OPERATIONS RESEARCH IN MANUFACTURING

Charles J. McCallum, Jr.

***Charles J. McCallum, Jr.**, is head of the Operations Research department at AT&T Bell Laboratories in Holmdel, New Jersey. He is responsible for operations research theory and practice, with emphasis on developing new operations research methodology and applying existing and new operations research techniques to the product realization process and other corporate efforts. He joined AT&T in 1970. He received an S.B. degree in mathematics from Massachusetts Institute of Technology and an M.S. in statistics and Ph.D. in operations research from Stanford University.*

Operations research models, although common in industry, government, and education, have not been widely used in manufacturing. That is changing, however, as operations research is brought to bear on today's extremely complex problems. Manufacturing managers can no longer depend solely on experience, but need the help of quantitative decision support tools. This article illustrates the importance of operations research in manufacturing. It provides examples of the use of operations research in factories and in manufacturing planning and support. It discusses future trends of operations research in manufacturing.

The Role of Operations Research

Operations research models, algorithms, and computer tools have played an impressive and essential role in leading to improved decision-making in the industrial world. They have brought a substantial economic savings by allowing extensive quantitative analysis on mathematical models rather than the costly development of physical prototypes. The savings in time has often been even more important than direct dollar savings.

Models can often be constructed, and revised, much more quickly than would be possible for physical prototypes. Moreover, by using sensitivity analysis on the mathematical models, it is possible to scrutinize more alternatives than would otherwise be the case. "What if" questions can be asked and answered, allowing the examination of a far richer range of possible scenarios before crucial decisions are made. In the manufacturing arena, in particular, operations research has been successfully and effectively applied in such areas as production scheduling, inventory management and control, resource allocation, and production layout and design.

The mathematical abstractions that are a key ingredient in operations research model building often lead to generic solution approaches that can reasonably be applied, with minor change, in a variety of different areas. Complex problems that look different close-up often are similar when the fundamental characteristics and issues are abstracted during the modeling process. Indeed, the building of mathematical models in itself often leads to invaluable insight, for example, regarding a separation of significant issues from unimportant details. In addition, preparing for the quantitative analysis of mathematical models requires measurability; objectives must be stated in quantitative terms that lead to a more formal examination of decision alternatives.

The operations research approach also leads to a system-oriented view of problems. Small issues are less likely to be tackled in isolation. Rather, there is a forced tendency to seek global rather than myopic answers—answers that are best for the whole instead of a piece of the whole. Mathematical modeling in operations research also leads to pursuit of what is best (or optimal) rather than what merely works (and hence is feasible).

Operations research models are common in industry, government, and education. Nevertheless, in the manufacturing arena they have, until recently, been underused. Today's competitive environment, however, will no longer allow that. Manufacturing problems today are exceedingly complex. Products of a greater variety are being demanded more quickly than ever before. And these products must be of higher quality and lower cost than those in the past. The manufacturing manager of today is no longer able to make decisions based solely on experience but needs the added benefit of quantitative decision support tools. Mathematical models and algorithms, and the analytical computer tools that follow naturally from them, must be developed for frequent reuse in seemingly diverse situations.

Although operations research is in part an art, it is primarily a scientific discipline. In abstracting a mathematical model from a real-world problem, the modeler must know what to include and what to exclude. Fundamental issues must be represented mathematically and unimportant details must be ignored. The model must capture the essence of the problem, but the need for mathematical tractability must be kept in mind. The algorithmic solution of well-formulated mathematical models is a highly technical task requiring considerable mathematical expertise in the techniques of operations research. In addition, efficient computer implementation of appropriate mathematical algorithms is not a trivial undertaking. Moreover, current state-of-the-art knowledge is not enough. It is necessary to advance the state of the art through basic research work. This is required both to permit solution of today's complex problems and associated models and to search successfully for generic models, techniques, and computer analysis tools that will allow the solution of a multiplicity of seemingly diverse problems in a cost-effective manner. The importance of such research work cannot be overstated; the technical expertise necessary to develop and apply appropriate operations research models, algorithms, and tools represents an important competitive advantage.

The purpose of this paper is to illustrate the importance of operations research in manufacturing. Examples are given and trends are noted. Clearly, the models and tools reported here have application potential more general than the examples given.

It is useful in what follows to note that operations research models are usually separated into two major classes: deterministic, where it can reasonably be assumed that model data is known or can be approximated, and stochastic, where uncertainty plays a central role. (See References 1 and 2 for more information.)

The primary area of deterministic operations research involves the field known as mathematical programming. Mathematical programming problems (or mathematical programs as they are often called) involve problems that require the determination of certain variables, called decision variables, in such a way that some quantitatively stated objective function, in terms of those variables, is optimized (maximized or minimized) while satisfying certain restrictions, or constraints, on algebraic combinations of the variables. Mathematical programs are called linear or nonlinear programs according to whether their objective functions and constraints are linear or nonlinear. The theory and practice of linear programming is well developed; this is less true of nonlinear programming. Other deterministic models and techniques related to mathematical programming in general include the subjects known as dynamic programming, combinatorial optimization, network flows, facility location, capacity expansion, and scheduling theory. Each of these topics has significant applications in manufacturing.

Important stochastic models include elements of probability, statistics, simulation, stochastic processes, queueing theory, inventory theory, and reliability. The range of applications for stochastic models is as rich as in the deterministic case. Moreover, once random phenomena are modeled, optimization methods can often be applied to help make decisions that explicitly consider the uncertainties involved.

Direct Factory Applications

The range of applications of operations research in manufacturing operations is broad; see, for example, References 3 to 6. This section gives representative examples that represent recent work in six main areas.

Scheduling in Assembly Line Factories. The planning of circuit pack production is a complicated task. Good planning must take into account such factors as the requirements (hard orders from customers, forecasted demand, and needed spare parts), current inventory and work in progress, and current and future supplies of components. The production capacity of the assembly line must also be considered. Furthermore, circuit packs are often subassemblies of final products and it is necessary to produce the correct distribution (or "mix") of circuit pack codes.

