Managing Increasing Product Variety at Integrated Steel Mills

Brian Denton • Diwakar Gupta • Keith Jawahir
IBM Microelectronics Division, 1000 River Road, Essex Junction, Vermont 05452
Department of Mechanical Engineering, University of Minnesota, Minneapolis, Minnesota 55455
Dofasco Inc., PO Box 2460, Hamilton, Ontario, Canada L8N 3J5
bdenton@us.ibm.com • guptad@me.umn.edu • keith_jawahir@dofasco.ca

This paper was refereed.

Intense market competition in recent years has made it increasingly important for integrated steel mills (ISMs) to differentiate themselves from competitors based on customer service, two key attributes of which are the duration and the reliability of order-fulfillment time. To improve responsiveness, some ISMs are shifting from a pure make-to-order system toward a hybrid make-to-stock, make-to-order system. They can then match certain customer orders to existing semifinished inventory, thereby reducing the time it takes to fill those orders. However, choosing which semifinished products to make to stock and how to manage their inventory are difficult problems. We developed an optimization model that one ISM implemented as a decision-support tool for choosing the designs of made-for-stock (MFS) slabs. Use of the model has reduced the number of MFS slab designs and increased the proportion of orders covered by those designs.
(*Industries: mining, metals. Inventory production: applications.*)

In North America, more than 100 million tons of steel are produced annually with an estimated value of over 50 billion dollars. Steel is an essential raw material for buildings, automobiles, household appliances, and many other consumer products. For many countries, the steel industry is vital to their global economic competitiveness. It is also a mature industry, often the quintessential example of the old economy. Yet changes in production technology in recent years have lowered the barrier to market entry and intensified competition. For example, minimills use newer electric arc furnace (EAF) technology to process scrap steel. A typical minimill consists of a scrap storage area, an EAF, and a continuous casting machine. Minimills produce between 300,000 and one million tons of steel annually and have capital investments measured in tens of millions of dollars. Integrated steel manufacturers (ISMs), on the other hand, carry out all of the processes necessary to convert iron ore into finished products. They have dozens of semifabrication processes, typically produce three to four million tons of steel annually, and have several billion dollars in capital investment.

Minimills are cost-efficient but restricted in the variety of steel grades they can produce. Nevertheless, they have generated unprecedented competition in the market for plain carbon steels. In response to this competitive pressure, some ISMs that have the technology to produce exotic grades and to customize finishing operations have positioned themselves in the high-end markets for exotic or custom-finished steel products. However, customers in these markets demand unique products and deliveries synchronized with their own production processes. Thus, ISMs are under pressure to increase the variety of products they produce and to improve their responsiveness to market demand. Even when their product portfolios have not grown in size, the composition of the portfolios is turning over much more rapidly than in the past. For example, more than 50 percent of the items in one steel manufacturer's portfolio, which contains thousands of unique end products, were introduced in the last 10 years. While in certain industries (for example, semiconductor manufacturing) this turnover would not be considered high, for steel manufacturing it is a sharp increase over historical trends. We also have anecdotal evidence that products requiring more finishing operations or exotic grades deliver greater contribution margins due to the lack of significant competition in these markets.

Managing variety has become the key to profitability for many ISMs. Whereas product proliferation is a common problem facing many industries, it poses a particularly difficult challenge for ISMs that have long operated in the make-to-order (MTO) production mode. Their production processes are designed to make steel in high volumes to minimize setup costs.

Managing variety has become the key to profitability for many ISMs.

Thus, invariably, order fulfillment times are long and range from 10 to 15 weeks. However, markets in which ISMs have greater price latitude demand custom products with shorter and reliable delivery lead times, in the range of five to six weeks. These requirements are not consistent with the assumptions of high-volume production with infrequent changeovers, upon which ISM production processes were built. As a result, where management intervention has been slow, the result of increased product variety has been capacity shortages as well as exploding inventory of semifinished and finished goods.

ISM managers see strategic inventory management as a challenge as well as an opportunity to improve operations. Strategically placed inventories of the right semifinished products in the right quantities can be used to achieve shorter, reliable deliveries while at the same time preserving production efficiencies. In effect, this changes the pure MTO architecture into a hybrid make-to-stock (MTS)–MTO architecture in which a portion of the finished products are made from existing stock of semifinished products. However, deciding which products to keep in stock and how to manage their inventories is not easy. These decisions are complicated by capacity, yield, demand uncertainty, process- and efficiency-related constraints, and the fact that the production process allows for an infinite range of semifinished products.

The Steel-Making Process

Steel making is a few-to-many industry. It uses a few raw materials to produce a variety of finished products. Product differentiation increases as raw material proceeds on its journey toward finished product. ISMs produce a variety of finished products, most commonly in the form of *flat rolled steel coils*, or *band* for short. Production consists of two basic stages: primary production in which raw materials (iron ore, coke, and limestone) are converted into band and finishing operations in which surface and structural modifications are made to the band to achieve customer specifications on an order (for example, tin planting, chromium-coating, and painting).

A typical ISM has a plant with the following primary production operations: coke ovens, a blast furnace, melt shop, ladle metallurgy, a continuous caster, and a hot strip mill (Figure 1). The first step in steel production is iron making. This process involves the separation of iron from iron ore made possible by a series of exothermic chemical reactions in a blast furnace. Next, the liquid iron, together with additional scrap steel, catalysts, and purifying fluxes are reduced in an oxygen furnace and transferred to a ladle. The ladle is then transferred to ladle metallurgy and vacuum degasing. At this stage, various alloying elements may be added to the ladle to modify the chemistry, purification processes are carried out, and additional processing is done to ensure a homogeneous chemistry throughout the ladle.

