Computer Developments for Mining Robotics

M.L. Maher, I.J. Oppenheim, D.R. Rehak

Department of Civil Engineering
Carnegie-Mellon University
Pittsburgh, PA, 15213

Introduction

Deployment of robots in the unstructured environment engages a set of problems which are absent in conventional factory robotics. Those new problems include two which are partly addressed in this paper:

- The geometrical interaction of a robot with a workspace (a domain) and with work objects. This problem requires a domain model to prevent robots from colliding with other solid objects, and to control their movements.
- The control of machinery in the presence of forces and displacement conditions fed back to it by the work process. Presently it is the human operator who controls machine positions, speed, force (and so on) so as to avoid machine damage, jamming, etc. Robotization of such equipment will require control procedures which can replace the operator.

Extensive research is demanded to solve these problems, and a large portion of that research is already under way. Research activities extend to sensor research, software engineering, theoretical/analytical studies, robot demonstrations, and so on. This paper reports on two efforts, both representative of computer methods studies, which have been directed at the two problem areas noted above and which are associated with the application area of underground mining. One is the development of a prototype data and command system (a software package) for operating a robotic roof bolter. The second is a study of using an expert system to replace the human operator in one narrow aspect of continuous miner operation. Both studies are small efforts, the first in each case, and represent starting points for further work.

1. A Prototype Data/Command System for Roof Bolting

Problem Statement

A software package was written to control the actions of a hypothetical roof bolter represented by the object shown in Figure 1. It has a base, 2m x 2m x 1m, and a manipulator modelled as a cylinder (0.15m diameter) which extends normal to the base over an extension range of 0 to 5m. "Bolting" is represented as positioning the robot in plan and extending the manipultor to contact the roof. The bolter has three degrees of freedom: two of base movement and one of manipulator extension. The bolter should be imagined to have treads, permitting it to make a general x-y move through a sequence of rotations and translations.

The mine is represented by a general collection of connected rooms. While each room may tend to be prismatic, the program was developed to permit considerable dimensional irregularity on any surface. For instance, floors may have sloping sections and walls or roofs may be irregular in section, containing non-planar facets, and so on. Regular geometry can, of course, be modelled easily using this more general capability.

The program was developed as a graduate project undertaken by three students. It was written in PASCAL on a PERQ Systems personal engineering workstation. One objective was to make the program user-friendly, envisioning it as a console from which an operator could control a mine area. As such, the program operates with screen graphics input-output and the robot commands are presented by menu to the user. Nonetheless, the program was equally designed to be driven by a higher-level program, in which case the command set forms the interface between the two programs. This paper addresses the user mode of operation, as it is more demonstrative.

Modelling

Figure 2 shows one domain segment (one room) modelled in a boundary element representation as used for this study. That representation is expedient if, as in this case, the surfaces can be pictured as a set of vertices on a planar grid, supplemented by elevations normal to the plane. This may be considered a 2 1/2-D (as opposed to 3-D) representation. It permits local irregularities to be modelled to any accuracy while not necessitating extensive storage for a detailed model of the entire mine map (often several square miles). A rectangular or trapezoidal wall can be defined using only 4 corner points. In Figure 2 the roof and right-hand wall are represented on a finer grid because they contain out-of-plane features modelled as faceted surfaces, as shown.

A data structure was created using a small set of different domain segments (channel, tee, cross, boxend, elbow, object) which connect (at vertices) by simple rules. The basic mine is represented by a directed graph, with each node being a domain element and the links representing the faced of the elements. A general mine layout can be built out of this set. An individual domain segment can be represented only by corner points (if prismatic). Vertex data for each face is stored in an array. Thus, the data structure can model accurately while providing fast access to the data. Since extensive memory management was not implemented, the prototype program is limited in the size and detail which can be represented. In the first phase demonstration that limitation was accepted.

A number of primitive functions were built into the program and they are the central features needed for creating a robotic command system. The most important is interference checking, to determine if the robot bolter interferes with the solid environment or with other objects within the space. The interference check operates during robot moves to signal a collision as a robot moves along a path. A move which does not trip an interference flag is one which can indeed be accomplished. If interference is found, it shows the point in the intended path at which the interference occurred. For user convenience in assembling sample problems, the interference checking can be turned off: this permits the user to lift a robot and place it elsewhere in the space without having to route it through the actual passages.

