Production and Sales Planning with Limited Shared Tooling at the Key Operation

Gerald G. Brown; Arthur M. Geoffrion; Gordon H. Bradley

Management Science, Vol. 27, No. 3 (Mar., 1981), 247-259.

PRODUCTION AND SALES PLANNING WITH LIMITED SHARED TOOLING AT THE KEY OPERATION*

GERALD G. BROWN, † ARTHUR M. GEOFFRION‡ AND GORDON H. BRADLEY†

The focus of this paper is multiperiod production and sales planning when there is a single dominant production operation for which tooling (dies, molds, etc.) can be shared among parts and is limited in availability. Our interest in such problems grew out of management issues confronting an injection molding manufacturer of plastic pipes and fittings for the building and chemical industries, but similar problems abound in the manufacture of many other cast, extruded, molded, pressed, or stamped products. We describe the development and successful application of a planning model and an associated computational approach for this class of problems.

The problem is modeled as a mixed integer linear program. Lagrangean relaxation is applied so as to exploit the availability of highly efficient techniques for minimum cost network flow problems and for single-item dynamic lot-sizing type problems. For the practical application at hand, provably good solutions are routinely being obtained in modest computing time to problems far beyond the capabilities of available mathematical programming systems.

(INVENTORY/PRODUCTION APPLICATIONS; PRODUCTION—FLOW SHOP; INTEGER PROGRAMMING APPLICATIONS)

1. Introduction

This paper addresses production and sales planning in a seasonal industry with a single principal production operation for which tooling can be shared among parts and is limited in availability.

The specific context of our experience is the production of injection molded plastic pipes and fittings for the building and chemical industries. The principal production operation is injection molding and the tooling consists of mold bases used to adapt the injection molding machines to the molds proper. Mold bases typically require 4-6 calendar months to obtain at a cost which can approach the cost of the molding machine itself, so their availability is limited and good utilization is important.

Other possible domains of application include production facilities based on casting, molding, stamping, extrusion, or pressing of finished or nearly finished products. Dies and molds and associated adaptive tooling are usually expensive and often designed for use with more than one end product. Machine tooling with elaborate jigs and fixtures constitutes another large area of potential application. Although there can be just one principal operation to be modeled, there may also be preparatory and final finishing operations.

An informal statement of the problem treated is as follows. A facility produces many different parts (products), each by a principal operation calling for a specific type of tool and any one of a number of machines compatible with the tool. Machines

247

are aggregated into machine groups and tools into tool types. Production and sales are to be planned for each part over a multiperiod horizon (typically monthly for a full year):

Determine
* how much of each part to produce in each time period
* how much of each part to sell in each time period
* how much of each part to carry forward as inventory from each time period into the next
* a tool/machine assignment schedule specifying, for each time period, the number of days of production of each tool type in conjunction with each compatible machine group

so as to satisfy all necessary constraints
* limited availability of tools in each time period
* limited availability of machines in each time period
* tool/machine compatibility restrictions
* for each part in each time period, sales cannot exceed forecast demand

and so as to satisfy desired managerial policy constraints
* for each part in each time period, sales must exceed a certain fraction of demand stipulated by management
* for each part, the ending inventory at the conclusion of the planning horizon must take on a stipulated value
* no planned backlogging (unfilled demand is lost)

in such a manner as to maximize total profits over all parts for the duration of the planning horizon, calculated according to
* incremental net profit contribution per unit sold
* less variable operating costs associated with production (by tool type and machine group)
* less fixed costs associated with production (by part, for each period with positive production)
* less part-specific inventory holding costs.

The problem as stated has elements in common with many familiar dynamic planning and resource allocation problems. It is more detailed than most seasonal planning problems in that discrete fixed costs are included and no aggregation is necessary over parts, yet it stops short of encompassing detailed scheduling because other aggregations are employed (tools → tool types, machines → machine groups, time → time periods). Related production planning and scheduling problems in the molding industry can be found in [3], [7], [9].

A proper mathematical formulation as a mixed integer linear program is given in §2. The next section presents a solution approach based on a particularly attractive Lagrangean relaxation and sketches our full scale computational implementation. §4 describes computational experience with the injection molding application mentioned earlier. For this application, solutions well within 2% of optimum are routinely produced in about 3 minutes of IBM 370/168 time for mixed integer linear programs on the order of 12,000 binary variables, 40,000 continuous variables, and 26,000 constraints.