At times, the planning of circuit pack production may be further complicated by shortages of components. During shortages, the interactions between circuit pack codes must be considered. For example, if there is an insufficient quantity of a component that is used in several circuit pack codes, then components must be allocated among the codes. The tradeoffs between production levels of different codes depend on the requirements, existing inventory, planned shipments of the component, and the desired mix of codes. Often, if one component is scarce, then several more are likely to be scarce, further complicating the planning.

The production planning becomes even more complex if the demand occasionally exceeds the capacity of the assembly line. During slack periods (when demand is less than capacity), enough inventory to satisfy future demands must be produced, or else backordering will result. Occasionally, backorders are inevitable if demand greatly exceeds capacity, or if component shortages are insurmountable. Other times, inventory must be stored for long periods to avoid too many backorders. However, it is important not to produce too much inventory because the cost of carrying inventory can be quite high.

To assist in the planning process, especially when there are shortages of components or demand exceeds capacity, a computer analysis tool based on a linear programming model can be used.[7] This tool determines the "best" allocation of components and the quantities of each circuit pack code to be produced over a several month horizon. Best is defined in terms of minimizing circuit pack inventory and backorder costs. The tool uses a specialized version of the Karmarkar linear programming algorithm[8] developed recently at AT&T Bell Laboratories. This version is fast enough to be implemented on a microcomputer. The tool is currently being evaluated for use at AT&T's Oklahoma City Works.

Microscheduling in a Continuous Manufacturing Environment. The coordination of the stranding and sheathing operations in a copper cable manufacturing facility represents a difficult scheduling problem. Each cable reel may undergo many stages of processing, and there may be several different machine groups at any given stage. Each cable code has its own processing route and preferred machine groups. Determining the optimal schedule for jobs on a shop floor is a combinatorial optimization problem. Depending on the number of machines on which each job must be processed and the number of different routes that can be chosen for this processing, this can be a massive problem that can only be addressed efficiently by heuristics. Indeed, it is a difficult problem to even determine whether a feasible schedule (one that meets all due dates) exists.

One approach taken to produce a schedule is a variation of "bin-packing," where the machines correspond to the bins and the jobs to the objects being packed in the bins, with each job having a feasible time window resulting from the time it becomes available and its due date. By changing parameters used in the bin-packing algorithm, several different schedules can be quickly produced and compared for desirability. Such a mathematical model, solution approach, and computer analysis tool have been developed.[9] They are being tested in collaboration with AT&T's Atlanta Works. Application at additional cable manufacturing facilities is expected. Similar approaches may be possible in other continuous manufacturing operations.

Manufacturing Line Analysis Tools. Manufacturing lines are created to produce products in a suitable mix (the product mix) at a suitable rate (the production rate, units per time) in suitable time (the production interval, time per unit) at suitable cost (including the costs of labor, equipment, and work-in-process inventory). New lines must be designed and existing lines must be managed to meet these important objectives.

Operations research models, algorithms, and software can help solve these design and management problems. In particular, existing manufacturing lines and potential new ones can often be well represented as networks of queues (see Figure 1). The circles correspond to various workstations or machines in the factory and represent the queues, or service facilities, in the mathematical model. The products being produced in the factory are "customers" or jobs flowing through the network. The products typically follow a prescribed route through the network, receiving service at designated queues, and waiting in buffers, when the servers (facilities or operators) are occupied. Different codes or products in the product mix are represented by different classes of customers, each of which typically follows its own route through the network. Analysis of the queueing network models helps predict and control the product mix, production rate, production interval, inventory levels, and use of labor and equipment.

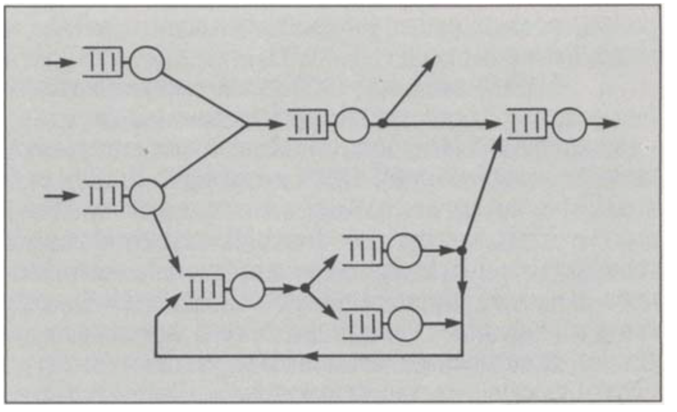

Figure 1. A network of queues.

Two operations research tools help analyze these queueing network models of manufacturing lines. The first is computer simulation, which is an experimental approach: The evolution of the system over time is generated on the computer, so that something close to the real performance of the manufacturing line over time can be observed. Simulation is being successfully applied to analyze many AT&T manufacturing lines.[10]

The second operations research tool for analyzing a queueing network model is mathematical analysis. The mathematical analysis is used to "solve" the mathematical model and describe the behavior of the model by formulas or algorithms implemented on the computer. Typically, the mathematical analysis describes only the long-term average behavior of the system. Such summary descriptions are much less detailed than the simulation output, but they often tell the essential story. Moreover, in some cases it takes much longer (weeks versus hours) to build and analyze a model by simulation. If possible, it is desirable to use the quicker mathematical analysis in the early phases of a

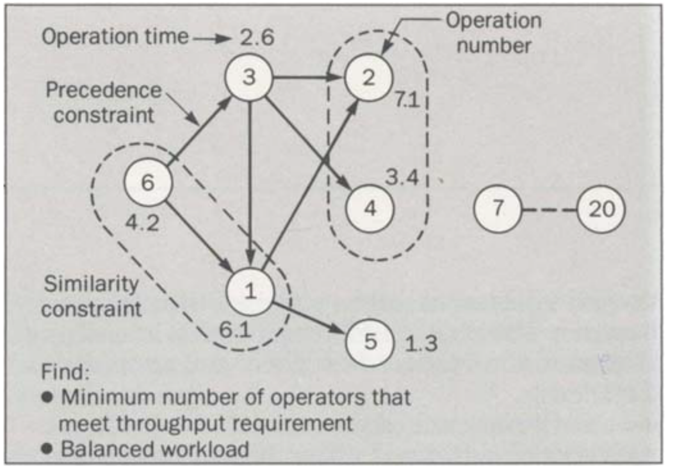

Figure 2. Optimal line balancing problem.

project and then confirm the results by a more detailed simulation study.

