Periodic Production Scheduling at a Fastener Manufacturer

Published in International Journal of Production Economics

by
Victor E. McGee and David F. Pyke
Amos Tuck School of Business Administration
Dartmouth College
Hanover, NH 03755

TEL: (603) 646-2517 FAX: (603) 646-1308

We want to acknowledge the excellent contributions of one of our former MBA students, Peter Stahl, who introduced us to the problem and worked diligently with us until his graduation. We also want to thank the many people in the company who engaged in the joint endeavor which led to the results reported here. Finally, we gratefully acknowledge the helpful comments of three anonymous referees.

Revised November 20, 1995

Abstract

This paper presents the application of periodic production scheduling at a metal-fastener manufacturing company. The problem was initially presented as a mandate to reduce finished goods inventory by 33%. Preliminary investigations, however, revealed that any effort to reduce finished goods inventory necessarily involved two major endeavors: (1) significant improvements to forecasting and (2) a periodic production scheduling system (PPS). The forecasting system provided the input to the PPS system. The PPS required initial procedures for sorting parts into families, which, as far as possible, shared the same setup patterns, engineering specifications, raw material requirements, etc. The PPS itself was developed on a spreadsheet, and the output from this package fed into an interactive PC package. In addition to graphic presentation of production scheduling, the PC software package also provided manpower implications of specific PPS plans for tooling and supervision during runs, on the basis of the periodic schedule. The paper concludes with lessons learned from attempts to finalize implementation.

1. Introduction

This paper presents the application of periodic production scheduling at a metal-fastener manufacturing company which we shall call Farnsworth-Shaw (F-S).[1] The firm produced more than 3600 stock keeping units (SKUs), using more than 40 cold nut forming (CNF) machines, and a number of tapping, heat treating, plating and cleaning machines. The parts produced include standard nuts, multipiece nuts and specialty fasteners. For instance, F-S made many standard hex-nuts that one would purchase at a local hardware store. One would more often see the parts, however, on automobiles produced in the U.S. (even if made by Japanese firms). The problem was initially presented to the authors as a mandate to reduce finished goods inventory by 33%, but it became evident immediately that production scheduling and forecasting issues needed to be addressed. The choices for production scheduling at F-S were many, but the complexity of the problem, and the lack of computer literacy on the part of the production scheduling personnel, narrowed the options considerably.

F-S received some forecasts from its automotive customers, who represented 70% of sales; but non-automotive customers ordered irregularly and in large quantities, creating lumpy demand and a scheduling nightmare. Even the automotive forecasts, which were provided for 8 weeks into the future, were firm for only 2 weeks hence; beyond that the forecast reliability degraded significantly. Schedulers did not believe any forecast from automotive firms past two weeks ahead. The response at F-S had been to build large finished goods inventories, and to *react* to demand as it occurred by initiating production orders in response to customer requests for parts. This system had worked quite well for many years, but recent events had changed their world dramatically.

Foreign competitors had entered the marketplace, taking market share in standard nuts by offering significantly lower cost. In addition, the owner of the firm had determined that inventories were 33% too high, and had asked factory management to adjust them accordingly. But managers knew that a blind reduction in inventory would be disastrous

---

[1] This is a fictitious name. The data and names reported in this paper have been disguised but the essential data relationships have been preserved intact to register the size and complexity of the problem. Some features of this project have been written up in somewhat simplified form as “Farnsworth-Shaw: a Case Study and Teaching Note on the Application of Periodic Review Scheduling Techniques in Manufacturing” and presented to the Eighth Working Seminar on Production Economics, Innsbruck, Austria, February 22-25, 1994.

for customer service. The primary source of the problems was that one year prior to this consulting assignment, the firm had purchased a competitor, nearly doubling the number of SKUs in the catalog. Now with 3600 SKUs, the reactive production scheduling system was falling apart. While efforts had been made at a more proactive system, they had reached their limit.[2]

This paper describes the systems developed for Farnsworth-Shaw, as well as the theoretical basis for these systems. First, however, in section 2 we describe F-S and the forecasting and production scheduling problems in more detail. We then present a brief review of relevant academic literature in section 3. Section 4 deals with the data sets available for analysis and gives some indication of the size and complexity of the task. The forecasting segment of the analysis is described in section 5. Section 6 presents the details of the production scheduling system. Section 7 describes the software developed to facilitate implementation. Finally, section 8 discusses implementation issues and ends with lessons learned.

2. The Company and the Problem

Farnsworth-Shaw serviced automotive manufacturers, automotive sub-suppliers, other original equipment manufacturers (OEMs) and a wide variety of distributors. Even though F-S served a cross-section of industry, 70 percent of sales came from the automotive industry. Indeed, the automotive industry had dominated Farnsworth-Shaw's focus ever since they moved away from the railroad industry in the early part of this century. That focus, however, led to poor service for the distribution side of the business.

The firm manufactured a complete line of stock nuts, many of which were patented innovations that revolutionized joint fastening. In recent years, as patents expired and international competitors picked off the high volume parts for manufacture on high-speed dedicated equipment, F-S shifted its focus to more specialized and complicated, often multipiece, fasteners. This business involved sophisticated selling to the end user, including communicating uses for the end product as well as extensive engineering and

---

[2] This is a case involving "Competent and dedicated people caught up in the never-ending eddies of daily production schedules, past-due orders, the wrong mix of raw-material inventory, hard-to-control WIP inventories, large finished goods inventories but not necessarily what is needed for Monday's shipment, second-class service to non-auto customers, lumpy demand, maintenance problems on the machines, etc." Quote from an internal memo.

long development times. For automotive companies, the development time often exceeded three years.

