NISTIR 88-3865

Real-Time Optimization in the Automated Manufacturing Research Facility

Wayne J. Davis

Professor of General Engineering
University of Illinois
Urbana, Illinois 61801

Richard H. F. Jackson
Albert T. Jones

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
(Formerly National Bureau of Standards)
Center for Manufacturing Engineering
Gaithersburg, MD 20899

October 1988

Issued February 1989

NISTIR 88-3865

Real-Time Optimization in the Automated Manufacturing Research Facility

Wayne J. Davis

Professor of General Engineering
University of Illinois
Urbana, Illinois 61801

Richard H. F. Jackson
Albert T. Jones

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
(Formerly National Bureau of Standards)
Center for Manufacturing Engineering
Gaithersburg, MD 20899

October 1988

Issued February 1989

National Bureau of Standards became the National Institute of Standards and Technology on August 23, 1988, when the Omnibus Trade and Competitiveness Act was signed. NIST retains all NBS functions. Its new programs will encourage improved use of technology by U.S. industry.

REAL-TIME OPTIMIZATION IN THE AUTOMATED MANUFACTURING RESEARCH FACILITY

Wayne J. Davis
Professor of General Engineering
University of Illinois
Urbana, IL 61801

Richard H. F. Jackson
Associate Program Manager, MTC
National Institute of Standards and Technology
Gaithersburg, MD 20899

Albert T. Jones
Deputy Project Manager, AMRF
National Institute of Standards and Technology
Gaithersburg, MD 20899

ABSTRACT

A major manufacturing research facility has been established at the National Institute of Standards and Technology. 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 planning/control architecture is under development to manage production and support activities. A three layer architecture is being developed to manage the data requirements of the modules within that hierarchy. Each of these architectures contain functions that require the solution to one or more optimization problems. This paper describes both the production planning/control and the data management architectures being developed at NBS. It emphasizes the optimization problems contained within those architectures. It also discusses the work underway at NBS to address some of those problems.

KEYWORDS: Automated Manufacturing, Data Administration, Flexible Manufacturing, Hierarchical Control, Real-time Optimization, Scheduling, Routing.

1. INTRODUCTION

Manufacturing plants typically contain various combinations of people, computers, and manufacturing equipment, working together to maximize corporate profits from the goods they produce. Many of these plants are plagued by large work-in-process inventories, low utilization of equipment, insufficient throughput, and excessive delays. All of these problems tend to decrease profits.

Hopes for alleviating these problems were raised when computer-controlled robots, machine tools, and transporters became commercially available. Many companies made large investments in purchasing, integrating, and learning to use this new equipment. The resulting Computer Integrated Manufacturing (CIM) systems were expected to increase quality and profits and lead to larger shares of the world markets.

In general, this has not happened. In fact, introducing CIM into an existing factory 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 originally anticipated. Second, the continued use of existing planning and scheduling strategies often exacerbated the problems mentioned above. Finally, existing data management and communication strategies are inadequate to handle the increased dependency on "data" in these CIM environments.

This paper describes some of the work being done at the National Institute of Standards and Technology's Automated Manufacturing Research Facility (AMRF) to address these issues. 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. The work done to address some of those problems is discussed in section 5.

2. THE AMRF

2.1 Overview

The National Bureau of Standards (NBS) has established an experimental test bed, the Automated Manufacturing Research Facility (AMRF) to address measurement and standards problems in CIM systems [SIM82]. Industry, academia, and other government agencies have played active roles 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. These individual hardware and software components have been successfully integrated into a small CIM system.

2.2 Planning/Control Philosophy

The AMRF is implementing a five level hierarchical planning/control architecture (see Figure 1). It is

- o partitioned into a temporal/spatial hierarchy in which manufacturing functions are decomposed into five levels,
- o intended to respond in real-time to feedback data obtained from machines equipped with a variety of sophisticated sensors,
- o implemented in a distributed computing environment using a variety of hardware platforms and programming languages, and
- o designed to be completely data-driven but separate from the data administration system.

Each module in the AMRF control hierarchy 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 [ALB81]. 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.

2.3 Functional Decomposition

The following sections provide a brief description of the five level AMRF control hierarchy. This hierarchy represents a temporal decomposition of manufacturing functions since the planning horizon and control cycle for each level decreases as one goes down the hierarchy. It is also a spatial decomposition since the AMRF is constructed around the notion of workstations and group technology cells [JON86].

2.3.1 Facility Level. Functions at this level can be broken into two major areas: business and engineering. Engineering functions include computer-aided design (CAD), Group Technology Classification, process planning, and quality performance tracking. Business functions include order tracking, sales, marketing, identifying production resource requirements, and initiating additional capital investments.

2.3.2 Shop Level. Functions at this level can be grouped into two categories: task management and resource management. Task management includes capacity planning, grouping orders into batches, assigning and releasing batch jobs to cells, and tracking individual orders to completion. Resource management includes 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 [JON84] is envisioned as a group technology cell which 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.

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 control computers. 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.

2.4 Data Management Philosophy

Although the flow of control in the AMRF 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. The management of that data is a key ingredient in the AMRF. The data management system attempts to provide shared data to all manufacturing processes in a timely, accurate, and completely transparent manner. That is, the requestor should not have to know where or how the data he needs is stored. Achieving this goal 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 capabilities, data formats, data types, and data structures.

NBS researchers have developed an architecture [LIB88] called IMDAS--Integrated Manufacturing Data Administration System--to address these issues. IMDAS is completely separate from the control hierarchy, and has been specifically designed to meet the manufacturing control requirements described above. It contains three levels of data management services: the Basic (BDAS), the Distributed (DDAS), and the Master (MDAS) Data Administration Service modules. The major functions in these modules are described below.

2.4.1 Basic Data Administration Service—BDAS. A BDAS exists on every computer system within the AMRF. It provides the services needed to interface local data repositories and the rest of the IMDAS. Those services include interprocess communication, data and command translation, and data manipulation. Interprocess communication is achieved by using a global shared memory scheme. In this scheme, data stored in a local shared memory is replicated into the shared memory areas on remote components which require a copy of that data. Each BDAS is also capable of translating from its own representation to an IMDAS-defined common representation, and vice versa. This translation includes type, syntax, structure, and format. Since IMDAS has a global data manipulation language for making database queries, 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 architecture is the Distributed Data Administration Service (DDAS). It integrates all assigned BDASs into a segment of the global database. After receiving a query from a user, the DDAS parses it 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 Mapper" which decomposes and restructures each query into one or more queries to be executed by subordinate BDASs. 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.

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 have partitioned these problems to match the control hierarchy described above. As one moves down this decision-making/control 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.

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 [EIL86]. 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.

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 transporter 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. As a result, 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. It is becoming increasingly more important to integrate many of these decisions with those discussed in the previous section. They 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 level of production planning/control hierarchy 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. Each module within the IMDAS architecture 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. Although techniques are available for completing this query decomposition within a centralized data administration system [MOH84,CHU86], little is known about approaches to optimizing the decomposition 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 (and must be integrated with) those described in the preceding sections on production scheduling. 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). In some CIM systems this function will be part of the Data Administration system. In others, it will be designed and managed separately. It is the DCS's responsibility to deliver all information to those processes that require it, at the time they require it. The collection of standard protocols for accomplishing this data transfer are being specified in the Open Systems Interconnection standards [DAP81].

Several optimization problems must be solved to design and manage these communication systems. The design issues involve the physical media and bandwidth to be used; the topology of the underlying network; and packetizing, queuing, and protocol strategies. The primary real-time management problems involve routing and configuration strategies. It is likely that many of the techniques described in [KLE76] can be used to resolve many of these issues.

5. CURRENT WORK

In this section, we discuss some of the solution techniques under development at the AMRF.

5.1 Job Scheduling

The earliest AMRF work on job scheduling is described in [JAC86]. The goal at that time was to ensure a basic scheduling/routing capability for the AMRF that was flexible and modular enough to allow incorporation of additional results from new research already underway. In [JAC86], the authors document the scheduler subsystem used in the June 1985 realization of the AMRF. This subsystem is responsible for managing the queues at each workstation, including material handling, and monitoring the completion of the tasks assigned there. The algorithms used were simple - SPT, FIFO, and LIFO. The system was capable of switching among those algorithms in real-time.

Raman et al. [RAM86] looked at the dynamic scheduling of a single workstation. The dynamic scheduling problem was treated as a series of static problems which are solved on a rolling-horizon basis. Characteristics of the optimal solutions to the static mean flow time and mean tardiness problems are developed and an implicit enumeration approach to the mean tardiness problem was also developed. These results are extended to drive dispatching procedures for the dynamic case.