Mathematical analysis is an extremely important tool for analyzing queueing network models, but there is a major stumbling block. Many of the important features in models of complex systems such as manufacturing lines are still beyond known capabilities for exact mathematical analysis. Thus, successful mathematical analysis of these complex models depends largely on appropriate approximations. Hence, an important research activity is developing new approximations. In particular, work is underway to develop approximations that provide satisfactory descriptions of the complex manufacturing lines. Examples of such research are in References 11 to 14. Some of the approximations have been implemented in the Queueing Network Analyzer (QNA) software package.[15,16] The QNA package is now also being used to analyze manufacturing lines at several AT&T factories, including Allentown, Merrimack Valley, and Reading.

The first version of QNA analyzes networks of multiserver queues with the first-come, first-served discipline and no capacity constraints. An important feature is that the external arrival processes need not follow a Poisson distribution and the service-time distributions need not be exponential, as usually required for exact mathematical analysis. Treating other kinds of variability is important, for example, with manufacturing networks to describe the congestion resulting from batchy arrivals of jobs and the nearly constant service times of production. The general approach in QNA is to characterize the arrival processes by two parameters and then analyze the individual queues separately. The queues are analyzed as standard multiserver queues partially characterized by the first two moments of the interarrival-time and service-time distributions. Congestion measures for the network as a whole are obtained by assuming as an approximation that the nodes are stochastically independent given the approximate flow parameters. Enhanced versions of QNA have also been developed to account for partial yields and rework in the manufacturing process.

Manufacturing Line Design. Line balancing tools are used to plan the efficient use of production facilities in an assembly line. For example, one section of a circuit pack assembly line consists of workstations at which people or flexible robots insert components that cannot be inserted by high-speed insertion machines. Each circuit pack moves from one station to the next. In each workstation certain components are inserted.

Optimal line balancing involves grouping insertion operations, characterized by insertion times, into a minimum number of workstations so that an externally imposed throughput requirement is met. Workloads on the workstations are then balanced. Additional problem complexities include precedence constraints and similarity constraints (similar-looking but different components can easily be confused and hence, in a manual environment, should be handled at different workstations). The model is illustrated in Figure 2. Only the first six components are modeled in detail. As can be seen, component 1 can only be inserted after components 3 and 6 have been inserted. Also, components 1 and 6 must be inserted by different workstations.

The modeling involved in this optimal line balancing work comes from an area known as combinatorial optimization that refers to a kind of mathematical programming problem in which the variables are restricted to discrete values. This class of problems is mathematically difficult to solve and typically requires sophisticated heuristics even to get near-optimal solutions. In the line balancing case, efficient heuristic techniques, known as "bin-packing heuristics," were developed to quickly generate good solutions to the problem. The computer implementation is currently being used at AT&T's North Carolina Works, on the automated in-line manufacturing (AIM) line there, as well as at AT&T factories in Denver and Little Rock. The model and solution approach are general and should find applicability in still other production lines.

Automation in Manufacturing. In the production of circuit packs, some manufacturing operations, including printed wire board drilling and component insertion, are performed automatically by numerically controlled or programmable machines. The problem of optimizing the sequence of machine movements to reduce the processing times required by these operations can be modeled as a classical operations research problem known as the traveling salesman problem (TSP).[17] In the TSP, a salesman must visit each city in a given set of cities, and seeks to determine the sequence to visit these cities that will minimize his total traveling distance (Figure 3). This problem is clearly analogous to the problems of minimizing circuit board drilling and component insertion times, where the "cities" to be visited are the drilling and insertion sites.

To develop sequencing methods and software, the machine movements are first modeled mathematically to allow calculation of the travel times (distances) between each pair of "cities." While no good algorithm exists for solving large TSPs exactly, many efficient heuristic programs have been developed for finding good solutions quickly. Software tools developed at Bell Laboratories use two such heuristics: a farthest-insertion-first algorithm for constructing a reasonably good initial tour,[18] and then an interchange heuristic for improving this tour.[19]

In on-line production tests at the AT&T Columbus Works, use of the optimized sequences produced by these tools yielded 7 to 30 percent reductions in the drilling times required for printed-wire boards. To date, software tools for optimizing printed-wire board drilling, metal punching, and certain component insertion processes have been provided to the AT&T Columbus, Montgomery, and Merrimack Valley Works, and AT&T Teletype Corporation. In addition, a heuristic sequencing method for very-large-scale TSPs (with several thousand cities) has been developed and deployed at the AT&T Richmond Works.[20]

Factory Resource Allocation. Manufacturing resources planning systems[21-24] such as that found in AT&T's Business Resources Planning System (BRPS) invoke a hierarchical, disciplined approach to the planning and execution of manufacturing functions (Figure 4). In such systems, marketing and manufacturing managers jointly establish production goals for broad product families over a long planning horizon. This plan in turn drives the manufacturing master production schedule, which specifies when production of each item should begin, and how many should be produced in each period. This schedule, combined with a bill-of-materials explosion process, determines the factory resource (materials and work center capacity) requirements in each period in the planning horizon.

In a manufacturing environment as large and complex as AT&T's, it is often difficult for factory managers developing the production plan or the master production schedule to know in advance whether enough material and capacity resources will be available to support the plan. For this reason, the literature on manufacturing resources planning systems stresses the need for feedback loops when shortages are discovered. Operations research models and optimization methods can provide quantitative decision support tools to help managers generate, review, and evaluate alternative production and resource allocation plans that explicitly incorporate resource availability constraints. These tools can be useful either in support of the initial development of feasible, long-range production plans, or in helping factory managers determine the impacts of and optimal responses to short-range resource shortages as they occur.