2. Problem Formulation

This section formally defines and discusses the problem as a mixed integer linear program.

We adopt the following notation. Essential use is made of the concept of a *standard day*, which is a part-specific measure of quantity. It is, for a given part, the quantity that would be produced in one calendar day if a tool of the required type were operating normally on any compatible machine.

Indices

i indexes parts

j indexes tool types

k indexes machine groups

t indexes time periods, $t = 1, \ldots, T$

$\mathcal{I}(j)$ index set of the parts requiring tool type j (these index sets must be mutually exclusive and exhaustive)

$\mathcal{K}(j)$ index set of the machine groups compatible with tool type j

Given Data

a_{jt} days of availability of type j tools during period t

b_{kt} days of availability of machine group k during period t

c_{jkt} variable daily operating cost during period t of tool type j on machine group k, for compatible combinations of j and k

d_{it} demand forecast for part i in period t, in standard days

f_{it} fixed cost associated with the production of part i in period t

h_{it} cost of holding one standard day of part i for the duration of period t

I_0 initial inventory in period 1 of part i, in standard days (must be $\geqslant 0$)

I_{iT} terminal inventory desired for part i, in standard days, at the conclusion of the last period (must be $\geqslant 0$)

m_{it} maximum possible production of part i in period t, in standard days

p_{it} profit contribution associated with selling one standard day's worth of part i in period t, exclusive of the other costs included in the model

α_{it} minimum fraction of d_{it} which must be satisfied as a matter of marketing policy $(0 \leqslant \alpha_{it} \leqslant 1)$

Decision Variables

I_{it} planned inventory of part i at the end of period t, in standard days $(1 \leqslant t < T)$

S_{it} planned sales of part i in period t, in standard days

V_{jt} planned production days for tool type j in period t

W_{jkt} planned production days for tool type j on machine group k during period t, for compatible combinations of j and k

X_{it} planned production of part i during period t, in standard days

Y_{it} a binary variable indicating whether or not part i is produced during period t

Using this notation, the problem can be posed as the following mixed integer linear program.

$$\underset{I, S, V, W, X, Y}{\text{MINIMIZE}} - \sum_i \sum_t p_{it} S_{it} + \sum_j \sum_k \sum_t c_{jkt} W_{jkt}$$
$$+ \sum_i \sum_t h_{it} (I_{i,t-1} + I_{i,t})/2 + \sum_i \sum_t f_{it} Y_{it} \tag{1}$$

subject to

$$\sum_{k \in \mathcal{K}(j)} W_{jkt} = V_{jt}, \qquad \text{all} \quad jt, \qquad (2)$$

$$V_{jt} = \sum_{i \in \mathcal{I}(j)} X_{it}, \qquad \text{all} \quad jt, \qquad (3)$$

$$\sum_{j} W_{jkt} \leqslant b_{kt}, \qquad \text{all} \quad kt, \qquad (4)$$

$$I_{it} = I_{i,t-1} + X_{it} - S_{it}, \qquad \text{all} \quad it, \qquad (5)$$

$$\alpha_{it} d_{it} \leqslant S_{it} \leqslant d_{it}, \qquad \text{all} \quad it, \qquad (6)$$

$$0 \leqslant X_{it} \leqslant m_{it} Y_{it}, \qquad \text{all} \quad it, \qquad (7)$$

$$0 \leqslant V_{jt} \leqslant a_{jt}, \qquad \text{all} \quad jt, \qquad (8)$$

$$I_{it} \geqslant 0, \qquad \text{all} \quad it \, (1 \leqslant t < T), \qquad (9)$$

$$W_{jkt} \geqslant 0, \qquad \text{all} \quad jkt, \qquad (10)$$

$$Y_{it} = 0 \text{ or } 1, \qquad \text{all} \quad it. \qquad (11)$$

It is understood that any summations or constraint enumerations involving j and k together will run only over compatible combinations of j and k.

The objective function is the negative of total profit over the duration of the planning horizon. It is the negative of profit contribution associated with sales over the planning horizon, plus: machine operating costs, inventory holding costs (applied to a simple 2-point estimate of the average inventory level of each part in each period), and fixed costs.

