

AN ARCHITECTURE FOR REAL-TIME SENSORY- INTERACTIVE CONTROL OF ROBOTS IN A MANUFACTURING FACILITY

J. S. Albus, C. R. McLean, A. J. Barbera and M. L. Fitzgerald

*Industrial Systems Division, Center for Manufacturing Engineering,
National Bureau of Standards, Washington, DC, USA*

Abstract. A hierarchical architecture is described for a robot integrated into a real-time sensory interactive factory control system. In this architecture, high level goals are decomposed through a succession of levels, each producing strings of simpler commands to the next lower level. The bottom level generates the drive signals to the robot actuators. Each control level is a separate process with a limited scope of responsibility. Each performs the generic control function of sampling its input and generating appropriate outputs. The input is characterized by three types of data - a command from the next higher level, processed sensory data, and status feedback from the next lower level. The outputs are of three types - a command to the next lower level, a request for sensory information to the processing module at the same level, and a status feedback to the next higher level. This paper describes this generic control structure and its implementation in a real-time sensory-interactive control system for a manufacturing facility.

Keywords. Hierarchical systems; robots; control theory; sensors; industrial control.

INTRODUCTION

A hierarchical control system such as the one shown in Figure 1 is partitioned vertically into levels of control. The basic command and control structure is a tree, wherein each computational module has a single superior, and one or more subordinate modules. The top module is where the highest level decisions are made and the longest planning horizon exists. Goals and plans generated at this highest level are transmitted to the next lower level where they are decomposed into sequences of subgoals. In general, the decomposition at each level takes into account information derived from: (a) sensors that measure the state of the environment, (b) reports from lower control levels as to the state of the control hierarchy itself, and (c) predictions (or expectations) generated by models, knowledge bases, and inference engines.

At each level, input subgoals are decomposed into sequences of output sub-subgoals in the context of the state of the environment, the next lower level of control, and the internal store of knowledge. Finally at the lowest level, output drive signals are computed and sent to the physical actuators. Each path, or chain of command, from the bottom to the top of the hierarchical tree can be further segmented, as shown in Figure 2, into three separate hier-

archies: (1) a goal, or task decomposition, hierarchy (H); (2) a feedback processing hierarchy (G); and (3) a world model hierarchy (M). This has been discussed in a number of previous papers (Albus and co-workers, 1981a, 1981b, 1981c).

In general, each of the hierarchical levels shown in Figure 1 can be further partitioned into sublevels. For example the control hierarchy for a machining work station robot shown in Figure 2, decomposes the equipment level into three sublevels. An assembly robot might decompose the equipment level into four or more sublevels. The number of sublevels required depends on the complexity of the tasks that must be decomposed at that level.

At all levels, the H, G, and M modules are concurrent processes produced by real-time programs executing simultaneously in each module. Perhaps the simplest way to treat this mathematically is to implement each of the modules in the hierarchy as a finite-state automaton. The state-graph describing the activity of the entire hierarchy can then be described by a Petri diagram (Albus, Barbera, Fitzgerald, 1982).

For each module in this architecture there are three concepts of time: the planning horizon, the response time, and the cycle time. The planning horizon is the interval over which a control module plans into the

future. The response time is the delay between a change in a module's input and the generation of a new output. The cycle time is the period between sampling the input variables. In general, the response time will be slightly longer than the cycle time, and the planning horizon will be many times longer than the response time.

The response time of the finite-state automata at each level depends on the requirements for stability and dynamic response at the respective levels. The response time requirement is shorter at the lower levels, but the complexity of the control computations is less. The response time is longer at the higher levels, and the complexity of the computations is greater. Thus, the total computational power required at any level of the hierarchy is more or less constant.

Communication between the various modules in such a system can be accomplished by writing messages in a common data base. Each message area (or mailbox) within the data base can be restricted so that only one system may write into it, although many can read its contents. If the cycles of the state-clock at all levels are synchronized, information transfer into and out of the common data base will occur at predictable time increments and each message can carry a time tag.

The flow of commands and control define the structure of the hierarchical tree. Each branch, or subtree, corresponds to a functional unit. There may, however, be informational pathways which exist outside the chain of command. Information of many types (such as timing signals, and knowledge of certain important events) need not follow the strictly tree-like pathways of the control hierarchy. Such information transfer pathways may crisscross the formal structure of the control hierarchy.

