AN ANALYSIS OF SCHEDULING POLICIES IN MULTIECHELON PRODUCTION SYSTEMS

Candace Arai Yano
Member, IIE
Department of Industrial & Operations Engineering
University of Michigan, Ann Arbor

Robert C. Carlson
Senior Member, IIE
Department of Industrial Engineering and
Engineering Managment
Stanford University

Technical Report 84-15

Revised February 1985

AN ANALYSIS OF SCHEDULING POLICIES IN MULTIECHELON PRODUCTION SYSTEMS

Candace Arai Yano
Member, IIE
Department of Industrial and Operations Engineering
University of Michigan, Ann Arbor

Robert C. Carlson
Senior Member, IIE
Department of Industrial Engineering and
Engineering Management
Stanford University

ABSTRACT

Results of a simulation study of the economics of frequency of rescheduling Material Requirements Planning (MRP) systems are presented for a single-product, two-stage system in which demand is uncertain. The results indicate that for systems with moderate demand uncertainty, frequent rescheduling to maintain customer service may be uneconomical when compared with the alternative of more stable schedules in conjunction with safety stock. This result arises primarily because the cost of "emergency" production setups which occur when rescheduling is frequent exceeds the cost of safety stock required to "protect" stable schedules.

AN ANALYSIS OF SCHEDULING POLICIES IN MULTIECHELON PRODUCTION SYSTEMS

1.0 INTRODUCTION

The issue of whether and how often to reschedule in multiechelon production systems which use Material Requirements Planning (MRP) has been debated for the past decade. Mather [8] argues for the use of firm planned orders and safety stock rather than frequent rescheduling. Bannerjee [2] reports results from a simulation study which indicate that the choice of scheduling policy has a greater effect on service levels than does safety stock.

In this paper we analyze the effects of rescheduling frequency on costs when demand is uncertain, building on the papers of Mather and Bannerjee. The objective is to determine which types of policies minimize the cost of achieving a customer service level requirement under various conditions. Service level is measured as "percent of demand filled immediately from stock," also referred to as "fill-rate." Total cost is composed of production setup costs, and inventory holding costs charged on end-of-period inventory.

We examine two scheduling policies. In the first policy the timing of production runs and the placement of orders is fixed far in advance of the planned production or order setup. Therefore, emergency setups are not permitted. We refer to such a policy as a "fixed" scheduling policy. The second policy entails replanning each period, as commonly is done in MRP systems. We refer to this as a "flexible" scheduling policy.

These policies are extremes. Many MRP systems operate with policies somewhere between these extremes. However, the use of these extreme policies in this study allows us to examine the effects of rescheduling frequency in multiechelon production systems.

Both policies are implemented in a rolling schedule environment. A rolling production schedule or plan is composed of the first-period plan of each of a series of finite horizon plans. Operationally, one would establish a plan based on known demand requirements or forecasts for a finite horizon and implement the production plan for the first period in that horizon. One would then add a new requirement or forecast to the end of the previous horizon, update other requirements/forecasts as necessary, devise a new, finite-horizon production plan, and implement the first period of that plan. The process continues in this manner. Most MRP systems operate using rolling schedules.

Other work related to rolling schedules has addressed choice of lot-sizing technique [3], modification of existing techniques to make them more robust [3,7], appropriate horizon length [1,5] and development of techniques for multi-stage problems [4]. Nearly all of this work assumes that demand is known (even if randomly generated in advance). Some work has been done (e.g., [11]) on lot sizing for single stage systems under demand uncertainty.

In this paper, we are concerned about the effect of rescheduling frequency when demand is uncertain. If rescheduling is infrequent (as in our fixed schedules), safety stock must be provided (in appropriate quantities and locations) to insure that the fill-rate is achieved. On the other hand, if rescheduling is done frequently (as in our flexible schedules), it may be possible to schedule "emergency" setups (those earlier than planned originally) when shortages are imminent. Safety stock also may be desirable, both to maintain service levels and to present too-frequent emergency setups.

We have developed algorithms to determine near-optimal safety stock levels [6,12] for two-stage assembly systems under two different scheduling policies. These algorithms are analytical, cost-based heuristics which have the objective of minimizing total cost subject to a fill-rate constraint. The algorithms are used to determine cost-effective safety stock levels, thereby enabling us to isolate the effects of scheduling policy on cost.

2.0 PROBLEM ASSUMPTIONS

Analyses are done for a single product with a two-level product structure. Two components are purchased from outside vendors and are assembled into the finished product. The product structure is diagrammed in Figure 1.

FIGURE 1

We assume that there are no production capacity constraints, and that the only type of uncertainty in the system is demand uncertainty. We assume that the demand process is stationary (demand in each period independent and identically distributed), and that all unfilled finished product demand is backordered. In the event that an insufficient number of components is available to complete the desired assembly quantity, a production run is made with the quantity available.

We use the Wagner-Whitin lot-sizing algorithm [10] independently for each item in the product structure, beginning with the end-item and moving toward the raw materials. This lot-sizing algorithm is used in the context of a rolling schedule.

The system operates as follows. Since the demand process is stationary, forecasts are set equal to mean demand. Available stock and backorders, if any, are then used to calculate net requirements. The Wagner-Whitin algorithm is implemented using these net requirements. If a lot must be produced in the current period, safety stock in the amount of $k_i\sqrt{t_i + L_i}\,\sigma$ is added to the production quantity, where

k_i = safety stock multiplier for item i

t_i = time until next planned setup for item i

L_i = leadtime for item i

σ = standard deviation of end-item demand in a single period.

The sum of t_i and L_i is used in the safety stock calculation because the safety stock must protect against forecast errors from the time the order is placed until the next lot is available for use.

If a fixed schedule is used, the Wagner-Whitin algorithm will cause an order to be placed every T_i periods only, where T_i is the natural cycle (i.e., a cyclic schedule results). Coined by Baker [1], the natural cycle is equal to $\sqrt{2S/Dh}$ where S is the setup cost, D is the expected demand per period, and h is the unit (linear) holding cost per period. It can be viewed as the the expected time between order/production setups measured in number of periods. The natural cycles are achieved by varying the setup cost, S_i, so that T_i equals the desired value. Thus, when fixed scheduling is used, the steps described above are equivalent to implementation of an (R,T) system in which one orders up to a quantity R which is equal to mean demand during $T_i + L_i$ periods plus safety stock, every T_i periods.

On the other hand, if a flexible schedule is used, one will order whenever it is anticipated (on the basis of on-hand inventory backorders and forecasted demand) that a stockout would occur otherwise. Thus, a setup may occur even if the anticipated shortage quantity is small (as is done in standard MRP logic). It is important to note that the use of the Wagner-Whitin algorithm is not a critical element in this study. It was selected because it consistently produced cyclic schedules in the fixed scheduling environment, and thus permitted us to control the values of the natural cycles in a systematic manner.

We now turn to a description of the safety stock policies. Computational results using one of the algorithms ([12]) mentioned earlier indicate that generally only finished product safety stock should be used in the fixed scheduling scenario. Qualitatively similar results have been obtained by Muckstadt and Lambrecht [9]. Thus, in our study, only finished product safety stock will be used when fixed scheduling is used throughout the system.

