FILE COPY

AN ARCHITECTURE FOR DECISION-MAKING IN THE FACTORY OF THE FUTURE

Richard H. F. Jackson
Center for Applied Mathematics
National Bureau of Standards
Gaithersburg, MD 20899

and

Albert W. T. Jones
Center for Manufacturing Engineering
National Bureau of Standards
Gaithersburg, MD 20899

ABSTRACT

A major manufacturing research facility is being established at the National Bureau of Standards. The Automated Manufacturing Research Facility has been designed to address the standards and measurement needs for the factory of the future. A five-layer hierarchical control architecture is under development to control the various production and support activities needed to drive that factory. The proper execution of many of these activities requires the solution to one or more optimization problems. This paper proposes a decision-making hierarchy which parallels that control architecture, describes the problems that exist at each level within that hierarchy, and discusses the work underway at NBS to address some of those problems.

KEYWORDS:
Automated Manufacturing, Factory Model, Flexible Manufacturing, Hierarchical Control, Real-time, Optimization, Scheduling, Routing.

1. INTRODUCTION

Manufacturing plants typically contain various combinations of people, computers, and machines, working together to maximize corporate profits from the goods they produce. Many of these plants are not meeting this goal. They are plagued by large work-in-process inventories, low utilization of equipment, insufficient throughput, and excessive delays, all of which tend to decrease profits.

Hopes for alleviating these conditions were raised when computer-controlled robots, machine tools, and transporters became commercially available. Many companies made large capital investments in this new equipment. They also acquired the sophisticated computer and database support needed to control this equipment. The result of integrating these new technologies into existing plants is often referred to as CIM (Computer Integrated Manufacturing) or FMS (Flexible Manufacturing System). The effects of this investment and integration effort were expected to be increased profits and larger shares of the world markets.

In general, this has not happened. In fact, introducing these new CIM technologies has the potential for an even greater negative impact. There are three major reasons for this surprising phenomenon. First, integrating equipment from different vendors was far more difficult than ever anticipated. This demonstrated the pressing need for software and hardware interface standards. Second, the continued use of existing planning and scheduling strategies often exacerbated the problems mentioned above. This resulted in very expensive but idle equipment and made it impossible to increase profits or achieve the desired rate of return on capital investments. Finally, because of the increased dependency on "data" in these CIM environments, poor data management and communication strategies have also caused equipment to be idle.

This paper describes the approach under development at the National Bureau of Standards (NBS) to address these issues within the Automated Manufacturing Research Facility (AMRF) [Simpson, Hocken, and Albus, 1982].

The remainder of this paper is composed of six sections. After this introductory section, section 2 provides an overview of the AMRF, including its design philosophy, hierarchical control architecture, and data management system. In sections 3 and 4, we identify the production management and data management decision problems that exist in the AMRF. Our approach to solving the sequencing and scheduling problems is discussed in section 5. Conclusions are given and future work is outlined in section 6. References are provided in section 7.

2. THE AMRF

2.1 Overview

The National Bureau of Standards has a fundamental commitment to promote the development of standards for automated manufacturing systems and to transfer technology to American industry. To meet this responsibility, the Center for Manufacturing Engineering at NBS has established an experimental test bed, the Automated Manufacturing Research Facility (AMRF) [Simpson, Hocken, and Albus, 1982]. Industry, academia, and other government agencies have played an active role in this development effort through direct appropriations, equipment loans, and cooperative research programs.

Physically, the AMRF contains several robots, machine tools, storage and retrieval systems, two wire-guided vehicles, and numerous computers. This equipment includes donations and purchases from four different robot manufacturers, three machine tool vendors, and every major computer company. This diversity of suppliers has forced NBS researchers to focus on designing and testing uniform software and hardware interfaces and data exchange formats to address the problems involved in system integration. Factory control software, a manufacturing data preparation system, a distributed data management strategy, a communications system, and numerous sensors have been developed. These individual hardware and software components have been successfully integrated into a working, flexible, small batch manufacturing system (see Figure 1).

* * * Figure 1 about here. * * *

2.2 Design Philosophy

The AMRF is intended to exhibit a greater degree of flexibility and modularity than any currently available FMS. To achieve these goals, the AMRF has adopted the following design philosophies concerning its control architecture. It is:

- o partitioned into a functional hierarchy in which decision making and control functions reside at the lowest possible level,
- o intended to respond in real-time to performance data obtained from machines equipped with sensors,
- o implemented in a distributed computing environment using state-of-the-art techniques in software engineering and artificial intelligence, and

o designed so that control processes are completely data-driven and communicate via NBS-developed hardware and software interfaces which are uniform throughout the AMRF.

As noted, the AMRF control architecture is based on the classic hierarchical, or tree-shaped, command/feedback control structure (see Figure 2) typical of many complex organizations [Albus, et al. , 1984]. This approach ensures that the size, functionality, and complexity of individual control modules is limited. In addition, each control level is completely data-driven. That is, the data required to perform its functions is separated from the actual control code.

Each module decomposes input commands from its supervisor into procedures to be executed at that level and subcommands to be issued to one or more subordinate modules (see Figure 2). This decomposition process is repeated until, at the lowest level, a sequence of coordinated primitive actions is generated which actuates shop floor equipment [Albus, Barbera, and Nagel, 1981]. The status feedback that is provided to supervisors by their subordinates, is used to close the control loop and to support the adaptive, real-time, decision making discussed in sections 3 and 4.

* * * Figure 2 about here. * * *

Although the flow of control in this hierarchy is strictly vertical and between adjacent neighbors only, it is necessary and even desirable to share certain classes of data across one or more levels. All data is managed by a completely separate, distributed data administration system. The concept of "common memory" is used to transmit to and from control processes via a communication network.

The following sections provide a brief description of the functional decomposition and data management methodology employed in the AMRF. Details can be found in the references provided.

2.3 Functional Decomposition

An analysis of traditional small batch manufacturing systems provided the foundation for the decomposition [Mclean, Mitchell, and Barkmeyer, 1983] of the manufacturing functions into five levels: facility, shop, cell, workstation, and equipment (see Figure 3). A brief discussion of the responsibilities assigned to each of these levels is given below. More details can be found in [Jones and McLean, 1986].

* * * Figure 3 about here. * * *

2.3.1 Facility Level

This level is responsible for implementing the "front office" functions that are typically found in manufacturing facilities. Activities at this level are grouped into subsystems that fall into three major functional areas: manufacturing engineering, information management, and production management.

Manufacturing engineering functions are typically carried out with human involvement via user-data interfaces. This includes computer-aided design (CAD), Group Technology Classification, and process planning. The information management activities provide user-data interfaces to support necessary administrative or business management functions. Production management tracks major projects, generates long-range schedules, identifies production resource requirements, determines the need for additional capital investments to meet production goals, determines excess production capacity, and summarizes quality performance data.

2.3.2 Shop Level

This level is responsible for coordinating the production and support jobs on the shop floor. It is also responsible for the allocation of resources to those jobs. Two major component modules have been identified within shop control: a task manager and a resource manager.