Constraints (2) and (3) serve to interrelate machines, tools, and parts. Constraints (4) and (8) respectively enforce the availability limitations on machines and tools. Constraint (5) defines ending inventories in the standard way. Constraint (6) requires the planned sales to be between forecast demand and some specified fraction thereof. Constraint (7) keeps production within possible limits and also forces Y_{it} to be 1 when X_{it} is positive. Constraint (9) specifies that there be no planned backlogging. Constraints (10) and (11) require no comment.

Some additional comments are appropriate.

1. There can be more than one tool (resp. machine) available of a given type (resp. group). Such census information, along with downtime estimates, determines the a_{jt} (resp. b_{kt}) coefficients.

2. The index sets $\mathcal{I}(\cdot)$ specify a unique tool type for each part. Typically these index sets will not be singletons, so that a number of parts compete for the same tooling.

3. The fixed cost coefficients f_{it} are perhaps best interpreted as surrogates for detailed setup costs; f_{it} is incurred in the model when part i is produced in period t irrespective of whether this requires a tool changeover (part i's tool type may be common to the part run previously), and irrespective of whether more than one machine must simultaneously make part i in order to achieve the planned production X_{it}. To specify setup costs at a greater level of detail would require a major revision of the model that would transport it from the realm of planning to the realm of detailed scheduling. Yet setup costs cannot be ignored entirely because this tends to cause some of every part to be produced during every period, a situation clearly unacceptable from the production viewpoint. Our solution is to take the f_{it}'s as empirically weighted average setup costs.

4. The terminal inventory level is the only significant terminal condition of the model. It is known from studies of related models (e.g., [8]) that I_{iT} can have a significant effect on solution quality, and hence that it should be set at some estimate of what the optimal inventory would be for part i at the end of period T in a similar model with many more periods. In practice, this means drawing on historical operating experience, insights obtained previously with the help of the model, and managerial judgment.

5. The maximum possible production m_{it} is the smaller of two limits: the physical limit imposed by full utilization of all available tooling and machines, and the limit on the amount of production that could be absorbed considering demand over the planning horizon, specified terminal inventory, and initial inventory. The second limit is redundant in view of (5) and (6); it is incorporated only to tighten the standard LP relaxation to be discussed later. The first limit, however, may well be binding (as it was in our application owing to part-specific limitations not subsumed by (4) and (8)).

6. The policy parameters α_{it} can be used, for instance, to maintain desired product line breadth when profit considerations alone would tend to narrow excessively the range of products produced.

7. There are several alternative problem representations which are equivalent to (1)–(11) and just as natural. Some of these are more compact; for example, the V_{jt} variables can be eliminated. The representation given is designed, principally through the introduction of the V_{jt} variables and the use of the standard day concept, to render the algorithmic manipulations of the next section as transparent as possible.

Before turning to the question of solving (1)–(11), we pause to define

$$h'_{it} = \tfrac{1}{2}(h_{it} + h_{i,t+1}) \quad \text{for all} \quad i \quad \text{and} \quad 1 \leqslant t < T,$$

$$H = \tfrac{1}{2} \sum_i (h_{i1} I_{i0} + h_{iT} I_{iT})$$

so that the third term of (1) can be expressed equivalently as

$$\sum_i \sum_{t=1}^{T-1} h'_{it} I_{it} + H.$$

3.   Solution Via Lagrangean Relaxation

For the practical application at hand, there are approximately 1000 parts, 90 tool types, 15 machine groups, 12 time periods, and 480 compatible combinations of tool types and machine groups. This leads to dimensions of approximately

40,000   continuous variables (I, S, V, W, X)

12,000   integer variables (Y)

26,000   constraints of type   (2), (3), (4), (5), (7).

Problems of this magnitude are generally considered to be far beyond the current state-of-the-art of general mixed integer linear programming. Consequently, we sought a way to exploit the special structure of the problem.

The key was to recognize that (1)–(11) is a network flow problem with fixed charges for certain arcs, and that Lagrangean relaxation [4] [5] [10] with respect to (3) is an attractive way to generate lower bounds on the optimal value of (1)–(11). The resulting

Lagrangean subproblem separates into as many independent simple transportation problems in the W variables as there are time periods, and as many independent dynamic single-item lot-size problems as there are parts. The original monolith is thereby decomposed into manageable fragments.