Another function is distance measurement along any ray from the robot. For instance, this command could be issued normal to the robot base (locally upwards) to determine the distance at which the roof is encountered. A bolting command would then order the manipulator to extend that distance. Similarly, if bolting is to be performed at a preset spacing, distance measurement to the closest wall would be useful in determining whether another bolt placement is indicated.

Another primitive function is to place or remove additional objects into the space. Face removal can be modelled in this manner, as can the introduction of another equipment into the workspace.

The program has other useful features which are not actually primitive functions but are more related to I/O capabilities. For instance, it is possible to preview an operation. While the robot remains at its present location, the operator can also picture its presence at a second location to study whether its movement there would be practical or not. Similarly, the graphic display can change its view between plan, longitudinal section, and transverse section; all views are accompanied by windowing and zooming display capabilities. Sequences can be recorded for future replay, and if the environment has been altered (through removal of material) the updated domain can be stored.

Program Description

Figure 3 shows the PERQ screen in its default configuration. The screen is divided into four areas, which will be described in turn.

The upper left region contains a plan view of the overall mine on an x-y coordinate system as shown. A symbol (*) denotes the present location of the robot. The untypical mine shape in Figure 3 is modelled from eight domain segments, and was selected to demonstrate the capacity for modelling irregular volumes. It is possible to window and zoom on the plan view. A robot move is reflected as a translation of the cursor in the plan view. In ordering a move, it is possible to direct the robot by moving the cursor with the mouse.

The large lower region is a transverse section, taken through the robot, showing the tunnel cross-section and the robot. The faceted nature of the roof is seen in this section. The view in this main window can be changed from the transverse section (called simply <u>section</u>) to a longitudinal section (called an <u>elevation</u>). A sloping floor or roof would be evident in such an elevation view. Robot movement or manipulator extension is seen in this main window, which also serves as the window when a preview is requested. The plan view can be moved into this main window also, all under control of the operator.

The area on the right above the main window is the status window. It contains the global coordinates of the robot, the manipulator extension status, the tunnel dimensions, and the error messages, if any.

The area at the topmost right is the menu window, showing the main menu. Sub-menus appear if an appropriate command from the main menu is chosen. Menu choices can be made with a cursor and mouse (not shown). The MOVE choice brings up a sub-menu to direct the robot movement. The INTERFERER ON choice permits switching the interference checking function on or off. The EXT/RETRACT (<u>ext</u> end/<u>retract</u> manipulator) choice generates a request for a command and argument to govern manipulator movement. The CHANGE choice brings up a sub-menu to permit altering of the display, windowing and zooming.

A series of sample exercises were run, demonstrating all the program capabilities and the ease of use. While this constituted only a first phase software package, it did demonstrate the various elements needed for a data and command system for robotic applications.

