



A correlation and heuristic approach for obtaining production sequences requiring a minimum of tool replacements

Patrick R. McMullen*, Mark Clark, David Albritton, John Bell

*College of Business, Department of Management, Auburn University, 415 W. Magnolia, Suite 401,
Auburn, AL 36849-5241, USA*

Received 1 June 2001; received in revised form 1 December 2001

Abstract

This research presents a technique to obtain production sequences requiring minimal tooling replacements, via exploitation of statistical metrics along with a heuristic approach. Tool preservation has always been an important issue, and in many cases, since the production of some products tends to wear the tooling differently than others, the production sequence plays an important role in this preservation. Thus, if tool preservation is a priority then maintaining uniform tool-wear through careful product sequencing should be of concern. In some industries, minimizing tool replacement by increasing its useful life can result in large annual savings for manufacturing firms—due to increased tool life, reduction of unscheduled downtime, and increased flexibility of the machining center. This research presents sequencing techniques that attempt to minimize the number of tool replacements on a single machine over a given period of time (via sequences which have uniformity of tool wear). The sequences are obtained via simulated annealing, with a measurement criterion of three correlation-related statistics. Experimentation indicates that sequences obtained via the presented metrics and simulated annealing provide fewer tooling replacements as compared to more conventional sequencing methods.

Scope and purpose

Scheduling of production has many objectives—minimization of in-process inventory and minimization of idle time are two of the more popular objectives. Scheduling production in a way to minimize tool wear, and subsequent tooling replacement is a scheduling objective that has seen relatively little attention in the literature. As a result, the authors feel compelled to offer the presented methodology addressing real-world scheduling problems with the objective of minimizing tooling replacements. Minimization of tooling replacements is especially critical if the tooling medium is of high value (i.e., industrial grade diamonds,

* Corresponding author. Tel.: +1-334-844-6511; fax: +1-334-844-4071.

E-mail address: pmcmullen@business.auburn.edu (P.R. McMullen).

or carbide cutting tips), because the amount of money saved can be sizeable. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Production sequencing; Tool wear; Heuristics; Simulated annealing

1. Introduction

There are many objectives to production sequencing—minimization of overtime/labor cost, minimization of in-process inventory, and maximization of production flexibility are typically the most common objectives. The authors, however, have seen little work done with regard to the scheduling of systems when the minimization of tool replacement is a priority. A literature search has uncovered work involved with efficient tooling and machining [1,2], and with efficient replacement of tooling [3–5], but nothing involving scheduling with the intent of prolonging the life of the tool.

If production is scheduled in a way that results in prolonging the life of the tool, the manufacturing firm can avoid unscheduled downtime and the cost of the tool itself. In some manufacturing cases, there is a positive correlation between production flexibility and the life of the tool. If a tool has the capability to produce a wide range of products, but only one product (i.e., Product A) is produced on the tool, then the tool is limited in its ability to produce other products due to the wearing characteristics that Product A has on the tool. This is assuming that the other products have different wearing characteristics. Conversely, if the production sequence is such that there are different parts produced by a tool, then the tool's life is prolonged because flexibility is maintained.

There is also a positive correlation between uniform tool wear and the life of the tool. Consider the very common and simple situation of car tire rotation. It is widely accepted that in order to maintain tire performance and maximize tire life, periodic rotation is required. This scheduled rotation requires a small amount of scheduled downtime (or changeover time), but alleviates more serious and costly consequences that are certain to occur. This example carries the implicit assumption that all sources of wear (tires) are replaced simultaneously.

While this tire example is simple, it illustrates the concept of optimizing tool life via uniformly distributing wear. Unlike the tire example above, the research here is concerned with the replacement of tools only on an “as-needed” basis. Examples of this problem feature are in industries where machine capacities are measured, in part, by their width, such as paper machines, carpet machines, and plywood manufacturers. In the paper industry, paper travels over doctor blades and rollers; and over time the paper's abrasiveness will start to wear the blades and the rollers where the two come in contact. Conversely, there is no wear where the blade and the tool never touch. This obviously causes uneven wear on the blades and eventually will prohibit the production of a “wider-web” of paper because of the non-uniform wear of the tooling. Subsequently, it is important to find production sequences (for processes of this type) that uniformly distribute tool wear so that tool replacements are kept to a minimum.