The results in [6] indicate that when flexible scheduling is used, safety stock serves primarily to reduce the number of emergency setups. Finished product safety stock has a small effect and component safety stock has very little effect on the fill-rate in this environment. Therefore, when flexible scheduling is used, it is important to optimize the safety stock level so as to achieve minimum cost, even if the safety stock does not affect the fill-rate.

Optimal component safety stock quantities are not known for situations in which flexible scheduling is used on both levels. Evidence from a few test cases indicated that the role of component safety stock under this scheduling policy is very similar to its role under the policy in which flexible scheduling is used only on the second level. The safety stock primarily serves to reduce emergency setups, but may also help to avert a few end-item shortages. We assume, therefore, that average total cost is nearly convex in the component safety stock quantity and that as the safety stock quantity is varied, changes in the service level are small. This assumption allows us to design reasonable test cases, but is not critical to the actual analysis. We determine approximately optimal second-level safety stock quantities via simulation by increasing the safety stock level until the average total cost begins to rise.

Fortunately, near-optimal safety stock levels for components can be obtained when fixed scheduling is used on level 1 and flexible scheduling is used on level 2. Using the basic and the modified algorithms which we developed (see [6]), we determine upper and lower bounds on optimal safety stock levels for the items using flexible schedules. From a sample of safety stock levels between these bounds, we choose the one with the lowest average total cost.

In each period, a plan is determined for a finite horizon, current decisions are implemented, a realization of the random demand process occurs, and the finite horizon rolls forward. The process then repeats. This differs from most of the earlier rolling horizon literature because demand is stochastic (not simply time varying) and is not known in advance.

3.0 SIMULATION MODEL

We developed a simulation model to evaluate the performance of the two scheduling policies. In the simulation model, demand in each period is distributed normally with a mean of 200, and a standard deviation of 10, 30, or 50.

We use a planning window of 24 periods, which for all cases in this study, is three or more times the length of the natural cycle of the component with the largest natural cycle. The fixed schedule is achieved by fixing the timing (but not the quantity) of all orders for a period of time equal to the largest integer multiple of the natural cycle less than the length of the planning window.

There are several reasons for this approach. First, research by Baker [1], Blackburn and Millen [3], and Carlson, Beckman and Kropp [5] indicates that using a horizon equal to an integral multiple of the natural cycle is better than a non-integral multiple when the Wagner-Whitin algorithm is implemented in a rolling horizon environment. Second, fixing the timing but not the quantity of the orders provides some latitude for responding to demand fluctuations without changing the ordering interval. Third, this technique limits production schedule changes to the end of the horizon, thereby essentially eliminating "nervousness" in the system. For instance, an item with a natural cycle of 4 would have its production schedule fixed for (24/4 - 1) x 4 = 20 periods when the length of the planning window is 24 periods. This technique also serves to avoid scheduling setups whose timing may not be optimal because of end-of-horizon effects. The production schedule for the latter periods becomes fixed as the horizon rolls forward.

We simulate 50 problems, each with a 24-period horizon, using approximately optimal safety stock levels for the items with flexible scheduling. We compare these results with results obtained from simulation of the same problems (under the same stochastic conditions) under fixed scheduling. In the fixed scheduling simulations, component safety stock quantities are set to zero, while end-item safety stock levels are varied. This is consistent with findings (see [9,12]) that in all realistic situations, it is a optimal to carry no component safety stock when fixed scheduling is used on both levels and a fill-rate criterion is used. The end-item safety stock level is varied as necessary in order to achieve a fill-rate approximately equal to that obtained under flexible scheduling in the corresponding case. The desired fill-rate is set implicitly by setting the value of the end-item safety stock multiplier (k_1) for the situation with flexible scheduling.

Note that if flexible scheduling is used at a particular level in the product structure, it must be used at all stages earlier in the production process (predecessors) in order to permit adequate coordination. Otherwise, required components are not likely to be available for any attempted emergency setup.

4.0 EXPERIMENTAL DESIGN AND RESULTS

We study two situations:

(1) A two-level product structure with flexible scheduling on both levels, and

(2) a two-level product structure with fixed scheduling on level 1 and flexible scheduling on level 2.

Each of the above is compared with an equivalent product structure with fixed scheduling on both levels, in terms of cost and fill-rate. For each of the situations above, we select a few problems for detailed analyses. The problems are selected in a manner so as to allow flexible scheduling the greatest possible advantage relative to fixed scheduling.

This selection criterion leads to problems with the following characteristics:

(a) Short component leadtimes.

(b) Low component holding costs.

(c) High service level criterion.

(d) High demand variability.

(e) Natural cycle combination giving significant advantage to flexible scheduling.

When component leadtimes are short, small quantities of safety stock can reduce emergency setups dramatically. If, in addition, the holding cost associated with this safety stock is low, the savings from reduction of setups can far outweigh the cost of the safety stock. The high service level criterion clearly gives flexible scheduling an advantage since the policy is less constraining and more adaptive. This is particularly true when demand variability is high.

The final selection criterion is a natural cycle combination giving significant advantage to flexible scheduling. Two findings which we discuss more fully in another paper [13] facilitate this selection process. The relationships were found to be consistent across levels of demand uncertainty and natural cycle lengths. The first finding is that under fixed scheduling, for a given T_1 and finished product safety stock quantity, the service level achieved when $T_1 = T_2 = T_3$ is higher than for systems with all other combinations of natural cycles. Simulation results for a typical case are illustrated in Figure 2. For each natural cycle vector indicated, average achieved service level (fill-rate) is plotted as a function of "theoretical" service level that would be achieved for a one-stage system with the same cumulative leadtime. The safety stock factor is set so as to provide protection during the total cumulative leadtime plus T_1. In other words, finished product safety stock is set equal to

$$k\sqrt{T_1 + L_1 + \max\ (L2, L3)}\ \sigma$$

Component safety stock is set equal to zero to isolate the effect of the natural cycle combinations on the fill-rate.

FIGURE 2

The second finding is that the opposite relationship is true when flexible scheduling is used. Systems in which $T_2 > T_1$ and $T_3 > T_1$ achieve higher service levels than other natural cycle combinations. Typical results are shown in Figure 3.

FIGURE 3

We conclude, therefore, from these two findings that flexible scheduling has an advantage over fixed scheduling when $T_2 > T_1$ and $T_3 > T_1$. Given the particularly large and well-defined effect of natural cycles in these circumstances, we decided to use the natural cycle combinations $\underset{\sim}{T} = (2, 4, 4)$ and $\underset{\sim}{T} = (4, 8, 8)$ for our studies, where $\underset{\sim}{T} = (T_1, T_2, T_3)$.

We, therefore, in each case following, are comparing two systems: one uses flexible scheduling on one or both levels and the other uses only fixed scheduling. For the systems with some flexible scheduling, the finished product safety stock factor is set to 1.0, yielding fill-rates in excess of 97% (and generally larger). Component safety stock quantities are "optimized" using a heuristic ([6]) if applicable, and via simulation otherwise.

The systems with fixed scheduling throughout have only finished product safety stock (consistent with our results reported in [12]) and the safety stock factor is adjusted (necessarily slightly upward because of the limitations of the inflexible schedules) to achieve fill-rates approximating the comparable system with some flexibility in scheduling.

4.1 Results for Flexible Scheduling on Both Levels

We begin our analyses with the following situations with flexible scheduling on both levels. They are to be compared with equivalent product structures with fixed scheduling on both levels.