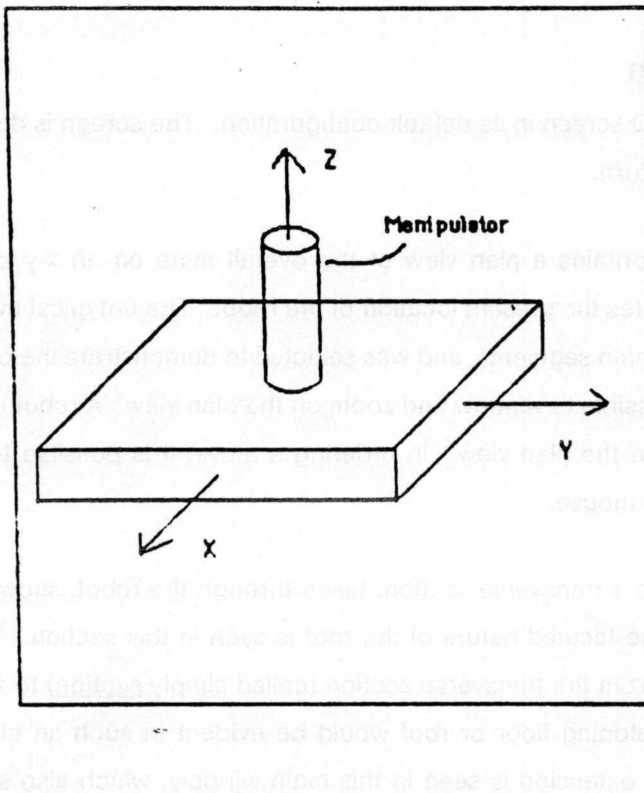

Figure 1: Robot Model and Coordinates

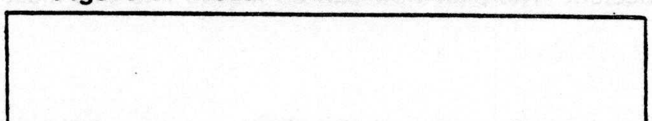

Figure 2: Domain Segment Example

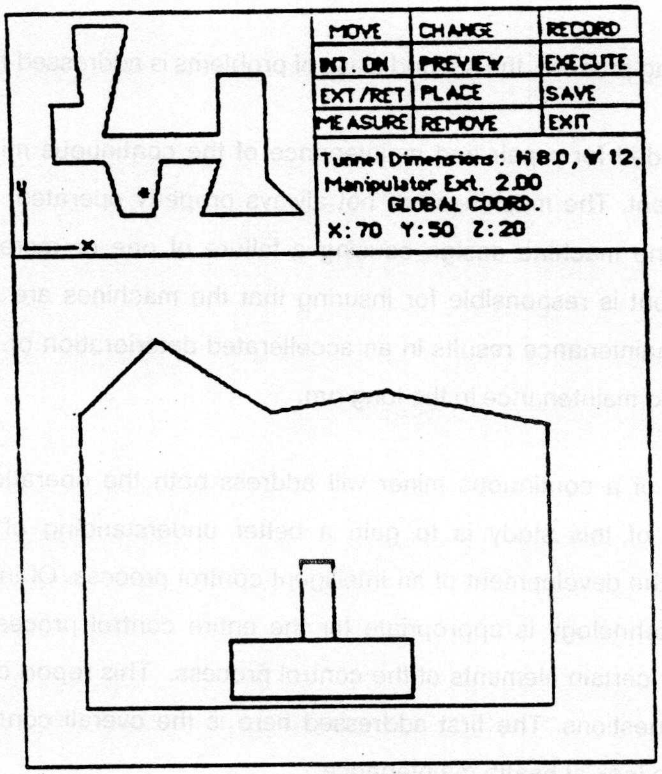

Figure 3: Typical Output Display

2. An Expert System for a Continuous Miner

Problem Background

The continuous miner, currently operated manually, has been identified as a potential machine to be used for the investigation of mining robots. The approach to robotic research in this case is an evolutionary one, where the initial mechanical configuration of the robot is modeled after an existing machine. The major difference in existing continuous miners and robotic continuous miners is that of intelligent automatic control.

The short term objective of this study is to improve the productivity of the continous miner. The productive time of the continuous miner is a relatively small percentage of the total time in a work shift. This small percentage is due to many factors which may be classified into two categories:

1. The method of mining coal when using a continuous miner causes logistical problems. For example, after a certain amount of coal has been cut, the roof support system must be extended into the most recently mined region requiring that machines be moved. This process detracts from the time a continuous miner could be cutting coal.
2. The repair and maintenance of the continuous miner takes up a considerable amount of time.

Since the objective of this study is to improve the productivity of the continuous miner without disturbing the entire mining process, the second class of problems is addressed here.

The extended time needed for repair and maintenance of the continuous miner has two causes: operation and management. The machines are not always properly operated; the human operator may push the limits of the machine design causing a failure of one or more components of the machine. The management is responsible for insuring that the machines are properly maintained. Improper or ineffective maintenance results in an accellerated deterioration of the machine thereby requiring more repairs and maintenance in the long run.

The intelligent control of a continuous miner will address both the operation and management problems. The purpose of this study is to gain a better understanding of how expert system technology could assist the development of an intelligent control process. Of interest to this study is whether expert system tehnology is appropriate for the entire control process or whether expert systems may be used for certain elements of the control process. This report discusses an ongoing effort to answer such questions. The first addressed here is the overall control process, and the second is a more in depth look at health maintenance.