The most recent work on job scheduling is described in [DAV88]. This paper proposed a decomposition of the production scheduling problem into two levels (see Figure 3). The authors made two important and realistic assumptions in developing this decomposition. First, decision makers at each level will behave in a cooperative fashion in solving their own problems. Second, the decision maker at the Process Coordinator level possess more detailed information about the variables and constraints associated with his decisions than the supremal. These assumptions result in a downward flow of authority and an upward flow of aggregated information about the state of the process and duration of activities.

The top level in this decomposition chooses the "best" among several candidate scheduling rules and determines the start and finish times of each JOB_j at each process P_n, E_{jn} and L_{jn} respectively (see Figure 4). The bottom level uses these bounds to determine the start and finish times for each of the tasks that make up JOB_j. This implies the ability to 1) quickly analyze alternatives at a given level, 2) perform contingency planning at each level, and 3) resolve conflicts between decisions at different levels. The foundations of this algorithm are two forms of simulation. The first consists of R on-line, concurrent evaluations of candidate scheduling rules. These are invoked whenever a new schedule needs to be generated. This occurs whenever a new job is ready for production or shop floor problems invalidate the existing schedule. These on-line simulations require each trial to be initialized from a known state, which is tied to the actual state of the manufacturing system, and that the initial state remains constant throughout the scheduling analysis. The second consists of continuously-running evaluations of potential scheduling rules to provide a contingency planning capability. Real-time simulation allows the initial state to change from one trial to another, based on the actual evolution of the manufacturing system. This approach creates several problems in output analysis not covered in the existing simulation literature [LAW82].

Although we are still in the model verification and time testing phase of the simulation analysis of the AMRF, our initial results are promising. To complete this phase, we are using a single scheduling rule and three different performance measures: time in system, productivity, and process utilization. Using a SIMAN[1] simulation package on an INTEL 80286 based personal computer, we can simulate the future response of the system at approximately 1000 times the speed at which a system emulation takes place. These numbers are based on results obtained from making 100 runs with each run scheduling 50 JOBs. We are in the process of analyzing the performance measures obtained from the various trials to verify the current AMRF SIMAN model.

5.2 Cart Scheduling

Material transfers are handled the same way as other jobs in the AMRF, in two phases. First, the location and pickup/delivery times are specified in the top level scheduler. Next, the process coordinator responsible for transporter scheduling uses these "due dates" to sequence material transfers and to assign a transporter to each of those transfers. The transporters in the AMRF are two bi-directional, Automatic Guided Vehicles (AGVs). Each AGV has two roller beds which means that it can transport two trays of materials simultaneously to any of the AMRF load/unload stations. The actual cart path is a wire which is taped to the floor. As shown in Figure 5, the path contains no loops. While this simplifies the routing, there is only one, it enhances the possibility of collisions and deadlocks.

To address this problem, the cart path has been divided into zones (see Figure 5) which interconnect the nodes denoting specific load/unload stations. From a scheduling perspective, these zones represent resources which must be allocated by the cart scheduler to a given AGV to permit a transfer from one node to the next. If a zone is already occupied by one AGV, then the other AGV must wait to traverse that zone. A Petri net [AGE79] was defined to formalize the rules for allocating and deallocating zones (see Figure 6). The tokens are used to denote ownership of zones and to detect potential deadlocks. It works as follows.

---

[1] Certain commercial equipment, instruments, or materials are identified in this paper. Such identification does not imply recommendation or indorsement by the National Bureau of Standards, nor does it imply that the material or equipment identified are necessarily the best available for the purpose.

Suppose an AGV wants to make a transfer from 8 to 9. First, there must be an open space in the buffer at node 9. If there is, it is reserved. Next, one token must be located at node 8, representing the requesting AGV, and another token must be at node 1, representing the availability of path 1. When both conditions are satisfied, the transition from node 8 to 9 can be made. Upon arrival at node 9, both tokens will reside at node 9, and the only subsequent permissible transition is to return to node 8 where ownership of path 1 will be relinquished. The reader will note that the transition from node 9 to node 10 is not permitted as it could lead to deadlock. That is, if an AGV is to travel from node 8 to 10, it must make this decision from the outset. As indicated, the transition from node 8 to 10 requires the AGV to first secure ownership of both paths 1 and 2. Upon reaching node 10, ownership of path 1 will be relinquished.

This logic has been incorporated into the AMRF simulation described above. This guarantees that the pickup and delivery times generated in the simulation are feasible and that a cart can be scheduled to carry out the required transfers.

5.3 Robot Path Planning

As noted above, obstacle avoidance and path planning are two of the most important optimization problems to be resolved in the control of robots. A path planning algorithm, which is guaranteed to be collision-free, has been the only moving object within the work volume and that this work space is modeled using an "oct-tree" representation [JAC80]. The output from this algorithm is a piece-wise linear, collision-free, 3-D path from the initial to the goal state.

There are three major search techniques used to find this path. A hill climbing technique, with the Euclidean distance as its objective function, is used to reach a local minimum. The A* best-first search technique [HER86], with h heuristic equal to the Euclidean distance from the current point to the goal, is used to move away from this local minimum. These two techniques are combined, possible many times, until the goal is reached. the resultant path is then checked, using a multi-resolution search, for collisions. This process is repeated until a satisfactory path has been found. This path is then passed to a trajectory planner where the velocity, acceleration, etc. required to move the robot are calculated.

5.4 Artificial Intelligence and Process Planning

As one would expect, automated manufacturing facilities are fertile areas for the application of artificial intelligence techniques. In [NAU86] an effort is described to apply these techniques to process planning and tool selection. Nau developed and implemented in the AMRF a Semi-Intelligent Process Selector (SIPS) which produces process plans for small set of machinable surfaces. Like most other reasoning systems, SIPS stores the data required by the solution procedure in frames which are manipulated by an inference engine, usually the rules in a rule-based reasoning system. However, unlike most other such systems, SIPS uses a technique called hierarchical knowledge clustering to manipulate the data and produce a process plan. This technique imposes a tree structure on the data frames which can be exploited to connect a sequence of frames to form a process plan. When cost information for each frame, or step in the process plan, is included, optimization techniques can be used in the production planning process to produce least cost plans. For example, SIPS uses a modified least-cost-first Branch and Bound procedure to find a least cost sequence of processes for making each machinable surface.

There are several advantages of this hierarchical approach to process planning. As with any hierarchical representation of data, it is easy to use, understand and exploit. Another advantage is the way in which the problem domain is automatically partitioned into regions associated with classes of machinable processes. This last, of course, provides much improvement in the speed of search procedures since the search need only be performed within one of the partitions.

An important aspect of this work is the way in which techniques of Operations Research and Artificial Intelligence are combined to produce a result that is more than just feasible, but is optimal with respect to some objective. Too often rule-based systems are proposed as the solution to a complicated problem before other analytic approaches have been considered. Too often system designers are content to settle for feasible solutions to a problem with little or no effort expended to find optimal or even improved solutions. This work is one attempt to improve the state of this art.

6. CONCLUSIONS AND FUTURE WORK

This paper has discussed the real-time optimization problems likely to be encountered in the Automated Manufacturing Research Facility. Two major areas were discussed: production planning and control, and data administration. We believe that the problems addressed herein are typical of any automated factory. We also provided a review of the recent efforts to solve some of these problems.

Future research will focus on two major areas. First, we will continue with the solution approaches already underway, and begin to focus on the development of solution techniques for the other decision problems described in the preceding sections. This research will be conducted in three concurrent phases. First, we will determine the information, both qualitative and quantitative, required to solve each problem. Next, we will find efficient structures for representing that information. Finally, we will attempt to marry techniques from Operations Research and Artificial Intelligence to solve each problem. Second, we will attempt to incorporate those techniques into an integrated decision-making and control architecture which manages both the fabrication of parts and the data needed to carry out that fabrication.

7. REFERENCES

[AGE79] Agerwala, T., "Putting Petri Nets to Work", **Computer**, 85- 93, December, 1979.

[ALB81] Albus, J., Barbera, A., and Nagel, N., "Theory and Practice of Hierarchical Control", **Proceedings of 23rd IEEE Computer Society International Conference**, September, 1981.

[CHU86] Chu, Wesley W., (ed), **Distributed Systems, Vol. II: Distributed Data Base Systems**, Artech House Inc., Deham Massachusetts, 1986

[DAP81] Data Processing - Open Systems Interconnection - Basic Reference Model, ISO Standard 7948, International Standards Organization, Geneva, 1981.