The task manager of the shop level system is responsible for capacity planning, grouping orders into batches, assigning and releasing batch jobs to cells, and tracking individual orders to completion. The resource manager is responsible for allocating the production resources to individual cells, managing the repair of existing resources, and ordering new resources.

2.3.3 Cell Level

At this level, batch jobs of similar parts are sequenced through workstations and supervision is provided for various other support services, such as material handling and calibration. The cell [Jones and McLean, 1984] brings some of the efficiency of a flow shop to small batch production by using a set of machine tools and shared job setups to produce a family of similar parts. The AMRF cells are dynamic production control structures which permit the time sharing of workstation level processing systems. This software structure was named the "virtual" cell to distinguish it from previous "real" manufacturing cells which are defined by fixed groupings of equipment or machinery on the shop floor. A detailed discussion of the virtual cell concept is found in [McLean, Hopp, and Bloom, 1982].

2.3.4 Workstation Level

The activities of small integrated physical groupings of shop floor equipment are directed and coordinated at the workstation level. A typical AMRF workstation consists of a robot, a machine tool, a material storage buffer and a control computer. Machining workstations process trays of parts that are delivered by the material handling system. The controller sequences equipment level subsystems through job setup, part fixturing, cutting processes, chip removal, in-process inspection, job takedown, and cleanup operations.

2.3.5 Equipment Level

These are "front-end" systems that are closely tied to commercial equipment or industrial machinery on the shop floor. Equipment controllers are required for robots, NC machine tools, coordinate measuring machines, delivery systems, and storage/retrieval devices. Equipment controllers perform two major functions: 1) translate workstation commands into a sequence of simple tasks that can be understood by the vendor-supplied controller, and 2) monitor the execution of these tasks via the sensors attached to the hardware. These controllers will be required for "off-the-shelf" equipment to provide extended functionality and compatibility with NBS control concepts, until higher level front-ends are incorporated by system vendors.

2.4 Data Management

As indicated in 2.2, the control modules described in the preceding section are completely data-driven. The management of that data is a key ingredient in the AMRF. The data management function is concerned with providing shared data to all manufacturing processes in a timely, accurate, and completely transparent manner. This function is complicated by both the manufacturing and computing environment in which it must be performed. The manufacturing environment requires dynamic and frequent updates to the data directory, data delivery paths which are separate from the existing control structure, and local but efficient storage of data for real-time operations. The computing environment consists of heterogeneous systems with different data manipulation languages, data management functions, formats, types, and structures. These constraints imply that data will, of necessity, be physically distributed around the factory.

NBS researchers have proposed an architecture [Barkmeyer et al., 1986] called IMDAS--Integrated Manufacturing Data Administration System--to manage this distributed data. IMDAS is completely separate from the control hierarchy, and has been specifically designed to meet the manufacturing and computing requirements described above. It consists of a three-level hierarchy of data management services: the Basic (BDAS), the Distributed (DDAS), and the Master (MDAS) Data Administration Service modules. The major functions of these modules are described below.

2.4.1 Basic Data Administration Service--BDAS

A BDAS exists on every component system within the AMRF, and it provides the services required to access data residing anywhere in the AMRF. These services include interprocess and network communication, data and command translation, and data management.

If required, interprocess communication is achieved by using shared memory. This approach permits data to be accessed by several different processes without any explicit action by the originator. Communication between component systems is achieved using the bottom four layers of the Open Systems Interconnect (OSI) seven-layer model [Data Processing, 1981]. A global shared memory scheme has been implemented in which data stored in a local shared memory is replicated into the shared memory areas on remote components which require a copy of that data.

Whenever data is moved from one system to another, it must be translated from the source representation to the target representation. Each BDAS is capable of translating from its own representation to an IMDAS-defined common representation, and vice versa. This translation includes type, syntax, structure, and format.

IMDAS also includes a global data manipulation language for making database queries. This implies that each BDAS must have a command translator to translate from this global language into the query language or access mechanism understood by the local physical data management tool. Typically, this tool will be either a simple file server, memory manager, or full database manager.

2.4.2 Distributed Data Administration--DDAS

The middle level in the IMDAS hierarchy is the Distributed Data Administration Service (DDAS). It is responsible for providing data management services to all processes residing on component systems assigned to it. Each DDAS has six major functional modules: distributed service executive, data manipulation language service, query mapping service, transaction manager, data dictionary service, and the data assembly service.

The Distributed Service Executive module provides the interface between the DDAS and the parent control system, local BDASs, and the MDAS (see below). It is also responsible for all initialization, coordination, and recovery procedures.

The Data Manipulation Language (DML) parses queries from the IMDAS Global DML into a tree of primitive operations. It then determines which of these operations it cannot perform. These are passed up to the MDAS. The remaining operations are then sent to the Query Mapping Service for eventual transfer to a local BDAS.

The Query Mapping Service decomposes and restructures each query into one or more queries to be executed by subordinate BDASs. This decomposition must also take into account the capabilities of the data server (DBMS, file servers, etc.) managing each BDAS database. Each of these new queries is sent to the transaction manager for execution.

The Transaction Manager (TM) is responsible for the control and management of distributed queries. In performing this function, the TM must also enforce integrity constraints, concurrence, consistency, replication, and recovery rules.

The Data Directory Service at the DDAS level integrates the directory information provided by the subordinate BDASs. This includes data location, structure, and delivery paths. The Data Assembly Service combines data received from multiple sources and formats the result in the GDML for transfer to a local BDAS.

2.4.3 Master Data Administration Service--MDAS

The Master Data Administration System (MDAS) coordinates the activities of multiple DDASs. This coordination includes managing the master data directory, directing query execution, resolving concurrence problems among DDASs, and controlling global initialization, integration, and recovery procedures. The internal functions of the MDAS are identical to those performed at each DDAS. It parses a query from a particular DDAS, decomposes that query into a tree of operations, determines which operations to route to the other DDASs, and manages the execution of those operations.

3. DECISION PROBLEMS IN PRODUCTION MANAGEMENT

In this section, we identify the decision problems that affect the actual production of parts on the shop floor. We believe that these problems exist in any automated manufacturing facility. Following the AMRF design philosophy described in 2.2, we have partitioned these decision-making problems to match the control hierarchy described above. The manufacturing data required to solve these problems--equipment, times, alternatives, and precedence relations--are contained in process plans.

As one moves down this decision-making hierarchy, several important observations can be made concerning the nature of these problems. First, each level must sequence through the list of jobs assigned by its supervisor, and develop a schedule of tasks for its subordinates. Second, there is a dramatic increase in the number of problems to be solved and the frequency with which they must be resolved. Third, there is a significant decrease in the time available to find solutions. Finally, the information used to solve them becomes more abundant, complete, and deterministic. These properties will have a tremendous impact on the techniques used to solve problems at different levels within this hierarchy.