Another common process that requires uniform tool wear for prolonged life is the hydraulic shearing press. In order to ensure a consistent cut across the width of the press, the die must be worn evenly from one side to the other, otherwise the die becomes inadequate and a replacement is

warranted. All firms that have machines with these characteristics are likely to be concerned about tool life, given the high-cost of tooling.

Because of the importance of this issue, the objective of this research is to provide an approach to production scheduling which attempts to prolong the life of the tool, subsequently minimizing tooling replacements. It is assumed throughout this research that individual tooling is replaced only when needed—when the level of tooling falls below some minimal threshold. The measurement of tool life is provided by three different statistical metrics: the multivariate statistic of canonical correlation, the univariate statistic of pairwise correlation, and a simple measurement of tool-wear “gap”. These three metrics are intended to serve as surrogate measures of sequence diversity with the ultimate intent of obtaining sequences requiring minimal tool replacements. Attainment of production sequences possessing the stated objective is obtained via an application of the optimization/artificial intelligence technique of simulated annealing.

The following sections of this paper outline the problem of maximizing tool life, illustrate the techniques dedicated to measuring uniformity of tool wear, describe the simulated annealing-based sequencing approach, present experimentation to assess the performance of both the uniformity metric and sequencing heuristic (via a tool wear simulation), and offer general observations.

2. Distribution of tool wear

To illustrate the issue of uniform tool wear, consider Fig. 1.

Fig. 1 represents a tool surface containing three areas—each area of the tool is a “source” of wear for the products in each product-set produced. For this example, we might consider a blanking operation, where there are three areas of the press. Each area can blank a product from the product-set simultaneously. Let us assume that the three products that make-up the product-set have similar size, weight, and shape characteristics, but vary in terms of their hardness, thus causing the tool to wear differently over time as a result of this characteristic. Again, we make the assumption that when the capability of one area of the tool drops below a minimum level, the entire tool must be replaced.

The rate of wear of each area is a direct result of the product-set that is processed. For purposes of this illustration, it is assumed that the wear of product type “1” is minimal wear, the wear of product type “2” is moderate wear, and a wear of product type “3” is heavy wear. Here, the products being processed by the outside areas (I and III) subject tool areas to heavy wear, while the product type in the middle area (II) subjects it to minimal wear. Assume that the very next product-set processed subjects the three areas to the exact same level of wear as that shown in Fig. 1. A sequence such




Fig. 1. Illustration of tool wear for a product-set.




Fig. 2. Processing of a product-set preferred to follow the product-set illustrated in Fig. 1.




Fig. 3. Uneven distribution of tool wear.




Fig. 4. Even distribution of tool wear.

as this is not “friendly” to prolonged tool life—this is because Areas I and III would wear much faster than Area II.

A more preferred product-set to follow the Fig. 1 illustration is one shown in Fig. 2.

Fig. 2 shows the wear of a product-set that would be reasonably appropriate to follow the product-set illustrated in Fig. 1, because it will better distribute the wearing of the tools. Each of the three areas of the tool will deteriorate at approximately the same rate, which will permit a prolonged life.

Fig. 3 shows an example of tooling (cutting tools here) which has been exposed to uneven distribution of tool wear—perhaps due to poor sequencing. Areas I and III have worn much faster than Area II. Subsequently, the entire set of tooling must be replaced.

Fig. 4, on the other hand, shows an example of tooling (again, cutting tools) where tool wear has occurred evenly over the three areas susceptible to wear—perhaps due to more responsible sequencing.

The reader must understand that the preceding examples are simplified, and are intended only to illustrate the concept of uniformity in tool wear. For more real-world problems, the working data will not be quite so trivial.

3. Measuring the life of the tool

Consider Table 1, and its associated data regarding tool wear for each of the nine product-sets to be processed.

