ENHANCING THE DECISION-MAKING PROCESS FOR PAPER MILL SCHEDULERS

SESH MURTHY, RAMA AKKIRAJU, RICHARD GOODWIN, PINAR KESKINOCAK, JOHN RACHLIN, FREDERICK WU, SANTHOSH KUMARAN, AND ROBERT DAIGLE

A NEW DECISION-SUPPORT SYSTEM FOR SCHEDULING PAPER MANUFACTURING AND DISTRIBUTION CONSIDERS MULTIPLE OBJECTIVES IN PROVIDING ALTERNATIVES.

PAPER COMPANIES NEED TO MAKE EFFICIENT USE OF THEIR resources to survive in a competitive market and make a return on their investment. Two approaches to increasing efficiency are to increase the efficiency of the underlying manufacturing process and to improve throughput through better scheduling and decision support. Historically paper manufacturers have focused their energy on improving the papermaking process through the use of improved manufacturing technologies and techniques. Paper companies are now beginning to focus their attention on improving efficiency and customer service and satisfaction through better scheduling and decision support software. To increase value to the customer and maximize profits, companies need to utilize their capacity efficiently, be highly responsive to customer demand, ensure on-time delivery, and provide short lead times. This requires scheduling to minimize production and distribution costs, inventory levels, and manufacturing disruptions. The relative importance of these competing objectives varies with the state of the production environment and market conditions. The complexity of the scheduling problem is compounded by process interactions wherein the scheduling of each stage of the production process affects downstream production or shipping. However, in contrast to the manufacturing technologies, the tools that are utilized in the area of production scheduling on a daily basis are the same ones that have been used for many years. The use of advanced planning and scheduling systems is generally in the early stages of use by most paper companies.

Currently, most paper companies develop a predetermined "block" manufacturing schedule that serves as the basis for generating detailed trim solutions and deciding whether it can accept new customer orders. Typically, the block schedule is generated manually, and separate software packages are used for trim scheduling and outbound logistics scheduling. Each of these packages focuses on a single process step and attempts to create an optimized schedule based on local objectives. Since there is no interaction between applications, the complete schedule obtained by combining the sub-schedules is usually of very low quality.

To make the process more manageable, companies adopt business rules, such as only taking orders in full vehicle loads and requiring trim schedules to produce exactly the amount ordered. This simplifies scheduling but reduces the ability to respond to customer requests, lowering customer satisfaction and reducing efficiency. Additionally, most scheduling packages combine multiple objectives into a single objective function and produce a single schedule that minimizes this composite objective function *(1)*. Since the importance of the underlying objectives may not be precisely known, this approach rarely generates a satisfactory solution *(2, 3)*. Furthermore, presenting schedulers with a single take-it-or-leave-it choice does not illustrate the tradeoffs between competing objectives needed to make an informed decision. To explore the solution space, schedulers must repeatedly modify the objective function and rerun the scheduling application to see other possible schedules *(4)*.

Given a set of firm orders and forecasted demand, the schedulers must decide how to:

- Allocate orders to mills and machines
- Sequence the orders on each machine and form a sequence of runs
- Trim the reels of paper to create rolls the right size for each order
- Load the rolls into vehicles for shipping to the customers and distribution centers.

Their goals are to:

- Maximize profit by minimizing trim loss, transportation costs and inventory carrying costs
- Maximize on-time delivery by minimizing both early and late delivery
- Minimize manufacturing disruptions
- Maximize product quality by minimizing customer preference violations.

Most problems in schedule generation are variations of well-known NP-complete problems. For example, run formation and sequencing is an extension of the multi-machine job sequencing problem with sequence dependent setup times, trimming is an extension of the cutting stock problem and loading falls into the domain of the multiple knapsack problem.

To assist the schedulers in this complex decision-making process, IBM has developed a decision-support system for mill scheduling. Based on a new multi-objective decision-support approach, this system produces multiple good alternative schedules that prescribe the operations for all the major steps in the manufacture and distribution of paper. The schedules are created by algorithms that take into account the interactions between processes and focus on enterprise-wide objectives, such as maximizing profitability, on-time delivery, product quality, and minimizing manufacturing disruptions. The algorithms use advanced optimization techniques such as linear programming, integer programming with and without randomized rounding, network flow and heuristic methods.