TASK DECOMPOSITION

In the robot control system architecture shown in Figure 2 the bottom (or first) level of the task decomposition hierarchy is where coordinate transforms and servo computations are made, and all joint motions are scaled to hardware limits on velocity and force.

At the second level, elemental movements (such as <REACH TO (A)>, <GRASP>, <LIFT>, <ORIENT ON (B)>, <MOVE TO (X)>, <RELEASE>, etc.) are decomposed into force and velocity trajectories in a convenient coordinate system. That coordinate system may be defined in the robot's work space, in the part, or in a coordinate frame in the robot's gripper.

At the third level, simple tasks (such as <FETCH (A)>, <MATE (B) TO (A)>, <LOAD TOOL (C) WITH PART (D)>, etc.) are decomposed

into elemental movements which can be interpreted by the second level.

In the Automated Manufacturing Research Facility (AMRF) currently under construction at the National Bureau of Standards (Simpson, Hocken, Albus, 1982), a machining workstation robot will receive input commands to the third level of its hierarchical control system from a WORKSTATION CONTROLLER, which is the fourth level in the robot's hierarchy, as indicated in Figure 2.

A typical machining workstation in the AMRF will consist of a robot, a machine tool, a work tray buffer, and several tools and sensors that the robot can manipulate. Trays of parts and tools will be delivered to the workstation by a robot cart. The workstation controller will be given commands by the cell controller consisting of lists of operations to be performed on the parts in the trays. It is the task of the workstation controller to generate a sequence of simple task commands to the robot, the machine tool, and any other systems under its control so that the set of operations specified by its input command list are carried out in an efficient sequence. For example, the workstation controller may generate a sequence of simple task commands to the robot to setup the clamping fixtures for the first part; to the machine tool to perform the specified machining operations; to the robot to modify the clamping fixtures for the next job; etc. The time horizon for the workstation may vary from several hours up to about a day, depending on the complexity and number of parts that are being processed.

The fifth level of the robot control hierarchy in Figure 2 is the CELL CONTROLLER which is responsible for managing the production of a batch of parts within a particular group technology part family. The task of the cell is to group parts in trays and route the trays from one workstation to another. The cell generates dispatching commands to the material transport workstation to deliver the required tools, fixtures, and materials to the proper machining workstations at the appropriate times. The cell must have planning and scheduling capabilities to analyze the process plans for each part, the tooling and fixturing requirements, and the machinability time estimates for each operation. It will use these capabilities to optimize the make-up of trays and their routing from workstation to workstation. The planning horizon for the cell will depend on the size and complexity of the batch of parts in process, but will be on the order of a week.

The sixth level in the robot control hierarchy is the SHOP CONTROLLER which performs long term production planning and scheduling. It also manages inventory, and

reports shortages in materials and tools to the next higher level (Facility Control) where orders are issued to outside vendors. The planning and scheduling functions are used to determine the workstation, robot, and material resources requirements for each cell. The shop then dynamically allocates workstations to, or reclaims them from the cells as necessary to meet the production schedule (McLean 1982). The control structure always satisfies the rules of a hierarchical tree at any instant in time, but the subtrees of the shop may shift from one moment to the next. For example, an individual robot may belong to Workstation #1 of Cell #26 at one moment, and to Workstation #16 of Cell #2 the next. The same logic can be applied to the passing of workstations between cells. This degree of flexibility becomes important in factories or construction sites where robots are mobile and rapidly move from one physical work site to another.

The seventh level is FACILITY CONTROL. It is at this level that engineering design is performed and the process plans for manufacturing each part, and assembling each system, are generated. Here also, management information is analyzed, materials requirements planning is done, and orders are processed for maintaining inventory. Because of the very long planning horizons at this level in the control hierarchy, the activities of the facility control module are not usually considered to be a part of a real-time control system. However, in the context of hierarchical control with exponentially increasing time horizons at each higher level, these facility control activities can be integrated into the real-time control hierarchy of the manufacturing system.