A batch of liquid steel, called a *heat*, typically varies in size between 100 and 300 tons and is transported between operations in a refractory lined container called a *ladle*. The grade of steel in a given batch is based on its chemical composition and the grade determines the physical properties of the eventual finished product. For example, grade often determines the ductility, tensile strength, and surface quality of the

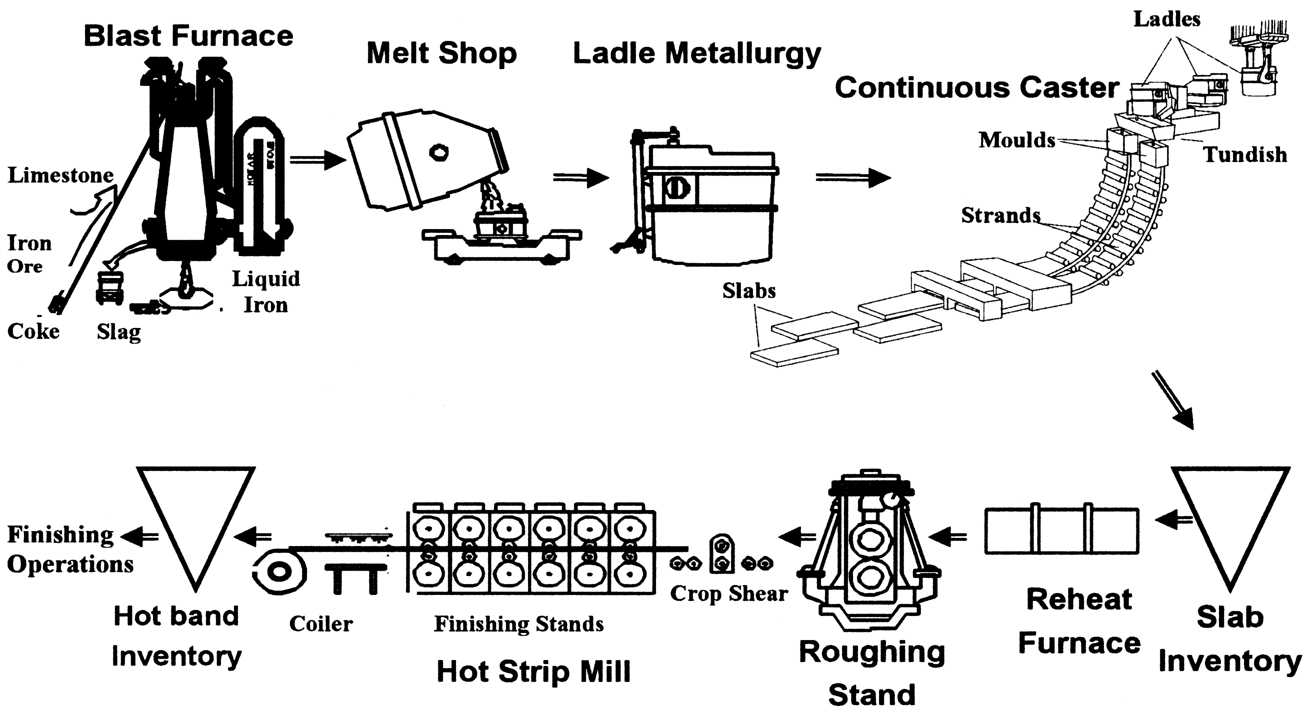

Figure 1: This diagram shows primary operations at a typical ISM for manufacturing steel coils from iron ore. Iron is separated from impurities in the ore, further purified, mixed with alloying elements, and then cast into slabs. Slab inventory is maintained at this stage. Next, slabs are reheated and hot-rolled into flat strips called coils. Coils are also usually stored in inventory prior to finishing operations. Finishing operations (not shown) consist of cold rolling and one of several different treatments, such as galvanizing, tinning, and painting.

product. Specifying the grade is the first step in customizing the finished product. After the chemistry requirements have been met, the next step is *casting*, which is a continuous process that transforms the steel from a liquid to solid form. The resulting slabs commonly have a fixed thickness but adjustable width. After the slab leaves the caster, it is cut with torches to a desired length, which is the parameter that controls the weight of the slab. Planners choose the width and weight of the slab to suit the dimensions of the finished coil ordered by the customer.

Theoretically, slabs can be cast in any width and cut to any length within the process capabilities. At one ISM slab widths ranged from about 30 to 65 inches, and slab lengths ranged from about 20 to 35 feet. However, large and rapid width changes to the mold during casting are expensive because they result in tapered slabs that have limited applicability to customer orders. Also, ISMs prefer to cast wide rather than narrow widths because wider slabs have higher throughput at the caster. Within order specifications, ISMs cut slabs as close to the specified weight as possible to avoid excess cropping downstream, which lowers slab weights and results in a revenue loss (because the price of the finished product is determined by its weight). Lower weights may also increase the number of pieces to be handled and processed at the hot mill if several smaller slabs are used to fill an order. The controllable attributes of a slab up to this point are the grade, width, weight, internal quality, and surface quality. It is typically not feasible to rework and correct deficiencies if any of these attributes are out of range with respect to the customer order or to the allowable hot-mill tolerances.

After casting, the slabs may either be labeled and sent to a slab-storage area or taken directly to the hot mill for processing. The *hot mill* is a flow line in which a slab is heated in a furnace to the desired temperature, moved on a conveyor through a system of rollers that are used to draw it out into a sheet, and sometimes reduced in width by roughing (applying pressure to its edges as it is rolled). The amount of width reduction that is possible depends on the metallurgical and process factors, and the internal and external quality specifications of the finished product. The steel sheet is subsequently spun into a coil (also referred to as *hot band*). It is then labeled, and sent to a coil-storage area where it cools and waits in inventory for further processing.

The three broad categories of inventory are slabs, hot band, and finished items. The slab stage is between the continuous caster and the hot mill, and the hot-band stage is just after the hot mill (Figure 1). Finished-goods inventory is further differentiated from hot band via one or more finishing operations. Slabs are the least differentiated and finished goods are fully differentiated products. In fact, the increase in the number of different types of items is an order of magnitude greater at the finishing stages than at the slab and hot-band stages. A typical ISM supplies thousands of unique finished products in response to tens of thousands of unique customer orders each year. However, within the slab and hot-band categories, ranges of potential specifications, called *designs*, may be used to fill each order. Planners must decide which of these are suitable candidates for MTS production.