Case 1:

$\underset{\sim}{h} = (1.0,0.1,0.1)$

$\underset{\sim}{T} = (2,4,4)$

$\underset{\sim}{L} = (1,1,1)$

$\sigma = 50$

$k_1 = 1.0$ when flexible scheduling is used on both levels

Case 2:

$\underset{\sim}{h} = (1.0,0.1,0.1)$

$\underset{\sim}{T} = (4,8,8)$

$\underset{\sim}{L} = (1,1,1)$

$\sigma = 50$

$k_1 = 1.0$ when flexible scheduling is used on both levels

For Case 1, approximately optimal second-level safety stock multipliers are

$$k_2 = k_3 \quad [0.8,1.0]$$

For Case 2, approximately optimal second-level safety stock multipliers are

$$k_2 = k_3 = 0.6$$

The average of 50 observations is plotted in Figure 4 for each of the following safety stock vectors in Case 1:

Flexible scheduling:	$\underset{\sim}{k} = (1.0,0.6,0.6)$
	$\underset{\sim}{k} = (1.0,0.8,0.8)$
	$\underset{\sim}{k} = (1.0,1.0,1.0)$
	$\underset{\sim}{k} = (1.0,1.2,1.2)$
Fixed scheduling:	$\underset{\sim}{k} = (1.0,0,0)$
	$\underset{\sim}{k} = (1.2,0,0)$
	$\underset{\sim}{k} = (1.4,0,0)$

and in Figure 5 for each of the following safety stock vectors in Case 2.

Flexible scheduling:	$\underset{\sim}{k} = (1.0,0.6,0.6)$
Fixed scheduling:	$\underset{\sim}{k} = (1.0,0,0)$
	$\underset{\sim}{k} = (1.2,0,0)$

FIGURES 4 and 5

Plots of average cost versus average service level (fill-rate) in Figures 4 and 5 indicate that fixed scheduling on both levels results in "better" performance than flexible scheduling on both levels. The output from the simulation is not amenable to statistical analysis because of its bivariate nature and the irregular pattern of costs and service levels when flexible scheduling is used on both levels. Therefore, we cannot be certain that the differences are statistically significant.

We can conclude, however, that the fixed scheduling policy performs at least as well, and perhaps with substantial cost savings, relative to the policy of flexible scheduling on both levels. Recall that only finished product safety stock is varied in the fixed scheduling scenarios. Therefore, although we have obtained a limited number of simulation results, it is reasonable to assume that the actual relationship between the fill-rate and total cost is a well-behaved concave, monotonically increasing function. Thus the straight lines in the figures represent "worst case" results and dominance of the fixed schedule is evident.

Recall also that the two cases analyzed here were chosen so as to afford advantage to the flexible scheduling policy. Flexible scheduling results in inferior performance even under opportune conditions, and will perform even worse in relative terms as parameter values become less advantageous.

4.2 Results for Fixed Scheduling on Level 1 and Flexible Scheduling on Level 2

We now turn to the comparison of the policies of fixed scheduling on the first level and flexible scheduling on the second level versus fixed scheduling on both levels. The following cases are studied:

Case 3: $\underset{\sim}{h} = (1,0.1,0.1)$

$\underset{\sim}{T} = (2,4,4)$

$\underset{\sim}{L} = (1,1,1)$

$\sigma = 50$

$k_1 = 1.0$ when flexible scheduling is used on the second level only

Case 4: $\underset{\sim}{h} = (1,0.1,0.1)$

$\underset{\sim}{T} = (4,8,8)$

$\underset{\sim}{L} = (1,1,1)$

$\sigma = 50$

$k_1 = 1.0$ when flexible scheduling is used on the second level only

For Case 3, using algorithms in [6] we determined that the approximately optimal second-level safety stock multiplier lies between 0.5 and 0.8. For Case 4, the approximately optimal second-level safety stock multiplier lies between 0.8 and 1.0.

The average of 50 observations is plotted in Figure 6 for each of the following safety stock vectors in Case 3:

Flexible scheduling
on level 2: $\underset{\sim}{k} = (1.0,0.6,0.6)$

$\underset{\sim}{k} = (1.0,0.7,0.7)$

Fixed scheduling:
$\underset{\sim}{k} = (1.2,0,0)$

$\underset{\sim}{k} = (1.4,0,0)$

and in Figure 7 for each of the following safety stock vectors in Case 4:

Flexible Scheduling
on level 2: $\underset{\sim}{k} = (1.0,0.9,0.9)$

Fixed scheduling: $\underset{\sim}{k} = (1.2,0,0)$

$\underset{\sim}{k} = (1.3,0,0)$

It is evident from these figures that the policy using fixed scheduling on level 1 and flexible scheduling on level 2 is also less preferred than the policy using fixed scheduling on both levels. These two figures illustrate plots of the average cost-average service level relationships for the two policies.

<u>FIGURES 6 and 7</u>

5.0 DISCUSSION

The cost penalty as a fraction of total system cost, computed as $(C_{flexible} - C_{fixed})/C_{fixed}$ from using flexible scheduling declines as fixed scheduling is used at levels deeper in the product structure. In the results presented here, the penalty from using flexible scheduling on both levels is 3% - 7%. The penalty from using a policy of fixed scheduling on level 1 and flexible scheduling on level 2 is approximately 1%. Since these particular situations were chosen so as to give significant advantage to the use of flexible scheduling, we can conclude that flexible scheduling may not be cost-effective. We want to emphasize the fact that the only costs included in these comparisons are setup costs and inventory holding costs. If costs associated with a change in the timing of a setup were included, the cost penalty from using flexible scheduling could increase dramatically.

Even if setup costs are considered not to be "real" but rather to be shadow prices for capacity constraints, a similar result could hold in many circumstances. Changes in planned setups can be costly even if capacity is not tight. Timing and quantity of production runs of predecessors and procurement plans must be revised. Setups which occur earlier than originally planned (as are most changes under flexible scheduling) typically necessitate expediting. Workforce plans may need to be revised and additional overtime employed. Whether some or all of these costs are relevant, it is possible that a relative ly cyclic schedule with safety stock as a buffer may be more economical or desirable than a policy with frequent rescheduling even when some attempt is made to optimize safety stock.

Conventional wisdom suggests that flexibility is desirable. However, that flexibility comes with a price. If "emergency" orders are permitted, as in the flexible scheduling environment, they <u>will</u> occur, incurring setup costs more frequently than may be desirable. This finding also has been made by Wemmerlov and Whybark [11]. A significant amount of safety stock is required to prevent most of these emergency setups from occurring, and indeed, it is not economical to attempt to prevent all emergency setups.

The instability of the "flexible" schedules can (and does) result in parts shortages (and therefore "short" batch regular and emergency production runs with whatever is available) as well as "unintended" work-in-process inventory at other times. Thus, not only are setup costs higher with flexible schedules than with fixed schedules, but inventory holdings costs may be greater as well.

