not as easily read as the above translation, the actual rule is still readable, and easily modified.

```
(p prim←ex←plan::top←slice
          (state ↑name Prim←Ex←Plan ↑status active)
          {(subgoal ↑name plan) <sub←goal>}
          {(working←domain ↑location top←slice
                       ↑status nil
                       ↑min←x <wd←minx> ↑max←x <wd←maxx>
                       ↑min←y <wd←miny> ↑max←y <wd←maxy>
                       ↑min←z <wd←minz>
                       ↑max←z {<wd←maxz> > 18.0} ) <wd>)
          (prim←ex←range ↑size←x <prim←ex←x>
                       ↑size←y <prim←ex←y>
                       ↑size←z <prim←ex←z>)
-->
          (bind <tolerance> 6.0)
          (make operating←domain ↑name 1 ↑status active
                    ↑min←x <wd←minx> ↑max←x <wd←maxx>
                    ↑min←y <wd←miny> ↑max←y <wd←maxy>
                    ↑min←z <wd←minz>
                    ↑max←z (compute( <wd←maxz> - <tolerance>))
                    ↑location TopSlice
          )
          (make tool ↑heading +z
                    ↑x←size <prim←ex←x>
                    ↑y←size <prim←ex←y>
                    ↑z←size <prim←ex←z>)
          (modify <wd> ↑status active)
          (modify <sub←goal> ↑name finished)
          (write Plan top cut as depth to closest object is > 18)
)
```

If it were desired to excavate the top slice only when its depth is greater than 24 inches, then the 18 would just be replaced. If some other condition is necessary before excavating the top slice, it could be added to the left side of this rule.

Evaluation Knowledge. Knowledge for the evaluation states is dependent on the same conditions as the planning knowledge with several additional constraints. This knowledge is also dependent on the current active operating domain, and on the objects returned from the graphic simulation. With this information, the rules contained in the evaluation states evaluate the operation that took place, and then choose the next operation.

After the preplan state has been executed, a scanning operation is always performed. This knowledge, although a decision choosing the next state, is placed within the preplan operation, as it always occurs (i.e., it is the only path out of preplan). Other decisions determining the next state of operation are based on the premise that all of the primary excavation should be accomplished before commencing any secondary excavation. This premise is based on two assumptions. The first is that the secondary excavator, as presently conceived, does not provide any method to remove the soil it dislodges; thus a sump has to be provided so that gravity can provide the mechanism to carry the soil away from the working area. Secondly, it is assumed that it is more expedient to use one tool first, then store it and begin operating with a second tool, rather than switching tools. By using one tool for as much work as possible, less time is involved deploying and storing tools between operations.

The evaluation of an activity's process entails updating object locations, and updating the locations and regions of the working domains. The present expert system always updates this information when reading the file returned by the graphic simulation.

An example of a translated evaluation rule is given below. This rule exhibits the dependency on the current active state and active operating domain.

```
IF
state is evaluate primary excavator
and top slice is active
and corner0←0 has not been excavated
THEN
make top slice completed
and make path to next state = scan plan
```

The actual rule associated with the translation is the following. Again, it is easy to read the actual rule which allows easy modification to the knowledge base.

```
(p prim←ex←evaluate::top←slice←excavated
    (state ↑name Prim←Ex←Evaluate ↑status active)
    {(subgoal ↑name target←check ↑status active) <sub←goal>}
    {(working←domain ↑location top←slice ↑status active) <wd>}
    (working←domain ↑location corner0←0 ↑status {<> completed})
    {(valid←transition ↑current←state Prim←Ex←Evaluate
                 ↑next←state Scan←Plan) <path>}
-->
    (modify <path> ↑enable TRUE)
    (modify <sub←goal> ↑status satisfied)
    (modify <wd> ↑status completed)
    (write Top Slice done so rescan top surface (crlf))
)
```

Currently, knowledge representing the performance of a given tool after an operation of an activity is not modeled, as a clear definition of the tools and how they work is not defined. However, this can be included in this process when that information is known.