What follows is, in a sense, a laundry list of problems, some classical and well-recognized, others new, arising from the introduction of new technologies into manufacturing systems. Indeed, some of the problems may seem insignificant now, but as our ability to understand and control these CIM systems increases, the marginal gain from having optimal solutions to these problems will also increase.

3.1 Facility Level

The facility level has sole responsibility for the business, and strategic planning functions which support the entire manufacturing enterprise. Better mathematical models are required to aid top management in assessing and justifying the potential benefits and costs of flexible automation. In addition, once the decision has been made to employ this technology, new techniques are needed in cost accounting, depreciation, capital investment strategies, and many other business functions [Eiler, 1986]. Existing methodologies are unable to measure the impacts of this flexibility in a meaningful way.

Another function performed at the facility level is the manufacturing data preparation crucial to the actual part production. Schedules must be generated for all of the activities required to complete this preparation. These schedules will include both new customer requests and revisions to existing data required by changing conditions on the shop floor. In addition, new methods are needed to aid in the classification and coding of parts from CAD data, geometric modeling, decomposition of complex geometries into primitive features that can be machined and inspected, and the design, revision, and verification of process plans.

3.2 Shop Level

The shop level receives a list of customer requests and any assigned priorities or due dates from the facility level. The shop level sequences through these requests, groups them into batches, and determines the order in which these batches will be released to the manufacturing cells on the shop floor. It then produces a schedule which indicates the cells to be used for each batch, estimated start and finish times at each cell, and the required material transfers among those cells. These plans must be updated any time a new request is issued, an existing request is cancelled or given a higher priority, or a significant problem occurs.

The shop also has overall responsibility for inventory control, tool management, capacity planning, and preventive maintenance for all equipment in the shop. These activities must be managed to support the schedules developed at this level.

An important issue to be resolved at the shop level is future use of existing techniques for Material Resource Planning and Master Production Scheduling. In an environment like the AMRF, in which decisions are pushed down to the lowest level, these global planning approaches may no longer be applicable. However, this is still an open question.

3.3 Cell Level

A cell controller must coordinate the activities of its subordinate workstations to complete the jobs assigned by the shop. Each job will require the services of one or more workstations including material handling and will usually have some due date and priority associated with it. The cell must sequence through these jobs and develop a schedule of anticipated start and finish times, and priorities for each job at each workstation. It must determine which workstations will be needed, and the order in which they will be needed. It must also arrange for the requisite material transfers in support of that schedule. When conflicts or delays are reported by a workstation controller, the cell must replan, reroute, and reschedule to overcome them.

Coordinating the activities at these workstations becomes more difficult when there exist shop-wide, shared resources like material transport devices. In addition, the introduction of "virtual cells" (see 2.2.3) will also complicate the problems both at the cell and the shop levels.

3.4 Workstation Level

As noted above, each workstation controller coordinates the activities of its subordinate equipment to execute a series of tasks assigned by a cell controller. Although the exact nature of the tasks are workstation-dependent, they typically consists of receiving materials, shipping materials, setup, takedown, and a list of features to be machined or inspected. The workstation controller must generate a sequence in which to perform these tasks and a schedule for each of its subordinates.

In addition to the aforementioned problems, the material handling workstation controller has several other problems that it must address. These special problems are directly related to its primary responsibility of planning and coordinating the activities required to move trays of materials around the factory. It must locate the material, assign a transportation device (or devices) to pickup and deliver that material, and determine the routes it will follow in executing the task. Further, all these activities must be coordinated and monitored for possible changes and updates.

Assigning trays to batches of parts must also be addressed. This problem is complicated in an environment in which a batch size of one or two is the rule rather than the exception. In this case, a single tray could contain several batches of parts, each having a different geometry. Further complications are that deliveries to more than one workstation may be combined on a single tray and that each transportation device may be capable of carrying more than one tray.

3.5 Equipment Level

The last level to be discussed is the equipment level, the lowest level in the hierarchy. There are three classes of equipment: stationary robots, machine tools, and material storage, retrieval, and transport devices. The mathematical decision problems to be solved by each equipment controller fall into two major categories. The first is sequencing and scheduling. Each controller must sequence through the current tasks assigned by its supervisory workstation. They may be rank-ordered, with expected completion times associated with each task. In addition, the controller must schedule and coordinate the activities required to execute these tasks. These activities will be performed by the subordinate systems to each particular controller (see below). The second set of problems is equipment-dependent, and discussed in more detail in the following sections.

3.5.1 Robots

Robots are used primarily to locate, move, and handle materials such as parts, tools, and fixtures. In addition, they perform housekeeping duties to remove chips during machining, and assemble and disassemble fixtures. Typical subsystems are vision, multiple hands and grippers, and other actuators. In addition to the sequencing and scheduling problems discussed above, robot controllers have several, more time-critical problems to solve. They include path generation, optimal routing for traversing parts, loading and unloading materials, and tray layout.

All robots are required to maneuver through three-dimensional space as part of their routine activities. This necessitates the generation of paths to allow the robot to move from one point to another. This problem is complicated by the fact that the robot's work space is filled with obstacles. If the position of these objects remains fixed, then this problem can be solved off-line, and to optimality. If, however, obstacles are constantly moving into and out of the work space, or changing position within the work space, then this becomes a real-time problem. In this case, it may be necessary, due to time constraints, to replace optimality with a sub-optimal, yet feasible and easily generated path.

Once the robot has reached its destination, it must then carry out some specified task. It may need to pick up a part, to place a part in a fixture, insert a tool into a tool drum, or any of a number of other similar activities. Each of tasks demands the "precise" positioning of the robot arm(s) before the activity can commence. The relative or absolute precision required will depend on the activity and the capabilities of the robot. For instance, a robot equipped with a vision system does not require the same precision as a robot without a vision system. This is an important problem and could be viewed as a solution to a nonlinear optimization problem in which the objective is to minimize the error in the actual or relative position.

Another area where optimization methods can be brought to bear is in the loading, unloading and layout of trays. In some respects, portions of the problems are scaled-down facility layout problems. Thus, some of the ideas from the facility layout and design literature could be useful. However, all of these problems can be complicated by the likelihood that multiple geometries may exist in the same confined space within a tray.

There is an interesting optimization problem concerned with finding optimal routes for traversing parts for inspection, cleaning, and deburring. These tasks usually require several different end-effectors such probes, deburring tools, etc. The objective would be to perform these activities in a way that is optimal with respect to some measure, perhaps time, number of two-handed moves, end-effector changes or part repositioning.

Pattern recognition for robot vision systems is another area where significant optimization problems appear. These range from simple nonlinear least squares problems that arise from attempting to match patterns, to more complicated nonlinear least squares problems that arise in combining small windows of bit patterns to form larger windows for faster scanning.