Table 1
Product-sets to be produced and their associated tool wear

Product-set	Area I	Area II	Area III
A	1	1	2
B	1	2	1
C	2	1	3
D	3	1	2
E	1	3	2
F	1	1	3
G	3	3	3
H	3	1	2
I	3	3	3

Table 2
Variable sets used for correlation calculations

Product-set	y_1	y_2	y_3	x_1	x_2	x_3
A	1	1	2			
B	1	2	1	1	1	2
C	2	1	3	1	2	1
D	3	1	2	2	1	3
E	1	3	2	3	1	2
F	1	1	3	1	3	2
G	3	3	3	1	1	3
H	3	1	2	3	3	3
I	3	3	3	3	1	2
				3	3	3

Here, there are nine product-sets requiring production (Product-sets A–I). Each product-set requires processing on three areas of the tool, and the amount of tool wear associated with each product-set is detailed. Like the previous example, “1” represents negligible wear, “2” represents moderate tool wear, and “3” represents heavy/severe tool wear. In this data, product-set H subjects tool Area I to heavy wear, tool Area II to minimal wear, and tool Area III to moderate wear. In order to have uniform tool wear, product-sets must be sequenced in such a way as to distribute the tool wear as evenly as possible. Prior to addressing the sequencing issue, however, a measure of “uniformity” of tool wear is presented. It is appropriate to mention at this time that the sequence of ABCDEFGHI is the processing sequence used for this particular example. The three areas of variation have been renamed y_1 , y_2 and y_3 for convenience. These three variables are then “lagged” once to create an additional set of variables. This new set of variables is called x_1 , x_2 and x_3 . The results of this matrix are shown in Table 2.

The lagging of the y_i variables to create the x_i variables are done to create a relationship between each product-set in the sequence and its adjacent product-set (this approach is analogous to autocorrelation in time series analysis). For example, the product-set G in the sequence has the y_i variables of 3, 3, and 3, and has the x_i variables of 1, 1, and 3. The x_i variables essentially come from the product-set in the sequence processed just prior to product-set G—product-set F. It

Table 3
Terms used for further analysis

Term	Description
R	Tooling replacements
x_i	The i th variable in the x variable set
y_i	The i th variable in the y variable set
a_{ij}	Members of the \mathbf{A} (tool wear) matrix
$\text{Corr}(x_i, y_i)$	The correlation between the x_i and y_i variables
n	Product-set (or jobs) requiring sequencing
m	The number of areas for tool wear
ρ	Correlation matrix
ρ_c	Characteristic matrix
Λ	Wilks' Lambda (multivariate correlation)
P	Pairwise correlation (univariate correlation)
G	Gap measure (univariate correlation)
$Tool_j$	Tooling remaining for area j
$Tool_j(max)$	Tooling capacity for area j
$Tool_j(min)$	Lower tooling threshold for area j
cv	Coefficient of variation

is desired to determine the level of correlation between the set of y_i variables and the set of x_i variables—specifically, it is desired to minimize the relationship between the x_i variables and the y_i variables. It is hypothesized here that less correlation between the x_i and y_i variable sets results in more uniformity of tool wear. The tool-wear matrix shown in Table 2 is referred to as the \mathbf{A} matrix hereafter and has n rows and $2m$ columns, and the x_i variables are determined by lagging the y_i variables once. Thus

$$a_{i,j+m} = a_{i-1,j}. \quad (1)$$

To clarify an understanding of the objective, definitions are given in Table 3.

The general objective is to minimize the number of tooling replacements occurring during the production run:

$$\text{Minimize: } \sum R. \quad (2)$$

During the production run, each tooling source experiences wear. When the level of tooling for *any* certain area falls below its threshold, the tooling is replaced, and the number of tooling replacements is incremented by one

$$\text{If: } Tool_j \leq Tool_j(min), \text{ then} \quad (3)$$

$$Tool_j = Tool_j(max), \text{ and} \quad (4)$$

$$R = R + 1. \quad (5)$$

Details of attempting to arrive at the objective function in (2) above are presented in Sections 3.1–3.4.

3.1. Canonical correlation and Wilks' Lambda

Canonical correlation is the exploration of the relationship between two sets of variables—each of the two sets of variables can contain multiple variables [6]. Canonical correlation is used here to explore the relationship between the set of tool-wear variables (y_i variables) and the lagged set of variables (x_i variables). Since Table 2 shows an unbalanced data set for the first and last observation as a result of the lagging, the rows with bold characters are omitted from further analysis.