[DAV88] Davis, W. J. and Jones, A. T., "A Real-time Production Scheduler for a Stochastic Manufacturing Environment", **International Journal of Computer Integrated Manufacturing**, Vol. 1, No. 2, 101-112, 1988.

[EIL86] Eiler, R., "Cost Accounting Faces Automation", **Managing Automation**, Vol. 1, No.1, pp. 25-29, 1986.

[HER86] Herman, M., "Fast, Three-Dimensional, Collision-Free Motion Planning", **Proceedings of the IEEE International Conference on Robotics and Automation**, San Francisco, CA, April 7-10, 1986.

[JAC80] Jackins, C. and Tanimoto, S., "Oct-Trees and Their Use in Representing Three Dimensional Objects", **Computer Graphics and Image Processing**, 14, 249-270, 1980.

[JAC85] Jackson, R., and Knapp-Cordes, J., "Guide to the June 1985 AMRF Scheduling Routines", National Bureau of Standards Internal Report, NBSIR 86-3338, Gaithersburg, Md., April, 1986.

[JON84] 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, August 12-15, 1984.

[JON86] 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.

[KLE76] Kleinrock, L., **Queuing Systems Volume II: Computer Applications**, John Wiley, New York, 1976.

[LAW82] Law, A. M. and Kelton, W. D., **Simulation, Modeling and Analysis**, McGraw-Hill, New York, 1982.

[LIB88] Libes, D. and Barkmeyer, E., "The Integrated Manufacturing Data Administration System (IMDAS)- An Overview", **International Journal of Computer Integrated Manufacturing**, Vol. 1, No. 1, 44-49, 1988.

[MCL83] McLean, C., Mitchell, M., Barkmeyer, E., "A Computing Architecture for Small Batch Manufacturing", **IEEE Spectrum**, May, 1983.

[MOH84] Mohan, C., **Tutorial: Recent Advances in Distributed Data Base Design**, IEEE Society Press, Los Angeles, California, 1984

[NAU86] Nau, D.S., Hierarchical abstraction of problem-solving knowledge, **Proceedings of 1986 ASME Winter National Meeting**, Anaheim, CA, December, 1986.

[RAM86] Raman, N., Talbot, F., and Rachamadugu, R., "Due-Date Based Scheduling in a Flexible Manufacturing System", National Bureau of Standards Internal Report, NBSIR 86-513, Gaithersburg, Md. April, 1986.

[SIM82] 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.