FEEDBACK PROCESSING

Each level of the task decomposition hierarchy is serviced by a feedback processing module which extracts the information needed for control decisions at that level from the sensory data stream and from the lower level control modules. The feedback processing modules at each level detect features, recognize patterns, correlate observations against expectations, and format the results to be used in the decisions and computational procedures of the task decomposition modules at that level.

At the lowest level of the robot hierarchy, the feedback processing modules extract and scale joint positions and force and torque data to be used by the servo and coordinate transformation computations.

At the second level, touch and proximity data, and simple visual measurements of distance and positions of grip points are extracted from the sensory input to be used in computing trajectory end points.

At the third level the three dimensional positions of visual features such as edges, corners, and holes are computed and combined to determine the position and orientation of surfaces and volumes of objects. Identities of objects may also need to be computed (or recognized) in order to generate the reaching and grasping commands at this level.

At the fourth (WORKSTATION) level, relationships between various objects need to be determined, in order to sequence simple task commands.

At the fifth (CELL) level, the location and composition of trays of parts and tools and the length of queues of parts need to be determined. This may be derived from sensors which read coded tags on trays, or may be inferred from sensory input from lower level sensors on the robot or in the workstation.

At the sixth (SHOP) level, the condition of machines, tools, and the amount of inventory on hand must be determined in order to generate schedules, allocate resources, and evaluate and set priorities for production.

At the seventh (FACILITY) level, the requirements for changes in part design, or in process plans need to be recognized in order to make engineering changes, or redesign parts or processes.

THE WORLD MODEL

The world model hierarchy, made up of M modules in Figure 2, consists of a knowledge base containing all the information currently known about the task, the parts, or the workplace. The M modules also contain procedures that allow them, based on the state of the task and other contextual information from various places in the hierarchy, to compute expectations and predictions about what the sensory data to the corresponding G module should be. This allows the G modules at each level to compare expectations with observations, and to measure both the degree of correlation and the degree of difference. A strong degree of correlation means that the proper model is being matched with the incoming sensory data. It means that the observed object or situation has been correctly recognized, and that information contained in the model can be safely used for decision making even though it may not be directly observable by the sensory system.

A large degree of difference between expectations generated by the model and observations derived from sensors means that either an incorrect choice of models has been made, or the model has not been correctly transformed spatially or temporally so as to generate the proper

set of expected feature relationships, or that the incoming sensory data is too noisy, or is being improperly processed and filtered. In this case, the computational problem for the task decomposition module is to decide which type of error is being encountered and what is required to remedy the discrepancy. In general, this type of problem can be solved either by a set of situation/action rules of an expert system, or a set of heuristic search procedures. At high levels the world model contains machinability data which enables it to predict tool wear, tool life, expected times to completion, etc.

At low levels, the world model contains dimensional information describing the shapes and sizes of parts. The M modules use this to generate expected positions of image features such as edges, corners, holes, and surfaces.

At the facility control level the model contains information about machining processes, material properties, shop processing capabilities, and expected lead times for procurements which can be used to compute estimated completion times for various production plans. At the shop level, the world model contains information about machine capabilities, tool life, and inventory levels and is able to simulate the performance of various cell configurations. At the cell level, the model contains information about workstation task times, and is able to simulate the expected performance of various hypothetical task sequences. At the workstation level, the world model contains knowledge of expected tray layouts including the names of parts and their approximate positions, orientations, and relationships. At the simple task level, the model is able to generate detailed positions and orientations of three dimensional objects in a robot or machine tool coordinate system. At the elemental move level, the model is able to generate expected positions and orientations of specific features of parts and tools, such as edges, corners, surfaces, holes, and slots. At the coordinate transformation and servo level, the model generates windows or filter functions that are used to screen and track the incoming raw data stream.

PROGRAMMING A HIERARCHICAL CONTROL SYSTEM

In a hierarchy of finite-state automata such as shown in Figure 2, a state-graph, or Petri diagram, is a useful program development tool. A state-graph is much like a flow chart for a procedural program, and the state-

transition table (which is the program for a finite-state automaton) is readily derived from the state-graph.

For example, the state-graph shown in Figure 3 is a representation of the set of operations required by the <FETCH (A)> task at the third level of the hierarchy shown in Figure 2. This state-graph shows several alternative actions to account for possible failure modes in the FETCH command. Figure 4 shows the state-transition table corresponding to the state-graph of Figure 3.

