Proceedings of the 1992 Winter Simulation Conference
ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson

TASK AND RESOURCE ALLOCATION VIA AUCTIONING

Dharmaraj Veeramani

Department of Industrial Engineering
University of Wisconsin-Madison
1513 University Avenue
Madison, WI 53706

ABSTRACT

This paper examines physical resource management in large computer-controlled manufacturing systems (LCMSs) under resource-constrained auction-based shop-floor control. This control scheme employs multi-level negotiation among autonomous system entities, such as workpieces, machine tools, and transport vehicles, as the mechanism for task and resource allocation. It was observed via simulation that a manufacturing system under auction-based control automatically adapts itself to different system loads. Other simulation results on the impact of various cutting-tool management policies on system performance are also reported.

1 INTRODUCTION

Computer-integrated manufacturing systems have traditionally been controlled in a hierarchical manner wherein the system is decomposed into several levels of authority in a pyramid-like structure. Hierarchically controlled systems are characterized by rigid master-slave relationships between adjacent levels in the system in which control commands progress in a top-down manner while status and sensory information progress in a bottom-up manner. The primary advantage of the hierarchical approach to manufacturing system control is that it mimics the hierarchy commonly found in human organizational systems, and this facilitates delineation of responsibilities and authority among people or control nodes in the system. Hierarchical control systems, however, grow in complexity with increasing size of the manufacturing system and suffer from limited extendibility and modifiability. In addition, failures at higher level control nodes can significantly cripple the operation of the manufacturing system. Hierarchical control schemes also require maintenance of global information. These limitations of hierarchical control systems have restricted their use to small and medium sized flexible manufacturing cells and systems having between two and ten machines.

The next century will see the emergence of a new generation of manufacturing systems that are considerably larger than extant computer-controlled manufacturing systems, and possibly consisting of 50 or more computer numerically-controlled (CNC) machines. The emergence of such Large Computer-Controlled Manufacturing Systems (LCMSs) would become necessary as large job-shops, each consisting of hundreds of stand-alone conventional machines, are modernized with CNC machines. For such large computerized manufacturing systems to become a reality, efficient shop-floor control schemes that have high reliability, fault tolerance, extendibility, reconfigurability, and adaptability need to be developed.

This paper presents a decentralized shop-floor control scheme that follows the paradigm of auction-based decision making among autonomous entities in the manufacturing system. No master/slave relationships exist among the control entities. Decisions are made locally by the entities on the basis of information exchanged amongst each other. The independence among the system entities makes the system modular in nature. This offers several advantages. In addition to reduction in the software complexity, the local autonomy of the system entities also enhances the fault-tolerance of the system and makes it more robust to failures. System modularity also facilitates easier expansion and reconfiguration of the system with minimal software changes.

In an auction-based control scheme, a workpiece, upon arrival to the manufacturing system, broadcasts its processing needs to the machines on the shop-floor. On the basis of a variety of factors, each machine entity that is interested in processing this workpiece constructs a bid and submits it to the workpiece. The workpiece then selects the machine with the most attractive bid as the winner of the auction. Thus, task allocation to machines under this scenario is achieved through the auctioning mechanism. Such a scheme assumes that manufacturing system entities have the capability to communicate with each other and are also sufficiently "intelligent" to make decisions locally. While none of the existing manufacturing systems utilize such decentralized schemes, the development of enabling technologies in future years will make this control paradigm increasingly attractive.

This idea of auction-based decentralized control has been studied by several researchers over the past decade (Lewis, 1981; Duffie and Piper, 1986; Maley, 1987; Parunak, 1988; Shaw and Whinston, 1985; Upton, 1988). However, such research has usually focused on the interaction among workpieces, machines, and workpiece delivery entities and has ignored cutting-tool and tool carrier entities. The distinctive aspect of the research discussed in this paper is the explicit consideration of cutting-tools and tool transport entities in the auction-based framework. This requires the consideration of cutting-tool availability, negotiation between machine and tool carrier entities, and inter-machine negotiation for cutting-tools. Thus, cutting-tool management policies directly impact the performance of the auction-based shop-floor control scheme.

This paper examines the performance of large computer-controlled manufacturing systems under resource-constrained auction-based control. In particular, the impact of alternative cutting-tool management policies are evaluated based on simulation experiments. Section 2 briefly discusses the results of a preliminary study on the impact of different system loads on the performance of a manufacturing system under auction-based control. Section 3 describes the protocols for multi-level negotiation and bid-construction and presents some simulation results that illustrate the impact of cutting-tool related issues in auction-based shop-floor control.