Slab Inventory and Storage

A typical ISM may carry semifinished inventory at the slab and hot-band stages. Throughout the remainder of the article, we concentrate specifically on issues surrounding the slab stage. It is important to understand that slab inventory is naturally present in ISMs. Slabs act as a buffer between two major production units, the continuous caster and the hot mill. Under current practice, some amount of slab inventory is unavoidable and some amount is planned. For example, unavoidable inventory is created by order cancelations. ISMs produce slabs after they receive initial orders but they confirm the orders later (typically five weeks prior to the delivery date) by contacting the customers to determine that their requirements have not changed. It is not uncommon for customers to cancel orders, or to change order sizes or specifications at the time of confirming the order. Usually, the ISM has produced the slabs for the original order specifications and must consign them to surplus inventory. The ISM's marketing department tries to find customers for such slabs, and its planners try to use such slabs for other orders as soon as possible.

Making some slab designs to stock provides some important strategic benefits. The production of slabs accounts for roughly half the time required to process an order. Thus, having slabs available for an order can potentially cut delivery lead times in half. Carrying the appropriate coil or finished goods inventory cuts lead times even more. However, slab inventory is much cheaper to carry because less value has been added at that stage, and because slabs have a lower rate of spoilage from surface corrosion. Also, ISMs have greater flexibility in matching slabs to orders than at later stages, which gives rise to significant risk-pooling benefits.

Steel making is a few-to-many industry.

There are two common ways of storing slabs. In the *random-access* area, slabs of different designs (grades, widths, weights, internal, and external qualities) are stacked in the same pile in random order. The piles are adjoining, and their heights are restricted to ensure structural stability. Slabs with identical dimensions may be stored at different locations. ISMs use random-access storage for low-volume slabs. They produce these slabs in small quantities but carry literally thousands of designs. Tracking and retrieval can be difficult in the random access area. In the *clone-bank* storage area, slabs of identical dimensions are stacked in piles in the slab-storage yard. The ISM stacks several piles next to each other and keeps the piles of identical slabs uniform in height by rotating the picking of slabs. It thus maintains stability in piles that may be five times the height of piles in random-access storage. Clone banks permit higher-density storage, simpler control and tracking, and shorter retrieval time than random-access storage. However, due to the large number of low-volume custom slabs that require random-access storage, ISMs have limited space for clone banks.

Order Matching Flexibility and Cold-Application Rules

As the ISM books orders, a planner determines whether an existing (cold) slab in inventory can be used to fill the order. This is done by checking if any of the slabs in inventory satisfy a set of rules, called the *cold-application rules,* for the order. If there is such a slab, it is assigned to the order; otherwise a custom-fitted slab design is included in the future production schedule.

Planners translate orders into finished hot band of the required dimensions and then into slabs of ranges of width and weight that can be applied to each order. They have some order-matching flexibility because ISMs can reduce slab width at the hot mill via roughing and thus use slabs slightly wider than the width specified for an order and because customers accept coils that fall within a range of weights around the aim weight they desire. Many customers permit weights somewhat lower than the aim weight but not higher because of constraints at their loading docks. Also, the reduction in width possible via roughing depends on such factors as the grade, width, cast duration, and gauge of the coil required for the order.

Our Approach and Related Literature

We developed and implemented a model for one ISM that it uses to choose which semifinished products to manufacture to stock. We first held discussions with senior planners in several different functional areas at a particular ISM. Participants have included, for example, inventory managers, purchasers, production planners, caster schedulers, and capacity managers. These meetings helped to underscore the key criteria behind the design of semifinished products as well as the business process used by the ISM for making such decisions. Motivated by our understanding of the problem, we proposed an optimization model suitable for making such decisions. The model was eventually implemented as a decision-support system (DSS), developed in C/C++, that runs on a PC and includes a user-friendly graphical interface. The DSS allows inventory planners to analyze various scenarios using heuristics to solve a mathematical optimization problem.

We proposed a two-step approach for managing product variety at the ISM. For the first step, we identified which slabs to produce for MTS. For the second step, we expect to develop methods for scheduling regular production or purchase of slabs chosen for each planning period. Overall, we expect our approach to reduce total inventory, improve delivery performance, and improve capacity utilization. So far, we have completed and implemented only the first of these two steps, which we describe in this article.

Several bodies of literature in manufacturing and operations management concern the problem we studied. Some researchers have focused on the steel industry and similar process industries (reviewed by Dutta and Fourer 1995). Recent such work falls into the following categories: strategic operations management and case studies (Bielefeld et al. 1986, Sinha et al. 1995), supply chain issues (Hafeez et al. 1996, Kisperska-Moron 1990), production planning and control issues (Boukas et al. 1990, Chen and Wang 1997, Lin and Moodie 1989, Sasidhar and Achary 1991, Vasko et al. 1991) caster and production scheduling decisions (Box and Herbe 1988, Diaz et al. 1991, Jacobs et al. 1988, Vonderembse and Haessler 1982), matching surplus inventory to orders (Kalagnanam et al. 2000), and cutting stock and ingot design (set-covering) problems (Vasko et al. 1992, Vasko et al. 1989, Vonderembse 1984).

Balakrishnan and Brown (1996) formulated the problem of choosing sizes of semifinished aluminum tubes (the bloom-sizing problem) from which tubes of different sizes can be drawn and finished to meet customer orders. For projected product mixes and volumes, they propose heuristics to find the best n tube sizes that minimize the overall drawing effort subject to an upper bound on the extrusion effort. This and the studies that deal with ingot design (Vasko et al. 1992, Vasko et al. 1989, Vonderembse 1984) are closely related to our problem. Unlike other researchers, we consider an infinite range of choices with respect to the width and weight of slabs, not a finite set of feasible designs. Our approach can approximate the best $k - x$ designs given that x designs are fixed (possibly by an earlier decision). This facilitates improved management of clone-bank inventories that are expected to change gradually over time as product portfolios change. Our model uses metallurgical and process constraints to determine which incoming orders match which designs. Its central focus is not on allocating existing surplus inventory to orders (a complicated problem in its own right). Thus, our work is distinct from work dealing with inventory-matching problems (Kalagnanam et al. 2000).