Each enterprise schedule contains an allocation of orders to machines, a sequencing of orders on the machines, a trim schedule for each machine and a loading schedule. The schedules are evaluated in terms of multiple objectives and the best schedules are presented to the scheduler. By examining these schedules and comparing the alternatives, schedulers gain an understanding of the tradeoffs that must be made and select a solution that strikes a suitable balance among the sometimes conflicting objectives. The system also allows schedulers to pose "what-if" questions and to work cooperatively with the software to improve schedule quality by injecting expertise and refining the alternatives. In addition, the system treats business rules as objectives rather than constraints. Some of the schedules presented by the system may violate the rules in order to substantially improve customer satisfaction or manufacturing efficiency. These schedules can be shared with schedulers responsible for different stages of production as well as with customer service representatives and serve as a basis for negotiating an exception to a business rule.

COOPERATIVE DECISION SUPPORT

Our approach to decision support is to augment the skills of the schedulers by offering an intelligent assistant capable of providing the information needed to make a good decision. In our decision-support approach, the scheduler plays two roles: participating in solution generation and deciding on the final schedule.

A good decision-support system should relieve the schedulers of mundane computational tasks and provide them with high quality solutions to allow them to focus on decision-making. The system should apply advanced solution methods to search systematically a larger space of possible solutions than a scheduler could consider. In addition, it should track and identify all constraint violations and perform extensive bookkeeping operations.

By combining the advanced solution techniques used by the system with the domain expertise of the schedulers, the system facilitates cooperation to produce a set of high quality alternative solutions for consideration. A premise of our decision-support approach is that the scheduler's expertise cannot be fully captured in software. As stated by Wagner *(5)*, "…an operations research model is never sufficient unto itself: it cannot become entirely independent of judgment supplied by knowledgeable managers… The question is not when to apply science and when to rely on intuition, but rather how to combine the two effectively." The system allows the scheduler to create new schedules or modify existing schedules at any level of detail in order to provide expert guidance. Unlike some earlier approaches *(4)* the iterations in the solution generation process are not only based on changing the objective weights, but also on more detailed input from the scheduler about the solutions. The schedules generated by the scheduler, like those generated by the system, serve as a basis for further search and refinement by the system. Again, this contrasts with earlier systems, where once a scheduler edits a solution, the system cannot further improve it. There are cooperative expert systems, such as Scheiker *(6)*, where the cooperation between the scheduler and the system is similar to our approach; however, expert sys-

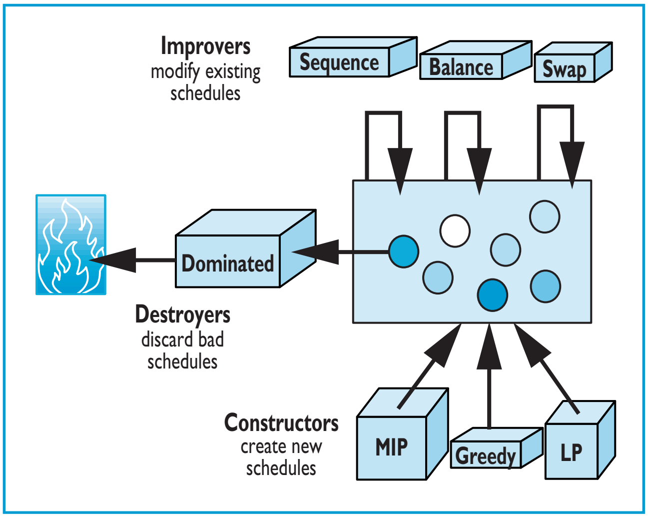

1. A-Team architecture

tem rules are typically far less capable of making useful improvements to scheduler generated solutions than the direct improvement methods we employ.

In the end, the final decision is in the hands of the scheduler who is responsible for the quality of the resulting schedule. We understand that a software system that offered a single take-it-or-leave-it solution for a complex multi-objective problem would not fulfill the information needs of the schedulers. It is difficult, if not impossible, for a software system to generate the "best" solution for a multi-objective problem, since the relative importance of the objectives is usually not known *a priori* and may change depending on the changing conditions *(3, 4, 7)*. Furthermore, a single solution does not provide any insight to the scheduler about the tradeoffs between the objectives and about how the solution compares to the other possible "good" solutions. On the other hand, a system that overwhelmed a scheduler with thousands of solutions would also be of little use. In order to present a reasonably sized set of alternatives that illustrate real tradeoffs, we evaluate each alternative in business terms with respect to key objectives and present the set of non-dominated solutions to the scheduler. A solution is called non-dominated with respect to a given set of solutions, if none of the other solutions in the set is better in every objective. The objectives used in the evaluation are profitability (including trim efficiency and transportation cost), customer satisfaction (measured by tardiness and other customer satisfaction measures, such as deviations from ordered quantities) and manufacturing disruptions (including setup time and number of short runs). These objectives represent the goals of the key constituencies of the enterprise. As the ultimate decision maker, the scheduler examines the tradeoffs offered by the system and selects the "best" schedule.