2 PRELIMINARY RESULTS ON AUCTION-BASED CONTROL

A pilot simulation study using SLAM was conducted on a model of a manufacturing system (consisting of six functionally identical machines) under a variety of scenarios having different system loads, different tool allocations, and different levels of similarity in tool requirements. This manufacturing system processes five different part types that arrive to the system with equal probability with exponentially distributed interarrival times. Processing time for each part type is deterministic and is the same on all machines. The preference for a particular machine by a part would arise primarily from the cutting-tool availability at the machine and the load on the machine. The machines are loaded with a set of *resident tools* that are permanently available to the machine. Additional tools can be borrowed from a central tool crib. However, a penalty in the form of a time-delay is associated with the transport of each tool borrowed from the central tool magazine. The borrowed tools are returned back to the central tool crib when they are no longer needed at the machine. Infinite tool lives are assumed. Inter-machine tool sharing is not considered in this preliminary model. In addition, transport systems are not explicitly incorporated in the model. Instead, time-delays associated with material handing is included. Task allocation is opportunistic and is based on a simple auction-based scheme.

Under the auction-based scheme considered in the preliminary model, a workpiece, upon arrival to the system, broadcasts its processing needs (tool requirement and processing time) to the system. The workpiece then joins the system input queue and awaits bids from the machines.

Machines whose input queues are not full and whose current tool magazine composition contains more than half the necessary tools are eligible for bidding on the workpiece. The machine constructs a bid of the sojourn time for the workpiece on the basis of (1) the processing times for the parts in its input queue, (2) the remaining processing time of the part being currently processed, and (3) the penalty associated with borrowing any tools from the central tool crib.

The machines submit the bids to the workpiece. After a predetermined deadline, the workpiece evaluates the bids it has received. The objective of the control scheme is to maximize throughput. Hence, the primary performance measure is average time in the system. Hence, the part selects the machine with the best bid for sojourn time. (If no bids have been received, the workpiece is reintroduced into the system after a delay.) The workpiece is then transported to the machine that won the auction. After undergoing processing at the machine, the workpiece subsequently leaves the system (the assumption being that a workpiece undergoes processing at only one machine every time it enters the system and requires re-fixturing before it can visit another machine).

After a machine finishes processing a workpiece, it returns to the tool crib all borrowed tools that are not required for processing any of the parts in the machine's input queue. If any of the borrowed tools are needed for another part at the machine, it is retained in the machine's tool magazine.

One of the primary findings of this preliminary study is that a manufacturing system under auction-based control automatically adapts itself to different system loads. Specifically, when part arrival intensities are high or low, the machines behave like dedicated machines and for medium arrival rates, the machines behave like general-purpose machines. (A dedicated machine can process a limited variety (possibly only one type) of parts whereas a general-purpose machine can process a large variety (possibly all types) of parts.) These different behaviors can be explained as follows.

Consider a production scenario in which parts arrive to the system at a high rate. Once a machine has won an auction for a particular part type, it would borrow all unavailable cutting-tools necessary to process this part type. Since the number of cutting-tools of each type is limited, other machines may not be able to bid on this part type in future auctions unless the machine returns the borrowed tools to the tool crib. However, due to high part arrival rates, the machines will continue to win more auctions for that part type, and, consequently, the input buffer at the machine is at full capacity for a high proportion of the time. Since most of the parts waiting at the machine might require tools that the machine has borrowed from the tool crib, fewer opportunities (if any) to return borrowed tools to the tool crib arise. This leads to a situation where the tool mix at each machine is relatively constant over time. Consequently, every part type develops a definite preference for some machine in the system. This means that the systems that consider tooling constraints in the auction process do indeed autonomously form virtual cells (as hypothesized in earlier research by Upton (1988)).

When part arrival intensities are low, all the machines are idle for most of the time. This provides many opportunities to return borrowed tools to the tool crib. Hence, when a part arrives to the system, the machine that has the most number of tools required by this part type is usually available and, therefore, wins the auction. If the allocation of resident tools is such that each machine is suited for a particular part type, then every part that comes into the system usually succeeds in selecting the machine that has the most suitable complement of resident tools. Hence, at low part arrival rates, each machine behaves like a dedicated machine.

At intermediate part arrival intensities, machines are not very heavily loaded. Therefore, it is possible for machines to return borrowed tools and also borrow other tools from the tool crib. The machine can, therefore, constantly change the contents of its tool magazine thereby allowing it to bid on a greater variety of part types. Thus, at medium part arrival rates, there exists a greater flexibility for task allocation to machines that allows the machines to behave like general-purpose machines.

The above explanation of the different behaviors of the manufacturing system under different part arrival intensities is also supported by its other characteristics. For instance, statistics on the number of tool transfers from the tool crib were collected. It was noted that the number of tool transfers is low at high system load. When part arrival rates decrease, the number of transfers first increases and then decreases. This is a reflection of the manner in which the tool magazine composition of the machines and thereby its processing flexibility changes with different part arrival rates.

In addition, the proportion of the part types that were processed at each machine for various part arrival intensities indicated that at high and low arrival rates, each machine produced a greater proportion of one part type than of others whereas at medium arrival rates, each machine produced a more balanced quantity of each part type.