Distributors, on the other hand, tended to sell only standard parts with little of the customer-specific design and service that F-S viewed as its specialty. These distributors placed orders at irregular intervals for sometimes large quantities of parts. Distributors and non-automotive OEM customers knew from experience that F-S's automotive business took priority. Delivery problems, difficulties with sales assistance, and extended delays for order errors to be resolved, all communicated to non-automotive customers F-S's focus on the automotive business. Nevertheless, F-S had maintained a strong name in the industry and was considered a high-quality supplier. Some OEMs, however, were beginning to shift their business to other manufacturers. Even some of F-S's largest distributors had started purchasing from other sources.

These other sources included domestic and foreign competitors. While it was difficult for foreign firms to take over the specialty fastener business because of its service and turnaround requirements, they had all but cornered the standard fastener business. A new type of domestic manufacturer evolved as well. These firms purchased blanks manufactured offshore and then tapped (cut threads) or added locking features in the U.S. F-S itself had dabbled in this business.

Customers. Automotive customers, the primary source of F-S's business, had changed their policies with regard to suppliers dramatically in the past few years. In fact, they essentially required the supplier to hold two weeks of inventory of every part. If the customer, say Ford, needed to dip into safety stock in an amount of less than two weeks' demand, F-S would provide those parts immediately. On the other hand, if Ford's forecasts were off by more than two weeks so that they requested significantly more parts, they would pay F-S to expedite shipment. The disruptions at the F-S factory were intense, even though F-S was paid by the customer to expedite shipping. In the previous year, it had become increasingly evident to managers that more than two weeks inventory may be required.

A recent survey of OEMs suggested that on-time delivery was their primary performance measure. It could be said that F-S generated business by having the engineering expertise to be able to build new and special parts for custom applications for new automotive products. High quality was an entry ticket to the marketplace. Then delivery and price were order winners. Distributors, on the other hand, viewed on-time delivery, defect free products and accuracy in ordering and billing as the top performance measures. Custom products were less of an issue for distributors since they stocked more standard parts. Hence, F-S was in the unenviable position of trying to provide customized engineering solutions to the automotive industry, high quality complex parts to OEMs and automotive customers, and high quality, low cost parts to distributors and some automotive customers. In the language of manufacturing strategy: cost, quality, delivery and flexibility were all important.

The Production Process. Metal fasteners at F-S were made in a multistage production process in which the first stage was cold nut forming. Long "wire" coils of steel, stainless steel, steel alloy, or aluminum, were fed into cold nut forming (CNF) machines. These wires ranged in diameter from 1/4 inch to nearly 1 inch. Within the cold nut forming equipment there were usually five stations. The first station cut a "slug" off the wire. The slug was a small disk of metal sliced from the raw material. Subsequent stations then formed the nut by high pressure punching of the slug with a punch and a die. Each station added new shape to the slug and the final station punched a hole. A conveyor dropped completed parts into a bin at the side of the machine. The completed parts looked exactly like finished nuts except that there were no threads on the interior. From CNF, parts moved to a number of other stages collectively called "secondary operations."

The second stage in the process was tapping. The tapping equipment cut the threads on the interior, and sometimes the exterior, of the nut. Changeovers (setup times) in tapping were minor compared with those in CNF and processing rates were faster as well. Parts then went to a washing station for cleaning. Occasionally parts were washed before tapping. After tapping and washing, some parts required heat treating, and/or shipment to a vendor for plating (adding some material to the exterior nut for rust prevention or appearance). After heat-treat and plating, if performed, nuts often went to other operations which added other parts such as washers, nylon inserts, and other pieces. These provided special features so that the nuts would perform unique locking or load spreading functions. Finally, they went to inspection and packaging. The finished goods warehouse at one end of the factory contained large quantities of many of the parts.

Factory personnel suggested that CNF was the clear bottleneck in the production process. Setup times were extremely long on CNF machines. When wire sizes were changed from one part to another, the setup took at least 8 hours. (Wire size changeovers were called *major setups*.) Complications could push the setup to 20 hours. When the wire size was the same the setup took up to 3 hours (termed *minor setups*), depending on how many stations in the CNF machine had to be changed. These long setup times, and the relatively slow processing times at CNF, gave clear indication that this stage was in fact the bottleneck. Production scheduling efforts, therefore, were focused on this stage.

Production Planning and Inventory Control. Automotive orders were known in advance since automotive customers provided weekly forecasts for eight weeks into the future. These 8-week projections were received by electronic mail at F-S. As noted above, the forecasts for the first two weeks were reasonably solid but beyond two weeks, orders could change significantly. Distributors, on the other hand, and non-automotive customers provided no forecast. They simply ordered at irregular intervals. In addition, prototypes, called “samples,” had to be produced on regular CNF machines. These often required long setups, and because they were innovative parts, they often had to be tested multiple times.

The reactive system of production scheduling caused major delays when the utilization of the CNF machines was high. It was not uncommon to observe two long setups when only one could have been performed: A part of a particular style would be produced, followed by a long setup to produce a very different part. While producing the second part, an order for a part similar to the first part would arrive. After completing the second part, another long setup had to be performed. Many hours of productive time would be lost. When F-S purchased the competitor and doubled the number of SKUs, utilization increased dramatically; and the reactive system fell apart. Delays were unacceptably long, customers were angry, and factory managers were extremely frustrated. Samples aggravated the problem, causing tension between production scheduling personnel who were being pushed to ship on time, and managers who knew that new product development was the source of future viability.