The fixed scheduling scenario, on the other hand, <u>forces</u> coordination among production stages. Therefore, each facility "knows" when to begin a production run and approximately how much will be needed far in advance. The demand uncertainty is dealt with using finished product safety stock and an adjustment of production quantities (but not production intervals). We have dealt with only a two-stage system. Although the results may not be generalizable to larger and more complex systems, it would appear that the level of coordination afforded by the fixed scheduling policy would be even more advantageous in many such situations. A significant amount of research in the areas of lot sizing and safety stocks in multi-stage systems remains to be done before the larger question can be answered conclusively. Nevertheless, the results here suggest that policies which encourage coordination may be preferable to those which do not.

The analyses were done with the assumption of moderately variable but stationary demand. If the demand process is highly erratic and nonstationary, however, some flexibility in scheduling may be desirable. Further research is necessary in order to evaluate various scheduling policies in such an environment.

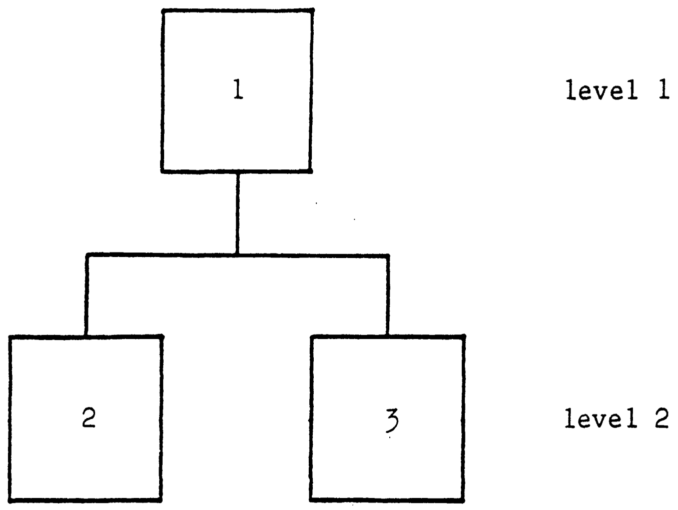

Figure 1

Two-Level Product Structure With Component Numbers

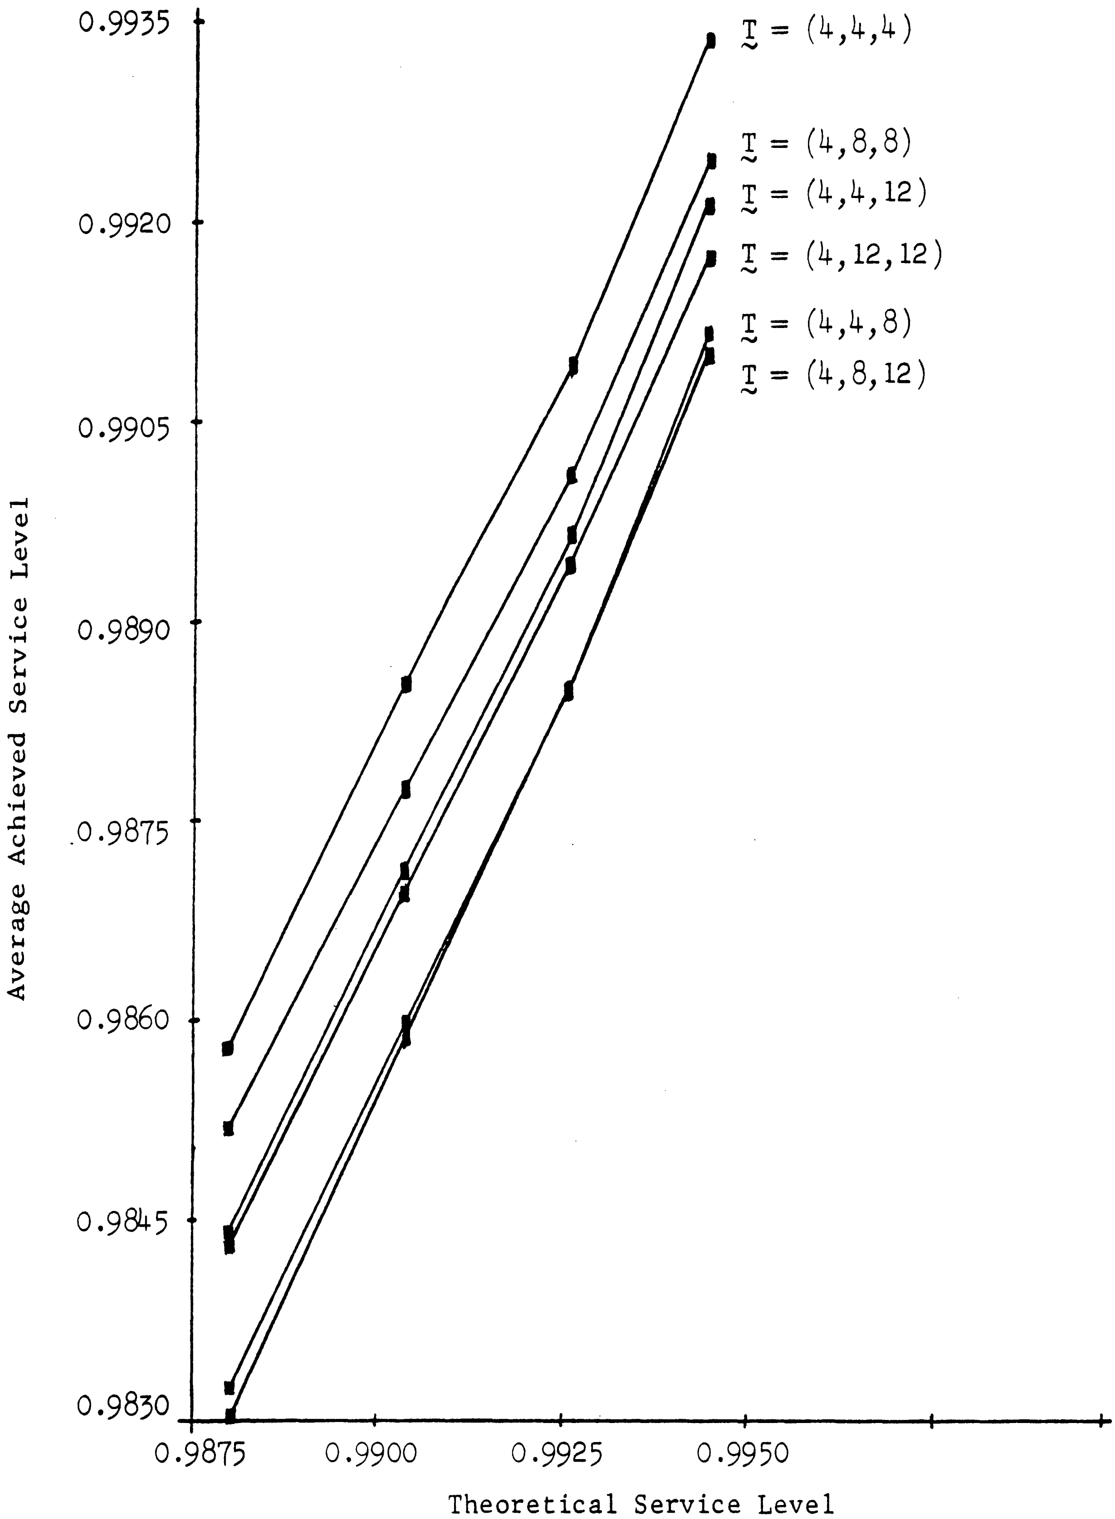

Figure 2

Average Achieved Service Level as a Function of Theoretical Service Level: $T_1 = 4$, $\sigma = 10$ With Fixed Scheduling