FIG 1. AMRF CONTROL HIERARCHY

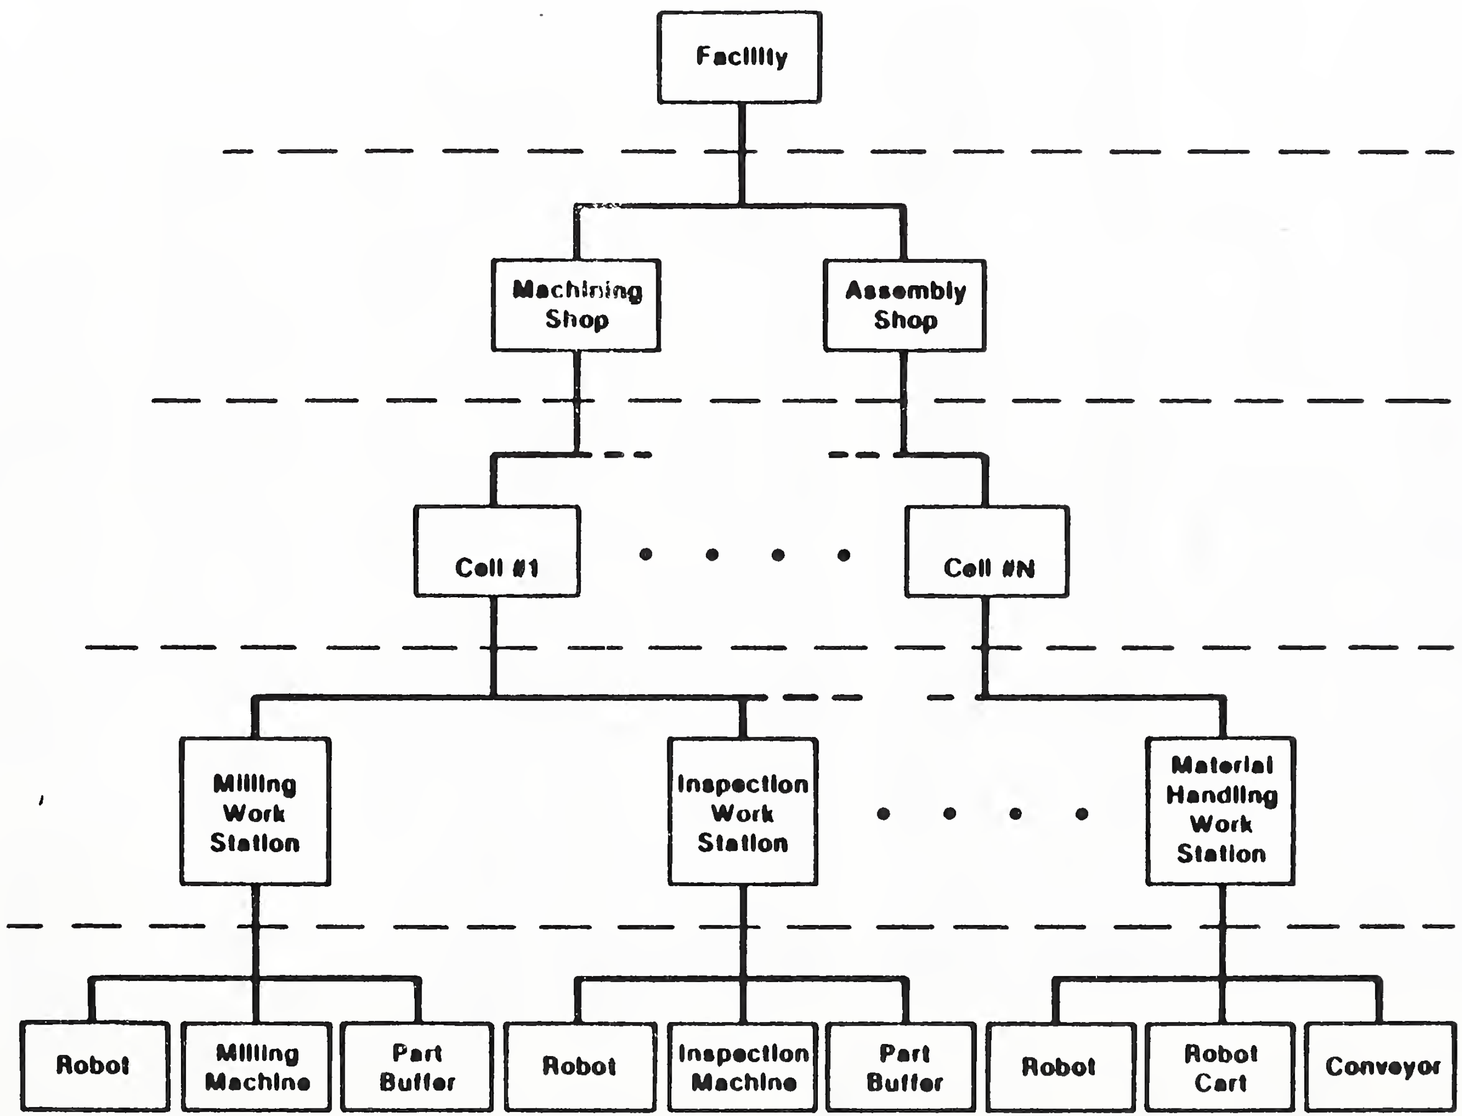

Fig. 2 Generic control module

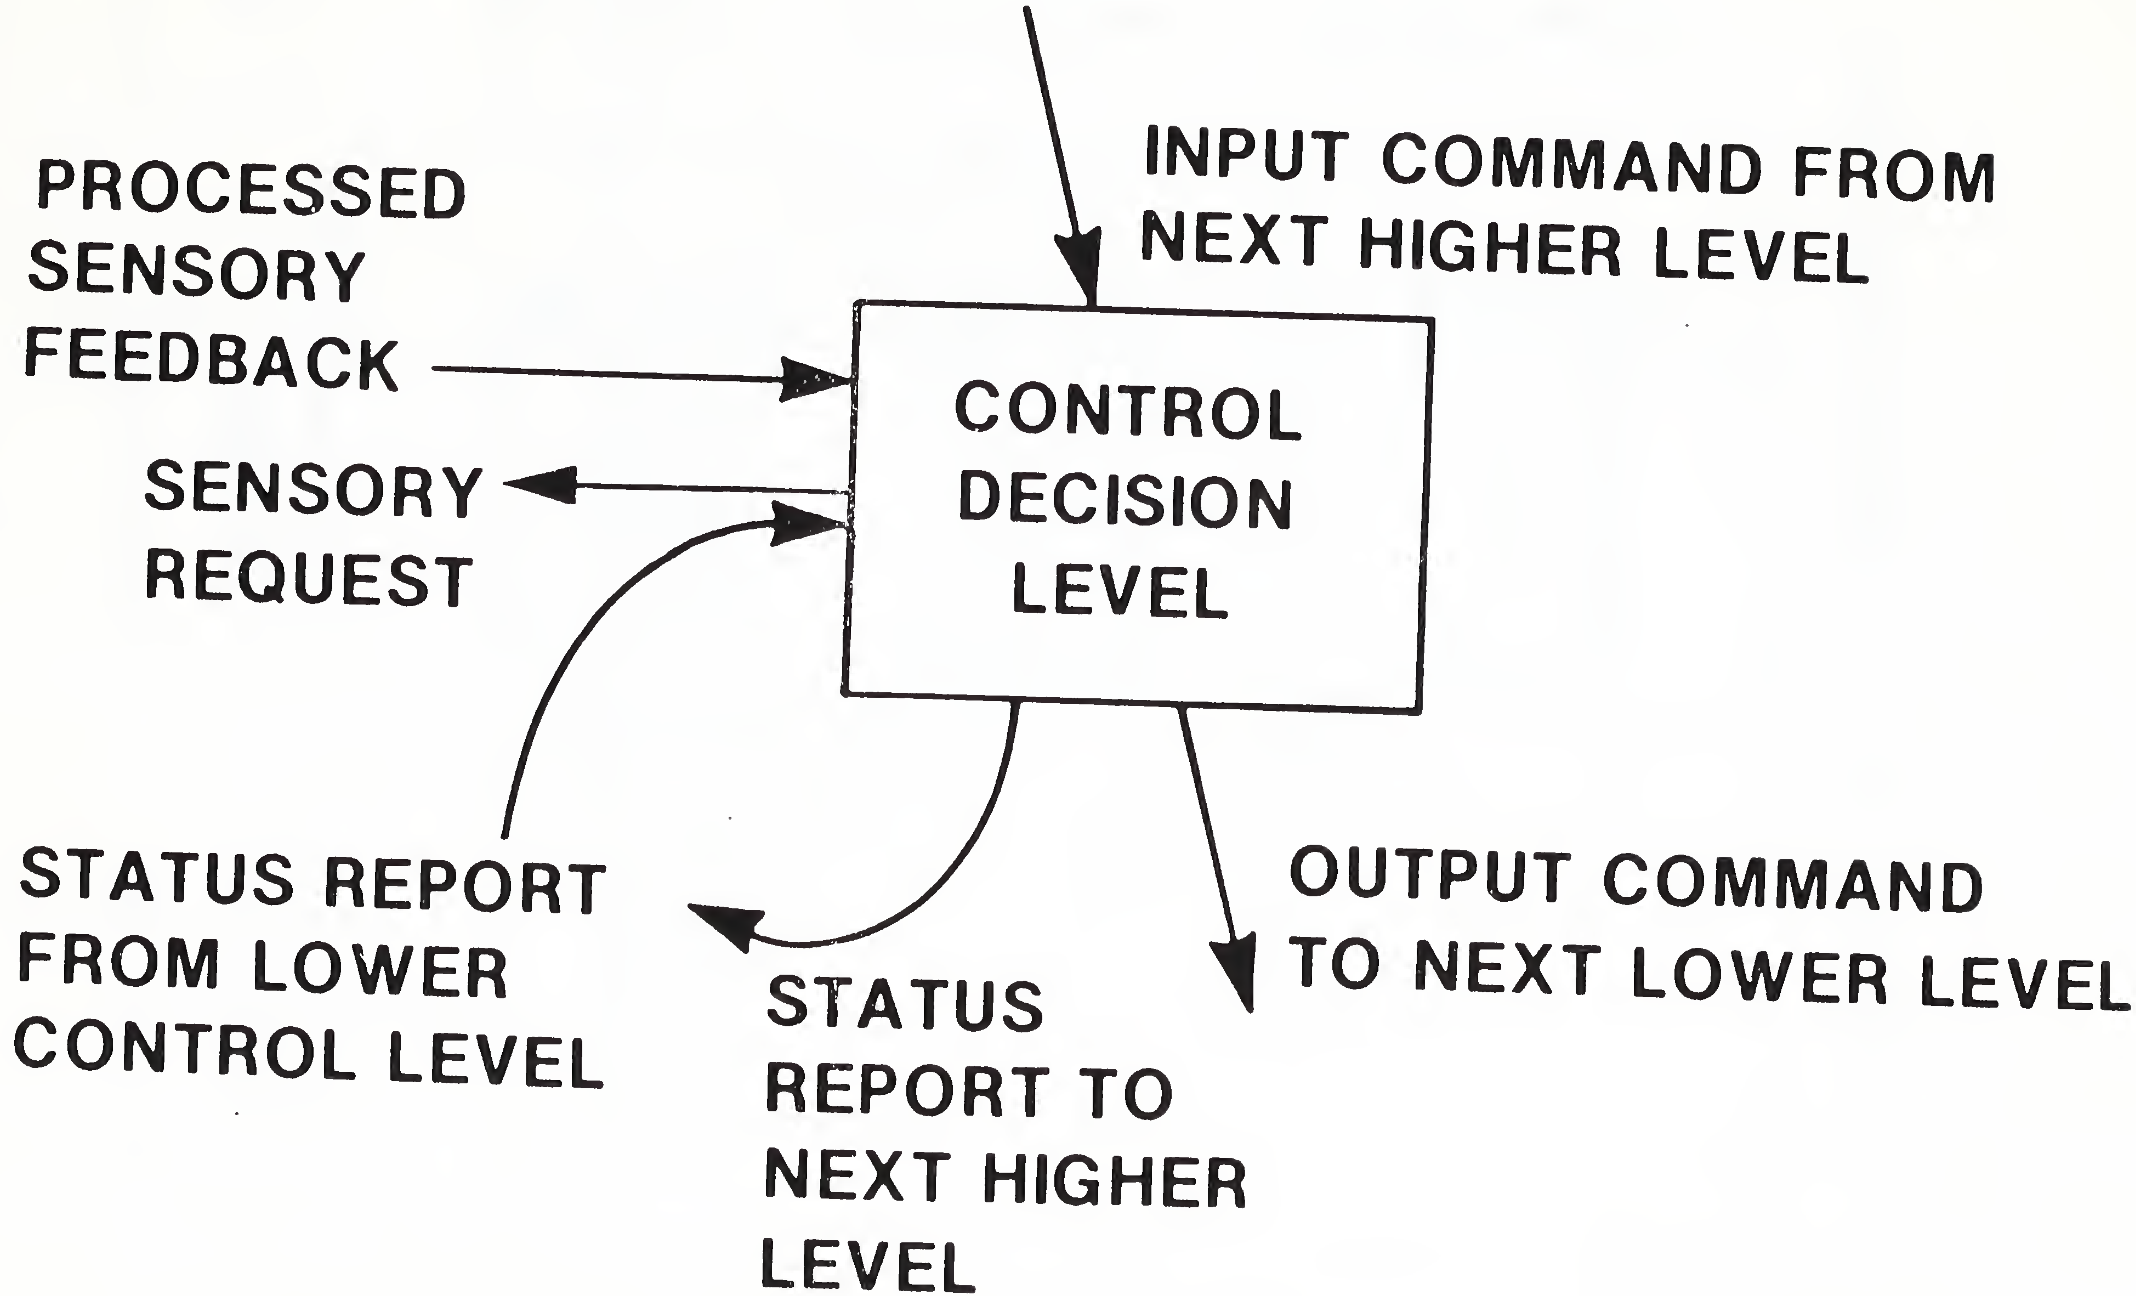

FIGURE 3 DECOMPOSITION

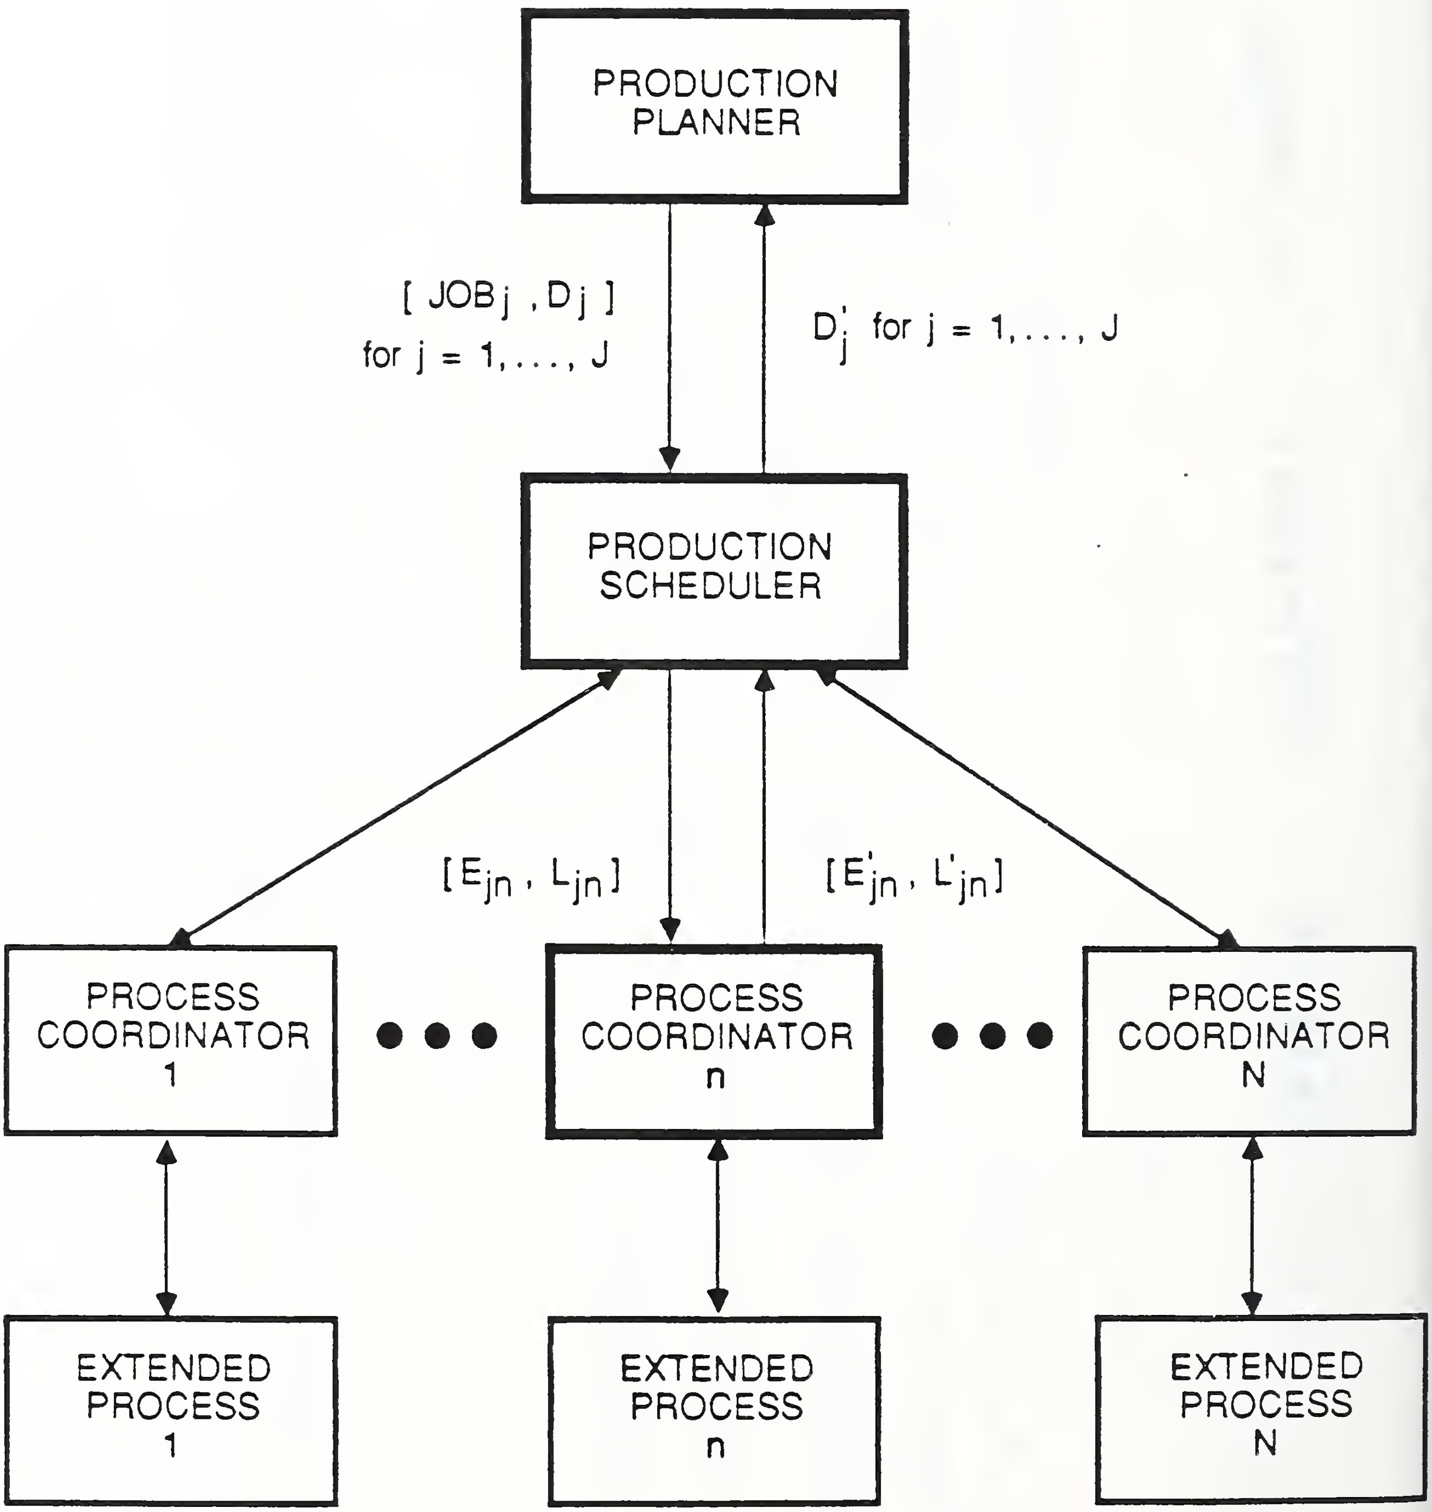

FIG. 4. SCHEDULING ALGORITHM

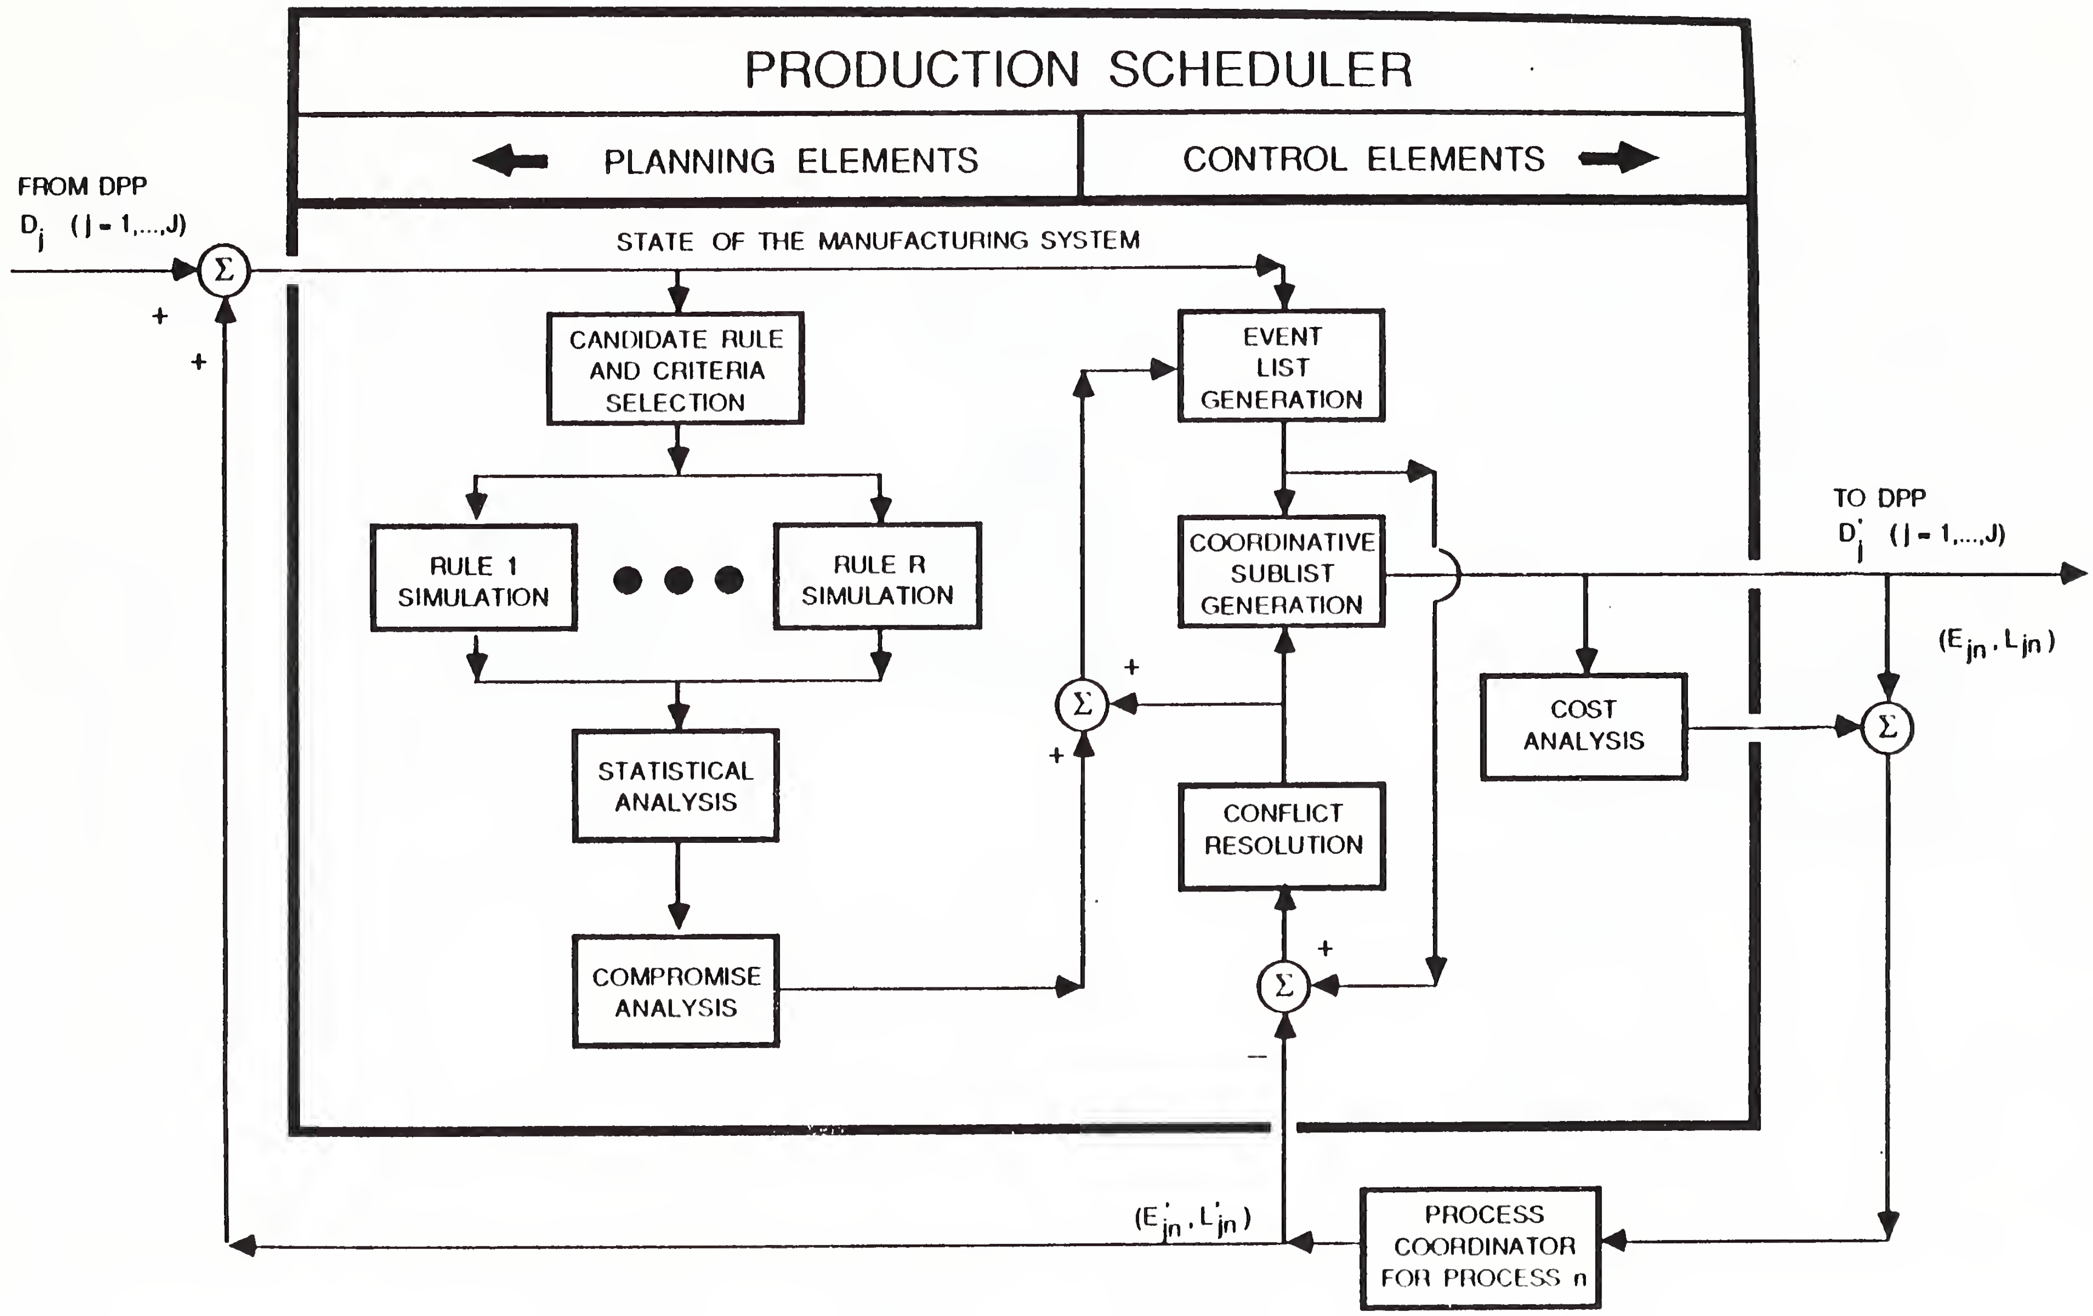

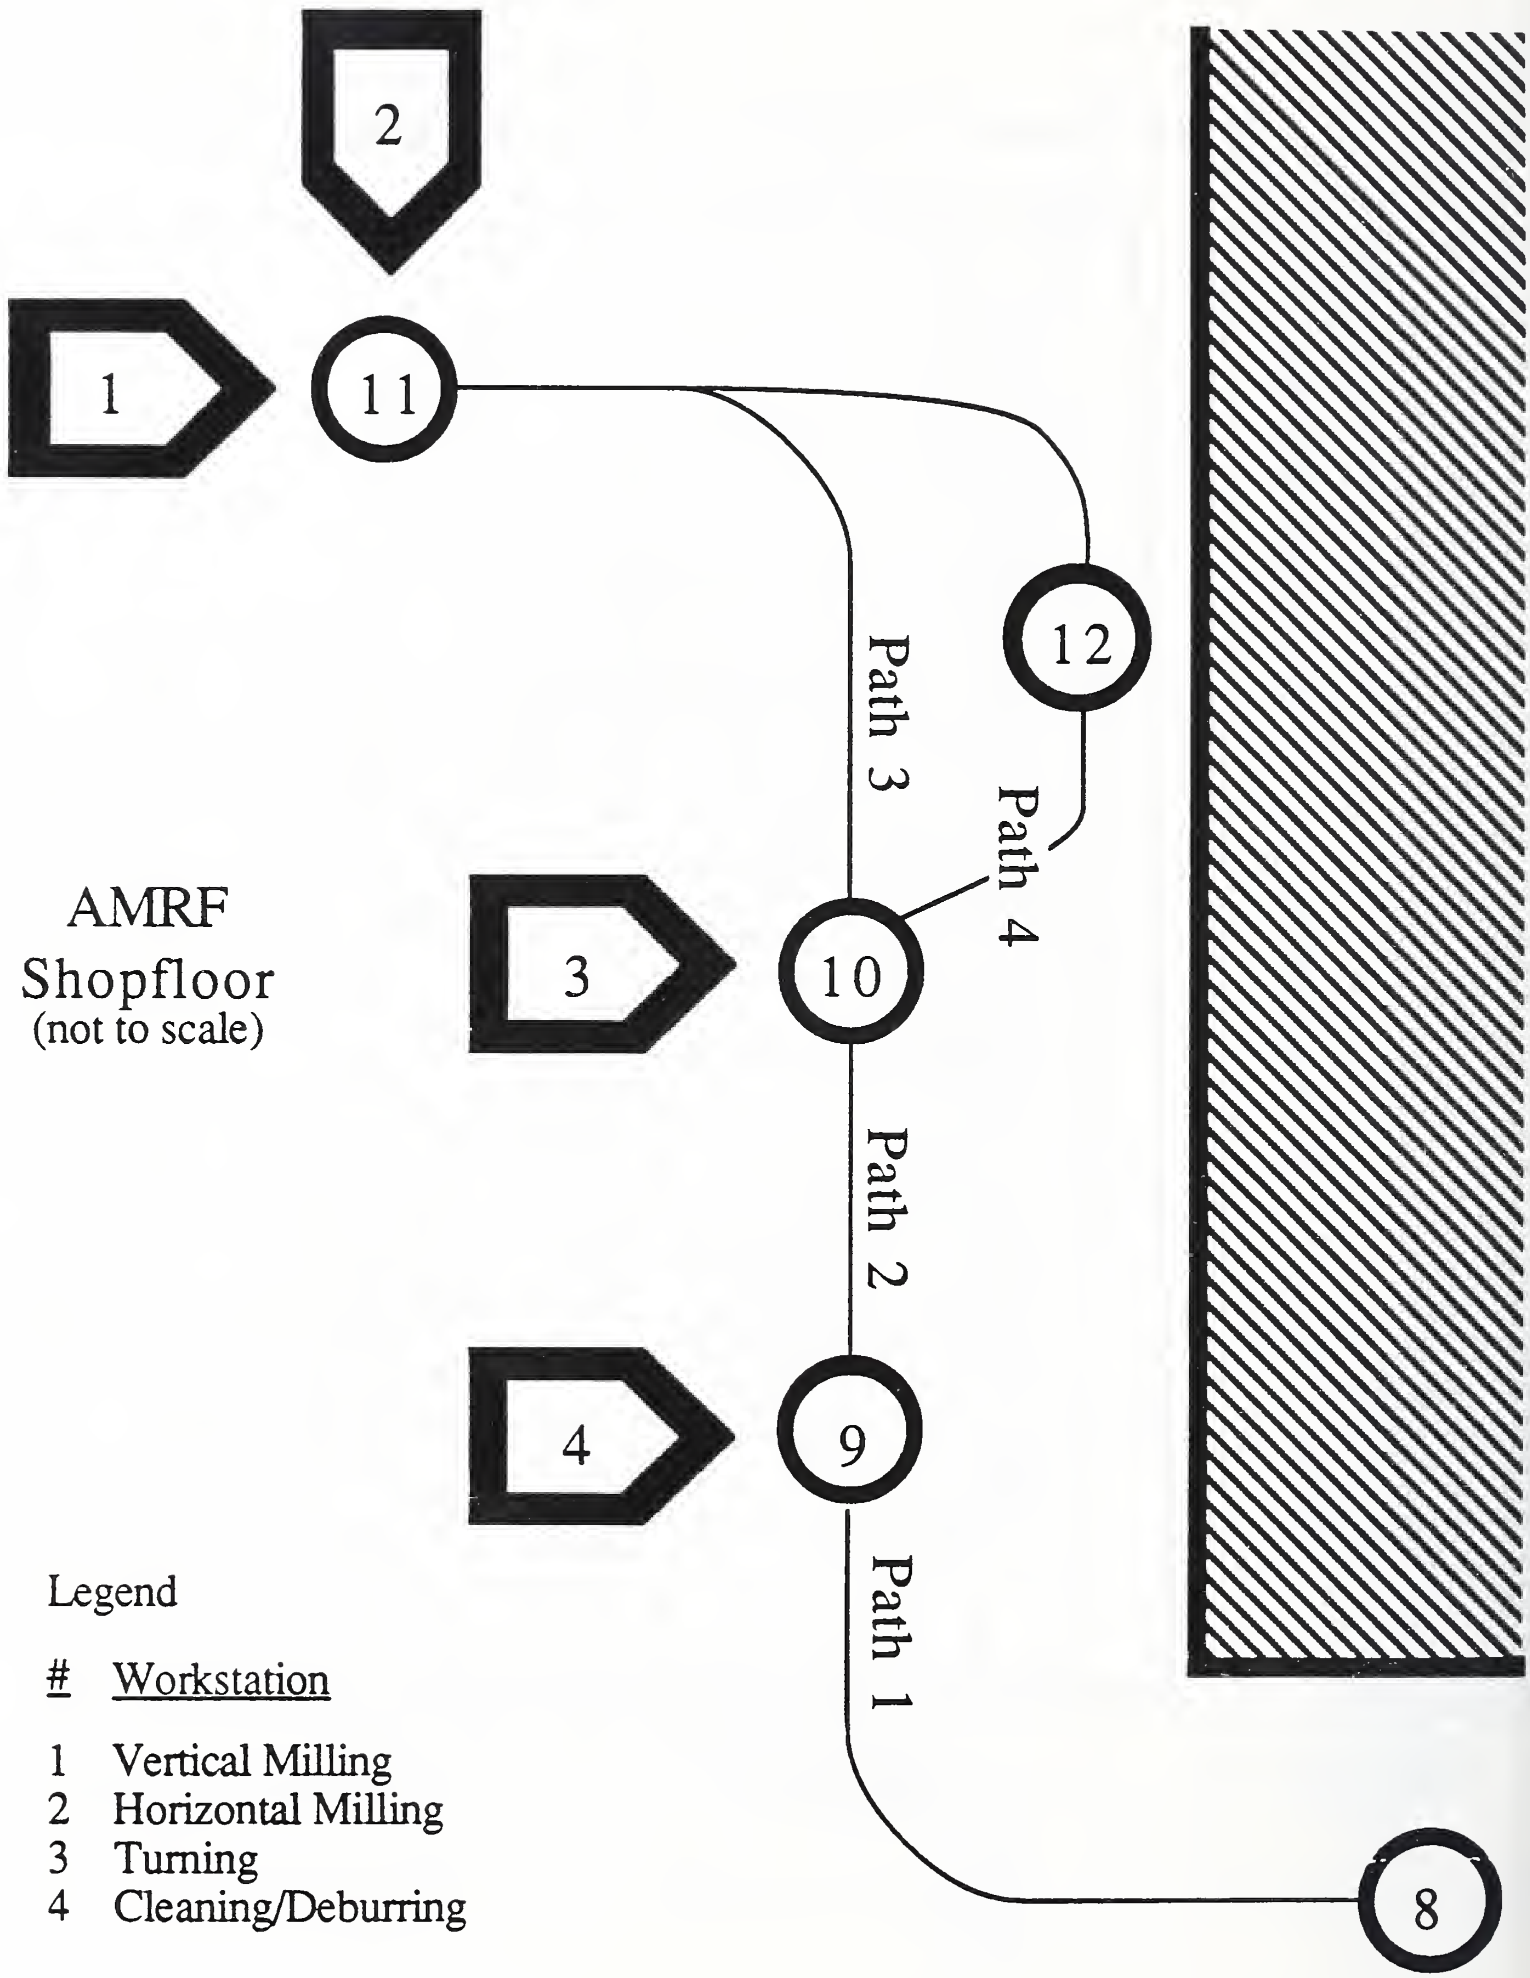

FIG 5. AMRF SHOP FLOOR

FIG 6. Petri Network for AMRF Cart Logic

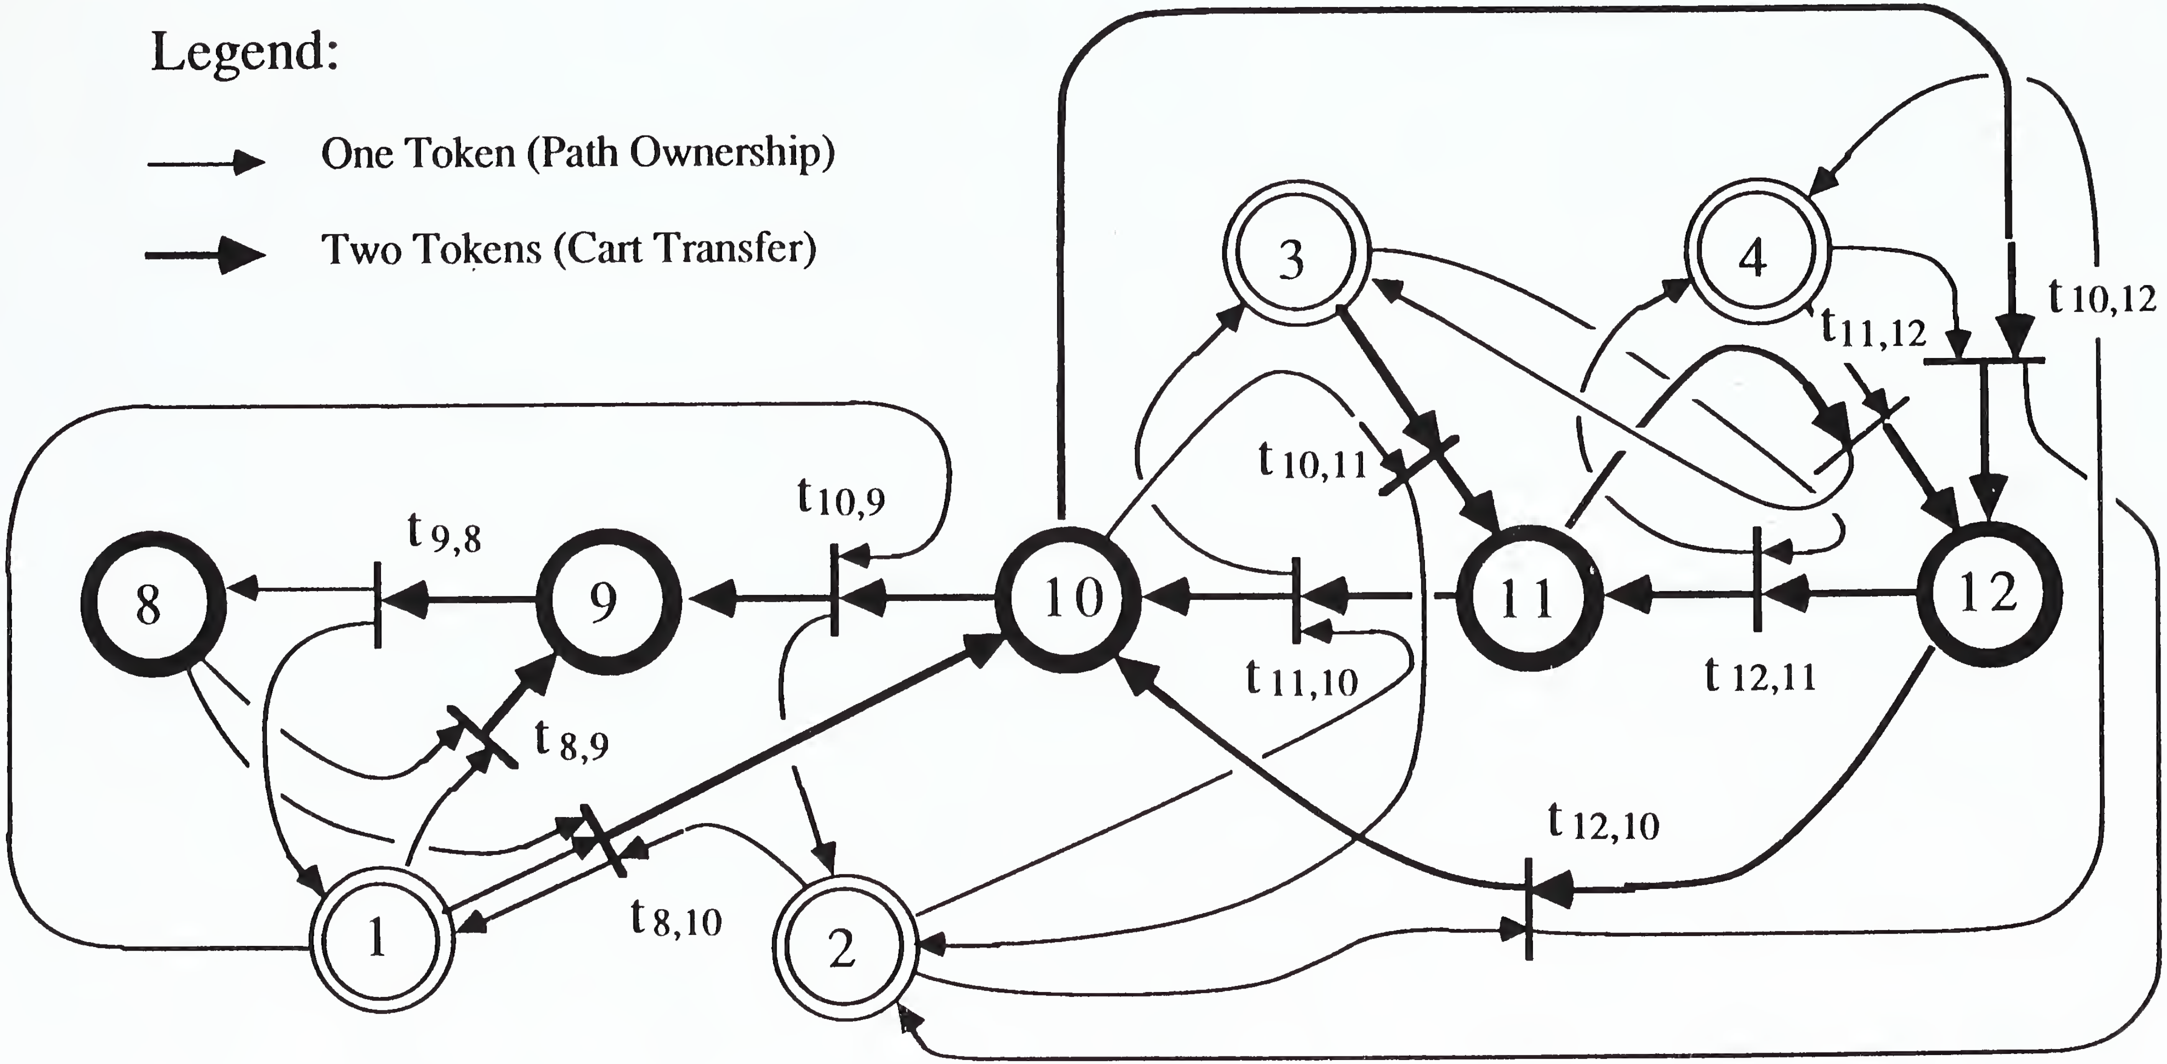

NBS-114A REV. 2-80

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET (See instructions)	1. PUBLICATION