Automatic Control

An overview of the automatic control issue is in progress. The overall control process has been mapped out providing some insight on where an expert system may be appropriate. Due to the large number of different configurations of the continuous miner, an effort is being made to separate the machine independent information from machine sepcific information. For example, the diagram in Figure 4 illustrates an inference network for determining if the machine may start. In this figure the machine independent information is represented in rectangles and the machine specific information is indicated by circles.

Although the intelligent control study is in the early stages, a few areas have been identified as appropriate for an expert system approach using current expert system technology. These areas are presented below.

- Assessment of Roof Conditions. In order for the continuous miner to safely continue to cut coal, it is necessary to assess the roof conditions. There is potential here for an expert system to draw inferences from accustic emissions and supply the control system with information on the current condition of the roof.

- Assesment of Floor Conditions. It is also necessary for the control system to assess floor conditions and recognize the potential for floor heave.

- Prediction of Explosive Seams. The control system must be able to characterize and predict potential suddenly explosive seams in order to determine the proper driving rate.

- Health Monitoring. The control system must be able to monitor the health of the machine and be able to discontinue any action that may result in catastrophic failure.

These items are not exhaustive, there are potentially many areas appropriate for an expert system approach.

The overall intelligent control issue is not readily addressed by current expert system technology. Some of the issues in developing such a system are real time porcessing of sensor data leading to the problem of real time control and feedback loops. The satisfaction of these issues requires pushing the limits of current expert system technology.

Health Maintenance

Computer assisted machine maintenance can take two approaches. In one approach, the computer indicates when regular maintenance should be scheduled. This is essentially a database problem where the database will serve as the storage mechanism for maintenance items, and an algorithmic program may be written to interface with this database.

The second approach requires the assessment of machine health and the recommendation of maintenance as indicated by the current status of the machine. This approach may recommend requirements that regularly scheduled maintenance may not predict. In many cases, the conditions of a particular mine may require that certain items on the miner be maintained by more frequent replacements and repairs and other items less frequently. The maintenance the continuous miner according to its health rather than according to a schedule derived from typical use of the miner may actually serve to decrease the amount of time and money spent on maintenance. This problem is appropriate for an expert system approach and will be discussed further.

In order to be consistent with intelligent, automatic control of the continuous miner, the monitoring system for health maintenance will be sensor based, eventually requiring no human interaction to assess the health state of the machine. Three systems have been identified for maintenance purposes: hydraulic, electrical, and mechanical. A depth first approach at this point would be more useful than a breadth first, indicating that, rather than develop an expert maintenance system for all three systems, we are concentrating on the knowledge engineering for one system, namely, the hydraulic system.

The development of an expert system for health maintenance monitoring requires the definition of an inference network defining the relationships between sensor data and maintenance requirements. The emphasis in this study is *expert system development*, assuming that once appropriate sensors have been identified, the details of hardware requirements for these sensors can be worked out. The