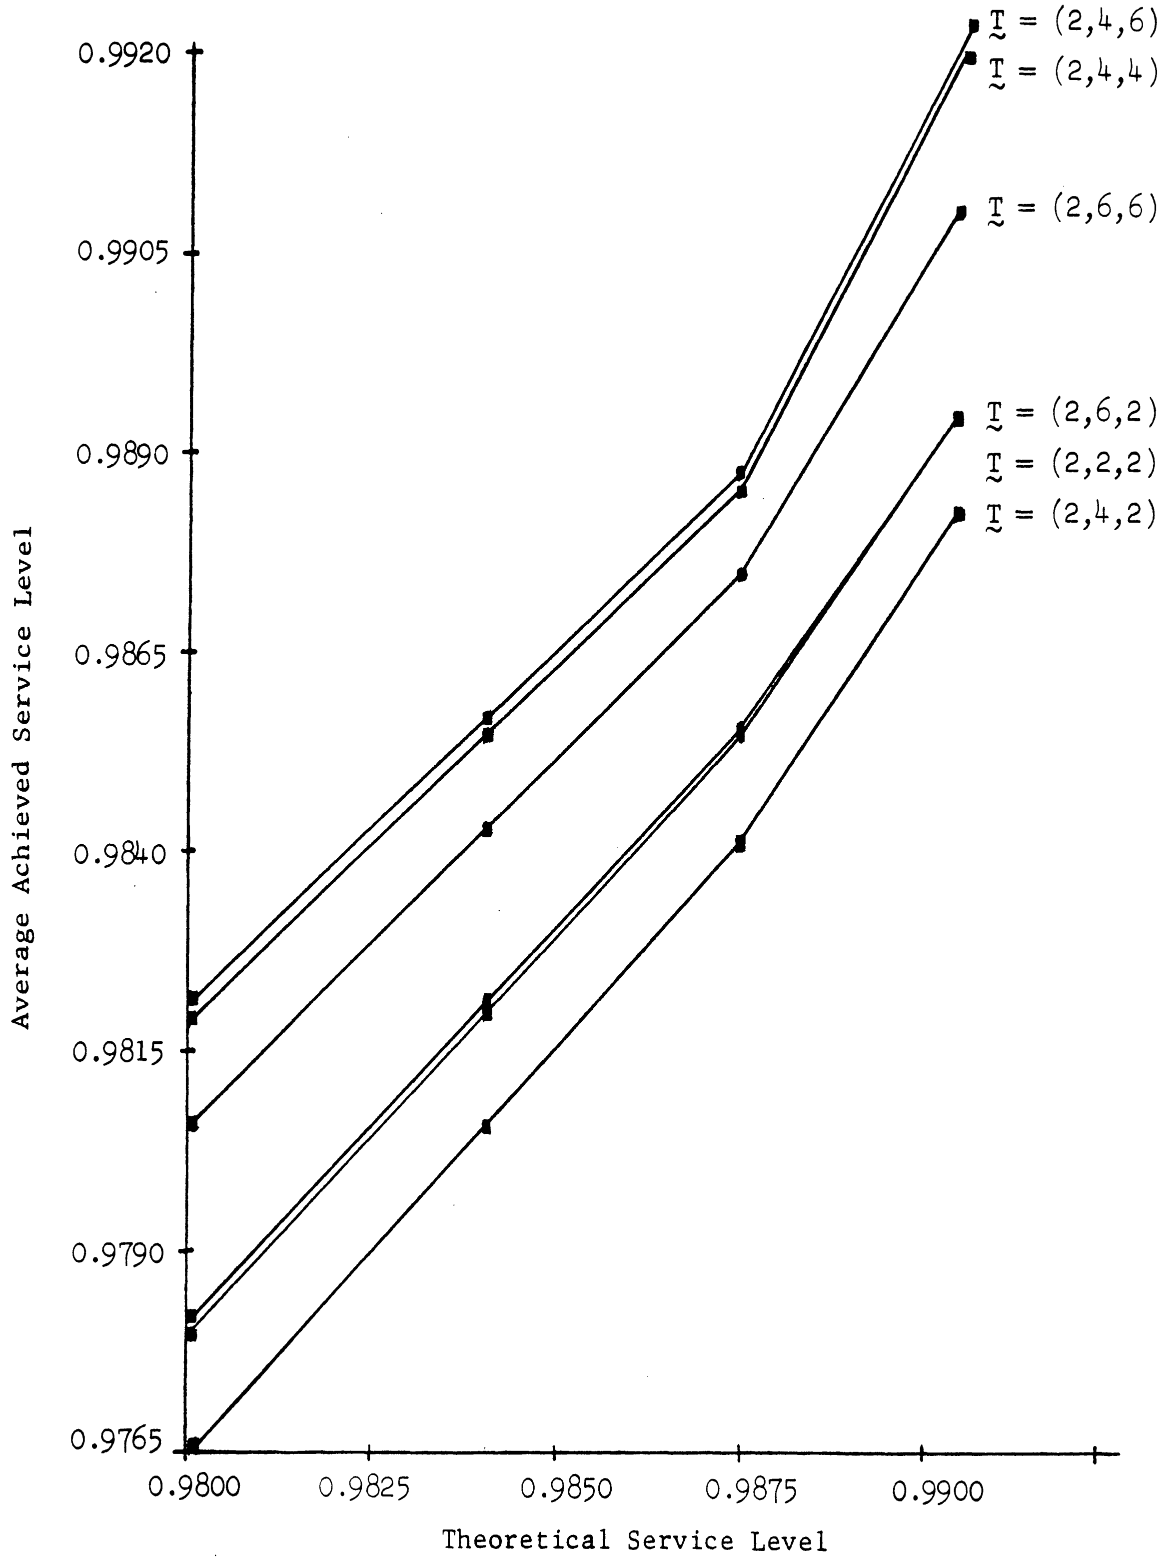

Figure 3

Average Achieved Service Level as a Function of Theoretical Service Level: $T_1 = 2$, $\sigma = 10$ With Fixed Scheduling on Level 1 and Flexible Scheduling on Level 2