It is easy to see that (1)–(11) is a fixed charge minimum cost network flow problem. See Figure 1 for an example with 3 parts, 2 tool types, 3 machine groups, 3 time periods, $\mathcal{I}(1) = \{1\}$, $\mathcal{I}(2) = \{2, 3\}$, $\mathcal{K}(1) = \{1, 2\}$, and $\mathcal{K}(2) = \{2, 3\}$. The *notational conventions* followed in Figure 1 and subsequent figures are: the minimand term corresponding to each arc is written *over* the arc, the upper capacity limit of each arc is written *under* it (omission implies infinite capacity), the constraint on the net outflow of each node is written under it (omission implies $= 0$, or strict conservation), and a fixed cost arc is drawn as a dashed line, with the amount of the fixed charge incurred for its use given as the first of the two annotations written over it. The curved arcs between the part nodes are not annotated for lack of room; the typical arc is:

Now consider the Lagrangean relaxation of (1)–(11) using multipliers λ_{jt} for (3): drop (3) and replace the objective function (1) by

$$\underset{I, S, V, W, X, Y}{\text{MINIMIZE}} - \sum_i \sum_t p_{it} S_{it} + \sum_j \sum_k \sum_t c_{jkt} W_{jkt} + \sum_i \sum_{t=1}^{T-1} h'_{it} I_{it}$$

$$+ \sum_i \sum_t f_{it} Y_{it} + \sum_j \sum_t \lambda_{jt} \left(\sum_{i \in \mathcal{I}(j)} X_{it} - V_{jt} \right) + H. \qquad (1 - \lambda)$$
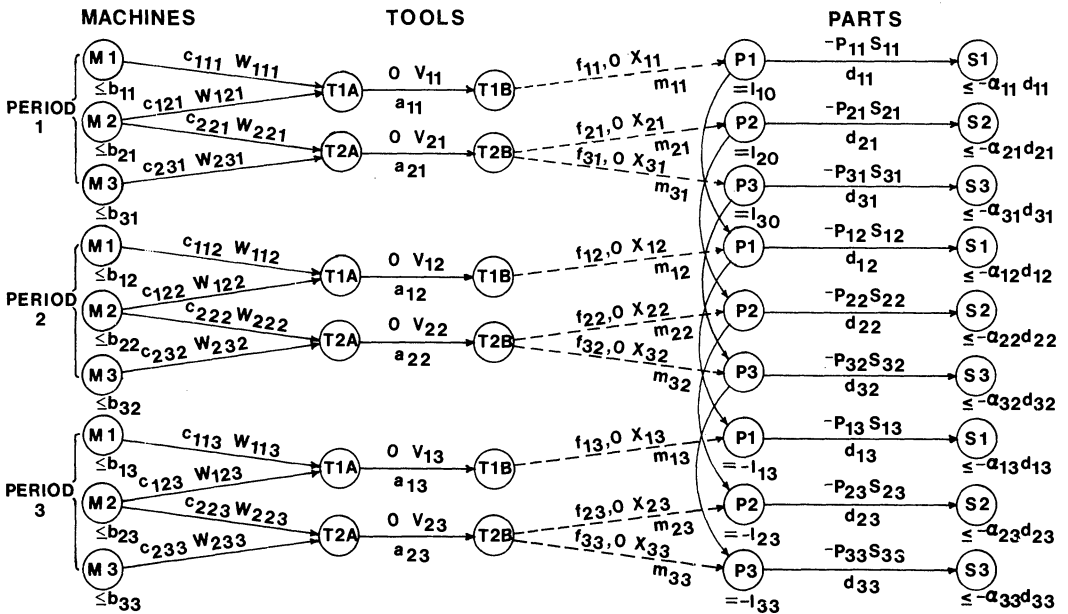

FIGURE 1.   Sample Equivalent Fixed-Charge Network for Problem (1)–(11).

Dropping (3) causes the variables $\{V, W\}$ to become completely decoupled from the variables $\{I, S, X, Y\}$. In fact, the decoupling extends even farther: the first set of variables is decoupled over t and the second set over i. Thus the indicated Lagrangean relaxation of (1)–(11) yields the following independent problems: for each t,

$$\underset{V, W}{\text{MINIMIZE}} \sum_j \sum_k \sum_t c_{jkt} W_{jkt} - \sum_j \sum_t \lambda_{jt} V_{jt}$$

subject to   (2), (4), (8), (10)   for fixed t (R^t)

and for each i,

$$\underset{I, S, X, Y}{\text{MINIMIZE}} - \sum_t p_{it} S_{it} + \sum_{t=1}^{T-1} h'_{it} I_{it} + \sum_t f_{it} Y_{it} + \sum_t \lambda_{j(i)t} X_{it}$$

subject to   (5), (6), (7), (9), (11)   for fixed i (R_i)

where $j(i)$ is the index of the tool type required by part i.