The robot carts that serve the workstations must address some of the same problems as the fixed-position robots; they may, however, take on a slightly different look. For example, path calculations for the robots become routing problems for the carts. The issue here is deciding which path to take to deliver or pick up trays from the workstations. If the cart can travel forward and backward, the problem becomes more complicated. The situation is further complicated by having multiple carts, although the coordination activity for this is performed at the next higher level. The layout of the wire-guided path is also a task that lends itself to mathematical analysis and could be studied to determine the best paths to lay down.

3.5.2 Machining Centers

The AMRF contains three CNC (Computer Numerically Controlled) machining centers: horizontal, vertical, and turning. They are capable of performing several metal removal operations, and limited, on-line inspection of parts and tools. In addition, the AMRF has a Coordinate Measuring Machine (CMM) to perform off-line inspection of machined parts. Typically, each machining center must coordinate the activities of a tool holder(s), part holder(s), spindle(s), and coolant sprayer(s). The CMM controls a rotary table, probes, and several other axes of motion. Each of these controllers is responsible for sequencing and scheduling assigned tasks. Examples of these tasks are tool and collet changes, remounting parts on pallets, chip removal, and the actual machining and inspection operations. These problems should be solved to optimality with respect to some performance measure, such as number of tool changes, number of refixturings, time in queue, or number of late tasks. Again, as noted with the robot controllers, these problems must be solved often and quickly.

Machining centers have several other problems related to the storage, selection, and use of tools. The storage problem is essentially a tool layout problem. The placement of tools in a drum (or other similar device) can impact the total time required to machine a set of features. Consequently, the exact arrangement of tools can be represented as an optimization problem in which the objective is to minimize the time required to access the tools required to perform a set of machining tasks. This assumes that the tools have already been selected, and the order in which they will be used is also known. The solutions to these two problems become constraints in the tool placement problem. Before the actual cutting can begin, a tool path, depth of cut, speed and feed must be generated. Finally, it is necessary to determine which tools will be kept for later jobs and which should be sent for storage or use elsewhere.

3.5.3 Automated Storage and Retrieval System

Automated storage and retrieval systems (AS/RS) are used to house raw, in-process, and finished parts, as well as robot end-effectors, fixtures, and tools. Basically, two decision problems must be addressed. The first is to determine the optimal size and location of these devices throughout the factory: this is typically an off-line problem. The second problem is concerned with the layout of the storage areas. One would like to store all of the materials required for a particular job in a contiguous area within a single AS/RS. But, since storage areas are assigned and released frequently, this may not be possible. Consequently, this becomes a dynamic storage allocation problem whose solution will have consequences for the time required to transfer these items to the required location for processing.

4. DECISION-MAKING PROBLEMS IN DATA MANAGEMENT

In this section, we identify the decisions involved in executing the data management functions for the AMRF. These decisions can be partitioned into three categories: storage, administration, and communication.

4.1 Data Storage Problems

Within the AMRF, data is physically stored on several different devices. The need to distribute data physically across the manufacturing facilities is motivated by the time-criticality factor involved in many data requests. This is especially true at the equipment levels of both the control and decision-making hierarchies described above. Several optimization problems arise as a result of this decision. First, there is the selection of the actual storage devices and their data management capabilities. In some cases, a simple file server will suffice: in others, a sophisticated data base management system will be required. Another set of problems are concerned with the location of data files: 1) how many copies are needed, 2) where are they stored, and 3) which is the master copy.

4.2 Data Administration

The distribution of data across a heterogenous collection of computer systems has a significant impact on the administration of that data. Typical administration functions include: 1) satisfying data requests, 2) ensuring the accuracy and consistency of the data itself and all data dictionaries, and 3) maintaining concurrence control over all duplicated data. IMDAS (see 2.4) is a three-level, distributed administration system designed to perform these functions. Each level manages a queue of data requests. Each request must be decomposed into a "query-tree" of more primitive database operations. These operations may be carried out at the same level or, possibly, by one or more modules at the next lower or next higher level. The fact that data is distributed complicates this decomposition process. Although techniques are available for completing this decomposition within a centralized administration system [Chu, 1986], little is known about approaches to solving this problem in an environment like IMDAS.

There are also sequencing and scheduling problems associated with managing these queues which contain both complex data requests and primitive database operations. These problems have similar characteristics to those described in the preceding section. However, they are complicated by the difficulty involved in 1) determining the time required to complete a task, 2) obtaining a "due date" for a given task, and, 3) coordinating the parallel activities at all three levels which may be involved in the completion of a single complex data request. Little is known about approaches to solving these problems.

4.3 Data Communication

The transfer of information between computer processes in an automated manufacturing environment will be managed by a Data Communication System (DCS). This implies that processes do not communicate directly with one another; rather, they simply make information available. It is the DCS's responsibility to deliver this information to those processes that require it, at the time they require it. The protocols for accomplishing this data transfer are being specified in the Open Systems Interconnection standards

[Data Processing, 1981] commonly referred to as MAP (Manufacturing Automation Protocol).

It is likely that many of the techniques used to design and manage computer networks [Kleinrock, 1976] the "virtual cell" notion (see 2.3.3) will have a significant impact on both the design and real-time administration of the DCS. The difficulties arise from the need to "disconnect and connect" processes, and therefore origin-destination pairs and paths, from one another in real-time. To date, these issues have not been addressed.

5. A DECOMPOSITION OF SEQUENCING AND SCHEDULING PROBLEMS

An important goal that remains for the successful completion of the AMRF is to expand the control hierarchy described above into a goal-directed hierarchy in which both planning and control functions are carried out at every level. In the preceding sections, the intent was to convey an appreciation of the benefits of hierarchical decomposition of the control problems within automated manufacturing facilities. The question addressed in this section is whether it is possible to mimic this hierarchical approach to control and decompose the planning process along the same lines described above for the facility as a whole.

5.1 Classical Approaches

Typically, a shop manager is responsible for sequencing and scheduling all jobs on the shop floor. The result of this effort is usually a GANTT chart [Baker, 1974] showing the start and finish times for each job on each machine. The manager is also required to update this chart frequently to account for changes in job and equipment status.

The literature abounds with mathematical programming, simulation, heuristic, and other techniques to aid the manager in solving these problems [Jackson and Jones, 1986; Graves, 1981; Raman, 1986; Sen and Gupta, 1984]. However, because of their computational requirements and restrictive assumptions, these approaches tend to have limited applicability in a real manufacturing environment. In applications such as this, a major constraining factor is the amount of computation that can be performed in real-time at each level. This is limited by the cycle time: the period of time over which each module is responsible for planning and updating local goals. While this horizon may be long at the highest levels in the hierarchy, it becomes very short at the lowest levels.

Researchers [Chiodini, 1986; Fox, 1983; Steffen and Greene, 1986; Wysk, Wu and Yang, 1986] have recently begun to build expert systems to address these problems. Unfortunately, this approach has met with little success. Expert systems are expensive and time consuming to build and they are still computationally inefficient. Furthermore, there are not enough "scheduling experts" around to create the required knowledge bases. Consequently, they are unable to provide useful solutions in a timely manner.