From a modeling perspective, the problem of choosing the optimal intermediate product design is related to the problem of determining the optimal point of differentiation, subject to a service level constraint (Lee 1996, Lee and Billington 1994, Lee and Tang 1997, Garg and Tang 1997, Swaminathan and Tayur 1998, Graman and Magazine 1998, Gupta and Benjaafar 2001). The key difference is that these models deal with assembled products. Finally, an extensive literature deals with inventory-placement issues in series production-inventory systems (Gallego and Zipkin 1999, Axsäter 2000, Chapter 5).

The ISM has improved its utilization of inventory space.

Model Formulation

Two key pieces of data are needed for formulating the slab-design-optimization problem: cold-application rules to determine what slab designs are feasible for each order and historical order-book data on the types and mix of orders the ISM typically produces. The time frame for the data set used in the model depends on the frequency with which the ISM selects slabs. We used six months because the ISM reevaluates its slab design decisions every six months.

We formulated the model with the objective of maximizing realized revenue when the maximum number of allowable designs is k. The ISM's inventory planners specify the number k. It cannot exceed the maximum number of cells dedicated to clone banks under existing constraints on space in the slab-storage yard. For each feasible application of a slab design to an order, we assume that the ISM earns a nonnegative reward. The reward represents a numerical measure of the benefits of satisfying the order with the particular slab. For applications of slabs to orders that are not feasible, this reward is zero. In the implementation of our model, we assumed the reward for each feasible application was equal to the size of the order in tons. However, in general the rewards could capture additional factors, such as variable production cost, order size (in tons), importance of the customer, and the width and weight discrepancy between the slab and the ideal width-weight combination for the order. We provide a mathematical description of the model in the Appendix.

A complication of the model is the continuous range of slab designs that are feasible for each order. This means that the set of potential slab designs is infinite. We tackle this problem by observing that for any subset of orders that can be satisfied by a range of slab widths and weights, it is possible to identify a unique optimal design. This holds as long as the reward associated with the application of a particular slab to a particular order is linear in the following factors: the amount of discrepancy between the ideal and actual slab dimensions applied to that order, the size of the order, and the importance of the customer. This implies that we can preprocess demand data to identify a finite set of slab designs and then search within that set to find the best k designs. The number of possible designs does not exceed c, the number of subsets of an order set of size c. Although this is potentially a very large number, in practice, cold application rules substantially limit the number of feasible designs.

The property described above parallels the node-optimality property of certain location problems (Mirchandani and Francis 1990, pp. 75–78). We demonstrate this property with the example in Figure 2 in which there are five orders. For a single order, we can assume there is a unique optimal design (width-weight combination in this example). Determining the optimal slab width for a given order depends largely on which production process is currently the bottleneck. For example, when caster capacity is a bottleneck, ISMs try

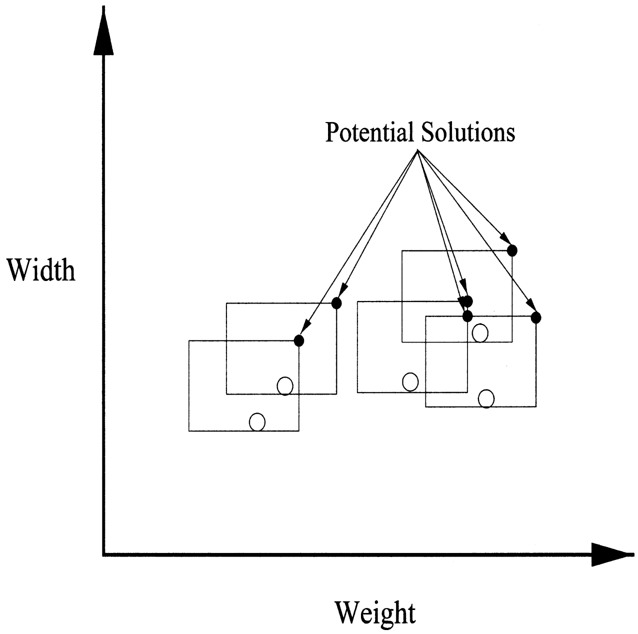

Figure 2: This figure illustrates the finiteness of the set of potential solutions. The open circles denote the aim width and weight for each order, the rectangles represent feasible width and weight ranges for each order, and the dark circles are the potential solutions. The subset of two orders on the left has two potential solutions, whereas the subset of three orders on the right has three potential solutions. No slab design can satisfy all five orders.

to cast slabs with the greatest possible width to maximize the total tonnage they can process. Thus, the ISM assumes the ordered weight is the maximum possible weight and uses this weight as the goal to simultaneously maximize revenues (because revenues are proportional to slab weight) and minimize the number of pieces handled and processed downstream. For illustration purposes, this example assumes that the caster is indeed the bottleneck and therefore maximum width and weight slab designs (top right corners) are preferred. In our example, no single slab design satisfies all five orders. If we consider the two subsets consisting of two and three orders shown in Figure 2, we notice that each of these subsets can be satisfied via common slab designs. Then, the set of designs that contains the complete set of optimal designs is the set of top right corners of all slab design sets that are feasible for at least one member of the original subset.