F-S’s production scheduling managers worked to create a new approach to production planning and control which aimed at minimizing changeovers on the cold nut forming machines. To some extent, it was possible to keep these setup times low, given the advance knowledge of automotive demands. However, samples, distributor orders and even some automotive orders could invade the schedule. When a distributor placed an order, a check was made to see if the part was in inventory and was not slated for delivery to an automotive customer. If it was available, a quick delivery was promised and an order was initiated at the shipping department. If it was not in stock, the head of non-automotive planning checked the production schedule and, through a negotiation process with the automotive scheduler, made a determination as to when the part could be inserted into the production schedule on the required CNF machine. Allowing one week for every secondary operation required, including the out-sourced plating operation, if necessary, a lead time was relayed to the customer. (This one-week-per-operation lead time was conservative, but not unrealistic. Queueing problems at secondary operations were minimal, so that loads were relatively light.)

In the process of gaining control over the schedule, production schedulers created a very large magnetic whiteboard that was mounted on a wall next to their desks. It was used as a large Gantt chart to control production on the CNF machines. Each machine was represented by a row on the board and time was represented horizontally. Order numbers, with information such as number of parts, were placed on the row according to the schedule of production for that CNF. This gave an immediate visual representation of the load on the factory at any given time and into the future. The net effect of using the Big Board was to decrease setup time substantially, but negotiations for inserting special runs, sudden distributor orders, samples, etc., continued to produce anguish and disarray in the schedule. At the time the authors were approached to help with the problem, the benefits from the Big Board had reached their limit. There were simply too many SKUs to keep track of, and an even more proactive approach was needed.

The Specific Problem. There were 46 cold nut forming machines at the time of the study and these were grouped into 9 classes -- referred to as the “half-inch” machines, the “nine-sixteenth” machines, and so on. There were roughly the same number of tapping machines, two washing stations, two heat treat ovens, and dozens of machines that added washers, nylon inserts, and so on. In spite of the fact that all CNF machines could handle a variety of parts, given appropriate setups, it was nevertheless true that most parts had a *preferred* CNF machine, a fact which had to be taken into account in production scheduling. The preferred CNF for each was known to engineers, operators, and production schedulers, and was based on previous experience running parts on various “identical” machines. Some parts simply had higher quality on some machines than on others, even though the machines were identical models.

In spite of the extra control exercised by the Big Board service was not good. Over a seven-month period that began nearly 10 months after the purchase of the competitor, on-time deliveries averaged around 80 percent. Close to 15 percent were over five days late. With roughly 1,000 parts to be produced each month "it is nearly impossible to create a rational schedule. Ray is constantly giving me (non-automotive) orders, and I understand. His customers are asking him for quick delivery -- or at least, certain lead times. However, our automotive representatives need to be 100 percent sure that their parts will be available for the automotive customers on time. I have to try to balance all those needs and keep setup times to a minimum."[3]

The authors also encountered issues regarding incentives and costing. F-S did not know in detail the cost to produce its products, e.g., the true cost of a setup. The practice was to use $40 per hour of setup time. That figure came from the labor associated with a setup, as well as several other assumed and real costs. The authors also deemed that it was entirely possible that the complexity introduced by the distributor business required so much time on the part of the non-automotive scheduler and so much interference in the production process that the products did not pay for themselves at the price charged. But F-S had little knowledge of the true costs of these activities. Also, at certain stages of the process, supervisors were rewarded on the basis of utilization of equipment and labor. It may have been preferable at secondary operations to incur more setup time in order to push orders through the process. Extra setup time, however, would reduce the performance of the supervisor as currently measured.

Detailed analysis of a 45-month history of production and shipment of some 3000 parts led to a statement of the company problem as a forecasting-production control issue. This approach was basically twofold: (1) a forecasting driver, which led to (2) a periodic production schedule (PPS). The front-end forecasting unit sorted parts into more-forecastable to less-forecastable groupings, anticipated the future demand and fed to the PPS an estimate of monthly demand and an estimate of the standard deviation of demand.

The choice was made to pursue a *periodic* schedule because it would be proactive and orderly. Because of sequence dependent setups, reactive systems would create enormous queueing delays. A proactive system, on the other hand, would allow schedulers to anticipate production of similar parts, and keep setup times to a minimum. Factors built into the PPS included reserved time for compulsory maintenance of machines, flexibility to juggle the placement of production runs (without destroying the periodic schedule), and

---

[3] Quote from the head of CNF scheduling.

automatic indication of manpower requirements for tooling and supervision during runs. The essence of the PPS was programmed for visual implementation and represented a computerized version of the Big Board. Details will be discussed in sections 6 and 7.

3. Literature Review

There has been considerable work on the deterministic version of the problem encountered by F-S. Silver & Peterson (1985), for example, address the choice of replenishment quantities for a family of items when demand is deterministic and there is a major setup cost for the family, and a minor setup cost for each item within the family. The familiar economic lot scheduling problem (ELSP) also considers deterministic demand, setup costs, and sequence independent setup times. See for example, Elmaghraby (1978), Dobson (1987), Roundy (1989), Zipkin (1991). Unfortunately, F-S faces stochastic demand, in some cases with high variability, and sequence dependent setup times. (Extensions to the ELSP have addressed sequence dependent setups (e.g., Dobson (1992)), but the issue of deterministic demand remains.)