Indeed, the behavior of a machine as a dedicated or a general-purpose one is dependent not only on the part arrival intensities but also on the tools that were made resident at each machine and on the similarity among the cutting-tool requirements for the various parts. Specifically, it was noted that when the resident tools at a machine were tailored towards the processing of a single part type, the machine was more likely to become dedicated to that part type. In addition, when the tool requirements for the various part types were very different, such focused resident tooling limited the variety of part types that a machine could bid on. This consequently resulted in an increase in the number of times that an eligible machine was unable to bid because of unavailability of tools. Indeed, with increasing system load, the machine whose resident tools consisted of a few tools needed for many part types was capable of achieving a higher utilization than other machines whose resident tooling enabled them to bid for only one part type.

3 AUCTION-BASED CONTROL OF LCMSs

Under an auction-based control scheme, no entity in the system has complete knowledge or control of the system. Hence, to prevent anarchy in the system, the autonomous entities must follow a set of protocols for negotiation-based decision making. Figures 1 and 2 respectively outline the workpiece entity protocol (similar to that of Upton (1988)) and machine entity protocols that govern the behavior of these two types of entities. Details of protocols for these entities and for the transport vehicle entities can be found in Veeramani (1991). We will now briefly clarify the following two key issues associated with the proposed auction scheme:

(1) How does a machine decide whether to bid on a part? How does a machine deal with pending auction results?

(2) How are bids constructed?

3.1 To Bid or Not to Bid

When a machine receives a request for a bid from a workpiece, it first checks its tool magazine for availability of cutting tools required to process the workpiece. If the number of such tools available at the machine is below a predetermined tool availability threshold, then the machine does not respond to the request for bids. Thus, the tool availability threshold constrains the number of participants in an auction. For instance, a threshold value of 50% would imply that for a workpiece initiating the negotiation, only the machines which have over half the required tools in their own tool magazines can consider placing a bid for that workpiece.

Clearly, as the tool threshold value increases, fewer machines will qualify for bidding in an auction. This can lead to an undesirable situation where some machines are heavily loaded and others are starved of jobs while a large number of workpieces await processing at the system input buffer. To prevent this situation from

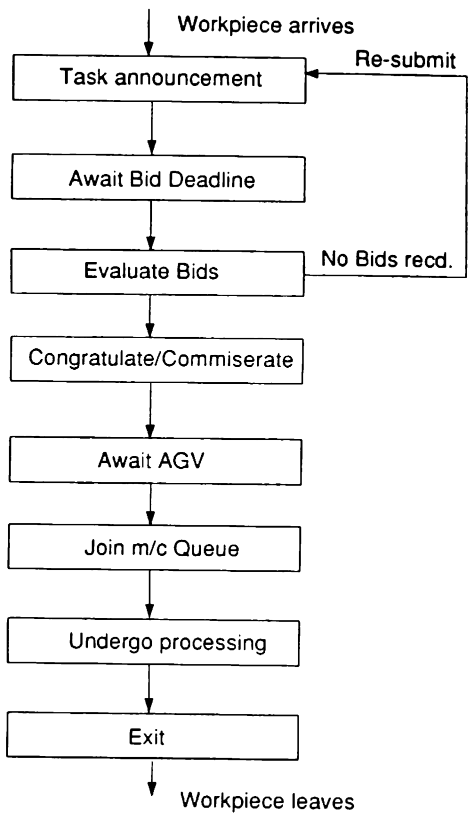

Figure 1: Workpiece Control Protocol

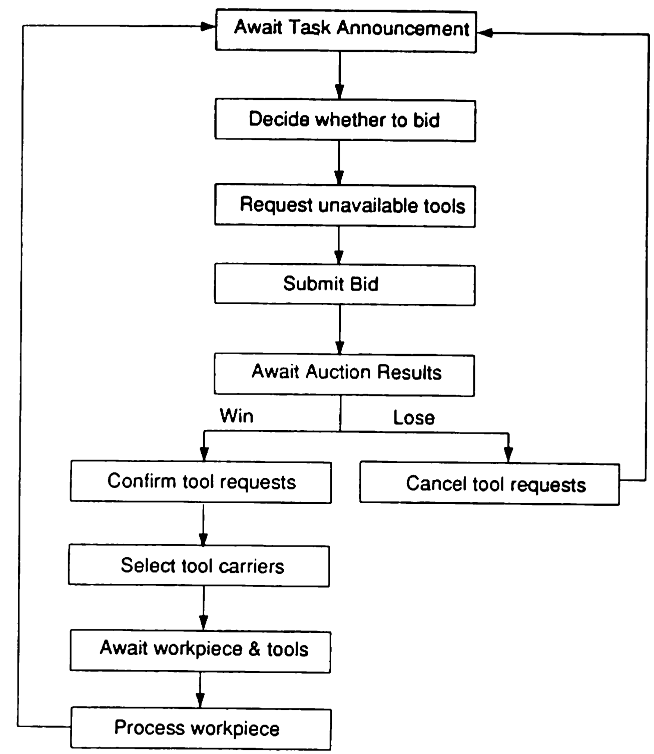

Figure 2: Machine Control Protocol

arising, when workpieces are re-introduced into the system due to lack of bids, their priority is increased. When a machine receives a task announcement from a priority job, it uses a lower threshold value in deciding whether to bid or not. Thus, machines use two different tool threshold values depending on whether the workpiece is a new workpiece or one which has been re-entered into the system with a higher priority.