Using $v(\cdot)$ to stand for the optimal value of problem $(\cdot)$, the optimal value of the full Lagrangean relaxation can be expressed as

$$\sum_t v(R^t) + \sum_i v(R_i) + H. \tag{12}$$

This quantity is a lower bound on the optimal value of (1)–(11) for any choice of λ.

Figures 2 and 3 portray (R^t) and (R_i) for the example illustrated in Figure 1. The same notational conventions apply. Observe that the network for (R^t) has been simplified by using (2) to eliminate the V-variables, and that the network for (R_i) has been simplified by tying the pure source nodes together. Notice that Lagrangean relaxation amounts to "scissoring" the B-type tool nodes in Figure 1 and placing a penalty or premium on the incident arc flows.

How good a bound is (12)? The answer depends on the choice of λ. It is known that a good choice is the optimal dual vector associated with (3) in the standard *LP* relaxation of (1)–(11), in which (11) is relaxed to $0 \leqslant Y_{it} \leqslant 1$ for all *it*. This choice of λ yields a value for (12) that is at least as high as the optimal value of the standard *LP* relaxation itself. Moreover, strict improvement is likely because the Integrality Property defined in [5] does not hold. Superior choices of λ may well exist, and might be
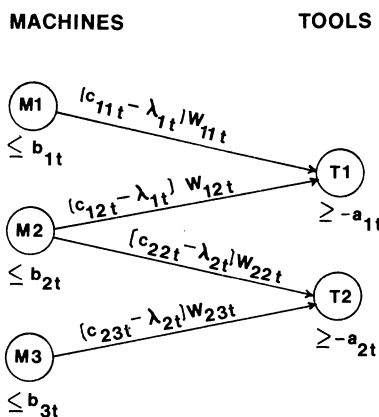
MACHINES   **TOOLS**
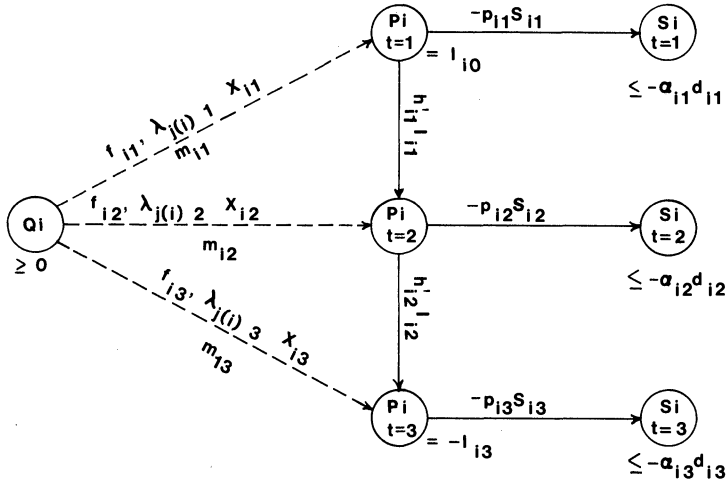

FIGURE 2.   Sample Equivalent Network for (R^t).

FIGURE 3.   Sample Equivalent Fixed Charge Network for (R_i).

found by auxiliary calculations other than the standard LP relaxation, but we do not pursue such possibilities here.[1]

One can build a branch-and-bound procedure around this Lagrangean relaxation. For the industrial application which stimulated this work, however, it has proven sufficient to generate a feasible solution to (1)–(11) based on the Lagrangean solution. The objective value of this solution has consistently been sufficiently close to the lower bound from Lagrangean relaxation that no further refinement has been needed.

A formal description of the solution procedure is now presented.

Step 1.   Solve the standard LP relaxation of (1)–(11) via an equivalent capacitated network formulation. Denote the associated dual variables for (3) by $\bar{\lambda}_{jt}$.