The present knowledge contained in the expert system also does not contain information as to how the tool should operate within the volume. Therefore, the knowledge represented closely resembles that knowledge required to operate at the global level of control. Presently, the knowledge to control the tools for local levels of control resides in the graphic simulator. If the Monitor/Controller is to have this knowledge to work at the local level of control, this knowledge will have to be represented in the expert system.

3.3 Example

Using the simulation presented in this report, an example excavation is discussed in detail. This example reflects several of the capabilities of the Monitor/Controller while maintaining a level of simplicity. Additionally, this example is representative of many excavations required in the field. Certainly there will be unique and geometrically complex cases arising in the field, but it is suggested that these situations will occur in a minority of cases.

The example excavation involves: the target object running in one direction, and an obstruction perpendicular to the target object located above it (Figure 3-1). The site is given as 5 feet to each side. The target object runs the length of the site in the x direction, at a depth of 4 feet, and 3 feet in the y direction. The target object is 4 inches square, and is preplanned. The obstruction object runs the length of the site in the y direction, at a depth of 2.5 feet, and 2.5 feet from the x axis. The obstruction object is also 4 inches square and is not preplanned. The two objects have been represented so that the simulated scanner will detect the objects' presence.

The simulation of this example provides the reader with examples of the reasoning used in the expert system, and with statistical data pertaining to the control level issue. The expert system will be discussed first and then the statistical results.

The following listing shows the excavation activities selected by the expert system for excavating this example. The reasons for each activity selected are also given.

1. **Start:** Initialize the control structure, and welcome the user to the expert system.

2. **Preplan:** Obtain the target object as it was preplanned.

3. **Scan:** Perform a scan of the top surface. The obstruction object is detected during this activity.

4. **Primary Excavation:** Perform a top cut over the site, as the object closest to the surface is at a depth greater than eighteen inches.

5. **Scan:** Scan the surface again so more accurate information can be gathered about the objects' locations.

6. **Primary Excavation:** Excavate along the target in the operating domain identified by the origin corner, to a depth below the target object.

7. **Primary Excavation:** Excavate along the target in the operating domain identified by the other corner along the x axis as the origin corner does not go to the end of the site since there is an obstruction object. This will be to a depth below the target object.

8. **Primary Excavation:** Excavate between the origin corner and the far corner along the x axis to a depth just above the obstruction object. This excavation removes soil between the corners.

9. **Secondary Excavation:** As the two corners and middle locations along the target object have been excavated with the primary excavator, start the secondary excavator. The first working domain is the origin corner, and the operating domain goes from the face of the present excavation and extends beyond the back face of the target object.

10. **Secondary Excavation:** Excavate around the target object using the far corner along the x axis as the work domain.

11. **Secondary Excavation:** Excavate around and beneath the obstruction between the corners in the x direction. The volume under the obstruction will be cleared, permitting the excavation of the target object under the obstruction.

12. **Secondary Excavation:** Excavate under the obstruction towards the target object.

13. **Stop:** As the excavations performed extend along the length of the object, the excavation is complete.