Some work has addressed lot sizing in a similar context to ours. Karmarkar (1987) points out that the setup cost is often a surrogate for the violation of capacity constraints and recommends focusing on setup time instead. Accounting for queueing effects, he shows that the optimal batch size is linear in the setup time, instead of the concave function captured by the usual EOQ formula. The results in Karmarkar's paper are not directly applicable because he assumes FCFS processing of jobs that arrive randomly to a shop, whereas F-S should necessarily focus on a fixed sequence to minimize changeover times. As we shall see below, however, our system follows Karmarkar's qualitative recommendations by allowing the manager to adjust the setup *cost* as a means of managing total *time* spent in setup. Therefore, although we begin with the (very soft) setup cost figure given by F-S, it is used simply to manage capacity utilization. Other work in this vein includes Karmarkar, Kekre, Kekre, & Freeman (1985), Weng (1994) and Benjaafar & Sheikhzadeh (1994).

Tang (1990) considers a problem in which major and minor setups are required between and within families, respectively. He provides a heuristic approach to solving for the production quantities and schedule of each part when the objective is to minimize the time at which the last part finishes processing. Processing times are very short relative to setup times and demand is given. F-S is faced with on-going random demand, and with batch processing times that can be very long, given the high demand.

One possible solution to the F-S problem is to employ "can-order" policies, of the (s, c, S) type, where s is the reorder point, c is the can-order point, and S is the order up to level. If the inventory position of a particular product reaches its reorder point, initiate production -- which in the F-S case means to put it in queue at its assigned machine. Any other parts that share similar setup characteristics to this part, and that are below their can-order points, should be put in queue next to the reordered part. See, for example, Silver & Peterson (1985) and Federgruen, Groenevelt, & Tijms (1984). Unfortunately, although this policy is intuitively appealing, optimal values of s, c, and S are difficult to find. Also, due to the reactive nature of the policy and therefore varying queue lengths, actual lead times are unknown. Computing the optimal policies is all the more complicated.

Several authors have extended the ELSP to include random demands. See for instance Gallego (1990), Federgruen & Katalan (1994a) and Federgruen & Katalan (1994b). Federgruen & Katalan (1994b) is particularly relevant, although we had completed our project prior to this paper appearing. Poisson or compound Poisson demand is assumed, and a base stock policy is employed. Idle time may be inserted between each production batch, but the search for optimal policies requires only finding the optimal total idle time in a cycle. Simple adjustments allow for sequence dependent setups. They also note a difference between two types of base-stock policies. Under a *pure* policy, an item's inventory is increased to an order-up-to level every R time periods. Under a base stock policy, production continues until the target level is reached, but the time between replenishments, R, is somewhat random. Our approach to the F-S problem includes the use of base stock policies, a fixed cycle for each product, and an implied adjustment of the total idle time allowed in a cycle. Our choice of a fixed cycle is based on the fact that F-S must schedule operators as well as batches on machines. Although the operator scheduling problem is minor compared to the problem of scheduling batches, a fixed cycle allows the foreman at CNF to generate a regular rotation of his operators.

In related work on the stochastic ELSP, Bourland & Yano (1994) examine the use of capacity slack and safety stock in the stochastic economic lot scheduling problem. They find that capacity slack is not as cost effective as safety stock. However, they assume that setups are sequence independent, and that backorders are not allowed. Thus, they employ a control policy that is essentially reactive -- a setup occurs when a part reaches its reorder point.

Given this literature, we chose to rely on work done by Atkins & Iyogun (1988), which assumes major and minor setup costs and shows that a simple periodic policy outperforms can-order policies. Their paper presents a procedure for allocating the major setup cost to several products. We employed this procedure, and then drew on the work of Karmarkar (1987) to manage capacity utilization by adjusting the major setup cost. Similarity with the work of Federgruen & Katalan (1994b) will also be noted, although as mentioned above, this work was not available at the time.

4. Initial Data Analysis

At the time of the consulting assignment, F-S was in the process of changing over to a new computer system and the initial data analysis was done on the basis of 45 months of past history on the old computer system. Two major files were made available. The first was a *year-to-date "sales" file* containing standard information for a particular product (some product identification, e.g., part number, product code, material code, heat treatment specification, finish code, style code; shipment information, e.g., quantity, customer, date; and cost information, e.g., cost per 1000 pieces, sales dollars, standard cost). The second dataset was the *engineering file*, containing such information as raw material code, thread size, thread pitch, slug weight, finished weight, hours to produce 1000 pieces, standard run size, and forming setup code. The forming setup code was a pointer to hard copy engineering files that detailed how a setup was to be performed. Two parts with identical forming setup codes were identical at the CNF stage, even if they became different parts at subsequent operations. The common elements in the two files allowed for conversion from "sales" to "production." The raw material code translated into *wire size* which was the diameter of the wire coil fed into the CNF machine. This, together with the forming setup code, was very important information, since setup time was a crucial consideration.

An analysis of the sales file indicated that over the 45 months 3,603 different part numbers had been produced and shipped. In any particular month some 1,000 different parts were produced and the number of new parts made in each month ranged from 4 to 41 in the previous twelve months. In the two months after the company had acquired the competitor, however, the number of new products produced jumped to 105 and 185, respectively, putting a severe strain on the system.

5. Forecastability

How forecastable were demands (shipments) for various part numbers? Textbook forecasting (see Makridakis, Wheelwright & McGee, 1983) usually means selecting from among a multiplicity of specific forecasting models (decomposition procedures, exponential smoothing models, Box-Jenkins ARIMA models, etc.) and making a choice based on a set of commonly used measures of accuracy (MAPE, MSE, Theil's U, etc.) for the *fit* of the model to historical data, and maybe also how well the model actually *forecasts*. This applies to a single time series (e.g., demand for one part number over time). For multiple time series, methods rapidly become esoteric (e.g., transfer function models) and if causal models are attempted in the form of econometric models the focus is on structural relations and not so much on ability to forecast. In the case of a company wanting to anticipate demand for 1,000 parts in a single month the problem cannot be handled in a simplistic manner by taking each part separately as a single time series, finding an acceptable model and forecasting over the next several months. There are several realities to contend with. In particular, even though the automotive customers sent their demand projections to F-S by electronic mail with 8-week projections, these projections were highly uncertain past two weeks ahead. On a weekly basis the computer projections were updated but changes could happen via a phone call at any time. Shipments were to take place on a regular day per week and customers wanted the total order shipped, not part of an order. In other words, even though the automotive customers were willing to accept partial shipments, they held it against F-S if the order was not complete.