At Bell Laboratories, a two-level hierarchical resource allocation model[25] has been developed with the AT&T North Carolina Works, and in a way consistent with the BRPS principles and procedures. The model helps managers optimize the allocation of resources across prod-

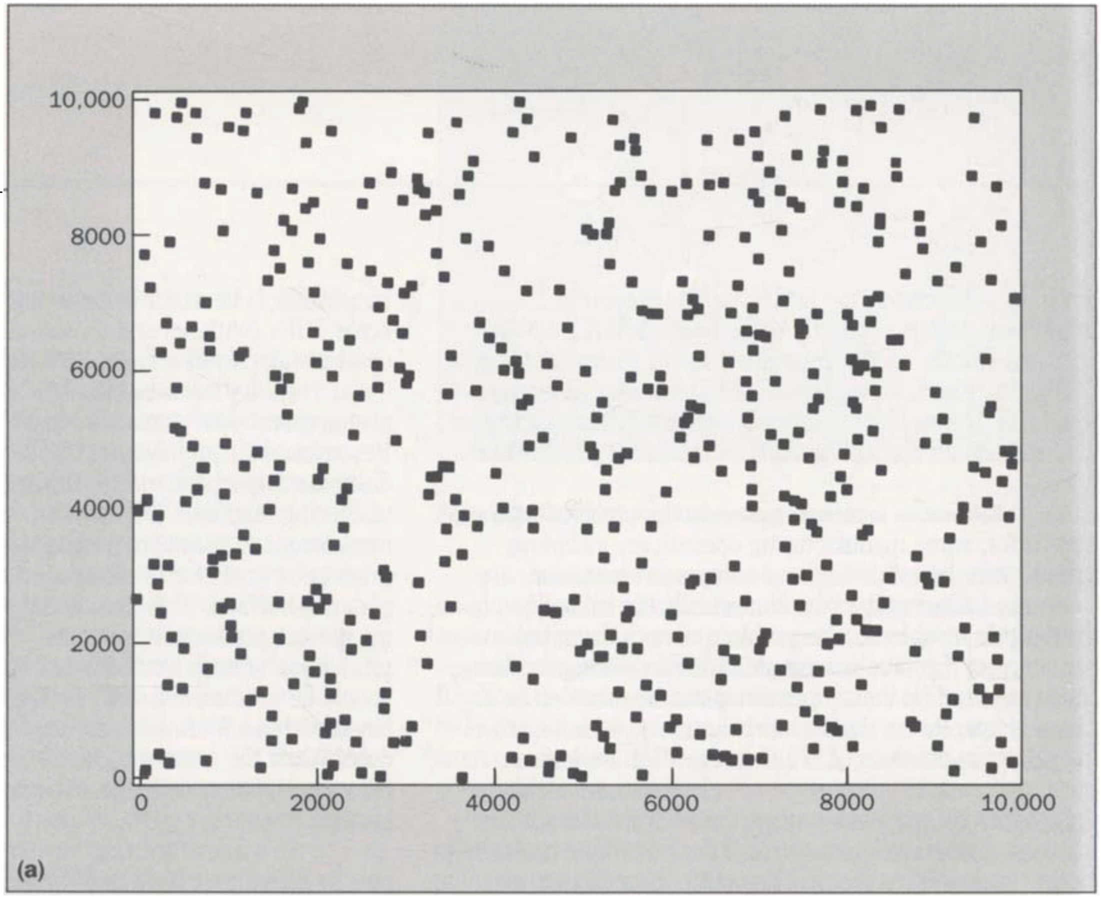

Figure 3. The traveling salesman problem: (a) 500 cities to be visited; (b) a possible solution. *(Source: D. S. Johnson)*

ucts within the factory. The objective of the optimization process is to determine scarce resource allocations that will minimize the deviations of production outputs from forecasted product demands or production targets. The model first determines optimal allocations of resources to broad product families over a long-range planning horizon. These results are then used to guide the development of detailed resource allocations to individual products in each time period. The use of weights W_{pt} for each product p in time period t allows managers to reflect product priorities; that is, managers may consider it more important for some products to meet production goals than for others. If no such preferences exist, then all the weights are simply set equal to 1.

Both the family and product level models are formulated as quadratic programs. A quadratic program is a variant of a linear program in which the objective function is a quadratic rather than a linear function of the decision variables. Use of quadratic functions penalizes large deviations from production targets more than small deviations and hence encourages some degree of balance across all products, while allowing a flexible, efficient allocation of resources. With proper formulation, large-scale quadratic programming models involving as many as several hundred product families and dozens of planning periods can be efficiently solved using state-of-the-art optimization codes such as MINOS.[26]

Manufacturing Planning and Support

Manufacturing operations require important support activities that indirectly affect the manufacturing process. Indeed they are critically involved in efficiently managing a complex manufacturing enterprise such as that at AT&T. And here, too, operations research can play a significant role.

Computer-Aided Design. Operations research methods have recently begun to be applied to problems arising in the computer-aided design (CAD) of electronic circuit

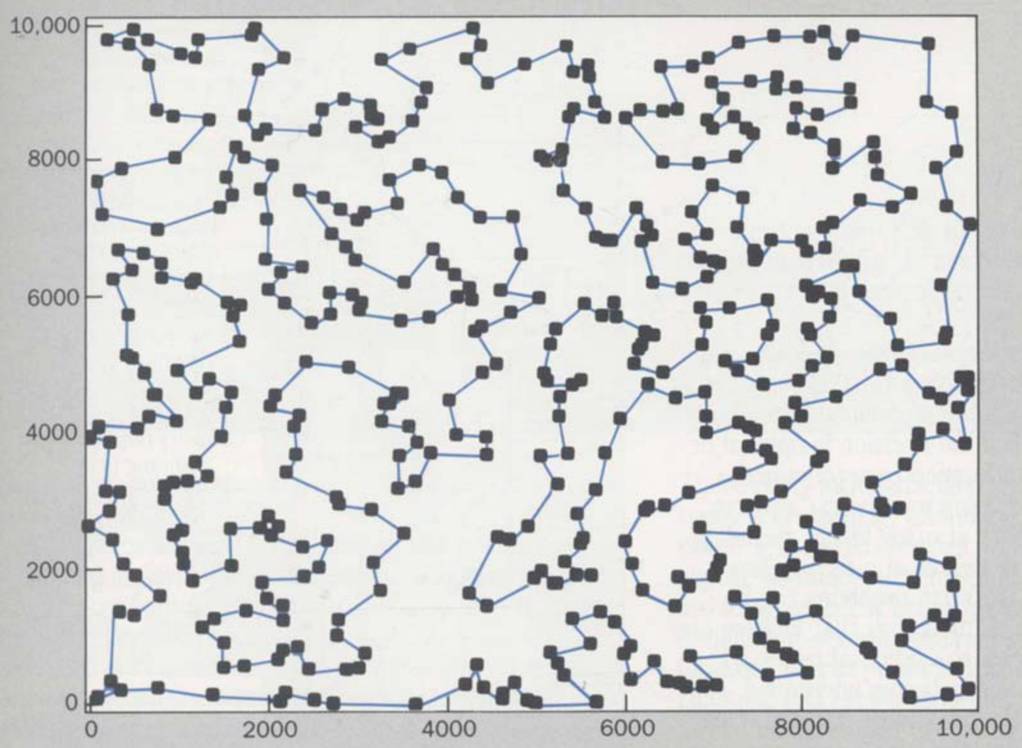

(b)