5.2 The Hierarchical Decomposition Approach

Another recent trend is to develop hierarchical planning and scheduling systems. Several researchers [Bitran, 1977; Davis, 1984; Gershwin, 1986; Lawrence and Morton, 1986] have proposed such systems. They are restricted in scope and limited to two or three levels. This paper also proposes a decision-making hierarchy for both production and data management. In both cases, sequencing and scheduling problems appear at every level. Each level is given a set of goals by its supervisor from which it defines a set of goals for its subordinates. These goals consist of an ordered list of jobs with proposed start and finish times. These new goals must be consistent with those set by the module's supervisor and they must commit the entire subordinate structure to a unified and coordinated course of action.

It is important to understand that, from a mathematical point of view, this is a new approach. Heretofore, scheduling was performed at one time for the whole system, and usually off-line, as mentioned above. This new approach proposes that each level be responsible for generating and maintaining its own schedules. These schedules should be created quickly and only as needed. They should also obey the constraints imposed by higher levels. These constraints are in the form of priorities among jobs, and start and finish times for each job.

5.3 A Proposed Modeling Methodology

As noted in the preceding section, each level within the decison-making hierarchy is responsible for generating schedules for the jobs assigned by its supervisor. We will model each of these jobs as a "project" comprised of several activities. These activities are related through a set of precedence constraints and resource requirements. Consequently, the scheduling problem at <u>each</u> level can be viewed as a multi-project, constrained resource problem.

Several researchers [Elsayed, 1982; Kutulus and Davis, 1982; Pritsker, Waters, and Wolfe, 1969] have developed techniques to solve multi-project, constrained resource problems. In particular, Pritsker <u>et al</u>. modeled this problem using zero-one decision variables, subscripted by project, task within a project, and time period, i.e, x_{ijt}. Using this idea, $x_{ijt} = 1$, if job j in project i was completed during time period t, and $x_{ijt} = 0$ otherwise. Objective functions are developed in their paper to minimize throughput time, makespan, and total lateness, or lateness penalty. Constraints are also developed to account for limited resources, precedence relations, job splitting, due dates, substitution of resources, and concurrent and nonconcurrent job performance requirements.

Direct application of this model to the five levels in the AMRF would require at least six subscripts, one for each level in the hierarchy, and one for the time period. (To be most accurate, more subscripts would be added for the bottommost level, to account for commands given to actuate shop floor equipment.) It should be easy to see that attempting to solve the scheduling problem at once in this way will result in a zero-one integer programming problem of tremendous proportions; certainly too large to be solved in real time in the average shop.

The approach we are pursuing is to look at x_{it} at each level, i.e, to decompose the scheduling problem along the same lines as the control problem is decomposed, and to schedule activities at each level, independent of the other level activities. The resulting solution yields a set of activities that are arranged for a subsequent level processor to decompose further and to schedule. In this way, the large problem is decomposed into smaller problems that can be solved as needed by each lower-level processor in real time.

The full mathematical implications of this approach are not yet understood. For example, a serious question arises regarding optimality of the resulting set of schedules. Moreover, it is entirely possible that, as this process evolves to ever increasing levels of complexity, and ever larger portions of the feasibility region are constrained, no even feasible solution will emerge at the bottommost level. These are serious problems that must be resolved.

We will be attacking this problem empirically. Our next step is to complete the development of the model proposed above, and to investigate the various solution techniques available in the literature. Then, we will implement it in the AMRF, and compare the results to the full model that incorporates all variables and constraints in the AMRF.

6. CONCLUSIONS AND FUTURE WORK

Two major areas related to real-time production planning for automated manufacturing systems have been addressed in this paper. First, these planning problems have been identified and partitioned into five layers to match the control hierarchy under development for the Automated Manufacturing Research Facility at NBS. Second, a review of the recent efforts to solve some of these problems has been included.

Future research will focus on two major areas. First, work will continue on the integrated planning and control architecture proposed in [20]. That framework consists of a generic production control module which can be used at every level in the hierarchy, a process planning system and command/feedback structures to provide data to those modules, and a data management system to store, update, and transfer that data in a timely and accurate manner. Second, we will focus on the development of solution techniques for the decision problems described in the preceding sections. This research will be conducted in three concurrent phases. First, we must determine the information, both qualitative and quantitative, required to solve each problem. Next, we must find efficient structures for representing that information. Finally, we will attempt to marry techniques from Operations Research and Artificial Intelligence to solve each problem.

7. REFERENCES

Albus, J., Barbera, A., and Nagel, N., "Theory and Practice of Hierarchical Control", Proceedings of 23rd IEEE Computer Society International Conference, September, 1981.

Albus, J., et al., "A Control System for an Automated Manufacturing Research Facility", Robots 8 Conference and Exposition, Detroit, MI, June 4, 1984.

Baker, K. R., Introduction to Sequencing and Scheduling, John Wiley, New York, 1974.

Barkmeyer, E., et al., "Distributed Data Management Architecture for Computer Integrated Manufacturing", Technical Report NBSIR 86-3312, National Bureau of Standards, Gaithersburg, MD, January 1986.

Bitran, G. R., Hax, A. C., "On the Design of Hierarchical Production Planning Systems", Decision Sciences, 8, pp. 28-54, 1977.

Chiodini, V., "A Knowledge Based System for Dynamic Manufacturing Replanning", Proceedings of the Symposium on Real-time Optimization in Automated Manufacturing Facilities, National Bureau of Standards Special Publication 724, Gaithersburg, MD, January 21-22, 1986.

Chu, Wesley W., (ed), "Distributed Systems, Vol. II: Distributed Data Base Systems", Artech House Inc., Deham Massachusetts, 1986

Data Processing - Open Systems Interconnection - Basic Reference Model, North Holland Publishing Company, 1981, pp. 81-118.

Davis, W. J., "Decision Making and Control Hierarchies for Production Systems", Technical Report No. 145, PLAIC, Purdue University, West Lafayette, IN, 1984.

Eiler, R., "Cost Accounting Faces Automation", Managing Automation, Vol. 1, No.1, pp. 25-29, 1986.

Elsayed, E. A., "Algorithms for Project Scheduling with Constrained Resources", International Journal of Production Research, 19, 1982.

Fox, M. Constraint-Directed Search: a case study in job shop scheduling, Doctoral Dissertation, Carnegie-Mellon University, 1983.

Gershwin, S., "An Approach to Hierarchical Production Planning and Scheduling", Proceedings of the Symposium on Real-time Optimization in Automated Manufacturing Facilities, National Bureau of Standards Special Publication 724, Gaithersburg, MD, January 21-22, 1986.

Graves S., "A Review of Production Scheduling", Operations Research, 29, 646-675, 1981.

Jackson, R., and Jones, A., (eds.), Proceedings of the Symposium on Real-time Optimization in Automated Manufacturing Facilities, National Bureau of Standards Special Publication 724, Gaithersburg, MD, January 21-22, 1986.