After obtaining the set S of candidate slab designs, we formulated a mathematical model of the slab-design-optimization problem (Appendix). This is a well-known mathematical model and arises in many different contexts. For example, the same mathematical formulation is obtained when a firm needs to choose at most k locations to maintain bank accounts for customer payments given fixed costs for opening facilities. In that context, the problem is called the lock-box problem (Cornuejols et al. 1977). A special case of the problem in which facilities are located at demand nodes and the fixed costs for opening facilities are zero is called the *k-median* problem in location theory (Mirchandani and Francis 1990, Chapter 2). Many operations researchers have developed heuristics that provide reasonably accurate solutions for this class of problems (Nemhauser and Wolsey 1999, pp. 495–512). For large problem instances, a combination of the greedy heuristic to obtain a good initial solution, and an interchange heuristic (Teitz and Bart 1968) to improve this solution yields good solutions (Cornuejols et al. 1977 give details). We chose this solution approach because it is fast (the model size for this application is computationally prohibitive) and because it requires no commercial software (for example, LP solver).

Implementation and Numerical Examples

An ISM applied our model and uses it as a decision-support tool for choosing which slab designs to store as clone-bank inventory and which slabs to purchase externally. We carried out the initial implementation and testing of the greedy-interchange heuristic on a Sun Ultra 10 Workstation. The algorithms were coded using C++ and tested using real data from current and previous years.

Figure 3 shows the different types of data required by the model. The model uses various sources of data at the ISM to formulate the optimization problem. We define an instance of the model using the set of rewards for applying orders to slabs. For each customer order, we obtained data on the ideal slab design and ranges of slab width and weight that constitute feasible

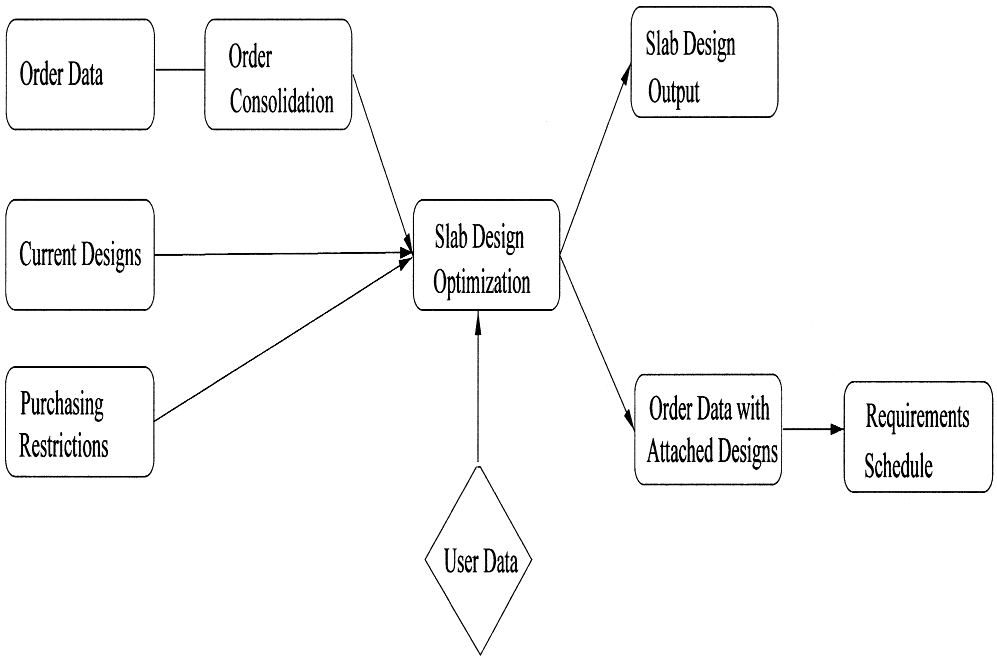

Figure 3: This overview shows the data inputs for the slab-design-optimization model and the outputs obtained from the model.

designs for that order. We then combined customer orders that had the same ideal slab design to obtain the demand volume for an order type. The set of all order types is denoted as C. To generate instances of the model corresponding to the needs of the decision makers, we made the following assumptions. First, we assumed that for each distinct subset of order types that could be covered by at least one common slab design, the optimal design was the maximum slab width and weight in the set of potential slab designs S. This is consistent with the assumption that the continuous caster is the bottleneck in the production process. We assumed that all of the distinct order types over a specified planning horizon (for example, three to six months) were rolled into a single time period. Also, we assumed rewards were equal to the estimates of the mean demand over the specified planning horizon for each order type.

To estimate the mean demand for each distinct order type, we could use historical order data, quantitative or subjective forecast of demand, or a mixture of the two. The advantage of setting rewards equal to mean demands is that mean demands are readily available data and we did not have to tune cost parameters to trade off the impact on different functional areas of the ISM. Furthermore, we based the objective function on a widely accepted measure of performance—the total tonnage of orders that can be covered with the chosen designs.

The objective is thus to maximize total demand covered by the k designs, given the secondary objective to maximize caster throughput. The basic structure of the algorithm used is as follows:

(1) Model generation: We compute a set of potential slab design choices based on cold-application rules in place at the ISM and numerical rewards, r_{ij}, are generated and stored using a sparse matrix storage scheme.

(2) Solution: We apply a fast greedy-interchange heuristic to generate a user-specified number of clone-bank designs.

(3) Reporting: We perform supplemental data analysis, producing a variety of output reports describing various properties of the proposed bank (for example,

factors affecting hot-mill scheduling, lists of final products, and customers that can be served by each clone bank).

An important step in implementing the model was validating it for the decision makers at the ISM. We did this by comparing decisions made using the model with those made in the previous year. Historically, decisions about which designs to keep in clone banks were assessed periodically. Approximately every quarter, about six inventory planners met to decide whether to add, remove, or switch existing designs. To validate the model, the planners used the model output to help them choose designs in a simulated decision process. Using several iterations they chose designs based on the model, reviewed the choices, and subsequently requested perturbations of the dimensions and analyzed the results. The perturbation analysis took the form of what-if type questions aimed at analyzing and capturing desirable attributes of designs not incorporated in the optimization. For example, in some cases, planners provided customer-specific information that resulted in changes to designs. Although such changes ultimately reduce total order coverage, they are able to better satisfy the preferred set of customers.