boards. Circuit board designs are becoming increasingly complex, and many use new technologies such as surface mounting of devices. These factors often make it impossible for existing CAD tools to route automatically all the interconnections required on the board. When automatic routers are blocked, several days or even weeks of manual routing efforts may be needed to complete the design of the board.

The set of possible interconnection tracks on a circuit board is naturally modeled as a rectilinear network, possibly with several layers to represent multilayer boards. Network optimization methods, including shortest path[27] and Steiner tree[28] techniques can then be applied to determine efficient routings for interconnecting each set of points as required. By applying these optimization techniques in a suitable iterative fashion, it is possible to construct routings that efficiently use the total available space on the board. This approach can avoid the pitfalls of "sequential" routers, which myopically route interconnections one at a time until any further interconnections are blocked.

Moreover, operations research methods have also been developed at AT&T Bell Laboratories to provide automatic rip-up procedures for removing (and later rerouting) some interconnections when blockages do occur. The problem of deciding which routings to rip up is modeled as a mathematical program. Efficient heuristic methods for solving this problem have been developed and implemented in prototype software; research on exact optimization methods is underway.

Tests to date of the iterative routing and rip-up methods have been very encouraging. The algorithms often achieve 100 percent complete routings automatically, and always give significantly more complete routings than do existing routers. Work has begun at Bell Laboratories on the development of an efficient coding of the algorithm for incorporation with other CAD tools. It is expected that this, and future work on the application of operations research methods to determine optimal component placements on new boards, will help to streamline the introduction of new board designs into manufacturing.

An Efficient Algorithm for Corporate Resource Allocation. The problem of allocating scarce corporate resources is an important and recurring decision in a manufacturing environment. The degree to which the decision is optimal or near-optimal can mean the difference between a firm being highly competitive and being merely an "also ran."

Allocation is necessary at many levels, including at a corporation-wide level as well as at a factory level. Typical corporate resource allocation problems can be exceedingly large. Hence, a methodology that enables one to solve large problems quickly is a practical necessity, particularly if an interactive application is envisioned.

Various objective functions can be pursued. One reasonable one consists of allocating scarce resources so as to balance (possibly weighted) deviations from production goals for all products considered. The weights represent the relative importance of products. This objective function is intuitively acceptable to production planners and avoids the "nervousness" (discontinuities) of the solutions based on a linear objective function.

A new efficient mathematical programming algorithm was developed that uses the previously described "minimax" or balanced objective function.[29] The time involved in using the algorithm is on the order of the time to determine the resource feasibility of a given production plan. Thus, it can solve very large problems extremely quickly, and can be employed when quadratic or linear programming approaches prove inadequate.

This algorithm has been coded in Fortran, and a demonstration model using it has been constructed using real data in a corporate setting which ties together multiple factory interactions. The initial reaction to the demonstration has been quite favorable.

Inventory Modelling. Inventories and the resulting turnover rates have a significant effect on a manufacturing firm's profitability. Inventories are obviously needed to facilitate smooth production operations. The key question is: how much? Appropriate inventory levels depend on the type of business (product types, production capacities, distribution network, and so forth). Hence, it is extremely important to model key inventory issues so that inventory levels can be quantified as a function of a set of parameters characterizing a specific firm.

Work is underway to better understand inventories in AT&T on a macro level. To this end, a high level model of inventory requirements for made-to-order products (for example, 5ESS™ switching machines) was developed. The model expresses required component inventories as a function of operational parameters such as

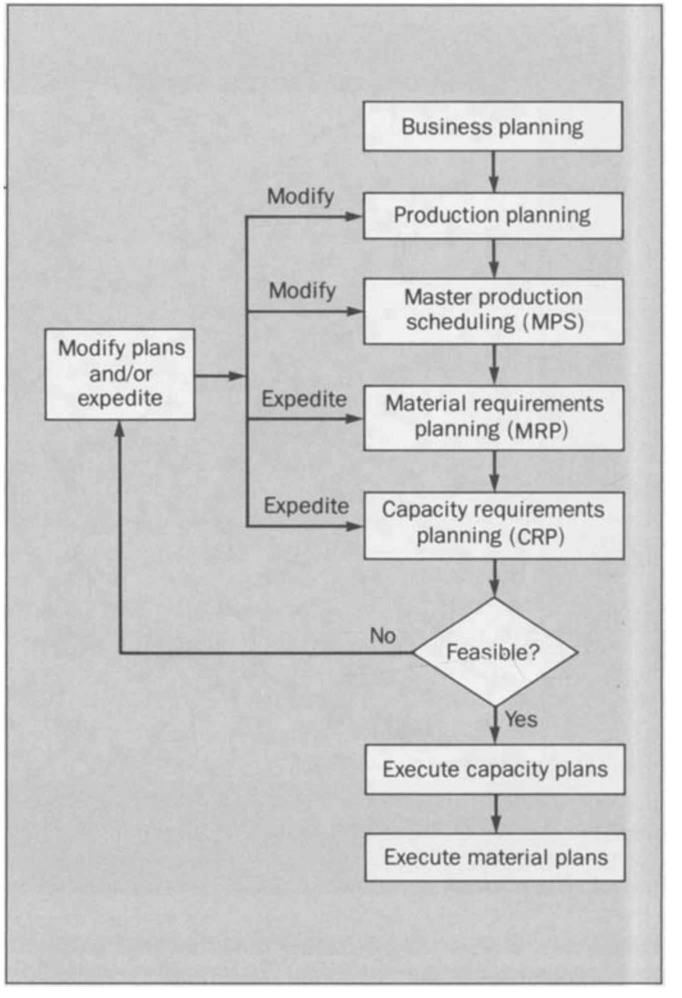

Figure 4. A view of the manufacturing resources planning process.