Step 2.   Form the Lagrangean relaxation of (1)–(11) with respect to (3) using the values of $\bar{\lambda}_{jt}$ from Step 1. Solve it via the independent subproblems (R^t) and (R_i). Denote the combined optimal solution to the full Lagrangean relaxation by $(I^0, S^0, V^0, W^0, X^0, Y^0)$ and its optimal objective value by LB.

Step 3.   Form the standard LP relaxation of (1)–(11) with f_{it} revised to 0 if $Y_{it}^0 = 1$ and augmented by a large positive constant if $Y_{it}^0 = 0$. Solve for an optimal solution (I', S', V', W', X', Y') via an equivalent capacitated network formulation. Let Y'' be Y' with all fractional components rounded up to unity. Form the revised solution with Y' replaced by Y'' and denote its objective value under (1) by UB. This solution is feasible in (1)–(11) and is within $UB-LB$ of being optimal. Stop.

We comment on the nature of the problems needing to be solved at each step.

Steps 1 and 3 each yield an equivalent capacitated minimum cost network flow problem, and hence can be executed efficiently using one of the modern primal simplex codes developed for such problems (e.g., [2]). It is evident that, when (11) is relaxed to $0 \leqslant Y_{it} \leqslant 1$ for all it, the right-hand inequality of (7) must hold with strict equality for all it at optimality. This permits elimination of Y, whereupon (7) and (11) can be replaced by

$$0 \leqslant X_{it} \leqslant m_{it}, \quad \text{all } it \tag{11}'$$

---

[1]See [4], [5], [10], and the new convergent procedure in [6].

and the $\sum_i \sum_t f_{it} Y_{it}$ term of (1) is replaced by

$$\sum_i \sum_t \frac{f_{it}}{m_{it}} X_{it}.$$

In other words, the standard LP relaxation of (1)–(11) has no fixed cost feature.

Problem (R') can be converted easily to a simple transportation problem. However, using LP duality theory it can be shown that the values of W_{jkt} found at Step 1 are also optimal in (R') for all t. Thus no work at all need be performed in connection with these subproblems.

Problem (R_i) is a mixed integer linear program with T equations, $3T - 1$ continuous variables and T binary variables. Tying the S_i nodes of Figure 3 together with the Q_i node and performing a change of variables $S_{it} = S'_{it} + \alpha_{it} d_{it}$ to eliminate the lower bounds on the sales variables yields an equivalent capacitated transshipment problem with fixed charges on T of the arcs. The transshipment problem has $T + 1$ nodes, T capacitated sales arcs (S'_{it}), $T - 1$ inventory arcs (I_{it}) and T capacitated production arcs (X_{it}) with fixed charges. Because only the minimum sales demand must be satisfied, this problem has T more variables than the closely related dynamic lot-size problem presented in [1]. This feature of our model allows "lost sales," which is an important aspect of the planning problem for the application at hand.

4. Application and Computational Results

The model and computational procedure just described have been under development and application since 1977 at plants of R & G Sloane Manufacturing Company of Sun Valley, California. The following identifications and specializations are appropriate.

General Model	Molding Application (main plant)
parts (i)	the top 1000 injection molded fittings (about 92% of all sales volume)
tool types (j)	about 90 types of interchangeable mold bases (total mold base census about 130)
machine groups (k)	about 15 groups of interchangeable injection molding machines (total machine census about 60)
tool/machine compatibility	about 480 jk combinations permissible
time periods (t)	typically the next 12 months
$c_{jkt}, f_{it}, \alpha_{it}, p_{it}$	taken as independent of t

The problem faced by R & G Sloane is a strongly seasonal one. Since the bulk of the company's business involves residential plumbing products, demand peaks along with residential construction in the summer months. Since the peak season demand rate exceeds the available capacity of mold bases and machines, constraints (4) and (8) tend to be binding at that time of year (typically, about 40% of the mold base

constraints and 80% of the machine constraints are binding 5-8 months of the year). Typical relative magnitudes of the major cost categories associated with an optimal solution are:

fixed costs	14.3
inventory carrying	16.9
variable operating	68.8
	100.0

Unfilled demand typically occurs for well under 10% of all parts.

Computational Implementation

A full scale computational implementation has been carried out for this application. The computer programs are in three modules:

1. data extraction and data base definition
2. problem preprocessing and diagnosis
3. optimization and report writing.