After being validated, the model was ready for use in planning future design choices. The model and greedy-interchange heuristic solution method were subsequently transferred to a Windows NT platform. Although the existing C++ program was easily transferred, we added a graphical user interface to make the decision-support system easy to use. Together with managers and inventory planners at the ISM, we developed an interface that allowed planners to conveniently make adjustments to the model's input (for example, perturbations for what-if type analysis) and to produce reports in formats consistent with other systems' reports.

Numerical Examples

We carried out a numerical analysis using the ISM's historical data. For reasons of confidentiality, we can give no specific data regarding order trends or chosen slab dimensions.

Initially, we analyzed the accuracy of a greedy-interchange heuristic using a computed Lagrangian dual upper bound (Nemhauser and Wolsey (1999) give details of the subgradient approach to solving a Lagrangian dual relaxation of the problem). Numerical experiments showed that the greedy-interchange heuristic typically finds a solution that is within two percent of optimality for the problems we considered. (We found that the interchange heuristic yielded small improvements to the greedy solution (less than three percent).) Typical running time was less that 10 minutes on a PC (366 MHz, 128 MB Ram). Thus, it provides a satisfactory solution method for the structure and size of the ISM's problems.

After applying the algorithm to one instance of the problem, that is, one set of six-month order data, we tested the robustness of the solution by calculating the percentage of orders (in tons) that could be covered by the same 50 designs (that are chosen optimally for the first data set) in different planning periods. We found that the percentage of orders covered by the chosen designs is quite robust with respect to short-term changes in order sets. We expect this not to hold for longer periods of time, because of portfolio turnover. Therefore, the software application of the slab-design-optimization problem has the functionality to start with some initial designs and add the best $k - x$ designs to the already existing x designs. Such flexibility accommodates the need for incremental changes to the portfolio over time.

Increasing the number of clone-bank positions provides diminishing returns (Figure 4). Using the model, we found that more than 50 percent of orders in 1999 could have been filled using 60 different designs, and thus these designs accounted for roughly half of the ISM's revenues. However, as the number of available clone-bank cells is increased, a long tail develops. Managing this tail is a problem all ISMs confront. The problem is that typically many orders require custom slab designs. This is consistent with the 80–20 rule or Pareto distribution of many naturally observable phenomena, including, for example, the ABC classification of inventory (Silver et al. 1998).

We contrasted the demand schedules for two slabs with similar geometries, weights, and quality attributes: a slab identified by the optimization model (Figure 5) and a slab that was kept in the clone bank prior

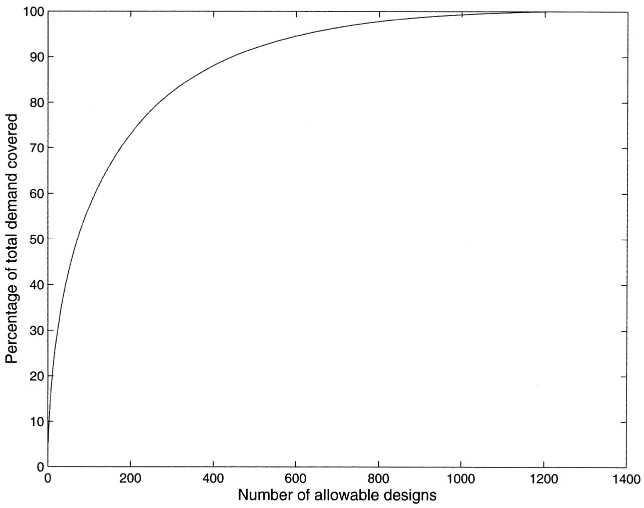

Figure 4: This figure shows the cumulative percent coverage of total demand for slab designs as a function of the number of allowable designs (obtained from using the greedy heuristic). A small number of designs account for a large percentage of total demand in tons.

to the implementation of the optimization model (Figure 6). To generate the schedules, we adjusted the due dates for orders applicable to each slab design according to production-routing-dependent processing times to determine the approximate week in which slabs for each order would have been required for processing at the hot mill. The weekly demand for the original slab design has a mean and standard deviation of 1,538 and 1,012 respectively. On the other hand, the mean and standard deviation of the weekly demand for the optimized slab design are 3,851 and 1,251 respectively. In general, the optimization model results in a significant increase in mean demand but only a small relative increase in standard deviation, giving the slab design a much smaller coefficient of variation (cv). From the point of view of inventory policy, a smaller cv implies lower safety-stock requirements per unit of demand. Thus, although demand variation is unavoidable, the applicability of a greater number of orders to an optimized slab design results in inventory pooling benefits.

Results

Since the DSS was implemented, it has led to the following significant improvements:

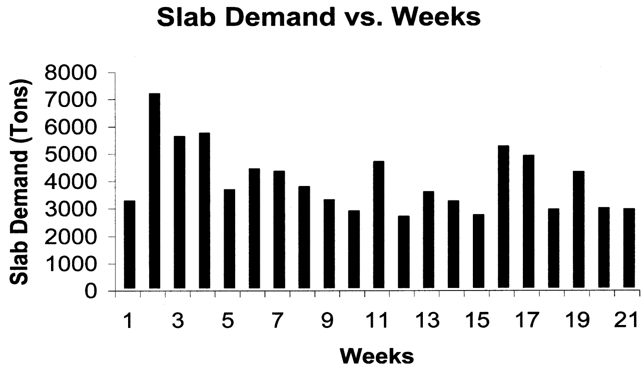

Figure 5: This example illustrates the demand pattern for a 21-week period in 1999 for a typical slab design generated after using the greedy heuristic.

—Total cycle time averaged across all product groups is 30 percent lower for orders sourced from slab stock versus those made from scratch during the year after we implemented the DSS.