ASYNCHRONOUS TEAM ARCHITECTURE

To implement our approach to decision support, we employ Asynchronous Team (A-Team) architecture. This architecture supports our aim of providing an intelligent assistant that can work cooperatively with schedulers to produce multiple solutions that illustrate tradeoffs *(8)*. In addition, the A-Team architecture allows us to combine multiple problem-solving approaches that increase the likelihood that good solutions can be found in a reasonable amount of time.

An asynchronous team is a team of software agents that cooperate to solve a problem by dynamically evolving a shared population of solutions *(9, 10)*. The agents work asynchronously and embody their own methods for deciding when to work, which solutions to work on, and how to generate new or improved solutions. Agents do not explicitly coordinate their activities with one another. Cooperation is achieved by allowing an agent to modify a solution in the population created by other agents. The solution alternatives are usually the result of contributions made by many agents *(11)*.

As shown in **Fig. 1,** there are three types of agents in an A-team: constructors, improvers, and destroyers. Constructors take only the problem definition and create new solutions. Deterministic constructors attempt to run only once, while constructors that use randomized algorithms may run many times to add a variety of solutions to the population. Improvers attempt to improve upon the current set of solutions in the population by modifying or combining existing solutions. Destroyers attempt to keep the population size in check and focus the efforts of the improvers by removing bad solutions from the population.

A scheduler can interact with the A-team, taking on the role of an agent by adding new solutions or by removing or improving existing solutions. Within the population, each solution is evaluated with respect to multiple objectives. Typically, solutions to scheduling problems are evaluated in terms of cost, timeliness, product quality, and manufacturing disruptions. As a group, the set of agents in an A-Team evolve the population of solutions to provide a set of good alternatives that illustrates important tradeoffs among the competing objectives.

For difficult optimization problems, any single algorithm is not likely to produce high quality solutions in a reasonable time for all problem instances. The combination of agents in an A-Team constitute a portfolio of algorithms which is more robust than any single algorithm *(10, 12)*. Beyond the benefits of using multiple algorithms independently, the A-Team architecture allows the agents to cooperate synergistically and leads to higher quality and more diverse solutions. Agents may focus on optimizing different objectives and explore different parts of the search space for feasible solutions. The A-Team architecture further enhances this synergistic cooperation by enabling the scheduler to function as an agent.

Agents in the A-Team

The following teams of agents constitute the A-Team:

- Agents that optimize the allocation of orders to machines across multiple mills in different geographical locations to balance capacity utilization, maximize manufacturing efficiency, and minimize shipping costs.
- Agents that form and sequence batches of similar types of paper on each paper machine to minimize waste and disruptions resulting from sequence-dependent setups while ensuring that product demand is fulfilled on time.
- Agents that generate detailed trim sheets, considering objectives such as minimizing waste and minimizing deviations from preferred order quantities.
- Agents that generate detailed load plans, which include schedules for shipping products to customers across multiple available modes of transportation to maximize on-time delivery and vehicle utilization while minimizing freight costs and satisfying vehicle availability constraints.

There are also several "improver" agents that modify the schedules generated by the constructor agents and create new schedules. The activities of these improver agents include moving orders between multiple runs, machines, and mills to improve load balance, transportation cost, and timeliness; resequencing orders on a given machine to ensure smoother transitions; and creating alternative groupings for orders to improve trim solutions and load plans. The tight integration of these multiple agents enables the system to perform strategic long-term planning in a globally optimal fashion, considering all significant upstream and downstream consequences. A detailed technical discussion on the agents can be found in Keskinocak et al. *(13).*

COOPERATION BETWEEN SCHEDULER AND SOFTWARE

To achieve improved solution quality and decision-making, the system cooperates with the scheduler in solution generation. The scheduler can suggest alternatives and modify solutions, as shown in **Fig. 2**. The system can take guidance from the scheduler, offer solution alternatives, and point out constraint violations. Schedulers can submit manually created or modified solutions back to the scheduling system for evaluation and comparison with other solutions, as well as for further improvement by the system. By working with the system to explore the solutions space, the schedulers gain insight that allows them to be more confident of their final decision.