Second, the job load on the machine is checked to ascertain if it can accommodate another workpiece in its input buffer. The machine considers the following issues in making this decision:

(1) Is the machine currently processing a part?
(If so, $w = 1$; 0 otherwise)
(2) Number of parts waiting in the machine's input buffer (x)
(3) Number of parts on their way to the machine (y)
(4) Number of auctions pending for which the results are not yet known (z)

There are two ways in which pending auction results can be accommodated by the machine. The first approach which shall be referred to as the *optimistic approach*, assumes that the machine wins all the pending auctions. This implies that when the machine calculates the buffer space availability for deciding whether to participate in an auction, it considers all the workpieces for which it has pending bids. In other words, if the buffer capacity of the machine is c, then the machine would consider bidding for a new part only if:

$$w + x + y + z \leq c + 1$$

In the second approach, referred to as the *pessimistic approach*, the machine assumes that it loses all the pending auctions. In other words, it ignores all pending bids while deciding whether to participate in an auction for a new workpiece. Under this scenario, the machine would qualify for bidding if:

$$w + x + y \leq c + 1$$

Upton (1988) considers both these options in his work but discards the latter approach as being "reckless" since it can result in a machine attracting more workpieces than its input buffer can handle. The pessimistic approach has, therefore, been suitably augmented (with additional acknowledgment message exchanges during auction result promulgation) to prevent such undesirable events from happening.

Thus, a machine decides whether it wants to participate in an auction for a new workpiece on the basis of availability of the necessary cutting-tools and buffer space at the machine.

3.2 Bid Construction Procedure

The bid that a machine submits to a workpiece is an estimate of the process completion time at the machine. The machine can begin processing the workpiece only when the workpiece, the machine, and the necessary tools are all available. Thus, the bid construction scheme considers three types of time delays. The first type of time delay is an estimate of the time at which the machine becomes available to process the workpiece. This quantity, denoted by α, depends on the amount of processing time remaining on the workpiece currently being machined at the machine (α_1), the cumulative processing times of the parts waiting in the machine's input buffer (α_2), and the processing times of parts with pending auction results (α_3). If an optimistic approach is adopted for dealing with pending auctions, then

$$\alpha = \alpha_1 + \alpha_2 + \alpha_3$$

If a pessimistic approach is used that ignores all outstanding bids, then

$$\alpha = \alpha_1 + \alpha_2$$

The second type of time delay is an estimate of the time by which the workpiece will be transported from the system input buffer to the machine. This corresponds to the transport-bid that the machine receives during the transport vehicle requisitioning process. The transport vehicle calculates this bid based on the information it maintains on its list of tasks and estimated completion times for each task. The transport time delay (denoted by β) can be estimated as follows:

$$\beta = \Sigma t_i + (d_{\phi,m} / v_1) + (d_{f,\phi} / v_1)$$

where Σt_i is the total time required for completing extant tasks, $d_{f,\phi}$ is the distance between the system input buffer and the location of the vehicle after completion of the last task, $d_{\phi,m}$ is the distance between the system input buffer and the machine, and v_1 is the average speed of the vehicle.

The third type of time delay is an estimate of the tool delivery time (denoted by γ). If tools need to be transported from more than one location, then the machine selects the maximum tool delivery time as the estimate of γ. The tool delivery time also accounts for the time required to exchange the tools from the tool magazine to the tool carrier and vice versa. Since the actual assignment of tool carriers to tasks is not done till the auction results have been announced, the estimate of tool delivery time is based solely on the distance between the source and destination of the tool transportation. (Information on tasks remaining for each tool carrier, and estimated transport time for each task, is used during the tool carrier selection process by the machine after it has won the auction.) Thus, for the bid construction scheme, the time delay associated with tool delivery is estimated as:

$$\gamma = \max_l (d_{l,r}) / v_2$$

where $d_{l,r}$ is the distance between the tool lending machine and the tool requesting machine, and v_2 is the average tool carrier speed. Although it is possible for the machines to communicate with the tool carriers to get accurate estimates of γ, the communication overhead involved makes this option unattractive. By allowing only the winning machine to communicate with the tool carriers, the level of message traffic in the communication network is significantly reduced.