An "order" from an automotive customer typically contained multiple (50 or more) parts. See Table 1, for example. Forecasting such a "customer order" (collection of parts) is not handled simply as a textbook problem. Regular weekly shipments took place on a Monday but each other day of the week shipments might be made (as in the week shown), mostly because of telephone requests from the customer, but occasionally from backorders. On a given day, the same part number (e.g., part number 35013) could appear several times because deliveries were made to different locations for the same customer. Altogether, in the example in Table 1 some 54 different parts were involved in the shipments to one customer in one week and quantities for individual parts varied considerably.

Another pragmatic consideration concerned the precise role that a forecast plays in a production-inventory system. Production scheduling requires the determination of the timing and the size of a production run (batch size) and this involves a demand forecast, a measure of variability, production lead-time, and safety stock considerations. Optimizing a specific forecasting model is not a critical issue. Clearly, if demand were statistically forecastable then unexplained variability would be decreased and this would help reduce the safety stock necessary to achieve a given service level. So a first step was to develop a quick indication of the forecastability of the parts. Later considerations involving the periodic scheduling system would render the precision of the forecasting model less important.

A first pass was made at identifying which parts, individually, were more forecastable and which ones were not. To do this the 45-month demand figures for each active part were examined as individual time series for any periodicities that could be counted on. A simple model was used, involving three harmonics (sine waves) of periodicities 3-, 6- and 12-months, and projections were made 6 months into the future. Figure 1 illustrates this procedure for four of the 57 parts that had activity in every one of the 45 months. They are all automotive parts and the dashed line shows the harmonic model projecting 6 months ahead. A rough estimate of the forecastability of the 4 parts is shown as an R^2 figure. A total of 961 parts was examined in this way and projections for the next 6 months were made along with a measure of standard error of forecast. Many of these time series were essentially impossible to forecast and so the standard deviation of actual demand over the last 12 months was taken to be the variability.

6. The Production Scheduling System

As indicated above, the clear choice for production scheduling at F-S was a periodic system. The key features of the system developed were as follows:

- Create families of SKUs and assign them to particular CNF machines

- Determine the cycle length for each family and CNF machine

- Create the planned sequence of SKUs for each CNF machine

- Check capacity and adjust cycle lengths accordingly

- Determine order-up-to levels for each SKU

The result of these activities was a regular periodic schedule for each CNF machine. For example, every 6 weeks, say, a given sequence of parts will be produced, in order, on a given machine. Some parts could be produced every other, or every fourth, cycle -- or even less frequently. The amount to produce will, on average, be 6 weeks of demand; but on a given run could be more or less, depending on available inventory and demand projections. We now discuss each feature of the system in turn. (These efforts were directed initially at one bank of machines -- the half-inch CNFs, which were highly utilized.)

Creating Families. The creation of families of parts was related to major and minor setup times on CNF machines. Recall that although CNF machines could handle a wide variety of parts many manufactured parts were identified with a preferred machine. A given CNF machine therefore was assigned all parts that "preferred" that machine. Then assignments of all other parts to CNFs were made to minimize the number of wire sizes on any machine. In some cases, only one part was assigned to a machine because of extremely high volume. In other cases, dozens of parts and up to 3 wire sizes were assigned. These assignments were checked by engineering and production scheduling personnel. We called a group of parts sharing one wire size on a CNF a "family" of parts.

The specific sequence of parts within families was determined later because the engineering staff indicated that the effort to create a "changeover matrix" of times required to set up part *i* after part *j*, for all parts would be prohibitive. Therefore, they recommended that as a first cut, we use a rough estimate for a minor setup and a fixed 10 hours for a major setup.

Determining Cycle Lengths. Once the families of parts had been identified for a given CNF machine the cycle lengths for a periodic production schedule could be determined. We refer to Table 2 to help illustrate this process.[4] Note that this table is a listing of a spreadsheet package for one CNF machine only. Twenty parts were to be made on this CNF machine and these parts have been arranged into three families according to wire size. The first family consisted of parts 1, 4, 6, 7, 8 and 13 which shared a common wire size. They were to be produced as a block, one after the other, in some defined order (to be determined later based on setup times) and produced every x weeks (where x is to be

---

[4] The data is this table is a realistic mockup. Actual datasets were sometimes simpler, sometimes more complicated than this illustrative example.

determined in this spreadsheet). The set of assumptions for this run is listed at the head of the spreadsheet, and of course, these values can be changed at will. Columns B and C list the expected value and the standard deviation of demand (which would come from the forecasting module). Column G records the minor setup times for these parts. (Note: for simplicity here it is assumed that going from part i to part j is the same as going from part j to part i within a family. Also note that when a setup time of 0.1 hours -- or 6 minutes -- appears, it means that there is no actual setup. The part is identical at the CNF stage to the one before it. Even though the specific sequence has yet to be determined, when parts shared the same forming setup code, we knew that only one part required even a minor setup. Others required only minor changes, such as brief paper work and changing the bin which caught the parts as they came off the machine. The values in Column G are representative of minor setup times *after* the actual sequence has been determined.)