—The ISM has dramatically improved its utilization of inventory space and its coverage of orders. Previously, 57 designs covered about 17 percent of total annual order volume. After implementation of the DSS, 50 designs cover about 50 percent of the total annual order volume. Increased coverage implies that the ISM has a better chance of satisfying orders by pulling slabs out of storage, which has lowered cycle times and improved on-time-delivery performance.

—By reducing the number of slab designs, the ISM

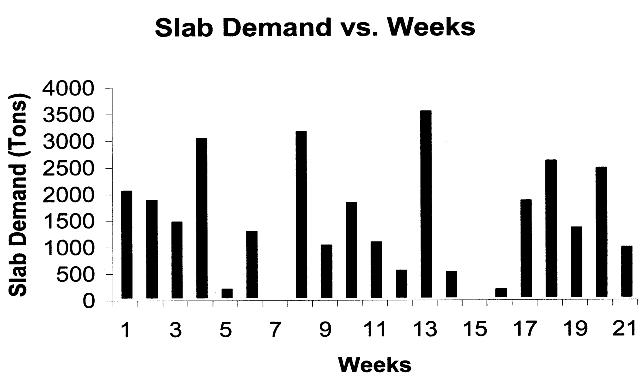

Figure 6: This example illustrates the demand pattern for a 21-week period in 1999 for a typical slab design stored in the clone bank prior to the implementation of the optimization model.

has effectively increased the capacity of the slab storage area. It is using the additional space to achieve greater efficiencies in the purchase of slabs.

—Using the DSS, the ISM has reduced risk (for example, order cancelation risk) by taking advantage of inventory pooling. In particular, it has increased the average number of orders that could use a particular design on average by a factor of five.

—Before we implemented the DSS a full-time senior planner spent about 50 percent of her time deciding what slab designs to carry in inventory. Now she spends about 10 percent of her time on such decisions.

Conclusions

The model we developed is a first cut at modeling the choice of clone-bank designs. It ignores capacity limitations at individual storage locations, that is, how many tons of slabs each location can hold, as well as decisions and policies regarding how clone-bank inventory is to be replenished over time. These factors may eventually affect the choice of clone-bank designs. Furthermore, the ISM could keep inventory at other staging points (for example, hot band) to improve its responsiveness to demand. We are currently working on developing an inventory management system that considers both multiple inventory staging points and inventory management policies for MFS slabs.

Appendix: Model Formulation and Solution Methodology

To express the problem mathematically, we let k denote the maximum number of designs that are to be made to stock. S denotes the set of potential slab designs, and $J = \{1, \ldots, k\}$ the index set of chosen slab designs with widths w_j, weights m_j, grades g_j, and qualities q_j. We define $C = \{1, \ldots, c\}$ to be the set of all orders within a historical data set. For each slab $j \in J$, and order i, we assume there is a nonnegative reward r_{ij}, which represents a numerical measure of the benefit of having a slab j to cover an order i. We assume the reward is nonzero only if the cold-application rules are satisfied for the slab-order pair and zero otherwise. Given r_{ij}'s, our problem is to choose the index set J with cardinality less than or equal to k that maximizes total reward. This can be written mathematically as follows:

$$\max_{J \subseteq S} \left\{ \sum_{i=1}^{c} \max_{j \in J} \{r_{ij}\} \mid \; |J| \leq k \right\}. \tag{1}$$

To formulate the problem as a mathematical program we define decision variables x_j's and y_{ij}'s. The objective is to maximize the total reward. That is, if

$$x_j = \begin{cases} 1 & \text{if } j \in S \text{ is chosen,} \\ 0 & \text{otherwise,} \end{cases}$$

$$y_{ij} = \begin{cases} 1 & \text{if order } i \text{ is assigned to slab } j, \\ 0 & \text{otherwise,} \end{cases}$$

and assuming that each order should be assigned to at most one chosen slab design, the problem can be formulated as

$$\max Z = \sum_{i=1}^{c} \sum_{j=1}^{s} r_{ij} y_{ij} \tag{2}$$

subject to

$$\sum_{j=1}^{s} y_{ij} \leq 1 \quad \forall i = 1, \ldots, c, \tag{3}$$

$$\sum_{j=1}^{s} x_j \leq k, \tag{4}$$

$$y_{ij} \leq x_j \leq 1, \tag{5}$$

$$y_{ij}, x_j \in \{0, 1\}, \quad j = 1, \ldots, s, \; i = 1, \ldots, c. \tag{6}$$

Whereas (2–6) is a well-known problem, it is also known to be NP-hard (Cornuejols et al. 1977). The size of the problem that is relevant for a typical ISM is indeed quite large. For example, a typical historical order set may contain tens of thousands of distinct orders. Similarly, the number of slab designs that we need to evaluate, albeit finite, may easily run into hundreds of thousands, and the clone banks may accommodate 50 different designs. Thus, there is little hope of solving a realistic instance of this problem exactly.

Acknowledgments

We thank two anonymous referees for their valuable comments that have helped to improve the manuscript. We are also grateful to many enthusiastic participants in this project, including Tom Braun, Maureen Crane, Robin Hallam, Ron Johnson, Mike Moore, and Jane Wood. This material is based in part upon work supported by the National Science Foundation under Grant No. 9988721. Additional support was received from the Natural Sciences and Engineering Research Council of Canada (Grant No. 45904), and Dofasco Inc. of Hamilton, Ontario, Canada through research grants to Diwakar

Gupta. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or others that have supported this endeavor.

References

Axšater, S. 2000. *Inventory Control*. Kluwer International Series in Operations Research and Management Science. Kluwer Academic Publishers, Boston, MA.

Balakrishnan, A., S. Brown. 1996. Process planning for aluminium tubes: An engineering-operations perspective. *Oper. Res.* **44**(1) 7–20.

Bielefeld, F. W., K. Walter, R. Wartmann. 1986. A computer-based strategic planning system for steel production. *Interfaces* **16**(4) 41–46.

Boukas, E. K., A. Haurie, F. Soumis. 1990. Hierarchical approach to steel production scheduling under a global energy constraint. *Ann. Oper. Res.* **26**(12) 289–311.