Jones, A., and McLean, C., "A Cell Control System for the AMRF", Proceedings of the 1984 ASME International Computers in Engineering Conference, Las Vegas, Nevada, 12-15 August, 1984.

Jones, A., and McLean, C.," A Proposed Hierarchical Control Model for Automated Manufacturing Systems", Journal of Manufacturing Systems, Volume 5, No. 1, 15-25, 1986.

Kleinrock, L., "Queueing Systems Volume II: Computer Applications", John Wiley, New York, 1976.

Kurtulus, I., and Davis, E. W., "Multi-Project Scheduling: Categorization of Heuristic Rules Performance", Management Science, 28, 161-172, 1982.

Lawrence, S. R., and Morton, T. R., "PATRIARCH: Hierarchical Production Scheduling", Proceedings of the Symposium on Real-time Optimization in Automated Manufacturing Facilities, National Bureau of Standards Special Publication 724, Gaithersburg, MD, January 21-22, 1986.

McLean, C., Bloom, H., Hopp, T. "The Virtual Manufacturing Cell", Proceedings of Fourth IFAC/IFIP Conference on Information Control Problems in Manufacturing Technology, NBS, October, 1982.

McLean, C., Mitchell, M., Barkmeyer, E., "A Computing Architecture for Small Batch Manufacturing, "IEEE Spectrum, May, 1983.

Mohan, C., "Tutorial: Recent Advances in Distributed Data Base Design", IEEE Society Press, Los Angeles, California, 1984

Pritsker, A. B., Walters, L. J., and Wolfe, P. M., "Multi-Project Scheduling with Limited Resources: A Zero-One Programming Approach", Management Science, 16, 1969.

Raman, N., "A Survey of the Literature on Production Scheduling as it Pertains to Flexible Manufacturing Systems", Technical Report NBS-GCR 85-499, National Bureau of Standards, Gaithersburg, MD, 1986.

Sen, T., and Gupta, S. K., "A State-of-the-Art Survey of Static Scheduling Research Involving Due Dates", Omega, 12, No. 1, pp. 63-76, 1984.

Simpson, J., Hocken, R., and Albus, J., "The Automated Manufacturing Research Facility of the National Bureau of Standards", Journal of Manufacturing Systems, Vol. 1, 1982, pp. 17-31.

Steffen, M. S., and Greene, T. J., "Hierarchies of Sub-Periods in Constraint Directed Scheduling", Proceedings of the Symposium on Real- time Optimization in Automated Manufacturing Facilities, National Bureau of Standards Special Publication 724, Gaithersburg, MD, January 21-22, 1986.

Wysk, R., Wu, S., Yang, N., "A Multi-Pass Expert Control System for Flexible Manufacturing Systems", Proceedings of the Symposium on Real- time Optimization in Automated Manufacturing Facilities, National Bureau of Standards Special Publication 724, Gaithersburg, MD, January 21-22, 1986.

* * *

The NBS Automated Manufacturing Research Facility is partially supported by the Navy Manufacturing Technology Program.