At any level in the task decomposition hierarchy such a state-transition table can be used by a computing structure such as shown in Figure 5 to generate the task decomposition function appropriate to that level. The left-hand side of the table consists of all the command, internal state, and feedback inputs that can be encountered at any tick of the state-clock. The right-hand side consists of an output command (and/or a pointer to a procedure which computes an argument which becomes part of the output command) to the next lower level, a next internal state, and a report to the next higher level or to other modules at the same level.

At the first hierarchical level, the number of lines in the state-transition table will be small. The left-hand side of the table consists simply of variables which select the type of coordinate transformation required and the type of servo computations needed.

At the second level, the left-hand side consists of variables which define the type of trajectories to be generated. The right-hand side contains pointers to procedures that compute forces, positions, accelerations, and velocities in the appropriate coordinate systems.

At the third level, the left-hand side consists of variables which specify the state of the environment as reported by sensors and the right-hand side the names of appropriate elemental movements to be made for each state. Pointers to procedures are used to compute arguments and modifiers.

At the higher levels, the state-tables may be compared to production rules in expert systems. Procedures that are invoked by these state-tables may consist of heuristic search algorithms or linear programming techniques for generating plans, schedules, etc.

CONCLUSION

There are a number of advantages of the table-driven hierarchical control system described above. The first is that it partitions the problem into simple, well-defined modules with clearly specified inputs, outputs, internal states, and rules for state-transitions. The control problem is partitioned vertically with respect to task complexity and abstraction, horizontally with respect to function (such as task decomposition, sensory processing, and world modeling), and along the time axis by the use of a state-clock. This simplifies the design and eases the synchronization of simultaneous processes in the many different computing modules.

The second advantage is that it facilitates the handling of error conditions. If additional states need to be defined to deal with unanticipated error conditions, they can simply be inserted into the state-graph and added to the state-transition table. There is very little interaction with other parts of the program code.

Third, this approach formalizes the control problem into a very orderly structure. Each line in the state-transition table for any module is an IF/THEN production rule. <IF (the command is such, and the state is so, and the feedback conditions are thus) / THEN (the output is whatever is stored on the right hand side of the table, and the system steps to the indicated next state)>. The addition of each node or edge to the state-graph, and the corresponding lines added to the state-transition table is the equivalent of the addition of a new chunk of knowledge about how to deal with a specific control situation at a particular point in a problem domain at a unique phase in the task execution. This system architecture thus bridges the gap between servomechanisms and finite-state automata at the lower levels and expert system technologies at the upper levels.

Fourth, the state table representation allows the easy insertion of conditional tests of new sensory or feedback data. For example, if a new touch sensor is added, it is possible to merely insert a new column in the feedback portion of the state-transition table, and a new line, or lines, in the state-transition table for each of the new edges to be added to the state-graph.

Fifth, it facilitates debugging. Each state of the system is formally identified and the set of conditions that lead to and from that state are clearly specified. Diagnostic routines can be used to read these states from common memory and stop the state-clock when necessary. This makes it easy to perform traces, to set break points, and to reason backwards from error states. The system is completely deterministic and errors in logic are simple to reconstruct. Program bugs are therefore relatively easy to locate and correct.

Sixth, it makes it possible to build teaching and learning capabilities into sensory-interactive robot control systems. Simple programs with few or no error conditions can be defined first by state-graph flowcharts. The system can then be run until it encounters unanticipated problems or undefined conditions which cause an error <STOP>. Then each problem condition can be dealt with specifically by adding lines to the state-transition table to address that particular problem state. Eventually the entire space of problem states will contain programmed solutions, and the frequency of error <STOP>'s will decline. The rule-based approach facilitates the integration of an expert system which would have the capability to use its higher "learning" rules to modify the world model, sensory processing, or decision structures, as experience on different problem states is acquired.

There are, of course, many unanswered questions regarding this approach. The hierarchical control system architecture described in this paper is still largely a theoretical construct. There are many problems that remain to be addressed and many of the details of this method will undoubtedly change as more experience is acquired during the construction and testing of the experimental system.