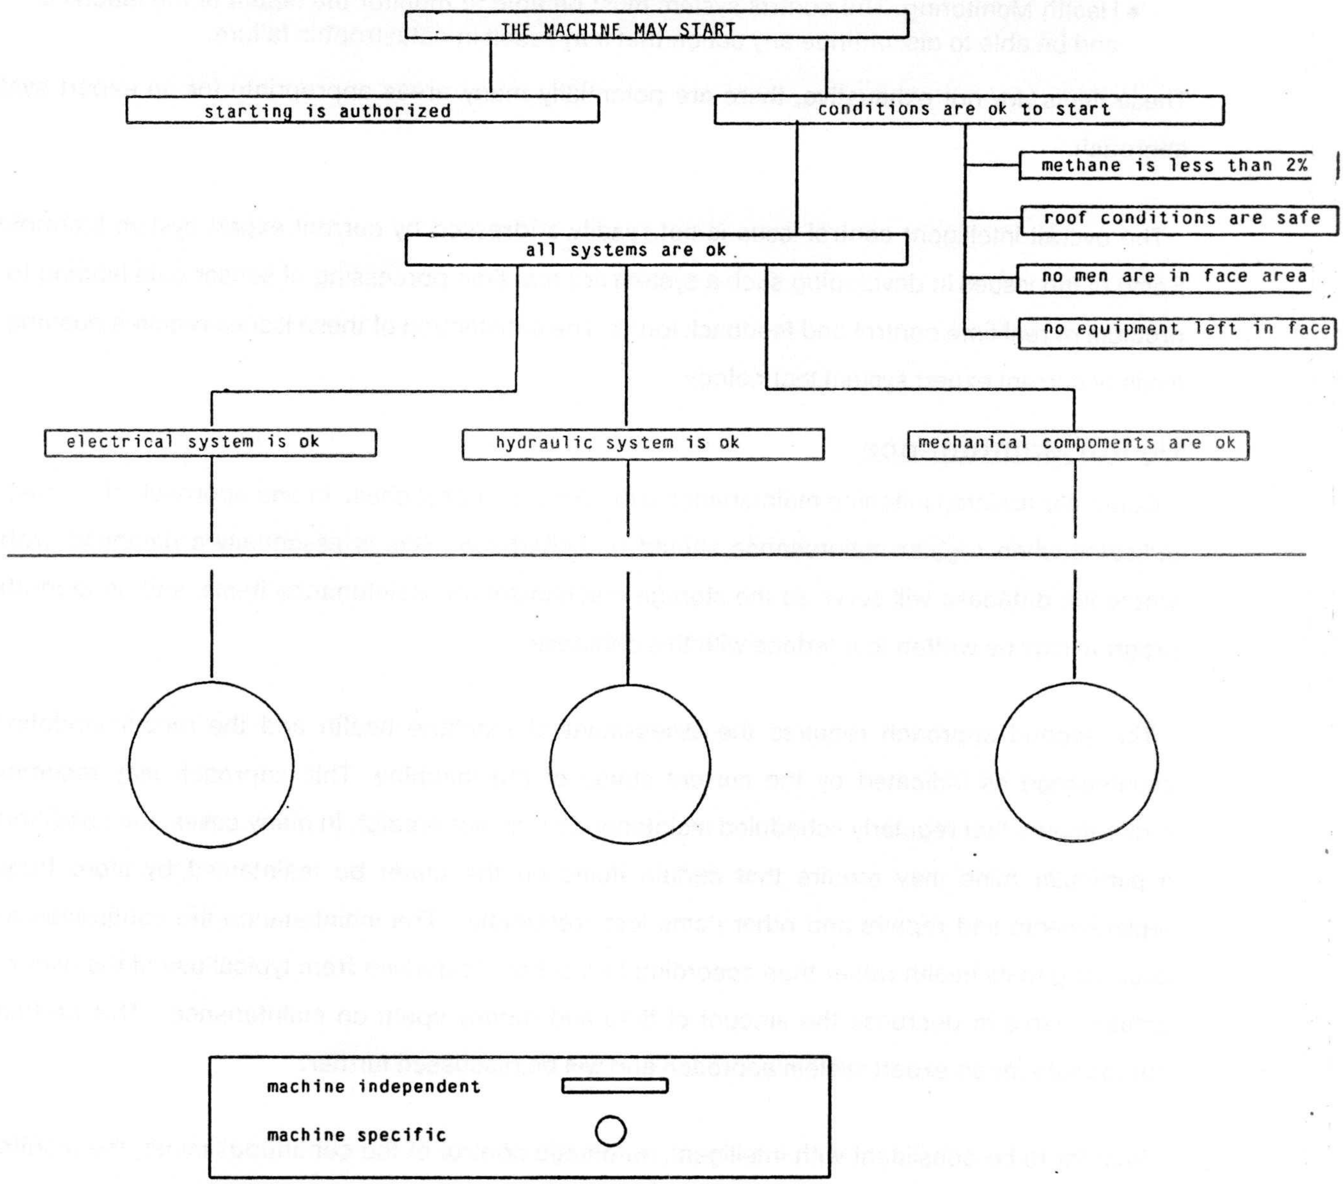

Figure 4: Inference Network For Start Conditions

difficulty in developing the expert system lies in formulating the expertise into an appropriate inference network that may be used as a model for the mechanical and electrical systems, where the inference network is a tree-like graph in which the leaf nodes represent input data and the root node(s) represent hypotheses (or in this case, maintenance requirements).