Data extraction primarily involves conversion of current production, marketing, and inventory control operating data to the form required by the model. The data base is organized and generated in sections:

· problem parameters and conditions
· machine group descriptions
· mold bases and their machine compatibility
· part descriptions and demand forecasts.

Preprocessing identifies structural and mathematical inconsistencies in the problem posed, and assists in preliminary diagnosis of critical shortages in equipment availability.

The optimization module solves the capacitated pure networks presented in Steps 1 and 3 with a GNET variant (XNET/Depth) [2]; an advanced starting solution is used which assumes high equipment utilization.

The fixed charge problems (R_i) are addressed with an enumeration algorithm that also employs GNET [2]. The overall algorithm applies implicit enumeration using a standard backtrack method to specify the binary variables. For each setting of the binary variables (or "case"), a capacitated transshipment problem is solved with a highly specialized version of GNET/DEPTH. The algorithm has several external parameters that permit tuning for high performance.

Extensive pre-solution analysis is performed to identify periods where production is mandatory (e.g., first period minimum demand greater than initial inventory) or precluded (e.g., large initial inventory). Additional dominance tests reduce the number of cases to be examined (a particularly effective technique involves keeping a running bound on the number of periods that can have positive production in an optimal solution).

Early analysis of this algorithm showed that a small number of problems ($\sim 10\%$) consumed a large fraction ($\sim 75\%$) of the total computer time in Step 2. Further analysis of the time-consuming problems showed that solution trajectories were characterized by frequent incumbent improvements in the initial cases followed by large numbers of cases with little or no improvement. This suggested a modification to

construct very good (but not necessarily optimal) solutions in significantly less time. A single parameter, JUMP, directs the enumeration to skip over a number of cases when the number of cases since the last improvement increases; the enumeration skips the integer part of the number of cases since last improvement divided by JUMP.

Solving Step 2 with the modified algorithm actually yields final solutions at Step 3 that are better than those from the exact algorithm. Although the exact and estimated bounds in Step 2 are nearly equal, the modified algorithm produces solutions with more setups than the exact algorithm. Up to a point, more setups improve the value of the final solution in Step 3. Since the final solution from Step 3 always has more setups than the optimal solutions from Step 2, it is better to allow Step 2 to construct good, but not necessarily optimal, solutions with more setups than to require Step 3 to insert the setups.

The parameter JUMP is a very effective control on the number of setups. As JUMP decreases from the value that yields an exact algorithm (2^T), the number of setups tends to increase. Experimentally, JUMP = 3 has been shown to produce superior final solution values from Step 3 for a wide variety of problems.

The bound produced from Step 2 is a valid lower bound only if the optimal solution is obtained for each (R_i) problem. However, the estimated bound has been repeatedly verified to be very close (less than 1%) to the correct bound. It would be possible to use the modified algorithm to construct solutions for Step 3 and the exact algorithm to construct a valid lower bound; however, we have chosen to utilize routinely just the modified algorithm, with occasional use of the exact algorithm to verify the quality of the estimated bound.

After Step 3, solution reports are presented at several levels of aggregation so as to facilitate managerial interpretation. They display all detailed solution features, estimated opportunity costs for critical mold bases and machines, and an overall analysis of profitability, turnover, and customer service.

Computational Results

Approximately 30 runs per year are carried out. Computational performance has exhibited a high degree of run-to-run stability in terms of solution quality and computer resources consumed.

Table 1 summarizes several aspects of performance for a recent typical run. The main storage requirement was about one megabyte. Notice that the bound produced

TABLE 1

Typical Computational Performance: A Problem with 953 Parts, 92 Tool Types, 16 Machine Groups

	IBM 370/168 CPU Seconds	Normalized Objective Value [a]	Pivots
Preprocessing	10.0		
Optimization			
Step 1 (*LP* Relaxation)	55.0	103.2	37, 933
Step 2 (Lagrangean Relaxation)	25.7	101.6	376, 241
Step 3 (Generate Feasible Solution)	54.2	100.0	36, 954
Report Writing	30.6		
	175.7		

[a]Normalization is relative to the value of the feasible solution obtained at Step 3. The Step 2 value is an estimated bound.

by the Lagrangean relaxation is significantly better than the standard linear programming relaxation bound. Notice also that the time in Step 2 is smaller than what one might expect; the 12-period fixed charge problems were processed in an average of only .027 seconds each (for comparison, the typical time quoted in [1] for exact solution of a proper subclass of these problems of the same size was 0.25 seconds on an IBM 370/158). For this run, 142 (resp. 10) of the 11,436 binary Y variables changed from value 0 (resp. 1) in Step 2 to value 1 (resp. 0) in Step 3. This shows that the solution to the Lagrangean relaxation of Step 2 required but minor adjustment with respect to the fixed charge arcs in order to yield the good feasible solution of Step 3.