The correlation matrix of the data is as shown in Table 4.

The correlation matrix (ρ) of the type shown in Table 4 breaks down into the following sub-matrices (a subscript of “1” in the matrices below relates to the “y” variables, “2” relates to the “x” variables):

$$\rho = \begin{bmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{bmatrix}. \tag{6}$$

A characteristic matrix (ρ_c) is defined as follows:

$$\rho_c = \rho_{11}^{-1/2} \rho_{12} \rho_{22}^{-1} \rho_{21} \rho_{11}^{-1/2}. \tag{7}$$

The eigenvalues of the above characteristic matrix are the canonical correlation coefficients for the model forming the correlation matrix. A singular, ubiquitous measure of canonical correlation for the entire model is known as Wilks' Lambda (Λ), and is as follows:

$$\Lambda = \prod_{i=1}^m (1 - \lambda_i), \tag{8}$$

where λ_i is the i th eigenvalue, and m represents the number of sources for tool wear in the model.

A high value of Wilks' Lambda shows a low level of correlation between the sets of variables, and is desired here—this would minimize the correlation between the two sets of variables (and subsequently evenly distribute tool wear). The value of Wilks' Lambda falls in the following range:

$$0 < \Lambda < 1. \tag{9}$$

In the context here, it is desirable to find a production sequence yielding a Wilks' Lambda value as close as possible to unity. It is asserted here that if the correlation between the set of variables representing wear on each area (y_i) with that of the lagged set of variables (x_i) is minimized, tool wear would be evenly distributed.

Table 4
Correlation matrix of data from Table 2

	y_1	y_2	y_3	x_1	x_2	x_3
y_1	1.000	0.018	0.315	0.309	-0.098	0.561
y_2	0.018	1.000	0.073	0.273	-0.688	-0.051
y_3	0.315	0.073	1.000	-0.121	0.236	-0.204
x_1	0.309	0.273	-0.121	1.000	-0.059	0.255
x_2	-0.098	-0.688	0.236	-0.059	1.000	-0.055
x_3	0.561	0.051	-0.204	0.255	-0.055	1.000

Using the data from Table 4, along with the methodology presented above, this example is further developed to eventually obtain the value of Wilks’ Lambda. The order of each of the matrices above is equal to the total number of areas for wear. In the case of this example, the order of the matrices is three ($m = 3$).

$$\rho_c = \begin{bmatrix} 1.04 & -0.0003 & -0.168 \\ -0.0003 & 1.002 & -0.037 \\ -0.168 & -0.037 & 1.04 \end{bmatrix} \begin{bmatrix} 0.309 & -0.098 & 0.561 \\ 0.273 & -0.688 & 0.051 \\ -0.121 & 0.236 & -0.204 \end{bmatrix} \begin{bmatrix} 1.072 & 0.048 & -0.271 \\ 0.048 & 1.005 & 0.043 \\ -0.271 & 0.043 & 1.071 \end{bmatrix} \\ \times \begin{bmatrix} 0.309 & 0.273 & -0.121 \\ -0.098 & -0.688 & 0.236 \\ 0.561 & 0.051 & -0.204 \end{bmatrix} \begin{bmatrix} 1.04 & -0.0003 & -0.168 \\ -0.0003 & 1.002 & -0.037 \\ -0.168 & -0.037 & 1.04 \end{bmatrix}.$$

The above computation reduces to the following:

$$\rho_c = \begin{bmatrix} 0.427717 & 0.155494 & -0.238629 \\ 0.155494 & 0.545187 & -0.23087 \\ -0.238629 & -0.23087 & 0.179011 \end{bmatrix}.$$

This matrix has the following eigenvalues: $\lambda_1 = 0.8198$, $\lambda_2 = 0.3289$, and $\lambda_3 = 0.00258$. From these canonical correlation coefficients, the Wilks’ Lambda measure is obtained

$$A = (1 - 0.8198)(1 - 0.3289)(1 - 0.00258) = 0.1206.$$

The value of Wilks’ Lambda is fairly low here, suggesting a high correlation between the y_i and x_i variables, and subsequently an uneven distribution of tool wear.