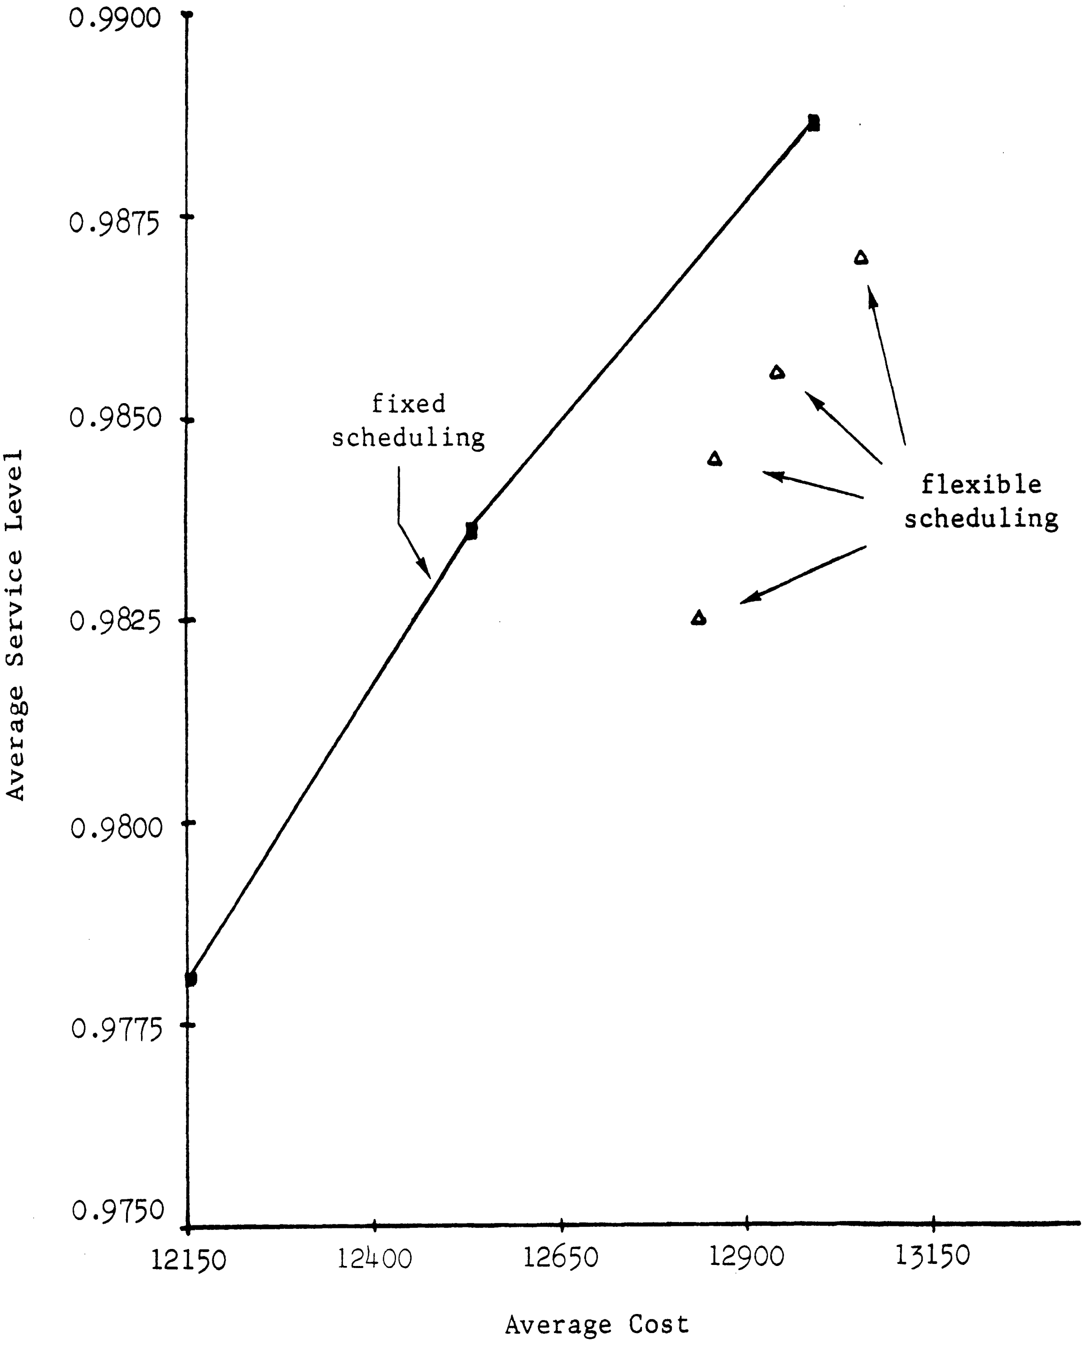

Figure 4

Comparison of Fixed and Flexible Scheduling: Case 1

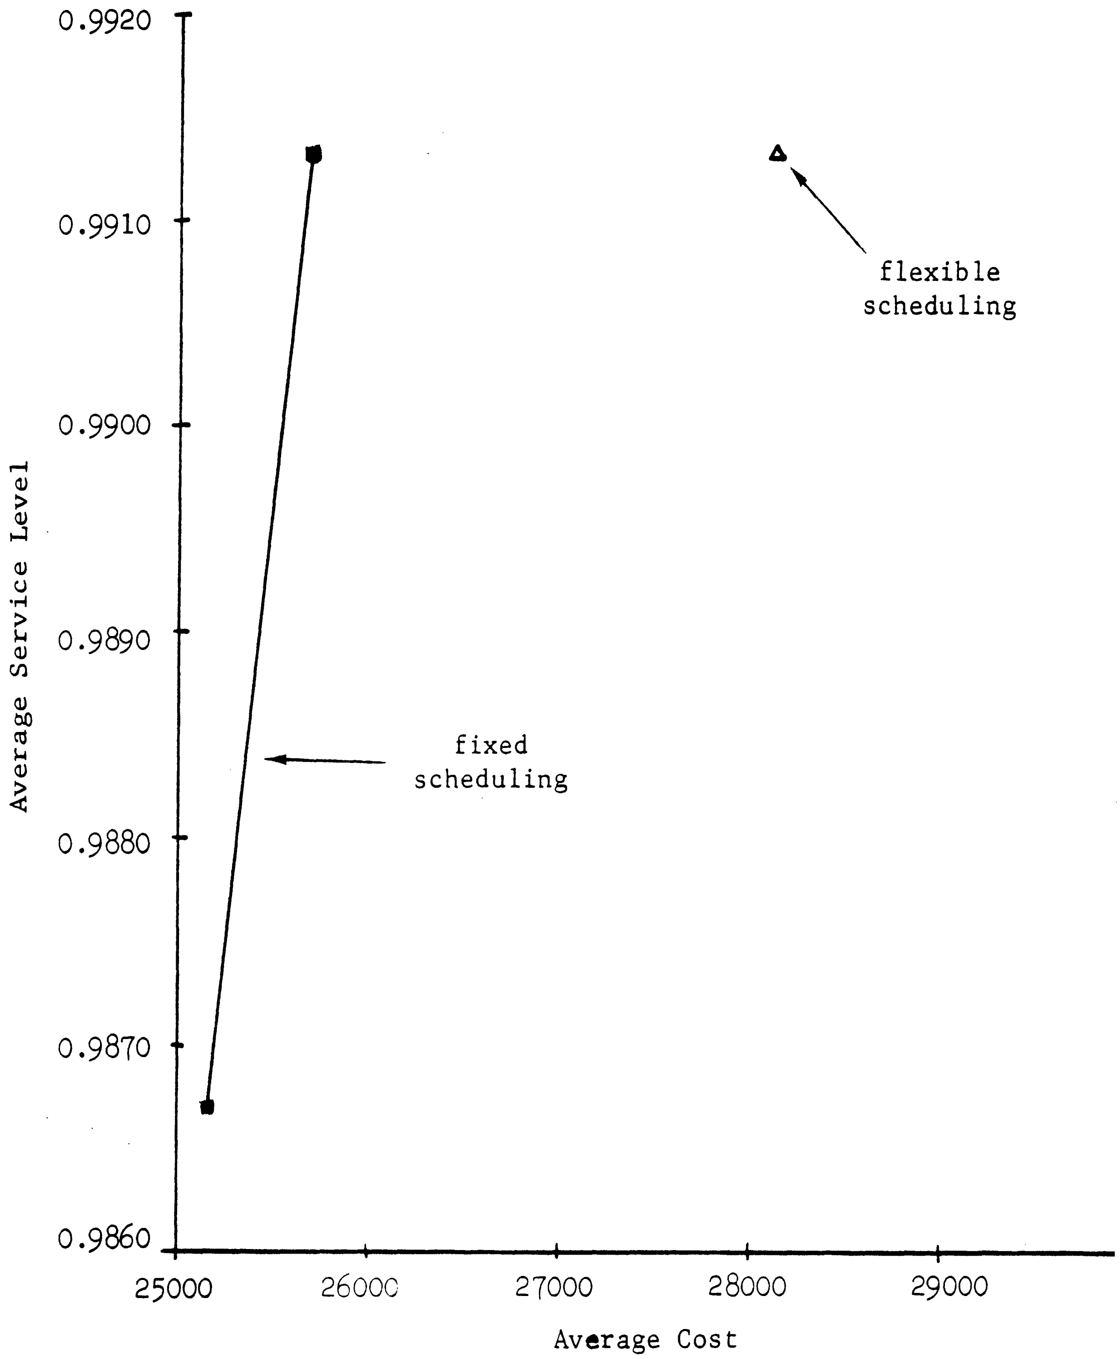

Figure 5

Comparison of Fixed and Flexible Scheduling: Case 2