To understand the development of a periodic schedule, focus attention on one CNF and one of the wire sizes assigned to it. Recall that F-S assumed a setup cost of $40 per hour. While it was clear that this was a soft number, it served as a starting point for the analysis. Using this setup cost and only minor setup times for each part, the spreadsheet computed the EOQ for each part, and translated the EOQ into a time supply (in weeks). The following data were required. We suppress the part subscript i for clarity:

C = setup cost per hour (in $)

m = minor setup time (assumed to be between 0.1 and 3 hours for all parts initially)

M = major setup time (assumed to be 10 hours for all parts initially)

r = carrying charge

v = unit cost of the part

D = annual demand in units for the part

The EOQ was then

$$EOQ = \sqrt{\frac{2CmD}{rv}}$$

And the desired time supply was EOQ$_i$(50)/D_i, for part *i* assuming a 50-week year. The result, as can be seen from column (J) in Table 2, was a set of time supplies for the parts in each family.

Following Atkins and Iyogun (1988), we then chose the part with the smallest time supply, which was the first part in the first family. (When working with the second family, this part would be the last part in that family.) A small portion, α_1, of the major setup cost was allocated to the first part of the first family. The basic idea was to eventually allocate the entire major setup cost to several of the parts in the family. (Recall that this cost was incurred whenever a changeover from one family to another was performed.) If all the cost were allocated to one part, its EOQ would increase significantly, thereby increasing its time supply. It would then be likely that this part would be produced less frequently than every cycle, and the setup cost (and time) would be incorrect. Hence, the algorithm calls for allocating the major setup cost to the most frequently produced parts.

Allocating α_i to a part increases its time supply, eventually to the point that it matches the time supply of the second most frequently produced part. For the first family in Table 2, the time supply of the first part would increase as α_1 increased, until it equaled 5.42, the time supply of the last part in that family. We then allocated α_1 and α_{13} to parts 1 and 13 until their time supplies equaled 6.73, the time supply of the second part in that family (part number 4). This process continued until the entire major setup was allocated -- that is, $\sum_i a_i = 1$.

If we express the total setup cost for an item as ($\alpha M + m$)C, it is easy to solve for α_i (including now the product subscript, *i*):

$$a_i = \frac{T_i^2 r_i v_i D_i}{2CM_i} - \frac{m_i}{M_i}$$

where T_i is the time supply, EOQ$_i$./D_i. This expression can be used with Goal Seek on EXCEL (Microsoft spreadsheet program) to very quickly find all values of α_i.

Once the α_i were computed, each part was assigned its minor setup and some portion (perhaps zero) of the major setup. A new EOQ and time supply were then computed. To lend order to the schedule, all time supplies were rounded to powers-of-2 multiples of the minimum time supply, which we called the “base period”. (See column (Q) of Table 2.) These values could then be adjusted by the user to account for time supplies that were greater than one year, for instance. Powers-of-2 multiples created a greater sense of order, allowing regular preventive maintenance, communication with non-automotive customers about production schedules, and easier scheduling of operators.

This process was repeated for each family on the CNF. If the base periods of different families on a CNF were different, the spreadsheet allowed for a fast search for the minimum cost, common base period across families. Otherwise, some parts would be on 6 week rotations, some on 9 weeks, some on 10 weeks, etc. A search might find a solution for this example with a base rotation of 6 weeks and all other cycles converted into powers-of-2 multiples of 6. Ten-week cycles would become 12; twenty-week cycles would become 24; and so on. Therefore, the machine would be scheduled on a regular cycle for all parts assigned to it. Our experience is that the small modifications of the base period required to achieve a common base period for all families on a CNF produce very small increases in cost. This observation is consistent with the fact that in the simple EOQ, relatively large deviations from the optimal lot size have a small influence on total costs.

Setting the Actual Sequence. Sequence dependent setups within and across families imply that one should solve a traveling salesman problem to minimize setup time. However, because finding all the possible setup times among parts would take an extremely long time, we worked with information from part numbers and forming setup codes to establish a very rough solution to the exact sequence of production. F-S engineers could then modify the sequence with their additional knowledge of the production process. In every case, however, as indicated in Table 2, we took advantage of parts that were identical at the CNF stage, placing them in order. Thus, relevant setup times are shown to be 0.1 hours.

Some parts were scheduled for production on alternate cycles. These were divided between the first and third, and the second and fourth, cycles to balance the load in every cycle as much as possible.

Checking Capacity. The remaining portion of the spreadsheet translated the cycle time (column Q) and the batch size (column R) into annual production (as a simple check), setups per year (column T), total time in setups (column U), total processing time (column V) and total annual capacity used (column W). For this illustrative example the assumed capacity was 6,000 hours and capacity used for all three families was 4,983 hours, before picking the same base period for all families.

The available capacity figure was based on number of shifts per year, reduced by time required for maintenance, prototype production, and unscheduled breakdowns. F-S suggested reserving more than 20% of available capacity for these activities. Our numbers were based on a conservative approach which would increase inventories over the optimal, but would provide a capacity buffer. We also checked capacity in each cycle to be sure there was available idle time for samples, demand surges, and so on.

If there was excess capacity after completing this schedule, we decreased the hourly setup cost and ran the algorithm again, thereby increasing the frequency of production cycles. More setup time would be consumed, but inventories would be reduced. If there was not enough capacity, the hourly setup cost was increased, decreasing the frequency of production cycles and setups.