3.2. Univariate pairwise correlation

A pairwise, univariate correlation measurement can also be employed to study the relationships between variables. This measure is expressed via the following relationship:

$$P = \sum_{i=1}^m Corr(x_i, y_i). \tag{10}$$

As with the multivariate approach detailed above, the intent of this univariate approach is also to minimize correlation between the x_i and y_i variables. Here, however, only pairwise correlations between the *same* tool areas are explored—the Wilks’ Lambda approach considers relationships across *differing* tool areas. This calculation is less CPU intensive than the Wilks’ Lambda approach, because with this pairwise correlation calculation, matrices need not be decomposed via eigenvalues.

Using the data from the same problem, the univariate pairwise correlation is as follows:

$$P = 0.309 - 0.688 - 0.204 = -0.583.$$

With these correlations, it is desired that this value of P be minimized—this suggests high levels of negative correlation, subsequently suggesting uniformity of tool wear. The value of P falls in the following range:

$$-m < P < m. \tag{11}$$

3.3. Gap measure of tool-wear uniformity

The third presented measure of tool-wear uniformity addresses the gap between tool-wear measures of adjacent product-sets (or jobs) in the sequence. This calculation utilizes the “unlagged” portion of the **A** matrix, and this gap measure (*G*) is calculated as follows:

$$G = \sum_{i=2}^n \sum_{j=1}^m |a_{i,j} - a_{i-1,j}|. \tag{12}$$

It is desired to maximize this *G* measure—this reflects large differences in tool wear for adjacent sequence members, thereby suggesting uniformity in tool wear. This measure of tool-wear uniformity differs from the Wilks’ Lambda and pairwise correlation in that the gap measure does not consider correlation between tool-wear areas, but only differences (quite literally). It is also worth noting that the gap measure calculation is much less CPU intensive than either of the other two. This value of *G* is > 0 and its upper-bound is determined by the values in the tool-wear matrix:

$$G > 0. \tag{13}$$

Applying Eq. (12) to the example problem data in Table 2, the *G* computation is as given in Table 5.

The sum of all values in the (*a_{i,j} - a_{i-1,j}*) columns represents the value of *G*. For this example, the value of *G* is 23.

3.4. Summary of tool-wear uniformity measurements

Three measurement approaches to tool-wear uniformity have been presented: Wilks’ Lambda, sum of pairwise correlation and a gap measurement of tool wear. It is desired to maximize the Wilks’ Lambda measure, minimize the pairwise correlation measure, and maximize the gap measure. Mathematically, these measures are as follows:

$$\text{Max: } \Lambda = \prod_{i=1}^m (1 - \lambda_i), \tag{14}$$

$$\text{Min: } P = \sum_{i=1}^m \text{Corr}(x_i, y_i), \tag{15}$$

Table 5
Calculation detail for tool-wear gap measure

Product-set	<i>a_{i,1}</i>	<i>a_{i,2}</i>	<i>a_{i,3}</i>	<i>a_{i,1} - a_{i-1,1}</i>	<i>a_{i,2} - a_{i-1,2}</i>	<i>a_{i,3} - a_{i-1,3}</i>
B	1	2	1	N/A	N/A	N/A
C	2	1	3	1	1	2
D	3	1	2	1	0	1
E	1	3	2	2	2	0
F	1	1	3	0	2	1
G	3	3	3	2	2	0
H	3	1	2	0	2	1
I	3	3	3	0	2	1

$$\text{Max: } G = \sum_{i=2}^n \sum_{j=1}^m |a_{i,j} - a_{i-1,j}|. \quad (16)$$

Each of these three objective functions is used in an attempt to find sequences offering uniform tool wear.