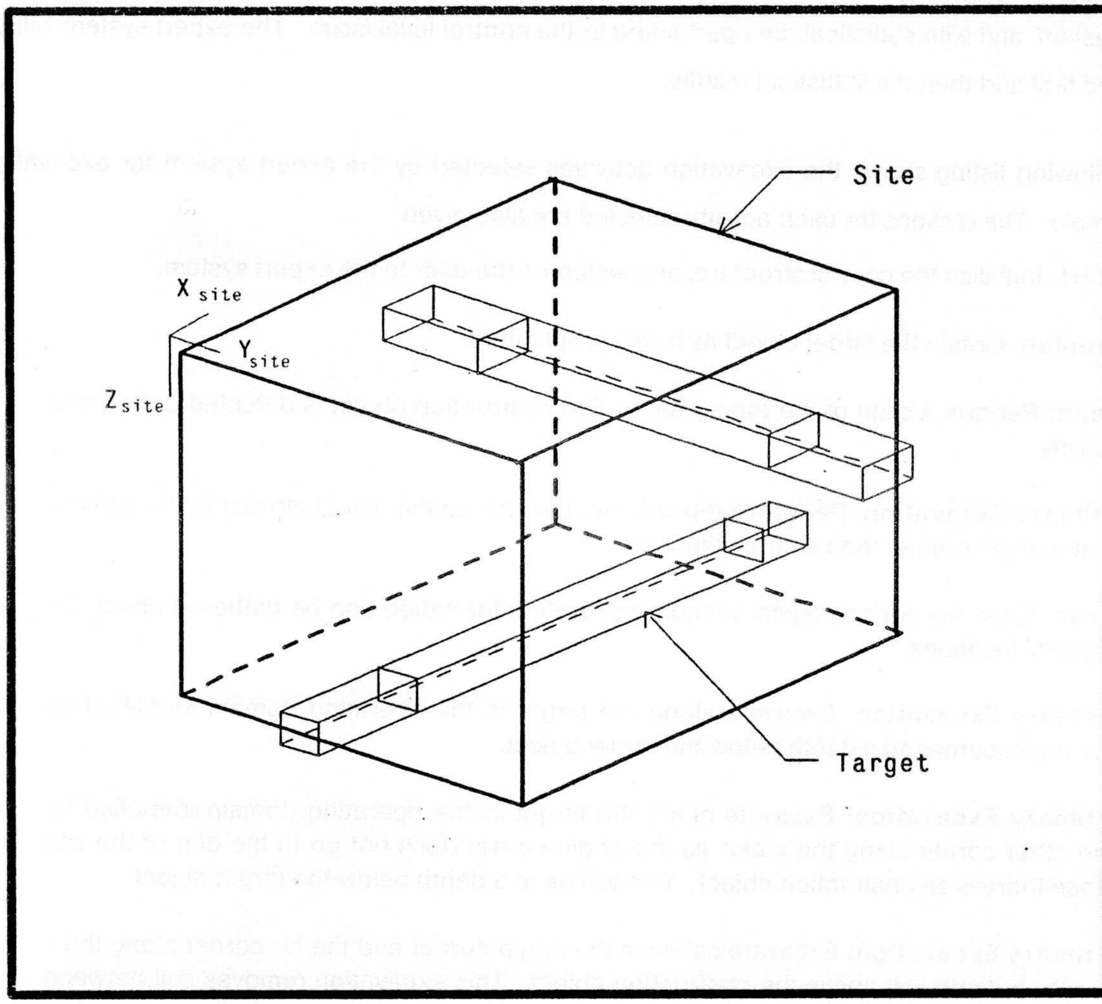

Type Control	# Clips	# Passed	# Extensions	% Excavated	% Primary Excavated
Element	535	396	3	79	72
Sequence	165	156	3	79	72
Volume	23	21	1	79	72

Figure 3-1: Example Data Run

The statistical results for this example are shown in Figure 3-1. The data shows that 79% of the site has been excavated. An even smaller percentage of the site could have been excavated if the top slice would have been made smaller. However, the thickness of the top slice excavation must be balanced with how much more information can be obtained in later activities (primarily scanning) since the tools will be closer to the buried objects.

The amount of the removed material performed with the primary excavator is 72% of the total excavated material. By knowing the percent of the site excavated, and what portion of the excavation was performed with the primary excavator, an idea of the time required to actually perform this excavation can be estimated if the rates of the tools are known. More important, is the high proportion of the excavation performed with the primary excavator. This reflects a time that will be shortest, as the primary excavator should remove soil faster than the secondary excavator.

As can be seen from the data, the amount of calculations performed by the Monitor/Controller (# Clips) is dramatically higher for the element control strategy than for the volume control strategy. This reflects the computation time required for the Monitor/Controller to determine the order and location of the elemental volumes for the tool.