Since all the machines are assumed to be functionally identical, the processing time for each workpiece is the same on each machine. The sojourn time for each workpiece is, therefore, determined by the total time delay involved before a machine can start processing it. Thus, the bid (B) that the machine submits to the workpiece is calculated as

$$B = \max(\alpha, \beta, \gamma)$$

3.3 Simulation Results for LCMSs

A detailed simulation using SLAM was conducted to study the performance of LCMSs under resource-constrained auction-based control. This model is larger and more detailed, and relaxes many of the assumptions that were made in the preliminary study. In particular, the system consists of 45 machines and processes 100 part types. The number of cutting-tool types that are considered is 500. The total tool inventory in the system is around 12,250. Each tool is assumed to have a finite tool life. Two types of material handling systems are included in the model: one for workpiece transport and the other for tool transport. The level of detail in the model allows the tracking of individual entities in the system including cutting-tools and tool carriers.

This section briefly reports the results of some simulation experiments that were conducted to gain a better understanding of the impact of the following factors on system performance:
(1) tool availability threshold,
(2) tool allocation to machines,
(3) how machines deal with pending auctions, and
(4) how machines deal with requests for tools from other machines.

The tool availability threshold is the proportion of requisite tools (expressed as a percentage) for a workpiece that a machine must have for it to consider participating in the auction for the workpiece. For a given tool allocation among machines, as the tool availability threshold increases, fewer machines will have tools in excess of the threshold value. This results in a decrease in the average number of received by each part and also an increase in the number of parts that do not get any bids.

Parts that do not receive any bids are re-entered into the system as priority jobs. When machines receive task announcements from priority jobs, a revised (lower) threshold value is used to determine whether the machine has sufficient tooling. Consequently, for re-entering jobs, more machines participate in each auction. Thus, as the threshold value increases, there is an initial decrease in the level of auction participation. However, as the threshold value continues to increase, the effect due to the increasing number of re-entering parts dominates that produced by the high threshold value for regular parts. In other words, increase in the number of priority jobs rejuvenates the auction process.

Consider, for instance, the number of tool request messages that are transmitted. A machine sends a tool request message to other machines and the tool crib to ascertain the availability of a specific tool that it needs to borrow. The tool request message traffic is, therefore, an indicator of the level of auction activity on the shop-floor. Figure 3 shows how the number of tool requests messages transmitted changes with the tool availability threshold. The initial drop in the tool request message traffic is due to decreasing number of candidate machines in each auction. However, as the number of priority jobs increases, more machines can participate in auctions, due to the lower tool availability threshold value for re-entering parts, resulting in an increase in the tool request message traffic.

To better understand the impact of tool allocation strategies on system performance, two kinds of allocation are considered. The first approach is referred to as a *random tool allocation* scheme. The second approach, referred to as a *clustered tool allocation* scheme, assigns sets of tools required by each part type to individual machines.

Consider the average number of bids per auction under each tool allocation scheme as a function of the tool availability threshold. At low threshold levels, machines need to have few tools for qualifying to participate in an auction. Under a random tool allocation scheme, more machines are likely to satisfy the tool availability threshold in comparison to the clustered approach. This is reflected in a higher average number of bids per auction under random tool allocation. However, at intermediate threshold values, the clustered approach performs better since sets of tools corresponding to various part types are allocated to the machines. At high threshold values, the number of re-entry parts increases. Since a lower threshold value is used for such priority jobs, the number of participants in the auctions also increases with increasing threshold value. Thus, at high threshold levels, the opposing effects of increasing threshold value and the increasing number of re-entry parts keeps the average number of bids comparable under both tool allocation schemes.

A comparison of the average time in the system indicates that overall the clustered tool allocation approach performs better than the random tool allocation approach (see Figure 4). This is because the clustered tool allocation scheme results in lower number of re-