4. Production sequencing and simulated annealing

With the three objective functions as surrogates for uniform (or even) tool wear presented, effort is now directed toward the issue of obtaining production sequences having an even distribution of tool wear. In the example above, there are nine unique product-sets requiring processing: A, B, C, D, E, F, G, H and I. The sequence ABCDEFGHI resulted in a Wilks Lambda of 0.1206, a pairwise correlation sum of -0.583 , and a gap measure of 23—not very good measures of tool wear distribution. Another sequence, such as CBAHEIDGF (chosen arbitrarily) would provide a Wilks' Lambda value of 0.3744, a pairwise correlation sum of -0.587 , and a gap measure of 21—improvements for Wilks' Lambda and the pairwise correlation sum, but not the gap measure. The issue then becomes one of deciding which sequence one should use to evenly distribute the tool wear (via the three presented metrics). For sequencing problems of this type, there are $9!$ ($n!$) or 362,880 possible sequences. If there are n unique product-sets requiring sequencing, there are then $n!$ possible permutations. If there are n_1 units of product-set 1, n_2 units of product-set 2, ... n_k units of product-set k , there are the following number of possible permutations [7]:

$$\text{Possible Permutations} = \frac{\left(\sum_{i=1}^k n_i\right)!}{\prod_{i=1}^k n_i!}. \quad (17)$$

One quickly realizes, then, that finding the best sequence for larger problems via enumeration becomes quite difficult if not outright impossible. Problems of this type are referred to as combinatorial optimization problems.

As a result of such difficulty, one could be motivated to find a desirable, if not necessarily optimal sequence in terms of the objective functions. Finding sequences that provide desirable, but not necessarily optimal solutions to *combinatorial* optimization problems can be done via search heuristics. There are several search heuristics that have been popular tools in the literature, such as genetic algorithms [8,9], tabu search [10–12] and ant-colony optimization heuristics [13,14]. This particular effort, however, exploits the popular search heuristic of simulated annealing [15,16].

Simulated annealing gets its name from the physical annealing of solids—heating a solid to a very high temperature and then slowly cooling the solid. During the cooling, the solid's particles rearrange themselves in such a way that one or more of the solid's physical properties are enhanced—such as strength or hardness. During this cooling process, particle rearrangement occurs at a diminishing rate. Enhancement of the solid's physical properties is the impetus behind the annealing.

Simulated annealing borrows from the physical annealing described above. With simulated annealing, a combinatorial optimization problem (analogous to a solid), such as a sequencing problem, has its feasible solutions rearranged (analogous to a solid's particles) until some stopping criteria are met. During the rearrangement of the feasible solutions, the objective function value of each feasible solution is determined and noted. The “best” solution encountered is presented at the end of

Table 6
User-defined parameters used for simulated annealing

Parameter	Symbol	Explanation
Initial temperature	T_1	Temperature used at start of search
Final temperature	T_F	Temperature used to stop search
Cooling rate	CR	Rate at which temperature (T) is lowered
Iterations	$Iter$	Number of user-defined iterations used for each temperature level
Boltzman constant	k_b	User-defined parameter associated with Metropolis Criterion

the simulated annealing process. Rearrangement of solutions is quite straightforward—components of the solution are randomly selected for pairwise swaps. Consider the following sequence which is used here as a feasible solution:

A B C D E F G H I (current solution).

The D and G components of this solution are underlined. This is because they have been randomly selected for a pairwise swap. The resulting sequence is as follows:

A B C G E F D H I (new solution).

During the rearrangement of the feasible solutions, a decision must be made regarding the acceptance of the newly created solution and subsequent replacement of the current solution. The rule for this “swap” exposes the very essence of simulated annealing. If the objective function value of the new solution is superior to that of the current solution, then the new solution replaces the current solution. If the objective function value of the new solution is inferior to that of the current solution, then the new solution will replace the current solution with the following probability of acceptance ($P(A)$):

$$P(A) = \exp(-\delta k_b/T), \quad (18)$$

Where δ is the percent difference between the objective function value of the current (Z_c) and new (Z_n) objective function values. Here, the objective function values are those of the three presented measurement values. Specifically

$$\delta = \frac{Z_n - Z_c}{Z_c}. \quad (19)$$

The value k_b is referred to as the Boltzman constant, which provides the user some degree of control regarding the probability of inferior solutions being accepted as current. The value T is referred to as the current temperature, and serves as a control parameter. The formula presented in (18) above is referred to as the Metropolis Criterion [17].