Setting Order-up-to Levels. The final step on the spreadsheet was to establish order-up-to levels for each part. Let

$\hat{x}_{R+L}$ = expected demand over the review period (R, or the cycle length for this part) plus the lead time (L) to process the parts,

s_{R+L} = standard deviation of demand over the review period (R, or the cycle length for this part) plus the lead time (L) to process the parts,

k = safety factor

Then, the order-up-to level, S_i, for part i, was

$$S_i = \hat{x}_{R+L} + k\, s_{R+L}$$

where k was based on the service or cost-based performance objective. In this case, a fill rate objective was used. (The fill rate is the percent demand met from the shelf.)

7. DO-IT: A Software Version of the "Big Board"

To maximize the likelihood of implementation, the periodic production system discussed above was prepared as a software package in a form that matched as closely as possible F-S's familiar Big Board scheduling procedure. In place of the magnets on the Big Board and the discussion that took place in front of the board, an interactive software package called DO-IT was prepared to follow on from the spreadsheet program.

Consider Table 3. The CNF machines were classified into nine types, of which one was called the "half-inch" type. There were eight machines of this type. Table 3 deals with these eight machines only and the part numbers that were allocated to each.[5] The first half-inch machine (1S-2) was used for just two parts, #25671 and #38365, and these were both to be produced on a 12-week cycle. Batch sizes were large and run times were long. ("Run Time" in Table 3 includes only processing time for this prototype version. Setup times are added in Figure 2.) Using this PPS for machine 1S-2 occupies a total of 19.75 weeks in the year. Much capacity was available for samples or rush orders that had not been scheduled.

Machine 20S-2 was scheduled to make 35 parts and the spreadsheet program yielded the cycle times indicated in column G of Table 3. Most of the parts were to be made on a 12-week cycle, some on a 24-week cycle, and a few on a 48-week cycle. Total production time on this machine was 21.16 weeks in the year, leaving lots of regulated slack time. Alternatively, production scheduling managers could lower the hourly setup cost as a means to increase the frequency of production.

The information in Table 3 was read into program DO-IT for interactive and graphical presentation of the periodic production schedule. Figure 2 converts the run time (column H of Table 3) into diagrammatic form and adds setup time, using a 48-week time line as one year, since the base cycle time (rotation period) was 6 weeks for all half-inch machines. For example, machine 1S-2 produced only two parts, numbers 25671 and 38365, and they were both produced on a 12-week cycle. The solid black bars indicate when each part was to be set up and produced and the "overall" line combines all activity on that machine. For machine 20S-2 there were 35 parts to be produced, variously on 12-

---

[5] The batch sizes in this table are based on actual data. Only the part numbers and the forming codes have been disguised.

, 24- and 48-weekly cycles. The diagram shows when a part was to be setup and how long it would occupy the machine. The “overall” line shows the more complicated scheduling required on this machine. Note that the third machine in the half-inch class, namely 25S1 was hardly used at all. This machine was due to be retired shortly and was therefore being phased out.

The graphical presentation of the PPS highlighted the organized slack time for compulsory maintenance, unpredictable orders, etc. Those responsible for machine maintenance were especially impressed with this feature since they argued that if a machine was not properly maintained and suddenly broke down, it would be down for several months. Managers who interacted with distributors were also pleased because they could go to their customers with the production schedule and suggest that orders be placed to match the schedule. Discounts could be offered to customers willing to adjust their ordering accordingly.

The DO-IT software package also translated the diagrammatic representation of PPS into manpower requirements. Key assumptions here were (a) that each setup hour required one manpower hour, and (b) that during batch run time one operator could monitor maybe three machines at a time. Some parts required more on line supervision than others but Figure 3 used these two assumptions for all parts on all machines. If the periodic schedules indicated in Figure 2 were put into effect, the manpower requirements after using machine 1S-2 were as shown in the first chart of Figure 3 under the caption “After 1-machine”. During setup time for part #25671 one man-hour was necessary and then during batch running time 1/3 man-hour was needed. When the second part, #38365, was setup one man-hour was needed and during batch running time for this second part, 1/3 man-hour was needed. Hence, the simple manpower chart after one machine. Of course, the order of these charts is arbitrary, and they are presented interactively to allow the scheduler to consider changes. When all eight half-inch machines have been scheduled as shown in Figure 2, the overall manpower requirements are as shown in the last of the eight diagrams of Figure 3.

8. Implementation and Lessons Learned

This paper has addressed the developed of systems designed to lend order to a firm that faced daily scheduling headaches. The systems were based on theoretical academic research, but were adjusted for the real problem, with all its complexity. F-S’s real problems encompassed the gamut of managerial issues. They were not simply inventory problems. There were issues about the appropriate marketplace for F-S. Should it continue to service distributors as well as automotive customers? Was it possible to provide the flexibility and customization, required by the automotive market (along with the low cost and delivery performance) while at the same time providing standard fasteners to the distributor market?

When the focus turned to reducing setups at the CNF machines, one option was to purchase new equipment that required vastly reduced setup times -- in the neighborhood of one-half hour even for major setups (wire-size changeovers). However, such machines cost close to a million dollars and by all measurements, F-S had enough machine capacity to handle its load. Operator availability was not so certain. Thus the immediate need was for a production scheduling system that recognized the current bounds of machine availability and machine operators. The periodic replenishment system devised for F-S was feasible, in that it could be accomplished in the context of the complete problem to be solved. It was understandable, largely because the schedulers' experience with the Big Board was an excellent base for discussing a strictly periodic production schedule. It helped focus activities on the part of the schedulers. For example, the search for adequate data for the change-over matrices was given a higher priority, and the creation of families of parts was taken very seriously.