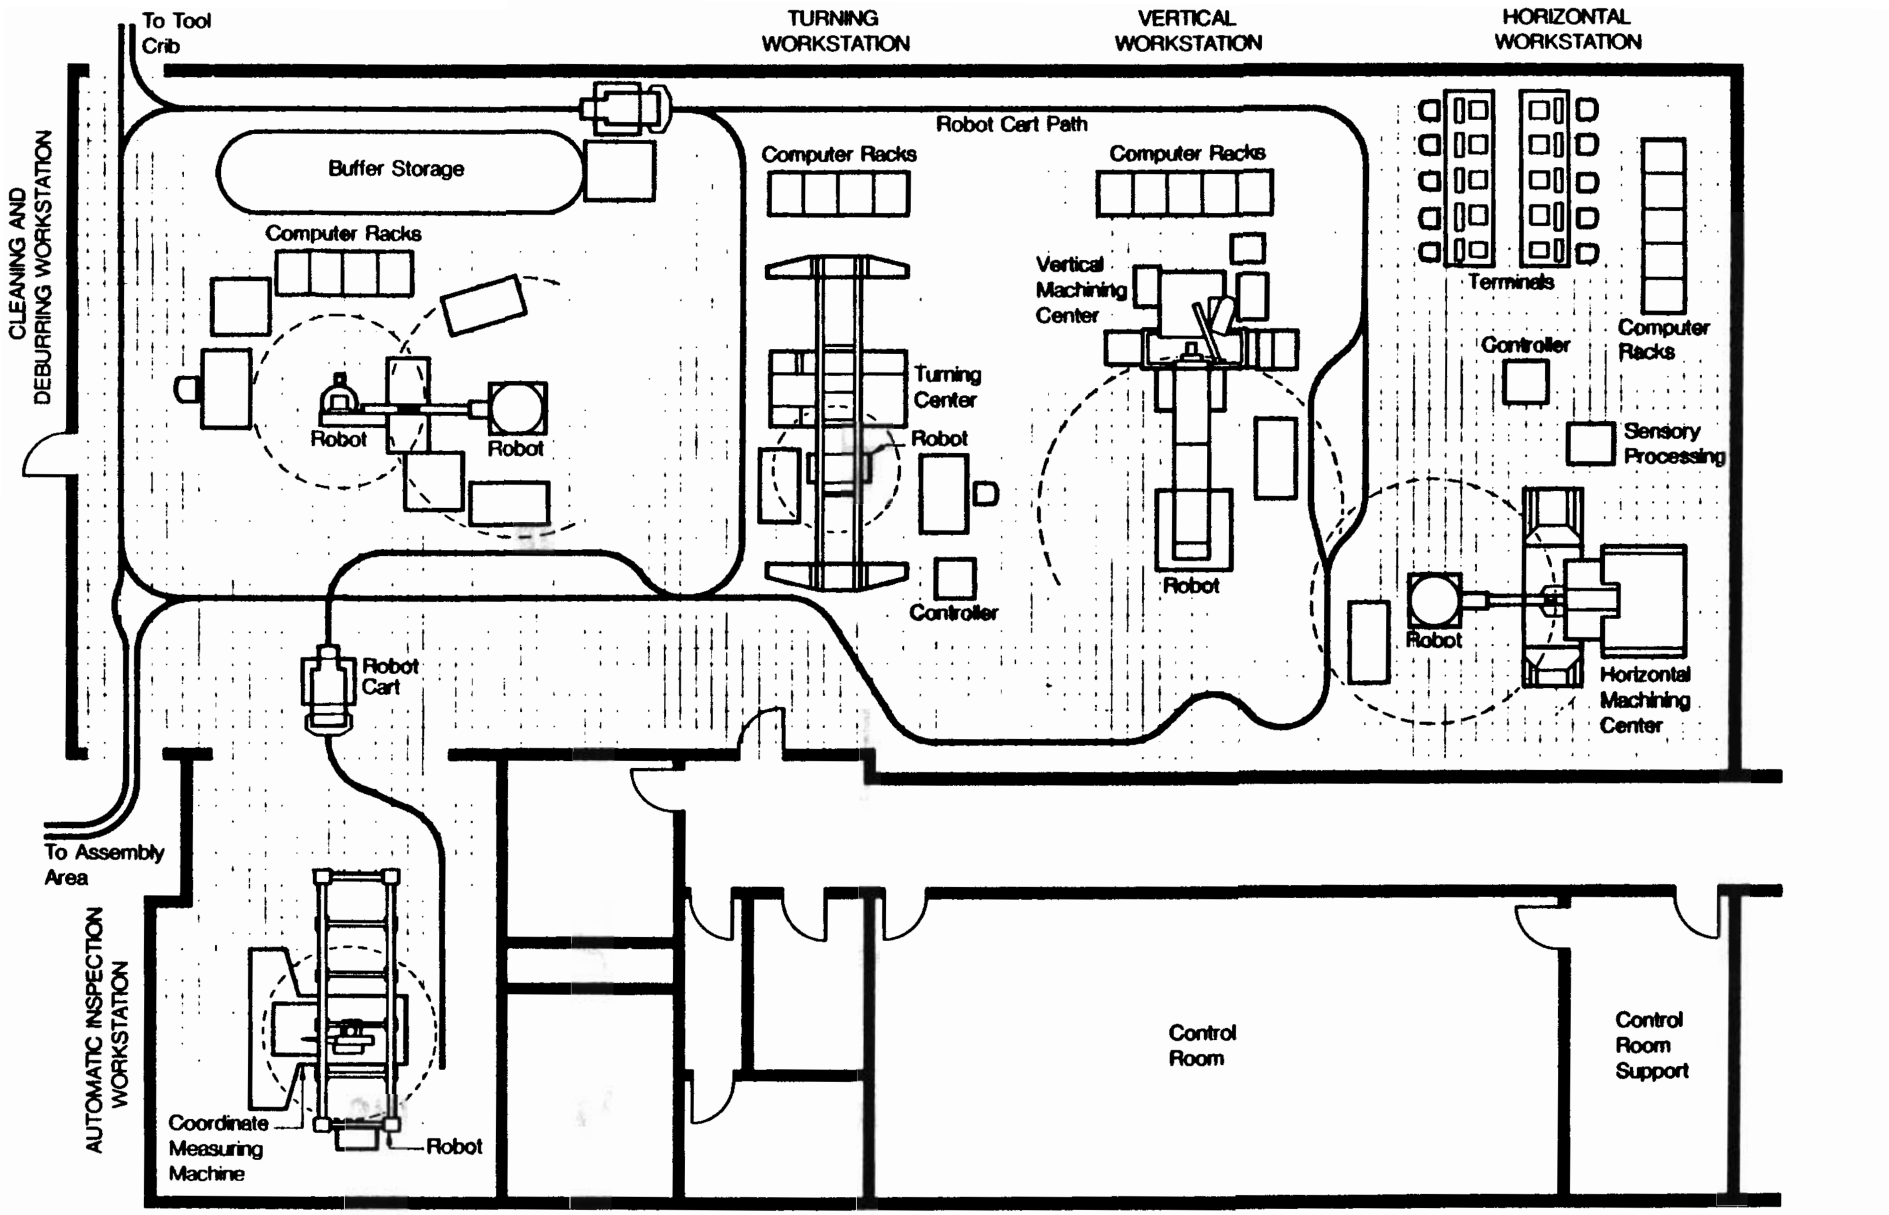

FIGURE 1. AUTOMATED MANUFACTURING RESEARCH FACILITY

CONTROL FLOW

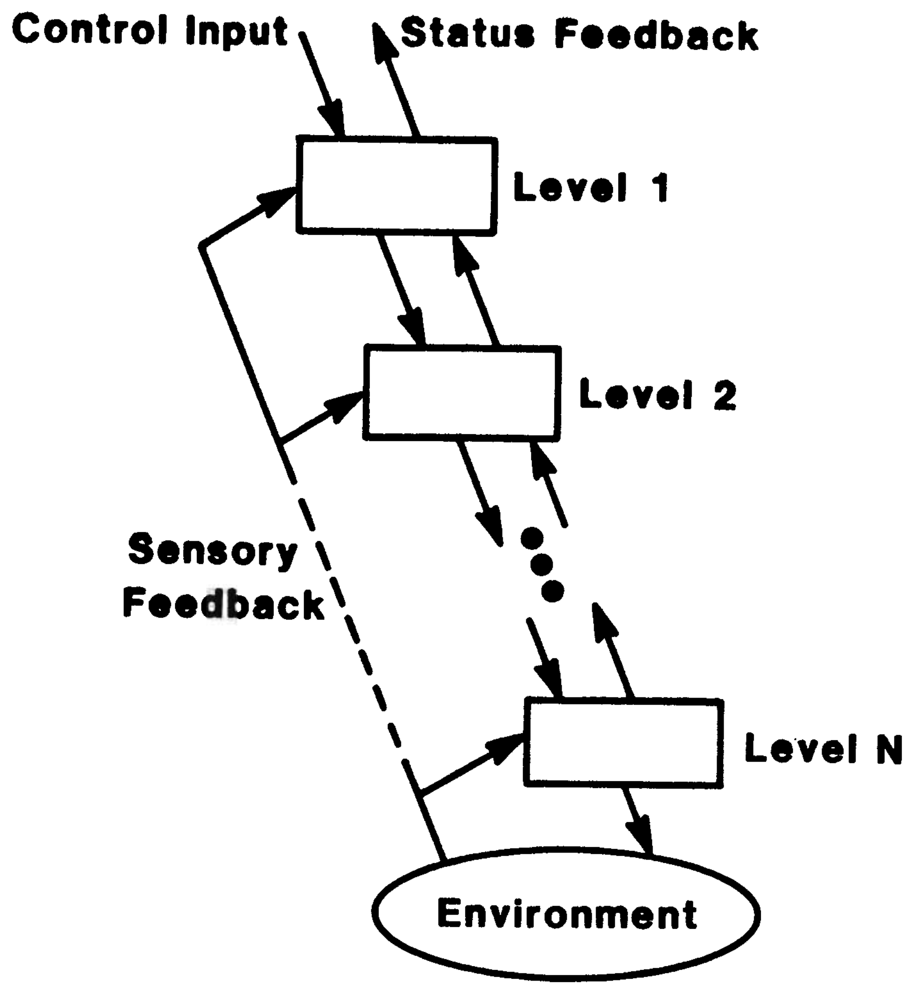