More generally, our experience has been that optimization CPU time for similar-sized problems seldom varies more than $\pm 10\%$. Computing time is very nearly proportional to the total number of parts. The final estimated optimality tolerance (which was 1.6% in the Table 1 run) tends to become smaller the more tightly capacitated tool and machine availability is; tolerances in the vicinity of 2/10 of 1% are commonly observed in the most tightly constrained situations. In no case has the tolerance ever exceeded 2%.

5. Conclusion

This paper has demonstrated the practical applicability of a procedure based on Lagrangean relaxation to a significant class of integrated production and sales planning models. The particular way in which this procedure is designed thoroughly exploits the recent major advances made for minimum cost network flow problems. Provably good solutions are routinely being obtained in modest computing time to mixed integer linear programs of a size far beyond the capabilities of generally available mathematical programming systems.

The system is used regularly at R & G Sloane Manufacturing Company. Day-to-day production scheduling is still performed manually, but with the benefit of the system's guidance and predictions of bottlenecks in the future. The integrated nature of the model has made the system valuable as a focal point for coordinating planning activities among the key functional areas of the firm: inventory control, finance, marketing, and production operations. Two specific illustrations are the evaluation of major capital expenditure and interplant equipment transfer opportunities. A recent *Business Week* article [11] featured this application, with the Vice President of Operations quoted as crediting the new system for an increase in total operating profits during a recent year of $500,000.[2]

References

1. BAKER, K. R., DIXON, P., MAGAZINE, M. J. AND SILVER, E. A., "An Algorithm for the Dynamic Lot-Size Problem with Time-Varying Production Capacity Constraints," *Management Sci.* Vol. 24, No. 16 (1978).
2. BRADLEY, G. H., BROWN, G. G. AND GRAVES, G. W., "Design and Implementation of Large Scale Primal Transshipment Algorithms," *Management Sci.*, Vol. 24, No. 1 (1977).

3. CAIE, J., LINDEN, J. AND MAXWELL, W., "Solution of a Single Stage Machine Load Planning Problem," *OMEGA*, Vol. 8, No. 3 (1980).
4. FISHER, M. L., "Lagrangean Relaxation Method for Solving Integer Programming Problems," *Management Sci.*, Vol. 27, No. 1 (1981).
5. GEOFFRION, A. M., "Lagrangean Relaxation for Integer Programming," *Mathematical Programming Study 2*, North-Holland Publishing Co., Amsterdam, 1974.
6. GRAVES, G. W. AND VAN ROY, T. J., "Decomposition for Large-Scale Linear and Mixed Integer Linear Programing," Working Paper, Graduate School of Management, UCLA, November 1979, *Mathematical Programming* (to appear).
7. LOVE, R. R. AND VEMUGANTI, R. R., "The Single-Plant Mold Allocation Problem with Capacity and Changeover Restrictions," *Operations Res.*, Vol. 26, No. 1 (1978).
8. MCCLAIN, J. O. AND THOMAS, J., "Horizon Effects in Aggregate Production Planning with Seasonal Demand," *Management Sci.*, Vol. 23, No. 7 (1977).
9. SALKIN, H. M. AND MORITO, S., "A Search Enumeration Algorithm for a Multiplant, Multiproduct Scheduling Problem," *Proceedings of the Bicentennial Conference on Mathematical Programming*, NBS, Gaithersburg, Maryland, December 1976.
10. SHAPIRO, J. F., "A Survey of Lagrangean Techniques for Discrete Optimization," in *Annals of Discrete Mathematics 5: Discrete Optimization*, P. L. Hammer, E. L. Johnson, and B. H. Korte (eds.), North Holland, Amsterdam, 1979.
11. "'What if' Help for Management," *Business Week*, 21 January 1980, pp. 73–74.

INSTITUTIONAL MEMBERS