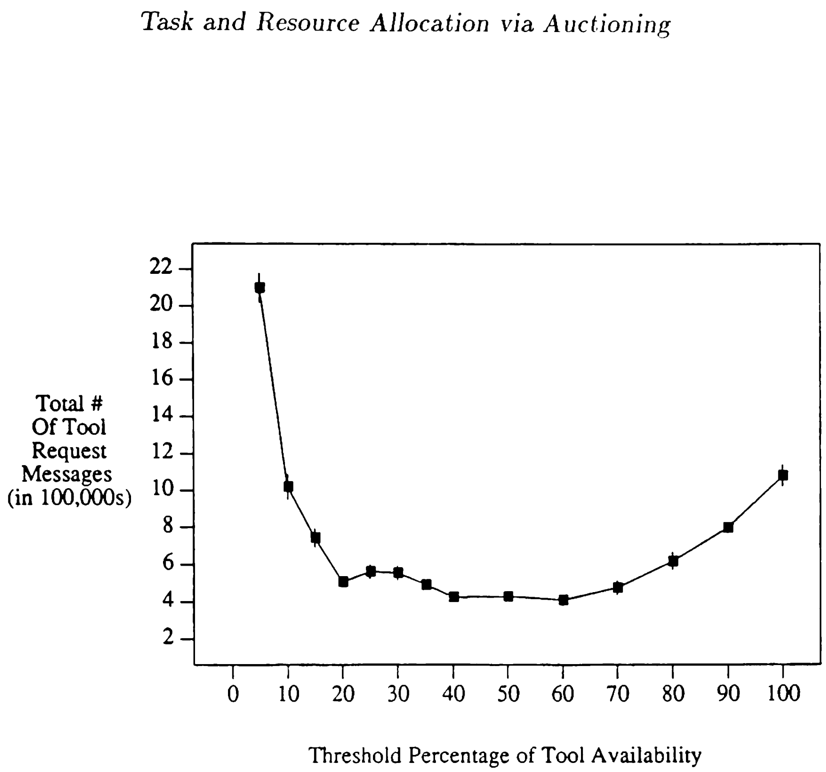

Figure 3: Tool Request Message Traffic Versus Tool Availability Threshold

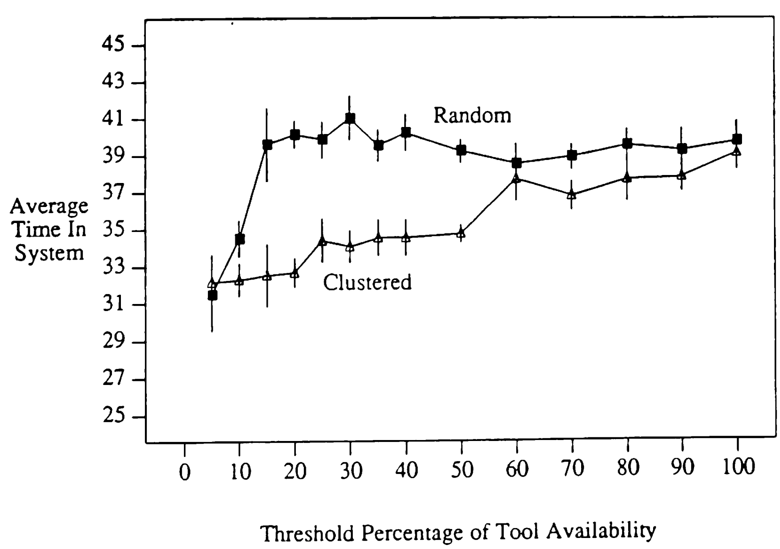

Figure 4: Average Time in System - Random vs. Clustered Allocation

entry parts for a given tool availability threshold value and also generates a higher average number of bids per auction (especially in the intermediate threshold range). As a result, each workpiece is capable of selecting from a larger number of candidate machines resulting in a better selection of the winning machine. At a low threshold value (5%), the random tool allocation appears to provide a slightly better performance in terms of the average time in the system. This can be attributed to the larger number of bids per auction that the random tool allocation scheme generates at low threshold values.

We now clarify the manner in which machines (tool donors) deal with tool requests. For each tool that a machine has in its tool magazine, the following information is maintained about its usage:

(1) Does the tool have a confirmed task at this machine?
(2) Does the tool have a confirmed tool request from another machine?
(3) Does the tool have a hold (unconfirmed reservation) on it from this machine? If so, for which part no.?
(4) Does the tool have a hold (unconfirmed reservation) on it from another machine? If so, for which part no.

For the tool crib, only (2) and (4) are relevant. The tool crib and machines that are considering tool requests from other machines shall be called *tool donors*.

When a tool donor receives a message from a machine seeking to borrow a tool from its tool magazine, it first checks whether the tool has any confirmed requests on it (based on (1) and (2)). If there are any confirmed requests, then the tool donor declines from responding to the tool request message. If not, then the donor looks at the unconfirmed tool requests (or tool holds) for that tool. If these unconfirmed tool requests are for a part number different from the one for which the tool is being requested, then the machine again ignores the tool request message.

When the unconfirmed requests are for the same (unique) part identification number that is under consideration, then the decision to respond to the tool request message can be based on two approaches. Under the first approach, referred to as the *conservative approach*, the tool donor ignores the tool request message. This implies that for any tool, the donor accepts only one unconfirmed request. If there already exists a request on a tool, then the donor accepts no further requests for it.

The other approach, referred to as the *liberal approach*, allows a donor to accept more than one request for a tool as long as all the requests are for the same part identification umber. Since only one machine will win the auction for the part, it does not lead to any conflicts or deadlocks if requests from more than one machine are accepted by the tool donor.

In general, it was found that when a part did not get any bids, there was at least one machine that did not bid because it was unable to find tool donors for necessary tools. Since, under the liberal approach, the number of tool donors for each tool is higher relative to the conservative approach, it is reasonable to expect that the liberal approach will perform better. Indeed, as a result of the higher number of tool donors, more machines can participate in each auction under the liberal approach. This leads to a higher number of bids per auction.