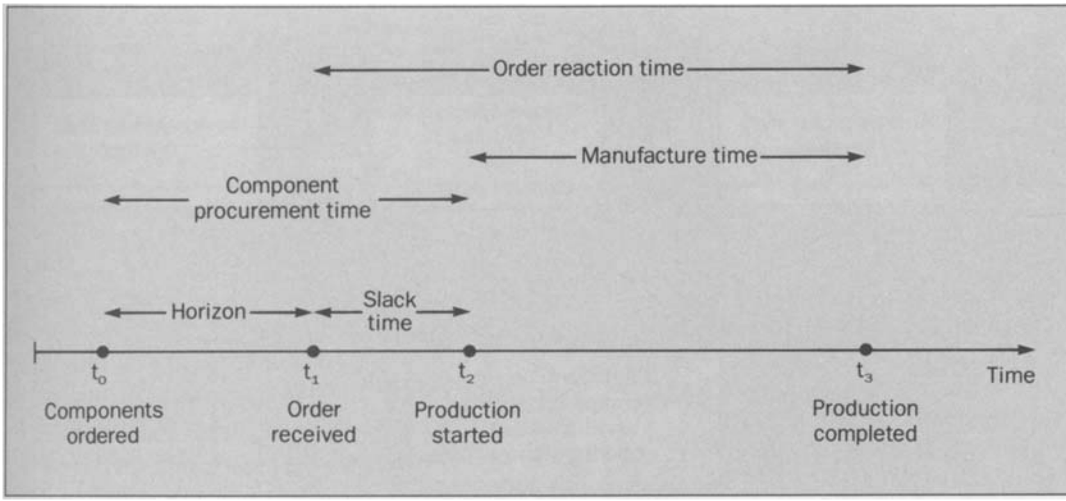

Figure 5. A production time line.

component procurement time, number and type of options offered, and so forth. Since components must be ordered before the configuration of a customized product is known, some component inventory is required; see Figure 5.

Other major factors affecting inventory are being studied (for example, interworks serial manufacturing, unreliable supply, pipeline, mismatching, and technological change), and models are being developed that quantify their effects. Once the models are completed, a "super model" will be developed that combines them under a unified umbrella. This will provide an invaluable strategic tool to analyze the effect of high-level policies on inventory requirements within the corporation.

Material Logistics Planning. The cost of establishing and operating a material logistics system is an important variable in the management of a manufacturing concern. Moreover, decisions about the number, size, and location of warehouse and distribution facilities are critically important to this cost. The Advanced Distribution and Location Model (ADLM), Figure 6, incorporates state-of-the-art mathematical programming algorithms in a software tool that helps to plan cost-effective logistics operations. It facilitates determination of the optimal locations and sizes of distribution centers (DCs) and local distribution centers (LDCs). It handles several product types characterized by different flow patterns from the sources to the demand areas, explicit capacity constraints, and concave operating cost functions.

Within ADLM, the number, size and location of DCs and LDCs are determined by various incurred costs. These costs include facility setup costs, facility closing costs, and shipping, inventory, handling, and operating costs. The model contains various features that do not appear simultaneously in standard facility location models, including nonlinear operating costs, capacity constraints, assignment of each customer to only one major stocking location, special products that are handled by only a few facilities, and establishment of LDCs where desirable. The algorithm developed to solve ADLM is highly complex.[30] The primary optimization routine determines what DCs should be opened and assigns demand areas to these DCs. Briefly stated:

- A capacitated facility location problem with nonlinear operating costs is solved to determine DC locations and demand area assignments (perhaps split among different DCs). This is done through the solution of a sequence of network optimization problems.
- Split demand area assignments are resolved using Lagrangian relaxation and branch-and-bound methods.

Additional heuristic routines determine: (1) a subset of open DCs that would handle special products; (2) LDC locations; and (3) appropriate sizes of DCs and LDCs.

The associated software tool has already been used to reconfigure parts of the AT&T distribution network.[31]

The Dynamic Distribution and Location Model (DDLM) is a large-scale mathematical programming model of the evolution of a material logistics system over a multiperiod time horizon. DDLM chooses a multiperiod schedule of openings, expansions and closings of DCs, and assigns demand areas to these DCs. The objective is to