Several inference networks are being derived and discarded as the study progresses. Due to the many varieties and manufacturers of continuous miners, an attempt is being made to keep the machine independent knowledge separate from the machine specific knowledge. The network in Figure 5 illustrates the formalization of machine independent information for the hydraulic system. The network represents the monitoring of the continuous miner for a global hydraulic system problem, where a circuit problem requires machine specific information.

In the preliminary stages of the expert system development it is useful to use the monitoring system in two modes: user supplies the data or sensors supply the data. The user supplies data mode is employed during development stages to facilitate the formalization of expertise. It is easier for a human expert to reason about the status of the continuous miner that he can check manually then for him to visualize automatic data acquisition from sensors that are not currently in place on the machine. The human's experience is derived from operating the machine using his own human sensors supplemented with measurements. In order to capture this expertise, an expert system is being developed that relies only on this kind of information.

The next stage of development is to determine the location of the actual sensors and formalize the knowledge on the interpretation of these sensors to provide the input information the human expert provides in the version described above. Once sensors have been identified, the sensor based monitoring system uses a mixed initiative control strategy. As described above, in an inference network the root of the tree indicates a hypothesis and the leaves represent input data. In a mixed initiative control strategy the expert system uses both a top down or bottom up strategy, selecting the strategy most appropriate for the current state of the problem.

The initial control strategy, a bottom up approach, requires that the sensors are montitored at predefined time intervals. Depending on the sensor values, the monitoring will continue or one or more hypotheses will be pursued. The sensor values will be interpreted and classified as being in one of three ranges:

1. Sensor value is normal.

2. Sensor value indicates that maintenance is needed.

3. Sensor value indicates that machine should be shut down.

A sensor value in the first range will cause the monitoring to continue. A sensor value in the third range will cause an emergency shut down flag to be sent to the operator. A sensor value in the second range will trigger a top down strategy by proposing a set of hypotheses to be pursued, indicating the order in which they should be checked. Once the top down strategy is invoked, only sensors that are relevant to the proposed hypothesis are checked.

Acknowledgements

Study number 1 was performed under the direction of the authors by three students: Mark Borland, Scott Swetz and Leslie Veach. We are indebted to them for their efforts. Facilities and computers were supplied by the Civil Engineering Construction Robotics Laboratory and several faculty colleagues shared their previous work with us. The study was enhanced greatly by the participation of Mr. Tom Fisher and others from the U.S. Bureau of Mines, Pittsburgh Research Center, serving as intended users for the student project effort.

Study number 2 is being carried out by two research assistants: Nitin Pandit from Carnegie-Mellon and Mark Levin from the Colorado School of Mines. Professor J. Kohler from the Pennsylvania State University serves as an expert consultant to this effort. It is sponsored by the U.S. Bureau of Mines under research contract HO358021. We appreciate the participation and guidance of Drs. G. Schnakenberg and J. Durkin from USBM.