REFERENCES

- Albus, J.S., Barbera, A.J., Fitzgerald, M.L. (1982). Programming a Hierarchical Robot Control System. Proc. 12th International Symposium on Industrial Robots, Paris, France, 9-11, June 1982.
- Albus, J.S., Barbera, A.J., Fitzgerald, M.L. (1981a). Hierarchical Control for Sensory Interactive Robots. Proc. 11th International Symposium on Industrial Robots, Tokyo, Japan, 7-9 October 1981.

- Albus, J.S., Barbera, A.J., Fitzgerald, M.L., Nashman, M. (1981b). Sensory Interactive Robots. Proc. of 31st General Assembly, International Institution for Production Engineering Research (CIRP), Toronto, Canada, September 1981.
- Albus, J.S., Barbera, A.J., Nagel, R.N., (1981c). Theory and Practice of Hierarchical Control. Proc. 23rd IEEE Computer Society International Conference, September 1981, 18-39.
- Barbera, A.J., Fitzgerald, M.L., Albus, J.S. (1982). Concepts for a Real-Time Sensory-Interactive Control System Architecture. Proc. 14th Southeastern Symposium on System Theory, Blacksburg, Virginia, April 1982, 121-126.
- McLean, C.R. (1982). The Virtual Manufacturing Cell. Proc. 4th IFAC/IFIP Symposium on Information Control Problems in Manufacturing Technology, Gaithersburg, Maryland, 26-28 October 1982.
- Simpson, J.A., Hocken, R.J., Albus, J.S. (1982). The Automated Manufacturing Research Facility of the National Bureau of Standards. Journal of Manufacturing Systems, Society of Manufacturing Engineers (in press).

This is to certify that this article was prepared by U.S. Government employees as part of their official duties and is therefore a work of the U.S. Government and not subject to copyright.