The higher level of participation in auctions under the liberal approach allows each workpiece to choose from a larger number of candidate machines. This leads to a better selection of machines by workpieces. In addition, the number of parts that do not get any bids is lower under the liberal approach. At low threshold values, the average number of bids under both approaches is comparable and there are few re-entry parts. As a result, the mean time in the system for both approaches is comparable at low threshold values. However, as the threshold value increases, the higher ability of machines to borrow tools and bid for parts under the liberal approach leads to a lower time in the system. With increase in the threshold value, the number of re-entry parts increases under both approaches and as a result the mean time in the system increases.

Consider the number of tools that are borrowed from the tool crib. Under the liberal approach, since machines have a greater ability to borrow tools from other machines, fewer tools are borrowed from the tool crib. This offers a couple of advantages. First, transporting tools from the tool crib can take more time since the tool crib is remotely located. This would tie up the tool carriers for longer periods of time thereby increasing the possibility of the tool handling system becoming a bottleneck. Second, since fewer tools are introduced into the system from the tool crib, the total number of tools at the machines is lower. The tool magazine capacity that is required at each machine will therefore be lower.

Figure 5 shows the number of tools borrowed from the tool crib as a function of tool availability threshold values. The tool allocation is based on the random tool allocation scheme. At low threshold values, a large number of machines qualify to participate in auctions. Therefore, there are a large number of machines seeking tools. Since the conservative approach allows each tool donor to accept only one request per tool, a large number of tool requests are submitted to the tool crib. At intermediate threshold levels, the average number of participants in auctions drops resulting in fewer requests for borrowed tools. This is reflected in a drop in the number of tools borrowed from the tool crib. At higher threshold values, the increasing number of re-entry jobs rejuvenates the auction and this results in an increase in the number of tool requests to the tool crib.

In the case of the liberal approach, the pattern of the number of tool requests to the tool crib is similar but less pronounced than that of the conservative approach.

Let us now consider the impact of how machines deal with pending auctions on the performance of the system. Under the proposed distributed shop-floor control scheme, several auctions are carried out simultaneously by several parts. This implies that a machine that is

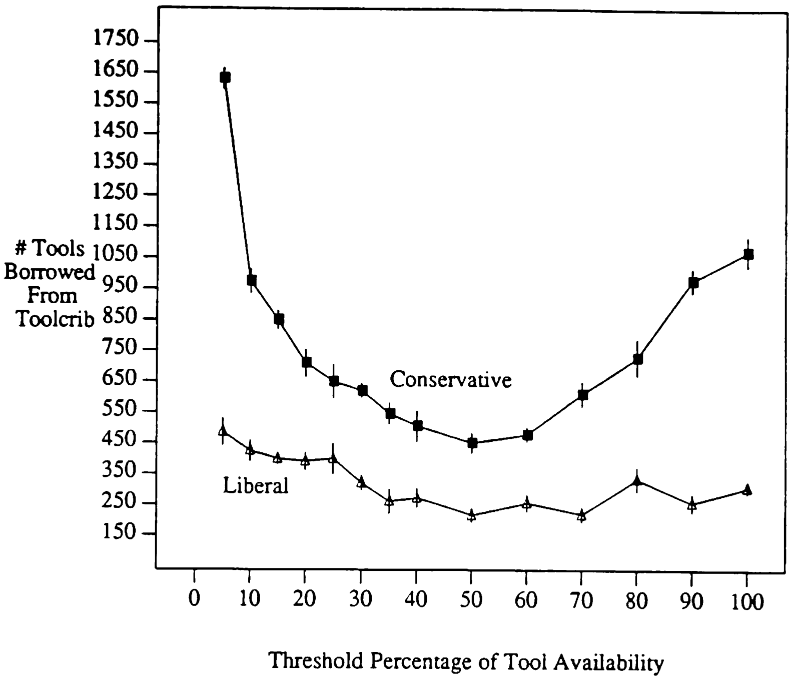

Figure 5: Number of Tools Borrowed from Tool Crib - Conservative vs. Liberal Approaches