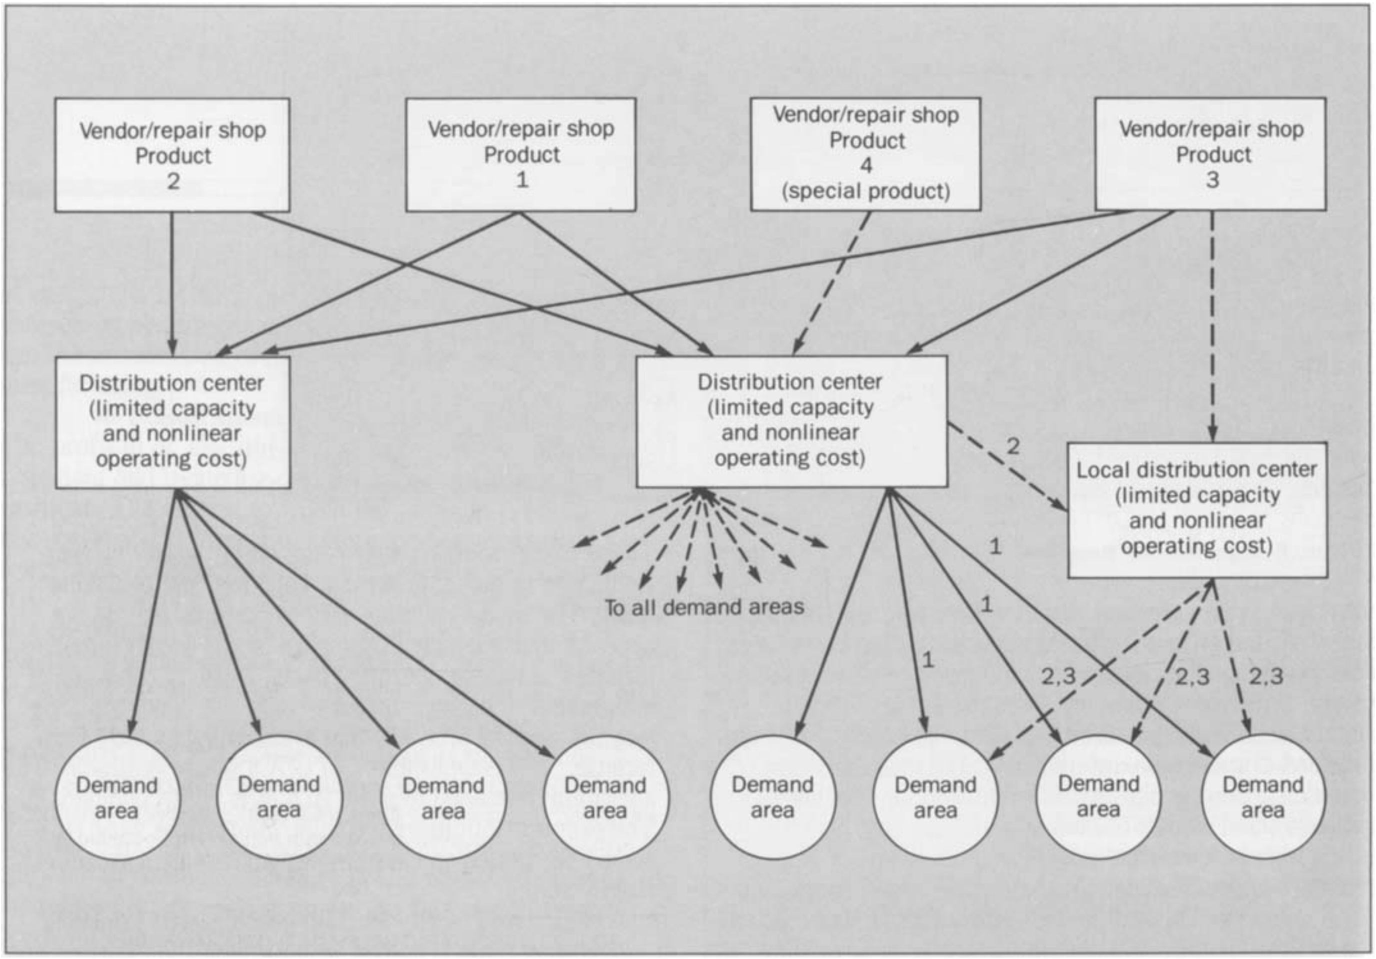

Figure 6. The Advanced Distribution and Location Model (ADLM).

minimize total discounted costs, including opening, expansion, and closing costs; shipping, inventory, handling, and operating costs; costs for reassigning demand areas; and holding costs for unused capacity.

DDLM is designed to complement ADLM. It was designed and implemented at the request of the Materials Management Services organization within AT&T Information Systems. DDLM will be used (jointly with ADLM) to enhance the distribution planning function for the entire AT&T corporation. Both ADLM and DDLM software can be used for large problems. Current implementations can handle up to 50 DC locations, 200 demand areas, and 10 product families.

Future Trends and Opportunities

As illustrated here, operations research models, algorithms, and computer tools are extremely important in the manufacturing area. Deterministic models and associated optimization techniques help to find optimal (or at least "good") solutions rather than merely feasible ones. Stochastic models help to understand random events so that they can be reasonably accounted for. Of course, once stochastic models are expressed in mathematical terms, optimal decisions can then be sought for them as well.

Over the years, new operations research methodologies have evolved as new technologies have emerged. Given the recent emphasis on factory automation, one can foresee fertile new areas for operations research in manufacturing applications. Obvious examples include flexible manufacturing systems (FMS) and robotics.[32-34] Not only will analytical tools help in designing efficient FMS and robots, but they will also help to control their daily operations.

One new effort just begun at Bell Laboratories focuses on disruptions management and control. This work addresses unplanned interruptions (such as design changes, material shortages, worker absences, machine breakdowns, and yield busts) to what should otherwise be a continuous clockwork-like manufacturing process. Whereas the standard approach is to treat these disruptions as inevitable and attempt to react optimally to them once they have occurred, the new effort involves forward-looking analytical studies to identify and quantify opportunities for streamlining the manufacturing process through elimination of the major causes of the disruptions. That is, the effort is aimed at treating the underlying disease rather than the symptoms. Insight gained from this preliminary work will provide a framework for focusing future efforts to prevent recurrent disruptions in those areas of greatest leverage.

AT&T is in the process of modernizing its information system network through such projects as Computer Integrated Manufacturing Architecture (CIMA),[35] the Business Resources Planning System (BRPS), and the Productivity Improvement System for Manufacturing (PRISM).[36] Having easy access to data and being able to exchange information among factories in real time is extremely important. As decision support tools based on operations research and other mathematical models and techniques become more commonly used, not only as stand-alone tools but as integral parts of these information systems, better decision making throughout all factories will become commonplace. This will not only facilitate better planning but will permit real-time scheduling as well.

As a final note, serious attention should be paid to new areas such as artificial intelligence and expert systems. It is expected that significant synergy will evolve between new accomplishments in these areas and new operations research endeavors. Through that synergy, one can expect more flexible and easy-to-use tools that solve larger and even more complex manufacturing problems.

Acknowledgments

The author is indebted to the members of the Operations Research department at Bell Laboratories, in particular, to those whose work is cited. Additional special thanks go to M. Farber, G. A. Kochman, E. Rosenberg, D. R. Smith, R. Steinberg, and C. S. Yu. The considerable efforts of these individuals and many others form the basis of this article.