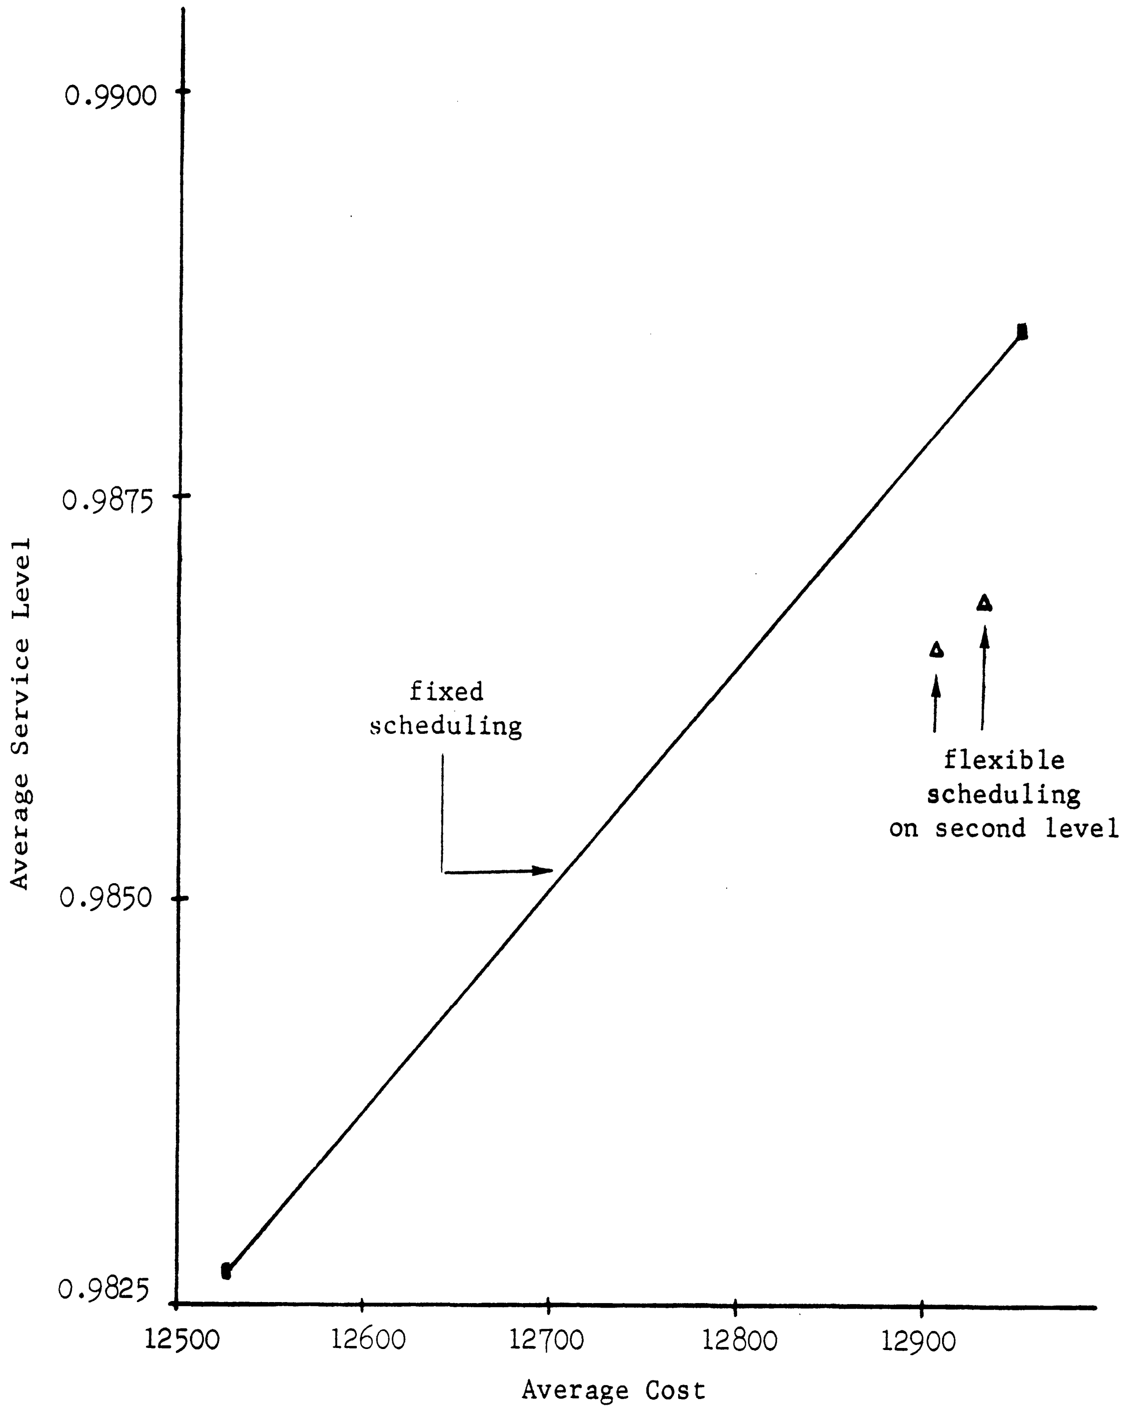

Figure 6

Comparison of Fixed Scheduling With Fixed Scheduling on Level 1
and Flexible Scheduling on Level 2: Case 3