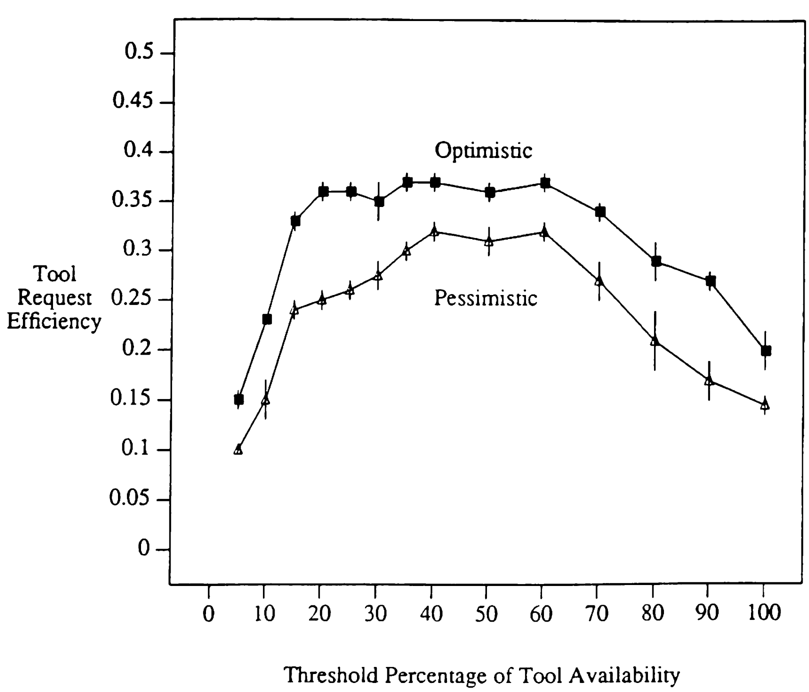

Figure 6: Efficiency of Tool Request Process - Optimistic vs. Pessimistic Approaches

constructing a bid for one auction may not be aware of the outcome of other auctions for which it has already submitted bids. As noted earlier, we consider two approaches, namely the optimistic and the pessimistic approaches, by which machines can deal with outstanding bids due to pending auction results. The optimistic approach assumes that the machine wins all pending auctions whereas the pessimistic approach assumes that the machine loses all the pending auctions. The auction process was modified by additional message exchange for the pessimistic approach to prevent machines from attracting more work than they could physically handle in their input buffer. In particular, when a part selects a machine as the winner of the auction, it first sends a congratulatory message to the machine. The machine, based on its current list of confirmed jobs, then decides whether it can accommodate this new job that it has won. If at that time the machine finds that it can no longer process the part, it sends back a message retracting its bid for the part. The part then broadcasts a congratulatory message to the next desirable machine from the auction results. The part eventually selects the machine that sends back an acknowledgment message confirming that the machine can indeed process the part. Thus, the pessimistic approach produces additional communication overhead but allows more machines to qualify for auctions. This results in a higher average number of bids per auction.

The higher level of participation in auctions under the pessimistic approach precipitates a higher traffic of messages in the communication network. For instance, Figure 6 (based on the random tool allocation scheme and the conservative approach to dealing with tool requests), indicates that the efficiency of the tool request process (that is, the ratio of number of tool confirmation messages to the number of tool hold messages) is lower in the case of the pessimistic approach than that of the optimistic approach. 10%-15% of the auctions required the selection of a runner-up machine. The additional message transmission associated with selecting the runner-up machine also adds to the communication traffic.

4 SUMMARY

This paper described a control scheme for task and resource allocation in Large Computer-Controlled Manufacturing Systems that is based on the paradigm of multi-level negotiation and resource-constrained auction-based decision making among autonomous entities in the manufacturing system. Cutting-tool management is an integral part of shop-floor control and has a significant impact on system performance. This paper has presented some simulation-based results that illustrate the impact of tool availability threshold levels, tool allocation schemes, tool sharing policies, etc.

ACKNOWLEDGMENTS

This research was supported by the National Science Foundation under grant CDR 8803017 to the Engineering Research Center for Intelligent Manufacturing Systems at Purdue University.

REFERENCES

Duffie, N.A. and R.S. Piper. 1986. Nonhierarchical control of manufacturing systems. *Journal of Manufacturing Systems* **5:2**: 137-139.

Lewis, W.C. 1981. Data flow architectures for distributed control of computer operated manufacturing systems: structure and simulated applications. Ph.D. Thesis, School of Industrial Engineering, Purdue University, West Lafayette, Indiana.

Maley, J.G. 1987. Part flow orchestration in distributed manufacturing processing. Ph.D. Thesis, School of Industrial Engineering, Purdue University, West Lafayette, Indiana.

Parunak, H.V.D. 1988. Manufacturing experience with the contract net. In: *Distributed Artificial Intelligence*, ed. M.N. Huhns, Morgan Kaufmann Publisher Inc., Los Altos, CA.

Shaw, M.J. and A.B. Whinston. 1985. Task bidding and distributed planning in flexible manufacturing. *Proceedings of the Second IEEE Conference on Artificial Intelligence*: 184-189.

Upton, D.M. 1988. The operation of large computer-controlled manufacturing systems. Ph.D. Thesis, School of Industrial Engineering, Purdue University, West Lafayette, Indiana.

Veeramani, D. 1991. Physical resource management in large computer-controlled manufacturing systems. Ph.D. Thesis, School of Industrial Engineering, Purdue University, West Lafayette, Indiana.

AUTHOR BIOGRAPHY

DHARMARAJ VEERAMANI is Assistant Professor of Industrial Engineering at the University of Wisconsin at Madison. He holds a Bachelor's degree in Mechanical Engineering from the Indian Institute of Technology, Madras, India. He received the MS and Ph.D. degrees in Industrial Engineering at Purdue University, West Lafayette, Indiana. His research interests include manufacturing systems control and concurrent engineering. He is a member of SME, ORSA, TIMS, and IIE. His membership in honor societies includes Tau Beta Pi, Alpha Pi Mu, Omega Rho, and Phi Kappa Phi.