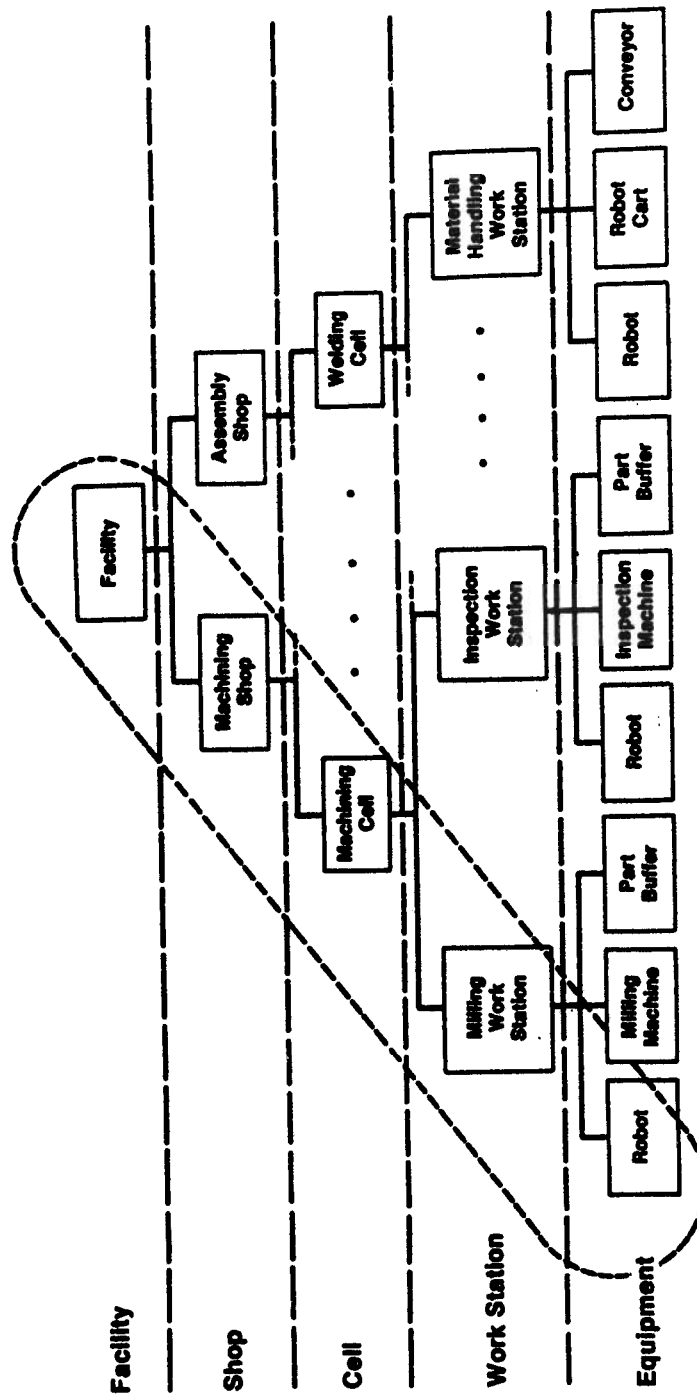


Fig. 1. A hierarchical control system architecture for integration of robots, machine tools, and material transport facilities into an automatic factory. The chain of command enclosed in the dotted lines is detailed in Fig. 2.

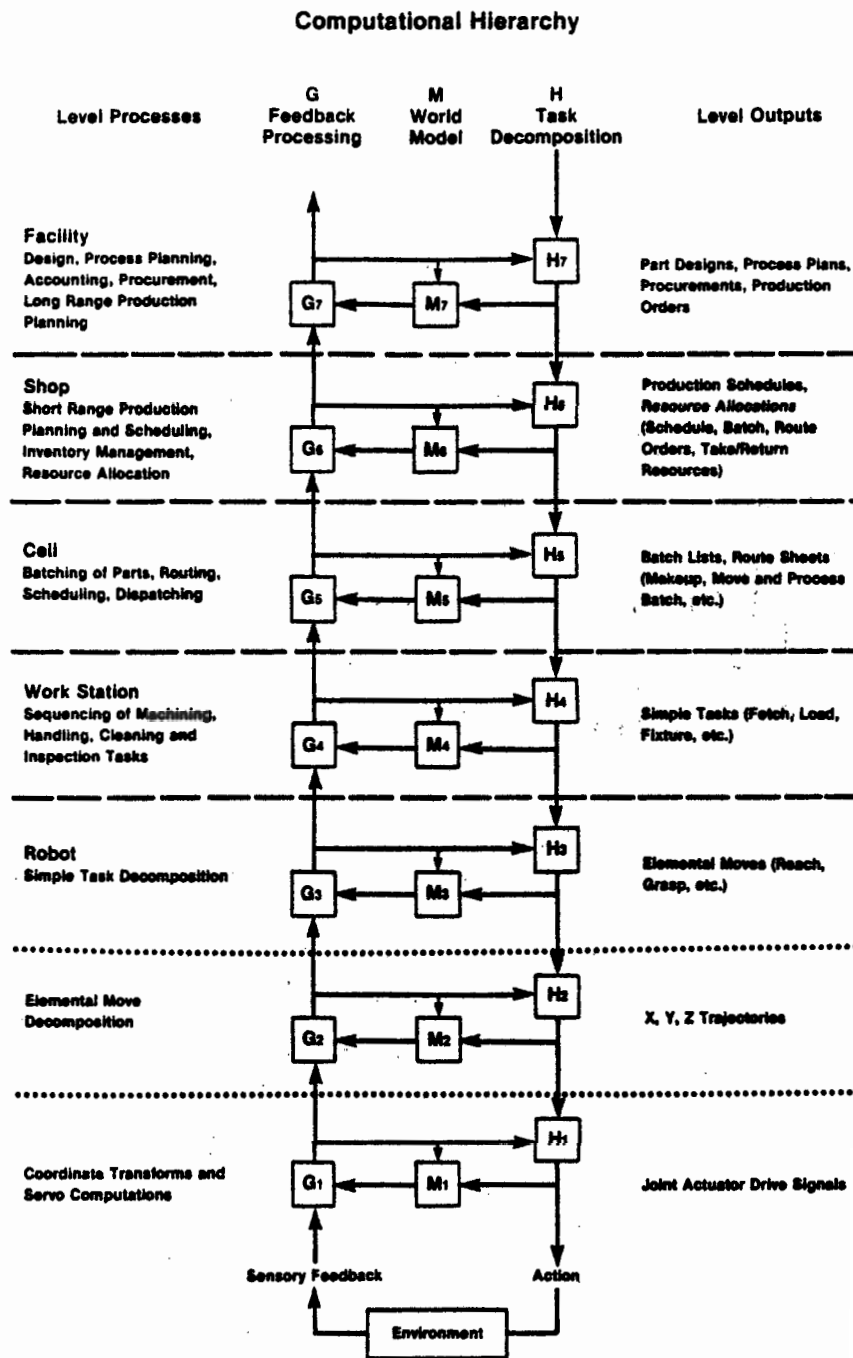


Fig. 2. The computational hierarchy for a robot in a machining work station. This hierarchy corresponds to the chain of command enclosed in dotted lines in Fig. 1.