The graphical user interface (GUI), developed specifically for paper mill scheduling, allows the scheduler and the system to communicate effectively. The interface facilitates communication by providing multiple views of the data at different levels of granularity and by using the scheduler's own graphical language as a basis for direct manipulation of a schedule. The views provided by the interface include the following:

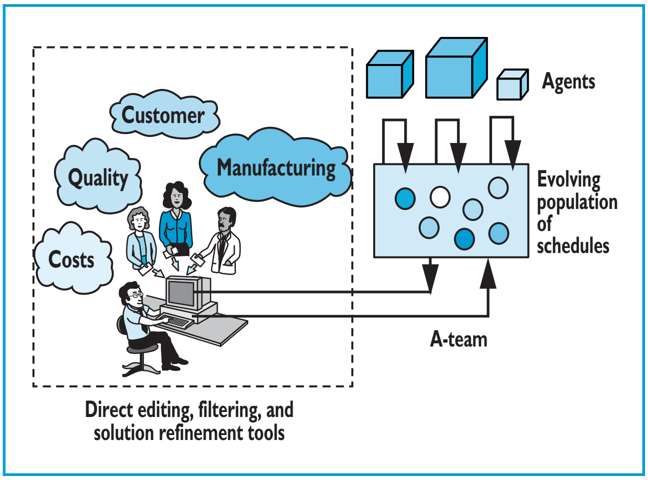

2. A-Teams provide the basis for building enhanced decision-support systems.

- Summary of results: presents the objective function evaluations for each enterprise schedule in the non-dominated set.
- Detailed view of each enterprise schedule: shows the duration of each run and the run sequence for all the machines in the schedule.
- Detailed view of each machine schedule in a given enterprise schedule: shows the duration and product type of each run and the run sequence on the machine.
- Detailed view of a run in a given machine schedule: shows the orders in the run, evaluations of alternative trim sheets and a selected trim sheet.
- Detailed view of a loading schedule: shows how orders are loaded onto vehicles.

Figure 3 shows the screen used for detailed trim scheduling. The grid on the upper right shows the specifications of the orders that need to be trimmed in this run. It includes roll size, diameter, and number of rolls for each order. The grid on the lower right is the trim sheet, which gives a graphical representation of a solution. Selecting an order in the top grid highlights its position in the trim sheet, as shown with order 3964-1. Finally, the grid on the left gives a summary view of the set of alternative trim solutions and their evaluations. The columns in the solution summary grid can be rearranged to show the subset of evaluations that the scheduler is interested in or can be scrolled to show the complete set.

Consider a typical scenario where a scheduler is trying to create a trim sheet for a production run and initially has the system generate a set of candidate solutions. Such is the case in Fig. 3, where the system has produced