The software packages (the PPS spreadsheet and DO-IT) offer a useful starting point for final scheduling, but one important task remains to be done. The realities of manpower availability (e.g., there are usually more operators available during the day shift than the night shift) mean that schedulers will not want all products to be started at the beginning of each cycle. It should also be remembered that Table 3, Figure 2 and Figure 3 are dealing with only one of the nine CNF machine classes. For each of the other CNF classes there are equivalent Tables and Figures. Thus when the total manpower implications of a particular PPS system are examined for the factory as a whole, it is inevitable that changes need to be made in the ordering and placement of parts within cycles, so that manpower requirements fit manpower availability. Our systems are flexible, easy to use, and fast. Thus, the inevitable changes can be accommodated relatively easily.

Managers expressed satisfaction with the ideas in the scheduling system, as well as with potential results. The proactive, orderly schedule promised multiple benefits for maintenance, customer relationships, operator scheduling, and a greater sense of control in the factory. An initial comparison of our system with the current procedures indicated that savings of 39% in setup time and 25% in inventory holding cost could be achieved. This comparison was necessarily rough because of limited data on current procedures. However, it was evident that the “Big Board” had its limits. Some parts we compared were scheduled for production identically as our system recommended; but others were scheduled too frequently, suggesting that the 39% savings in setup time could be realized.

At this writing, however, implementation is not complete. F-S is pursuing a new strategic initiative to eliminate up to 1000 slow moving parts from its catalog. Discussions with customers will determine whether these parts are truly necessary for customer relationships. One supposes that the true cost of producing extremely slow moving parts is not covered by the price F-S is able to charge. Hence, many parts may be dropped. To pursue a proactive scheduling system when the number of parts to be scheduled is so uncertain is not a wise use of time. It is our hope that F-S will finalize implementation of this, or a modified, system as the part count stabilizes.

References

Atkins, D., & Iyogun, P. (1988). Periodic Versus 'Can-Order' Policies for Coordinated Multi-Item Inventory Systems. Management Science, 34(6), 791-796.

Benjaafar, A., & Sheikhzadeh, M. (1994). On the Effect of Setup Minimization in Scheduling Policies. Department of Mechanical Engineering, University of Minnesota

Bourland, K., & Yano, C. (1994). The Strategic Use of Capacity Slack in the Economic Lot Scheduling Problem with Random Demand. Management Science, 40(12), 1690-1704.

Dobson, G. (1987). The Economic Lot-Scheduling Problem: Achieving Feasibility Using Time-Varying Lot Sizes. Operations Research, 35, 764-771.

Dobson, G. (1992). The Cyclic Lot Scheduling Problem with Sequence Dependent Setups. Operations Research, 40, 736-749.

Elmaghraby, S. (1978). The Economic Lot Scheduling Problem (ELSP): Review and Extensions. Management Science, 24(6), 587-598.

Federgruen, A., Groenevelt, H., & Tijms, H. C. (1984). Coordinated Replenishments in a Multi-Item Inventory System with Compound Poisson Demand. Management Science, 30(3), 344-357.

Federgruen, A., & Katalan, Z. (1994a). Make-to-Stock or Make-to-Order: That is the Question; Novel Answers to an Ancient Debate. The Wharton School, University of Pennsylvania.

Federgruen, A., & Katalan, Z. (1994b). The Stochastic Economic Lot Scheduling Problem: Cyclical Base-Stock Policies with Idle Times. The Wharton School, University of Pennsylvania.

Gallego, G. (1990). Scheduling the Production of Several Items with Random Demands in a Single Facility. Management Science, 36(12), 1579-1592.

Karmarkar, U. (1987). Lot Sizes, Lead Times and In-Process Inventories. Management Science, 33(3), 409-418.

Karmarkar, U. S., Kekre, S., Kekre, S., & Freeman, S. (1985). Lot Sizing and Lead Time Performance in a Manufacturing Cell. Interfaces, 15, 1-9.

Makridakis, S., Wheelwright, S. C. & McGee, V. E. (1983). Forecasting: Methods & Applications. Second Edition. New York: John Wiley & Sons.

Roundy, R. (1989). Rounding Off to Powers of Two in Continuous Relaxations of Capacitated Lot Sizing Problems. Management Science, 35, 1433-1442.

Silver, E. A., & Peterson, R. (1985). Decision Systems for Inventory Management and Production Planning. New York: John Wiley & Sons.

Tang, C. S. (1990). Scheduling Batches on Parallel Machines with Major and Minor Set-ups. European Journal of Operational Research, 46, 28-37.

Weng, Z. K. (1994). Manufacturing Lead Times in Make-to-Order Manufacturing Systems: A Model and Insights. School of Management, Georgia Institute of Technology

Zipkin, P. H. (1991). Computing Optimal Lot Sizes in the Economic Lot Scheduling Problem. Operations Research, 39(1), 56-63.

Figure 1 The first 4 of 57 Parts with Shipment Activity in each of the 45 Months

Here the actual shipments are shown as solid lines. The fitted 3-harmonic model (sine waves with 3-month, 6-month and 12-month periods) and a six-month forecast are shown as dashed lines.

Figure 2 PPS Output from the DO-IT Package

For each of the eight CNF machines in the “half-inch” class, the periodic production schedule is displayed graphically over a 48-week “year”. For example, on machine 1S-2 only two parts are scheduled, both on a 12-week cycle, and the “overall” line indicates total commitment on this machine over the 48-week year.

Figure 3 Manpower Requirements from the DO-IT Package

If each of the eight CNF machines produced according to the PPS indicated in Figure 2 then the manpower requirements are as indicated here. Assumptions are that each setup hour requires one man-hour and each production run hour requires 1/3 man-hour. The order of running the eight machines is arbitrary and the arrangements of the cyclic productions can be modified to control the manpower levels as needed.