Additionally, the element control strategy requires many more communications (# Passed) between the Monitor/Controller and the tool, than does the volume control. As more information is required to be transferred, the more complex the communication process becomes. Therefore, the volume strategy minimizes the transfer of information, enabling the communication to be as simple as possible.

The last statistical piece of information is the number of extensions required for each of the control structures. This number reflects the number of times the Monitor/Controller had to internally extend the object while performing its calculations. It is desired to reduce this number to save computations.

This evidence strongly suggests local planning should be performed at the tool level as much as possible, allowing the Monitor/Controller to make only the global decisions.

4. SECONDARY EXCAVATION LOCAL PLANNER

4.1 Purpose

The expert system **Secondary Excavator Local Planner**, (SELP) , is designed to serve as the Secondary Manipulator task developer for the Monitor/Controller global planner, M/C. Fundamentally, upon being directed to a volume operating domain by the M/C, SELP determines a plan of excavation for the Secondary Manipulator to follow in pursuing the immediate goal of clearing the target object.

SELP develops an excavation recommendation based on considerations of the location, size and orientation of the target and any inclusion objects present within the domain, and a limited set of kinematic constraints on the Secondary Manipulator.

4.2 Implementation

4.2.1 Context

The context includes three major classes of elements:

Excavation Objects consist of three working memory element types which hold size and position data for target and obstruction objects contained within the site which SELP must process.

These three elements, which constitute the communication area between M/C and SELP, are:

- *object*s, which contain excavation object descriptions "input" to SELP from M/C
- *excavation*s, intermediate constructs for SELP internal use, and
- *box*es, which contain excavation recommendations "returned" to the M/C.

Work Domain is a single element containing data about the current operating domain in which work is to be performed, and mechanical attributes of the Secondary Manipulator which may impact excavation.

Miscellaneous consists of data elements which serve utility functions such as counters, pointers, temporary variables, and process state management within SELP. These are used as required during the local planning processes.

All of the above elements are local private copies constructed for SELP with each invocation. Therefore, attributes of these elements can be and are modified or destroyed during a local planning process with no ill effects on global activity.

4.2.2 Control Structure

To implement the system, a forward chaining control format was utilized. There are six major process states employed in SELP, as follows:

- **Step 1** determines the local coordinate system to be used.
- **Step 2** expresses all objects in local coordinates.
- **Step 3** determines what objects are in the path to the target, and, of these, which are physically too close together to be considered as individual entities.

  Objects passing this proximity analysis are assigned a unique sequence number, a derived datum used in subsequent processing. Each non-target object is then compared with each *other* non-target object and combined, if required, to form a *new* object. If any combinations occur, object sequence numbers are adjusted to insure a monotonically ascending sequence.
- **Step 4** arranges any remaining non-target objects in proper ascending order by sequence number for calculation of bounding excavations.
- **Step 5**, calculates the recommended bounding excavations. Since Step 4 properly ordered the non-target objects, the sequence number attribute can be used to select the objects one by one and calculate the bounding excavations. A bounding excavation may be degenerate, i.e. have zero volume.
- **Step 6**, lastly, transforms the local coordinates of bounding excavations back to the original global system for M/C use, excluding any degenerate volumes.

4.2.3 Domain Knowledge

The organization of domain knowledge within the system follows a transition through the task environment from a state of zero knowledge, i.e. initial conditions, to a final state of knowledge, i.e. when excavation recommendations are returned to the M/C.

Fundamentally, this knowledge state transition involves:

- *interrogating* the environment to determine its condition
- *establishing* a starting reference frame from which to work
- *determining* the *path* to the target object and what obstacles may be in the way of achieving this object
- *deploying* a simple set of rules to traverse this *path* to the target
- *formulating* the excavation recommendations and returning them to the M/C

No specialized technical domain knowledge was involved in this effort; a straightforward common sense approach was taken to devise a path to the target. Several agreed-to conventions were observed and full advantage was taken of the geometrical properties of the constrained environment. Reference [Balash 85] offers a more detailed treatment of the features and characteristics of SELP.