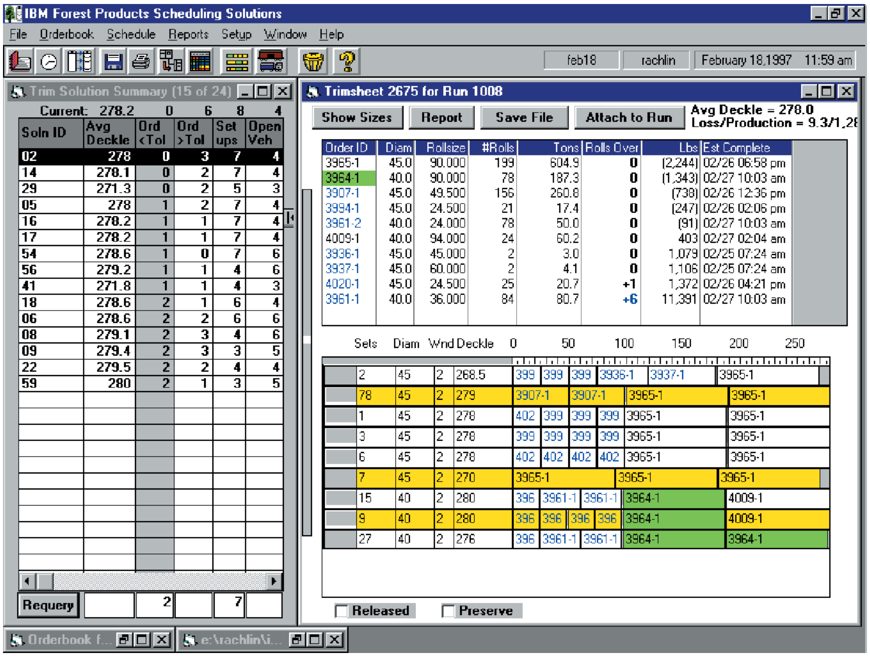

3. Screens from the graphical user-interface. The GUI enables effective cooperation between the scheduler and the agents in the A-Team.

multiple solutions and presented a set of 24 non-dominated solutions. Using the criteria boxes at the bottom of the grid, the scheduler has narrowed the set to view only solutions with no more than seven setups and no more than two orders under tolerance. Clicking on a column, the scheduler can sort the solutions by a particular evaluation. In this case, the scheduler has sorted by "orders under tolerance," which is the number of orders for which the quantity produced is under the minimum order quantity. Solutions that don't underrun orders are of particular interest if the mill has more capacity than it has orders for. Double clicking on a solution shows its details on the right-hand side and allows the scheduler to determine which orders are being overrun and by how much. In this case, orders 4020-1 and 3961-1 are being overrun by 1 and 6 rolls respectively. To improve the solutions and explore alternatives, the scheduler can directly manipulate the trim sheet by dragging rolls between patterns, reordering, or adding or deleting patterns. Once the scheduler finalizes a trim sheet, it can be released for production, using the check box at the bottom of the screen, and attached to the production run as its current trim sheet using the button at the top of the screen.

BUSINESS BENEFITS

The IBM mill scheduling system has been successfully deployed at several paper mills in North America. The system significantly improved the scheduling and the decision-making process of paper companies, giving them substantial monetary returns. One paper company reports a reduction in trim loss by 6 tons per day and a 6% reduction in freight costs *(14, 15)*. The benefits of the IBM scheduling system are an increase in manufacturing yield, a decrease in transportation costs, and improvement in customer service.

Improved manufacturing yield

Increases in the manufacturing yield are a result of scheduling production to make most efficient use of the installed manufacturing capacity. Yield increases can come from several sources, including improved trimming, smooth transitions, and fewer bottlenecks.

Improved trimming. It is important that the user can choose from multiple "candidate" trim sheets that examine the different tradeoffs between trim objectives prior to making the trim decision. This advanced optimization, combined with a graphical user interface which allows the user to see and adjust the trim sheet and its related patterns, provides the ability to improve trim yield. Improved trim also can come from the scheduling system's reacting to unplanned manufacturing disruptions, such as a linear defect caused by a hole in the wire or some other problem on the paper machine. In such a case, the location of the reel defect can be communicated to the trim package, along with the up-to-date status of roll production. The remaining production can be dynamically re-trimmed, producing a new trim sheet with only the rolls that have not been made.

"Smooth" paper machine transitions. The scheduling system is aware of the impact of various transitions on the paper machine. When generating the production schedule, the system will work to minimize transitions which produce excessive amounts of broke. If a paper company has a policy of selling the higher quality transition paper, the system will work to ensure that transitions between runs are smooth, and at worst, produce sellable transition paper.

Avoid downstream bottlenecks. In certain paper mills the production schedule can produce downstream bottlenecks that can affect yield. For example, if an excessive number of rolls are produced over a short time, there may be a backup at the winder. In the extreme, reel spools may become unavailable, causing the paper machine to produce paper directly into the broke pit. The system works to minimize this scenario by spreading out the production of narrow rolls over the duration of a run.

Improved transportation planning

Transportation optimization is considered from the very beginning of the scheduling process. During the formation and sequencing of the manufacturing runs, the scheduling system seeks to organize production so that products that are shipping to the same (or similar/drop shipment) locations are produced in the same time frame. The system considers such issues as delivery date, trim optimization, inventory requirements and availability, manufacturing smoothness objectives, and potential downstream bottlenecks.

Improved customer service

The system can be utilized as an "Available to Promise" and "Capacity to Promise" system. We have divided this "promise" capability into three categories.