A State-Graph Representation of Fetch (A)

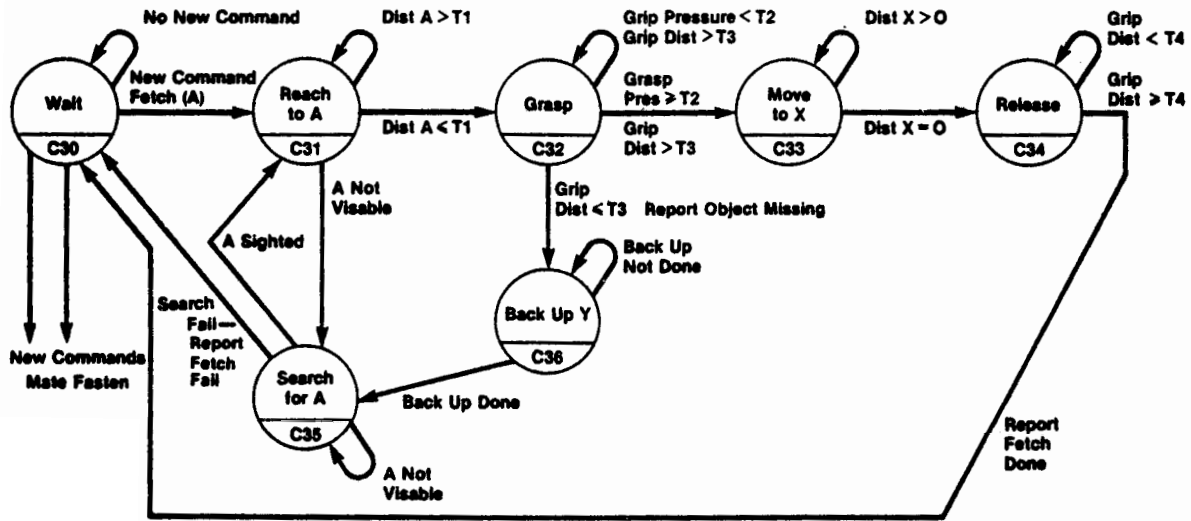


Fig. 3. A state-graph at hierarchical level 3 for a decomposition of the <Fetch (A)> command.

A State-Transition Table Representation of Fetch (A)

Command	State	Feedback	Next State	Output	Report
—	C30	No New Command	C30	Wait	—
Fetch (A)	C30	New Command	C31	Reach to (A)	—
"	C31	Distance to A > T1	C31	Reach to (A)	—
"	C31	Distance to A < T1	C32	Grasp (A)	—
"	C31	A Not Visible	C35	Search for (A)	—
"	C32	Grasp Pressure < T2 Grip Dist > T3	C32	Grasp (A)	—
"	C32	Grasp Pressure ≥ T2 Grip Dist > T3	C33	Move to (X)	—
"	C32	Grip Dist ≤ T3	C36	Back Up (Y)	Object Missing
"	C33	Distance to X > O	C33	Move to (X)	—
"	C33	Distance to X = O	C34	Release	—
"	C34	Grip Dist < T4	C34	Release	—
"	C34	Grip Dist ≥ T4	C30	Wait	Report Fetch Done
"	C35	A Not Visible	C35	Search for (A)	—
"	C35	A in Sight	C31	Reach to (A)	—
"	C35	Search Fail	C30	Wait	Report Fetch Fail
"	C36	Back Up Not Done	C36	Back Up (Y)	—
"	C36	Back Up Done	C35	Search for (A)	—

Fig. 4. A state-transition table representation of the state-graph shown in Fig. 3.

A Computing Structure Designed to Execute State-Transition Tables




Fig. 5. A computing structure designed to execute state-transition tables of the type shown in Fig. 4.