5. SUMMARY AND CONCLUSIONS

The simulation of the Monitor/Controller for REX has proven that the framework developed can perform the strategic planning and process evaluation for autonomous excavations.

The simulation has created a global process model for performing robotic excavation for buried objects. Using the control structure of the expert system, additional processes can be easily added. Excavation knowledge can also be added easily to decide what activity to perform next, and where to perform it. This simulation has created a tool for experimenting with different excavation strategies. New knowledge can be tested using the simulation, prior to incorporating it into a field version. Even in the present studies, as excavation cases have been executed with the simulation, additional knowledge has been gathered. It is clear that the simulation provides an excellent way of obtaining knowledge to control the excavation process. By simply creating a new simulation model, new requirements for knowledge not thought of before have been identified.

The expert system framework has shown that it is well suited for handling non-sequential, opportunistic processing of data. The flexible control structure implemented in this expert system helps provide this capability. The framework does have a deficiency when reasoning with geometric representations. This simulation was able to develop spatial relationships due to the limitations on object placement within the site. Future versions of the expert system will have to address the problem of spatial relationships if the limitations on object locations are to be relaxed. Primitives such as "on top of", or "next to", may have to be provided.

The present knowledge base demonstrates the capabilities for strategy planning and process evaluation. This knowledge must be supplemented with more realistic knowledge about the actual tools and their strategy options when they become defined. These strategies should reflect the underlying risk that can be safely accepted while performing autonomous excavations.

The most critical issue addressed from a global viewpoint has been the determination of control strategy location within the system. Issues of placing tool dependent knowledge in the Monitor/Controller versus placement of this knowledge within the tool's local controller/planner have been addressed. Concerns involving location of this knowledge deal with system modularity, so that new tools can be added fairly easily, and communication expense (both cost and time) involved in sharing information between the Monitor/Controller and the individual tools.

Simulations conducted so far strongly suggest that each of the tools comprising REX should have as much local knowledge as possible. This local knowledge should allow the tool to act within a large volume of material. Given suspected or known locations of objects, the tool should be able to move around these objects performing its task. It should then be able to describe to the Monitor/Controller the locations of the objects found during the operation. It should also indicate what portion of the volume it acted upon. With this breakdown of knowledge, the Monitor/Controller can concentrate on global issues of strategy planning and process evaluation, while remaining generic to new or modified tools.

Within the plan states of the scan, primary and secondary excavation activities, an additional responsibility could be provided. Another communication mechanism can be established to call a local planner, either an expert system or a procedural algorithm, to decide on local planning issues prior to the activity's operation. This concept has been demonstrated with the secondary excavation local planner.

Finally, the simulation addresses the issue of interfacing an expert system environment with a graphics package. Although the current implementation is slow, as the two different systems reside on two pieces of hardware, the approach shows promise for future graphics applications of expert systems. If the expert system environment used for this simulation could be implemented on a machine having graphic capabilities, a better system could be provided. The capabilities provided with this system configuration have enabled easy access to implemented strategy, and a valuable package to visualize the strategy. Whether the strategic planning and process evaluation remain in an expert system format for the prototype version of the Monitor/Controller is a future decision. However, the ability to provide and test new strategies easily with this environment suggests keeping this simulation concurrent with the design of REX. When acceptable levels of control strategy exist, they can be transferred to faster procedural code or possibly even placed on a chip if required. The simulation could then start work evaluating strategies for the next version of the Monitor/Controller.

REFERENCES

[Baker 85] Baker, Nelson C., and Fenves, S.J., *Simulation of the Monitor/Controller for a Robotic Excavator*, Technical Report R-85-149, Department of Civil Engineering, Carnegie-Mellon University, August 1985.

[Balash 85] Balash, John, *An OPS5 Strategic Planner for the REX Secondary Manipulator*, unpublished Final Report for Expert Systems in Civil Engineering, 12-743, Carnegie-Mellon University, May 1985.