Box, R. E., D. G. Herbe. 1988. A scheduling model for LTV Steel's Cleveland Works' twin strand continuous slab caster. *Interfaces* **18**(1) 42–56.

Chen, M., W. Wang. 1997. A linear programming model for integrated steel production and distribution planning. *Internat. J. Oper. Production Management* **17**(6) 592–610.

Cornuejols, G., M. L. Fisher, G. L. Nemhauser. 1977. Location of bank accounts to optimize float. *Management Sci.* **23**(8) 789–810.

Diaz, A., L. Sancho, R. Garcia, J. Larraneta. 1991. A dynamic scheduling and control algorithm in an ENSIDESA steel plant. *Interfaces* **21**(5) 53–62.

Dutta, G., R. Fourer. 1995. A survey of mathematical programming application in integrated steel plants. Presented at the INFORMS New Orleans meeting, New Orleans, LA, November 1995, http://iems.nwu.edu/~4er/steel/survey.html.

Gallego, G., P. Zipkin. 1999. Stock positioning and performance in serial production-transportation systems. *Manufacturing Service Oper. Management* **1**(1) 77–88.

Garg, A., C. S. Tang. 1997. On postponement strategies for product families with multiple points of differentiation. *IIE Trans.* **29**(1) 641–650.

Graman, G. A., M. J. Magazine. 1998. An analysis of packaging postponement. *Proc. 1998 MSOM Conf.* School of Business Administration, University of Washington, Seattle, WA, June 29–30, 67–72.

Gupta, D., S. Benjaafar. 2001. Make-to-order, make-to-stock, or delay product differentiation?—A common framework for modeling and analysis. Working paper, Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN.

Hafeez, K., M. Griffiths, J. Griffiths, M. Naim. 1996. Systems design of a two-echelon steel industry supply chain. *Internat. J. Production Econom.* **45**(1–3) 121–130.

Jacobs, T. L., J. R. Wright, A. E. Cobbs. 1988. Optimal inter-process steel production scheduling. *Comput. Oper. Res.* **15**(6) 497–507.

Kalagnanam, J. R., M. W. Dawande, M. Trumbo, H. S. Lee. 2000. The surplus inventory matching problem in the process industry. *Oper. Res.* **48**(4) 505–516.

Kisperska-Moron, D. 1990. Analysis of the logistic system of steel products in Poland. *Engrg. Costs Production Econom.* **19**(1–3) 47–55.

Lee, H. L. 1996. Effective inventory and service management through product and process redesign. *Oper. Res.* **44**(4) 151–159.

——, C. Billington. 1994. Designing products and processes for postponement. S. Dasu, C. Eastman, eds. *Management of Design: Engineering and Management Perspectives*. Kluwer Academic Publishers, Boston, MA, 105–122.

——, C. S. Tang. 1997. Modeling the costs and benefits of delayed product differentiation. *Management Sci.* **43**(1) 40–53.

Lin, C. W., C. L. Moodie. 1989. Hierarchical production planning for a modern steel manufacturing system. *Internat. J. Production Res.* **27**(4) 613–628.

Mirchandani, P. B., R. L. Francis. 1990. *Discrete Location Theory*. John Wiley and Sons, New York.

Nemhauser, G. L., L. A. Wolsey. 1999. *Integer and Combinatorial Optimization*. John Wiley and Sons, New York.

Sasidhar, B., K. K. Achary. 1991. A multiple arc network model of production planning in a steel mill. *Internat. J. Production Econom.* **22**(3) 195–202.

Silver, S. A., D. F. Pyke, R. Peterson. 1998. *Inventory Management and Production Planning and Scheduling*, 3rd ed. John Wiley and Sons, New York.

Sinha, G. P., B. S. Chandrasekaran, N. Mitter, G. Dutta, S. B. Singh, A. R. Choudhury, P. N. Roy. 1995. Strategic and operational management with optimization at Tata Steel. *Interfaces* **25**(1) 6–19.

Swaminathan, J. M., S. R. Tayur. 1998. Managing broader product lines through delayed differentiation using vanilla boxes. *Management Sci.* **44**(12) S161–S172.

Teitz, M. B., P. Bart 1968. Heuristic methods for estimating the generalized vertex median of a weighted graph. *Oper. Res.* **16** (September–October) 955–961.

Vasko, F. J., F. E. Wolf, J. A. Pflugrad. 1991. An efficient heuristic for planning mother plate requirements at Bethlehem Steel. *Interfaces* **21**(2) 1–7.

——, ——, K. L. Stott, O. Ehrsam. 1992. Bethlehem Steel combines cutting stock and set covering to enhance customer service. *Math. Comput. Modelling* **16**(1) 19–17.

——, ——, ——, J. W. Scheirer. 1989. Selecting optimal ingot sizes for Bethlehem Steel. *Interfaces* **19**(1) 68–84.

Vonderembse, M. A. 1984. Selecting master slab width(s) for continuous steel casting. *J. Oper. Management* **4**(3) 231–244.

——, R. W. Haessler. 1982. A mathematical programming approach to schedule master slab casters in the steel industry. *Management Sci.* **28**(12) 1450–1461.

Bruce Gabel, Assistant Director of Order Fulfillment, Dofasco Inc., P.O. Box 2460, Hamilton, Ontario, Canada L8N 3J5, writes: "The outcome of the project described in this manuscript has been of significant value to Dofasco. The project involved the development and analysis of a model for choosing which designs of steel slabs and hot band steel coils to carry in stock for the purpose of reducing cycle times. The resulting model has formed the basis of decision-support software that is used on a regular basis to aid in advance planning of inventory requirements.

"Use of a mathematical model as a basis for choosing designs has not only improved the quality of decisions, but also the speed and consistency with which they are made. To give a specific example of the improvements, the model showed the number of different designs of slabs could be reduced by nearly 50% without any negative effect on performance. This will lead to significant reduction in inventory."