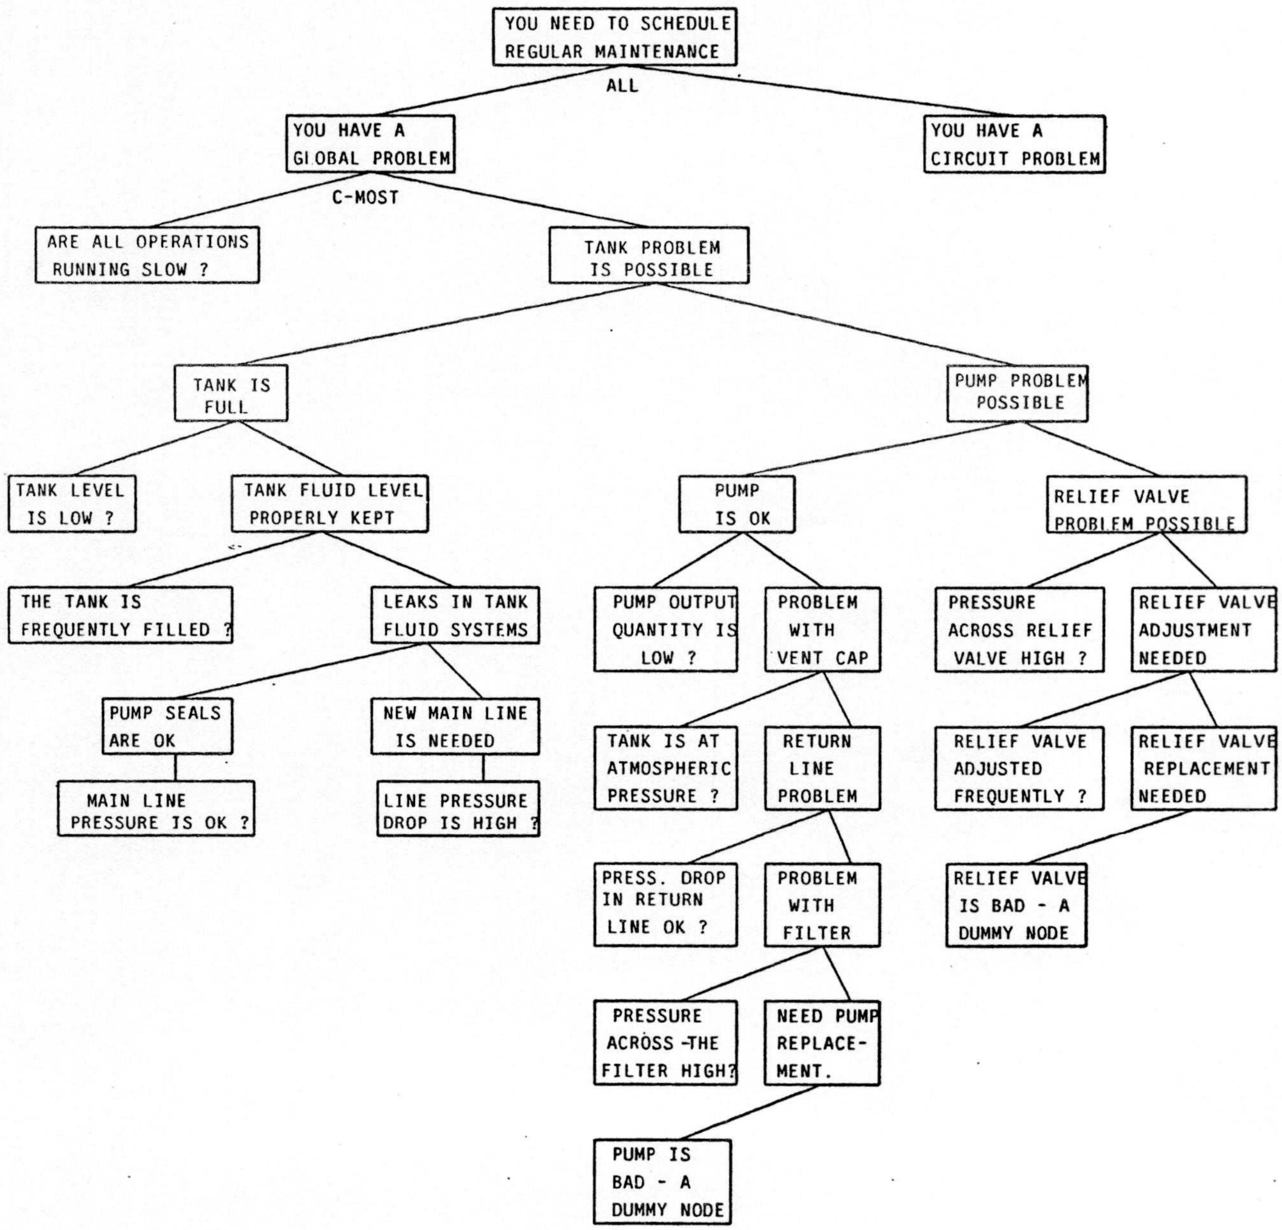

Figure 5: Inference Network For Hydraulic System Maintenance