FIGURE 2. AMRF CONTROL FLOW

AMRF HIERARCHY

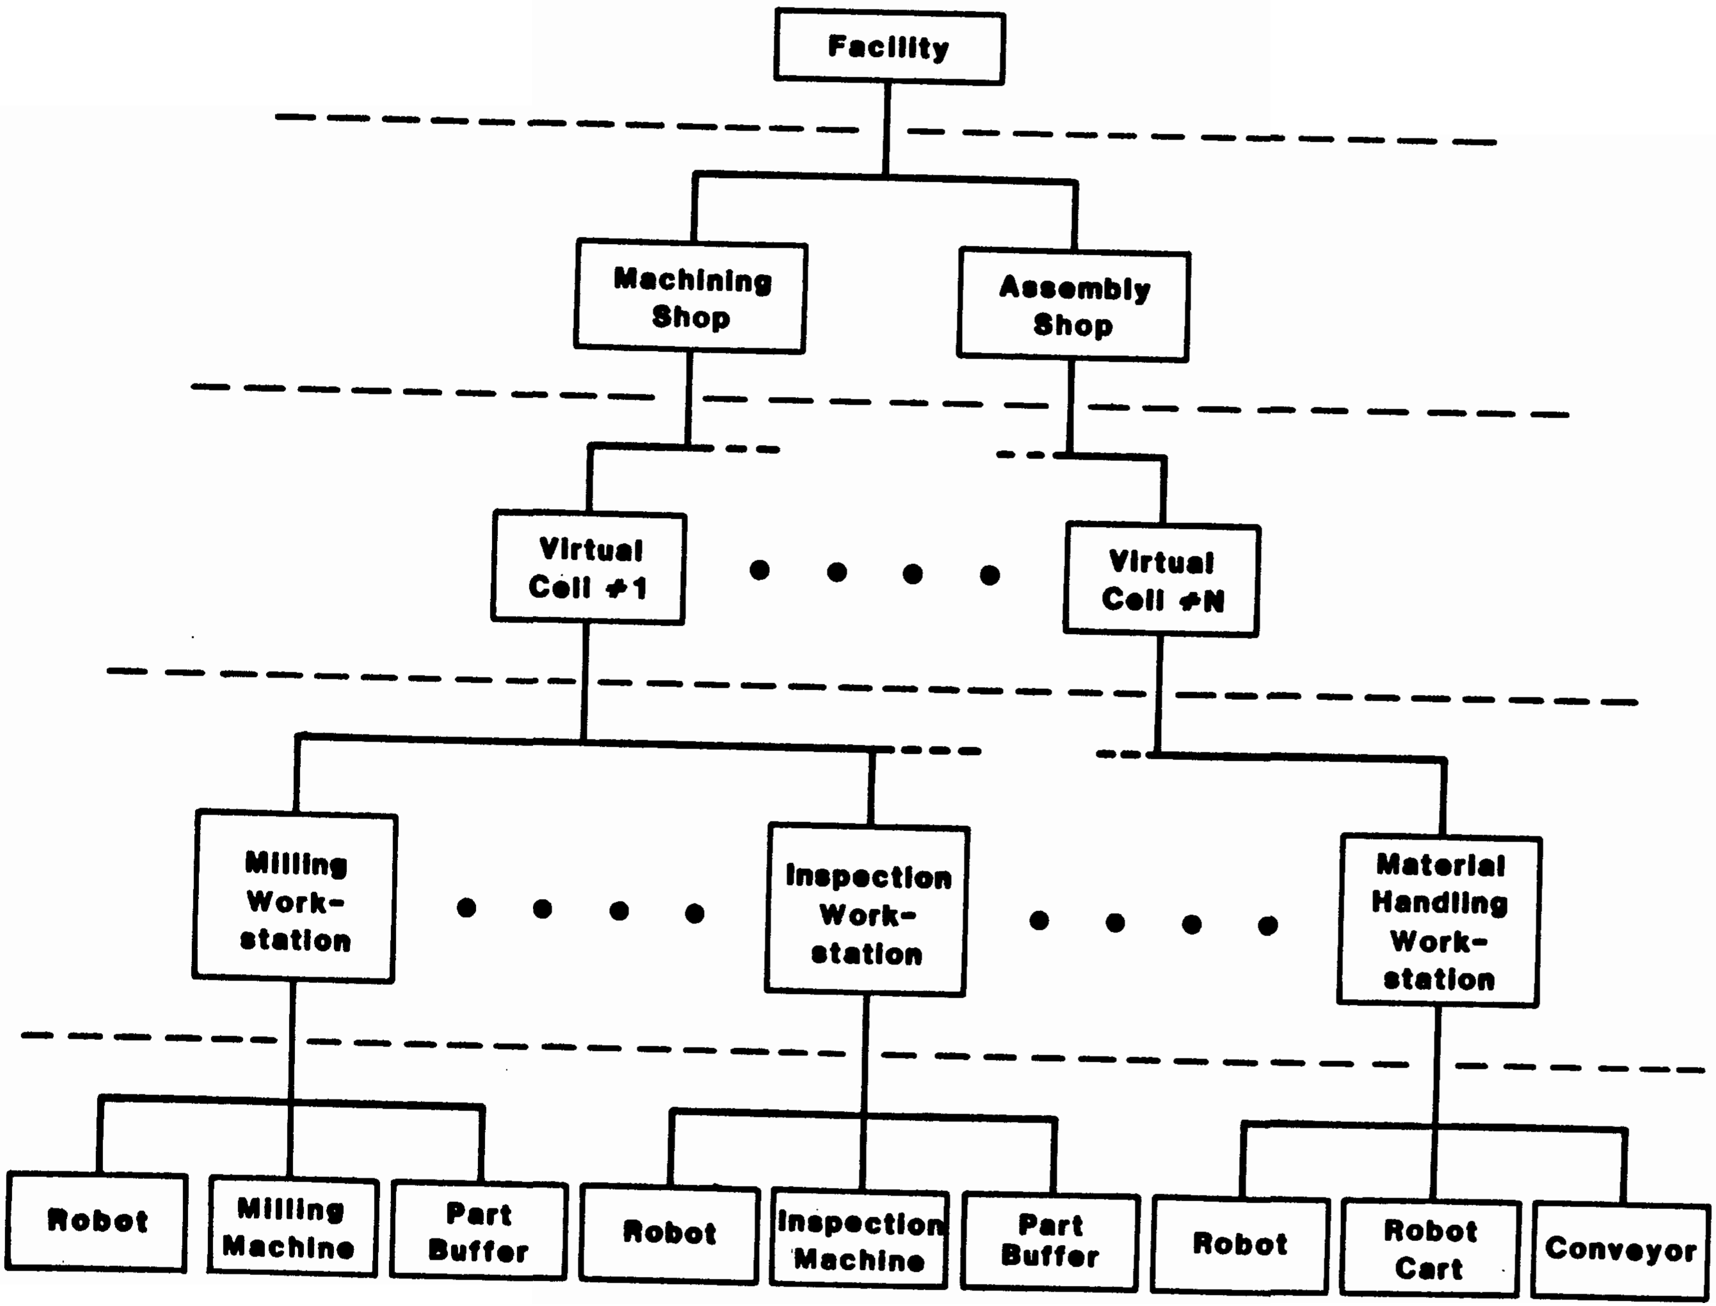

FIGURE 3. AMRF CONTROL