[Brownston 85] Brownston, Lee, Farrell, Robert, Kant, Elaine, and Martin, Nancy, *Programming Expert Systems in OPS%, An Introduction to Rule-Based Programming,* Addison-Wesley Publishing Company, Inc., 1985.

[Fenves 84a] Fenves, Steven J., and Baker, Nelson C., *REX Controller/Monitor Operation, Progress Report No. 1*, unpublished REX Progress Report, Civil Engineering & Construction Robotics Laboratory, Carnegie-Mellon University, September 1984.

[Fenves 84b] Fenves, Steven J., and Baker, Nelson C., *REX Controller/Monitor Task Interface, Progress Report No. 2*, unpublished REX Progress Report, Civil Engineering & Construction Robotics Laboratory, Carnegie-Mellon University, November 1984.

[Foderaro 83] Foderaro, John K., and Ericson, Lars W., *The FRANZ LISP Manual,* Department of Computer Science, 1983.

[Forgy 81] Forgy, Charles L., *OPS5 User's Manual,* Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213, July 1981.

[Hayes-Roth 83] Hayes-Roth, Frederick, Waterman, Donald A. and Lenat, Douglas B., *Building Expert Systems,* Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1983.

[IRIS 83] *IRIS User's Guide,* Version 1.0, Silicon Graphics Inc., 630 Clyde Court, Mountain View, California 94043, 1983.

[Kernighan 78] Kernighan, Brian W., and Ritchie, Dennis M., *The C Programming Language,* Prentice-Hall, Inc., Englewood Cliffs, New Jersey 07632, 1978.

[Kernighan 84] Kernighan, Brain W., and Pike, Rob, *The UNIX Programming Environment,* Prentice-Hall, Inc., Englewood Cliffs, New Jersey 07632, 1984.

[Kochan 83] Kochan, Stephen G., *Programming In C,* Hayden Book Company, Hasbrouck Heights, New Jersey, 1983.

[Newman 79] Newman, William M., and Sproull, Robert F., *Principles of Interactive Computer Graphics, Second Edition,* McGraw-Hill Book Company, 1979.

[Sriram 84] Sriram, D., Maher, M.L. and Fenves, S.J., "Knowledge-based Expert Systems in Structural Design," *Computers and Structures,* January 1984.

[Winston 84] Winston, Patrick Henry, and Horn, Barthold Klaus Paul, *LISP, Second Edition,* Addison-Wesley Publishing Company, 1984.

I. GLOSSARY

This brief glossary is intended to assist the reader with the terminology used in the discussion of the Monitor/Controller and the excavation process of REX.

- *activity* - a major component of REX excavation procedure, such as preplan, scan, primary excavation, or secondary excavation.
- *cuboid* - a regular rectangular prism (i.e. a three dimensional volume where all opposite faces are parallel, and all interior angles are ninety degrees).
- *domain* - a cuboid that represents a volume that the excavation process has or will occur within. See also working domain, and operating domain.
- *Monitor/Controller* - the central process that oversees the excavation process.
- *objects* - objects are represented as cuboids and are of two types: target and obstruction. The target object is the desired object to excavate and all others are obstructions which must be excavated around.
- *operating domain* - a cuboid that represents a volume that a given excavation activity is currently operating on.
- *primary excavation* - an activity that REX uses to remove soil at a rapid rate in a coarse manner.
- *REX* - the notation for Robotic EXcavator.
- *scanning* - an activity that REX uses to locate buried objects before excavating.
- *secondary excavation* - an activity that REX uses to remove soil at a slow rate in a fine manner.
- *simulator model* - the data structure containing the site size, and all objects locations within the site for which the excavation process simulates.
- *site* - the volume in which the excavation is confined
- *tools* - the hardware used by the activities to perform their operation.
- *working domain* - a cuboid that represents a volume that the excavation process can work within as it has already been cleared by the process.
- *world model* - the data structure containing the site, the domains, and all objects currently known to the Monitor/Controller.