1. Simple ATP (Available to Promise). The customer service representative (CSR) checks for available inventory. This inventory may be exactly what the customer desires or it may need to be re-wound in order to meet the customer's requirements. In either case, the answer to the inquiry can be made immediately.

2. CTP (Capacity to Promise). Production "blocks" have been created which reserve space on the paper machines for the production of specific paper types. The CSR will visually check the predetermined block schedule and determine if space is available.

3. Advanced CTP (Capacity to Promise). The CSR has determined that there is neither available inventory nor available room in the block schedule. The action at this point is to enter the customer inquiry into the scheduling system as a what if and determine whether the order can be accepted.

The scheduling system enables customer inquiries to be answered quickly and accurately. Additionally, the "pushing out" of customer orders to accommodate unanticipated or late customer orders can be understood.

One of the major considerations for the scheduling system is to keep the delivery promises that have been made. This applies not only to making the product early enough to be shipped on time (on the least expensive transportation mode, if possible), but also to make the product in the right weight, length, diameter, width, core type, and paper grade/basis weight.

CURRENT DEVELOPMENTS

We have been working to extend the capabilities of this system in several key areas. First, IBM has developed this system so that it can be efficiently configured both for the initial installation at a customer's location and, over time, as the manufacturing and transportation environment changes. This capability has reduced the time necessary to implement these products by approximately half and consequently has reduced the price of implementation and project risk. IBM is working on the inclusion of finishing scheduling capability to address the requirements of paper manufacturers who convert the rolls into other value-added paper products. **TJ**

Murthy, Akkiraju, Goodwin, Keskinocak, Rachlin, and Wu are researchers with IBM T.J. Watson Research Center, Route 134, Yorktown Heights, NY 10598. Kumaran is senior development manager with IBM Supply Chain Optimization Development, 3200 Windy Hill Rd., Atlanta, GA 30339-5610. Daigle is sales manager, supply chain optimization, with IBM, 3 Canal Plaza, Portland, ME.

LITERATURE CITED

1. Pickard, G.C., *Pulp & Paper* 71(9): 125(1997).
2. Goodwin, R.T., Rachlin, J., Murthy, S., et al., "Interactive decision support: Advantages of an incomplete utility model," AAAI Spring Symposium in Interactive and Mixed Initiative Decision-Theoretic Systems, 1998.
3. Steuer, R.E., *Multiple Criteria Optimization: Theory, Computation, and Application,* Robert E. Krieger Publishing, Malabar FL 1989.
4. Yu, P.L., *Multiple-Criteria Decision Making,* New York: Plenum Press, 1985.
5. Wagner, H.M., *Principles of Operations Research with Applications to Managerial Decisions,* Prentice-Hall, Englewood Cliffs, N.J., 1969.
6. Numao, M., and Morishita, S., "Cooperative scheduling and its application to steelmaking processes," *IEEE Trans. on Industrial Electronics* 38(2): 150 (1991).
7. Bennett, D.P., "Multicriteria Optimization," 1998, http://www.ieor.berkeley.edu/~bennett/multicrit.html.
8. Murthy S., Rachlin, J., Akkiraju, R., et al., "Agent-Based Cooperative Scheduling." In "Constraints & Agents," Technical Report WS 1997-97-05. AAAI Press, Menlo Park, CA, 1997.
9. Talukdar, S.N., Souza, P. de, and Murthy, S., "Organizations for computer-based agents," *Engineering Intelligent Systems,* 1(2): 1993.
10. Murthy, S., *Synergy in Cooperating Agents: Designing Manipulators from Task Specifications*, Ph.D. dissertation, Carnegie Mellon University, 1992.
11. Salman, F., Kalagnanam, J., and Murthy, S., "Cooperative Strategies for Solving the Bicriteria Sparse Multiple Knapsack Problem," IBM Research Report RC 21059 (94164), 1997.
12. Lee, H.S., Murthy, S., Haider, S.W., and Morse, D., "Primary production scheduling at steel-making industries," *IBM Journal of Research and Development,* 1996.
13. Keskinocak, P., Akkiraju, R., Goodwin, et al., "Scheduling solutions for the paper industry," IBM research report RC21297(94999)1OCT1998.
14. Hoffman, T., "A. I. Based Software Models Help Cut Production Cost," *Computer World,* September 2, 1996. http://www.computerworld.com/idx_usea.htm.
15. Shaw, M., "Madison streamlines business processes with integrated information system," *Pulp and Paper* 72(5): 1998.