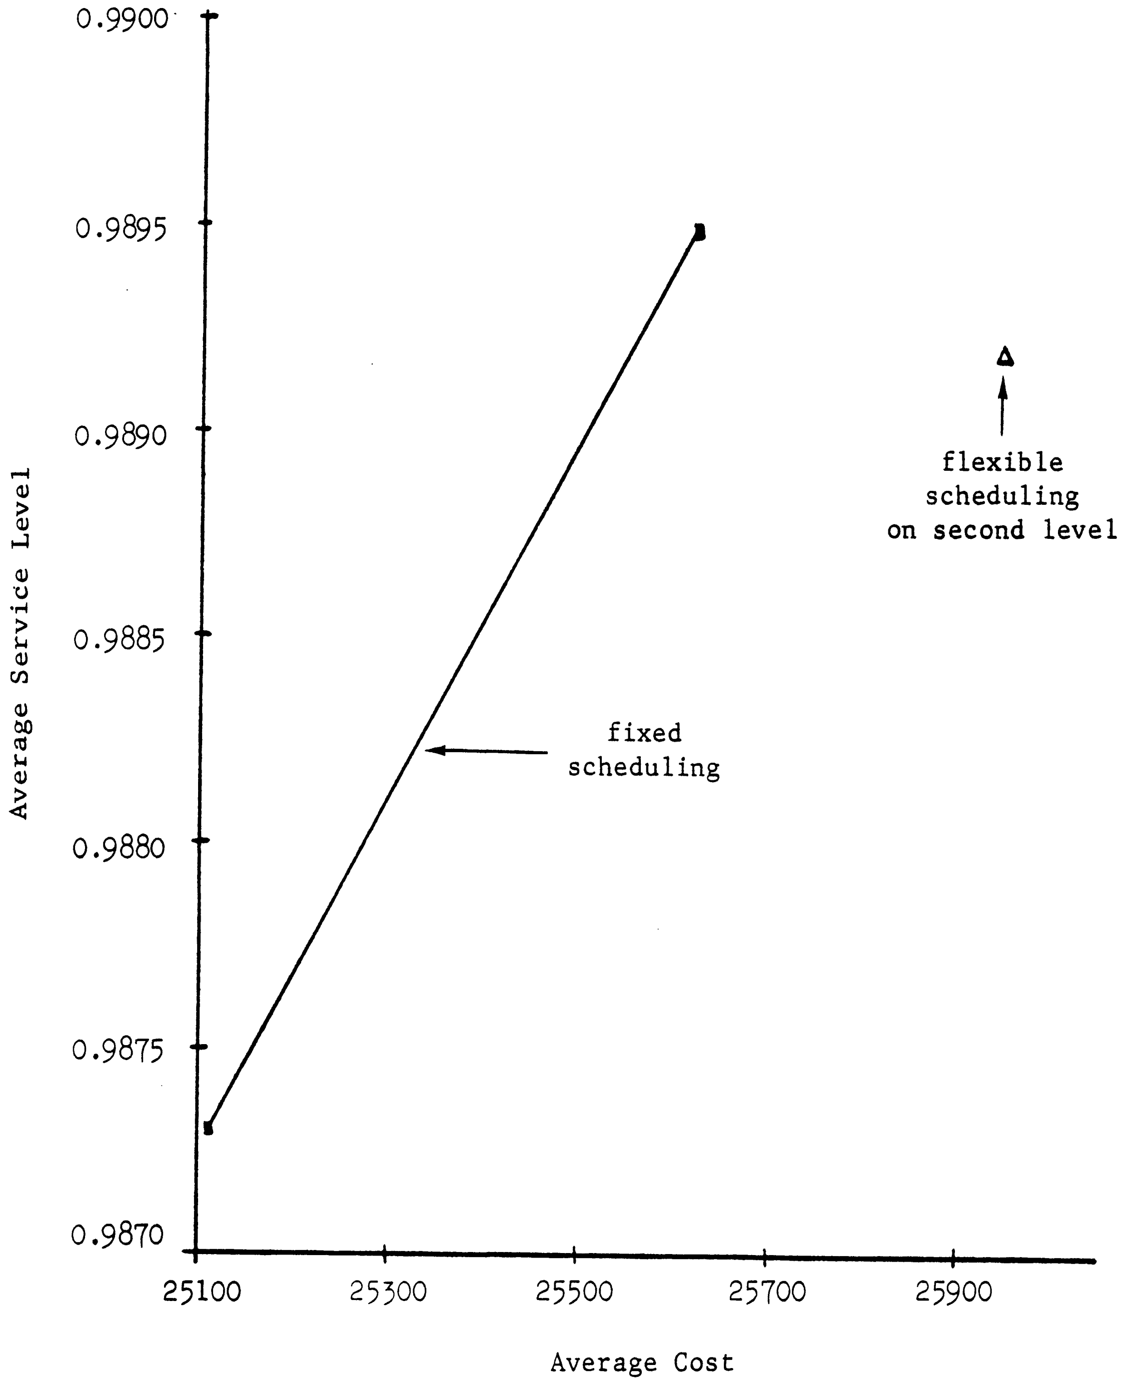

Figure 7

Comparison of Fixed Scheduling With Fixed Scheduling on Level 1
and Flexible Scheduling on Level 2: Case 4

REFERENCES

[1] Baker, K. R., "An Experimental Study of the Effectiveness of Rolling Schedules in Production Planning." Decision Sciences, 8 (1977), 17-27.

[2] Bannerjee, A., "Relative Importance of Scheduling and Safety Stocks in Material Requirements Planning Systems." AIDS Proceedings, New Orleans (1979), 251.

[3] Blackburn, J.D. and R. A. Millen, "Heuristic Lot-Sizing Performance in a Rolling-Schedule Environment." Decision Sciences, 11 (1980), 691-701.

[4] Blackburn, J.D. and R. A. Millen, "Improved Heuristics for Multi-Echelon Requirements Planning Systems." Management Science 28 (1982), 44-56.

[5] Carlson, R. C., S. L. Beckman and D. H. Kropp, "The Effectiveness of Extending the Horizon in Rolling Production Scheduling." Decision Sciences, 13 (1982), 129-146.

[6] Carlson, R. C. and C. A. Yano, "Buffering Against Demand Uncertainty in Material Requirements Planning Systems--Systems with Emergency Setups for Some Components." Technical Report 84-13, Department of Industrial and Operations Engineering, University of Michigan, 1984.

[7] Chand, S., "A Note on Dynamic Lot Sizing in a Rolling Horizon Environment," Decision Sciences 13 (1982), 113-118.

[8] Mather, H., "Reschedule the Reschedules You Just Rescheduled--A Way of Life for MRP?" Production and Inventory Management, 18 (1977), 60-79.

[9] Muckstadt, J. and M. Lambrecht, "Analysis of Safety Stock Policies for MRP Systems," Technical Report No. 562, School of Operations Research and Industrial Engineering, Cornell University, 1982.

[10] Wagner, H. H. and T. Whitin, "Dynamic Version of the Economic Lot Size Model." Management Science, 5 (1958), 89-96.

[11] Wemmerlov, U. and D. C. Whybark, "Lot-Sizing Under Uncertainty in a Rolling Schedule Environment," International Journal of Production Research 22 (1984), 467-484.

[12] Yano, C. A. and R. C. Carlson, "Buffering Against Demand Uncertainty in Material Requirements Planning Systems--Systems With No Emergency Setups." Technical Report 84-14, Department of Industrial and Operations Engineering, University of Michigan, 1984.

[13] Yano, C. A., and R. C. Carlson, "Impact of System Parameters and Scheduling Policies on System Performance in Material Requirements Planning Systems," Technical Report 84-20, Department of Industrial and Operations Engineering, University of Michigan, 1984.