OR REPORT NO. NISTIR 88-3865	2. Performing Organ. Report No.	3. Publication Date FEBRUARY 1989

4. TITLE AND SUBTITLE

Real-Time Optimization in the Automated Manufacturing Research Facility

5. AUTHOR(S)

Wayne Davis, Richard Jackson, Albert Jones

6. PERFORMING ORGANIZATION *(If joint or other than NBS, see instructions)*

NATIONAL BUREAU OF STANDARDS
U.S. DEPARTMENT OF COMMERCE
GAITHERSBURG, MD 20899

7. Contract/Grant No.

8. Type of Report & Period Covered

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS *(Street, City, State, ZIP)*

National Institute of Standards & Technology
Bldg. 220 Room A-127
Gaithersburg, MD 20899

10. SUPPLEMENTARY NOTES

☐ Document describes a computer program; SF-185, FIPS Software Summary, is attached.

11. ABSTRACT *(A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here)*

A major manufacturing research facility has been established at the National Institute of Standards & Technology. 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 planning/control architecture is under development to manage production and support activities. A three layer architecture is being developed to manage the data requirements of the modules within that hierarchy. Each of these architectures contain functions that require the solution to one or more optimization problems. This paper describes both the production/control and the data management architectures being developed at NIST. It emphasizes the optimization problems contained within those architectures. It also discussess the work underway at NIST to address some of those problems.

12. KEY WORDS *(Six to twelve entries; alphabetical order; capitalize only proper names; and separate key words by semicolons)*

Automated Manufacturing, Data Administration, Flexible Manufacturing, Hierarchical Control, Real-Time Optimization, Scheduling, Routing.

13. AVAILABILITY

☒ Unlimited
☐ For Official Distribution. Do Not Release to NTIS
☐ Order From Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402.
☒ Order From National Technical Information Service (NTIS), Springfield, VA. 22161

14. NO. OF PRINTED PAGES

26

15. Price

$12.95