References

1. F. S. Hillier and G. J. Lieberman, *Operations Research*, 4th ed., Holden-Day, San Francisco, 1986.
2. H. M. Wagner, *Principles of Operations Research*, 3rd ed., Prentice-Hall, Englewood Cliffs, N. J., 1975.
3. A. C. Hax and D. Candea, *Production and Inventory Management*, Prentice-Hall, Englewood Cliffs, N. J., 1984.
4. J. C. Mitchell, "Operations Research Methods in Production/ Operations Management: A Survey," AT&T Bell Laboratories, Holmdel, N.J., December 1984.
5. H. Luss, "Operations Research and Capacity Expansion Problems: A Survey," *Operations Research*, Vol. 30, No. 5, September-October 1982, pp. 907-947.
6. IBM, "Computers in Manufacturing," *IBM Journal of Research and Development*, Vol. 29, No. 4, July 1985, pp. 329-440.
7. M. S. Meketon, "A Production Scheduling Model and Algorithm for a Flexible Manufacturing Line," *Symposium on Real-Time Optimization in Automated Manufacturing Facilities*, Gaithersburg, Md., January 1986.
8. N. K. Karmarkar, "A New Polynomial-Time Algorithm for Linear Programming," *Combinatorica*, Vol. 4, No. 4, 1984, pp. 373-395. See also D. M. Gay, N. K. Karmarkar, and K. G. Ramakrishnan, "The Karmarkar Algorithm: Adding Wings to Linear Programming," *AT&T Bell Laboratories Record*, Vol. 64, No. 2, March 1986, pp. 4-10.
9. B. S. Munson, "Job Shop Scheduling Using a Bin-Packing Approach," ORSA/TIMS Joint National Meeting, Miami, October 1986.
10. D. Y. Burman et al., "Performance Analysis Techniques for IC Manufacturing Lines," *AT&T Technical Journal*, Vol. 65, No. 4, July/August 1986, pp. 46-57.
11. S. L. Albin, "Approximating a Point Process, II: Superposition Arrival Processes to Queues," *Operations Research*, Vol. 32, No.

5, September-October 1984, pp. 1133-1162.
12. D. Y. Burman and D. R. Smith, "An Asymptotic Analysis of a Queueing System with Markov-Modulated Arrivals," *Operations Research,* Vol. 34, No. 1, January-February 1986, pp. 105-119.
13. J. G. Klincewicz and W. Whitt, "On Approximations for Queues, II: Shape Constraints," *AT&T Bell Laboratories Technical Journal,* Vol. 63, No. 1, January 1984, pp. 139-161.
14. W. Whitt, "Approximations for Departure Processes and Queues in Series," *Naval Research Logistics Quarterly,* Vol. 31, No. 4, November 1984, pp. 499-521.
15. W. Whitt, "The Queueing Network Analyzer," *Bell System Technical Journal,* Vol. 62, No. 9, November 1983, pp. 2779-2815.
16. W. Whitt, "Performance of the Queueing Network Analyzer," *Bell System Technical Journal,* Vol. 62, No. 9, November 1983, pp. 2817-2843.
17. E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan, eds., *The Traveling Salesman Problem,* John Wiley and Sons, New York, 1985
18. A. Rosenkrantz, R. Stearns, and P. Lewis, II, "An Analysis of Several Heuristics for the Traveling Salesman Problem," *SIAM Journal of Computing,* Vol. 6, No. 3, September 1977, pp. 563-581.

19. S. Lin and B. W. Kernighan, "An Effective Heuristic for the Traveling Salesman Problem," *Operations Research,* Vol. 21, No. 2, March-April 1973, pp. 498-516.
20. L. W. Phillips, "Solution of Large-Scale Traveling Salesman Problems," *IE News: Operations Research,* Vol. 17, No. 3, 1982, pp. 1-3.
21. O. W. Wight, *MRP II: Unlocking America's Productivity Potential,* OWL Publications, Essex Junction, Vermont, 1981.
22. J. Orlicky, *Material Requirements Planning,* McGraw-Hill Book Company, New York, 1975.
23. S. C. Aggarwal, "MRP, JIT, OPT, FMS? Making Sense of Production Operations Systems," *Harvard Business Review,* Vol. 63, No. 5, September-October 1985, pp. 8-16.
24. F. R. Jacobs, "OPT Uncovered: Many Production Planning and Scheduling Concepts Can Be Applied with or without the Software," *Industrial Engineering,* Vol. 16, No. 10, October 1984, pp. 32-41.
25. J. C. Mitchell, "A Hierarchical Production Planning Model to Allocate Scarce Resources," AT&T Bell Laboratories, Holmdel, N. J., December 1985.
26. B. A. Murtagh and M. A. Saunders, "Large-Scale Linearly Constrained Optimization," *Mathematical Programming,* Vol. 14, No. 1, January 1978, pp. 41-72.
27. R. E. Tarjan, *Data Structures and Algorithms,* Society for Industrial and Applied Mathematics, Philadelphia, 1983.
28. R. T. Wang, "A Dual Ascent Approach for Steiner Tree Problems on a Directed Graph," *Mathematical Programming,* Vol. 28, No. 3, April 1984, pp. 271-284.
29. H. Luss and D. R. Smith, "A Minimax Mathematical Programming Algorithm," TIMS/ORSA Joint National Meeting, Los Angeles, April 1986.
30. J. G. Klincewicz, "A Large-Scale Distribution and Location Model," *AT&T Technical Journal,* Vol. 64, No. 7, September 1985, pp. 1705-1730.
31. T. Griggs, "ADLM System Puts Materials in Their Place," *Bell Labs News,* Vol. 25, No. 39, October 1985, p. 2.
32. K. E. Stecke and R. Suri, conference chairpersons, *Proceedings of the First ORSA/TIMS Special Interest Conference: Flexible Manufacturing Systems—1984,* Operations Research Society of America and The Institute of Management Sciences, Ann Arbor, Mich., August 1984.
33. *ORSA/TIMS Joint National Meeting Bulletin,* Atlanta, November 1985.
34. Richard M. Cyert, "Potential of Automation in Manufacturing," ORSA/TIMS Joint National Meeting, Atlanta, November 1985.
35. L. B. Boza, "CIMA—An Integrated Architecture for Product Realization," *AT&T Technical Journal,* Vol. 65, No. 4, July/August 1986, pp. 17-24.
36. T. Griggs, "PRISM Reflects Focus on Manufacturing," *Bell Labs News,* Vol. 26, No. 12, March 1986, p. 1.

(Manuscript received December 13, 1985)