The important thing to understand here is that inferior new solutions are sometimes accepted to replace the current solution. At first glance, this may seem destructive, but this is one of the features of simulated annealing which sets it apart from other local search heuristics. The logic here is that if inferior solutions are occasionally accepted to replace current solutions, the search has “migrated” to a new part of the feasibility continuum, with the hope of obtaining (or at least approaching) a global optimum, and avoid being trapped at local optima.

Table 6 shows some user-defined parameters (and explanations) necessary for execution of the simulated annealing heuristic.

It is also important to note that as the control parameter, T , decreases, it becomes less likely for an inferior solution to replace the current solution (analogous to the diminishing rate of particle rearrangement in physical annealing).

The following is pseudocode used to explain the simulated annealing search heuristic:

```

Initialize  $T_1$ ,  $T_F$ ,  $CR$ ,  $Iter$  and  $k_b$ 
 $T = T_1$ 
While ( $T > T_F$ )
  For  $i = 1$  to  $Iter$ 
    Generate_new_solution
    If ( $Z_n > Z_c$ ) then replace_current_solution
    Else {
      If ( $\text{rand}() < P(A)$ ) then replace_current_solution
      Else keep_current_solution}
    Next  $i$ 
   $T = T * CR$ 
End While
Report_best_solution

```

5. Experimentation

Experimentation is conducted to evaluate the performance of the three presented measurements of tool-wear uniformity: Wilks' Lambda, sum of pairwise correlation and the gap measure (via the simulated annealing approach) to obtain sequences having tolerable levels of tool-wear uniformity. Several test problems are used to assess this performance. For these test problems, sequences are obtained (via simulated annealing), for each of the presented objective functions, as well as with two additional approaches.

These two additional approaches do not involve the objective function measures presented above, nor do they require simulated annealing or any type of search, but instead obtain sequences via inspection, or application of a rule. The first of these two additional approaches is finding a sequence minimizing makespan, or the length of a production run. Consider a problem set where four units of product-set A are demanded, two units of product-set B are demanded, and a single unit of product-set C is demanded. A production sequence minimizing makespan is as follows:

A A A A B B C (minimized makespan).

This of course, can be done by inspection. The other of the two additional approaches to sequencing is the sequence that optimizes (minimizes) the stability of material usage rate [18]. This approach attempts to find sequences that maximize system flexibility and evenly distribute the requirements for the demand of raw materials. This approach to sequencing is very important to manufacturing firms having an interest in just-in-time (JIT) production. A sequence optimizing the usage rate [19] for this simple sequencing problem is as follows:

A B A C A B A (minimized usage rate).

Some consider minimization of this usage rate as a formal way to maximize the product “inter-mixing”—the authors *are* comfortable with this particular terminology.

Once sequences are found via the three objective functions and simulated annealing (or via the two additional approaches dedicated to minimization of makespan and usage rate), these sequences are simulated in an environment that determines how well they perform in terms of tool repair and/or replacement.

5.1. Problem sets

Table 7 shows the eleven problem sets used for experimentation.

The first column of Table 7 is self-explanatory. The second column details the number of areas for tool wear, while the third column details the specific product-mix. The fourth column (range of tool wear) shows the range of values taken on by the **A** matrix—these values were generated via a uniformly distributed random number generator (taking on the ranges detailed above). The number of solutions is also shown for the smaller problem sets (sets 1–6), complete enumeration was also performed to obtain sequences having optimal values of Wilks’ Lambda, sum of pairwise correlations and gap measures. The cooling rate was engineered to force the number of solutions evaluated via the simulated annealing approach to be proportional to the number of feasible solutions for each problem set. For the largest problem sets (sets 7–11), the cooling rate was set to 0.9999. For the objective functions addressing minimization of Wilks’ Lambda and the sum of pairwise correlations, the Boltzman Constant (k_b) was kept at 1.0000. For the objective functions addressing maximization of the gap measure, the Boltzman Constant was kept at 346.57 (this provides a 50% probability of acceptance when the new solution is 5% inferior to that of the current solution when the value of T is at 25). The Boltzman Constant values used here were the result of experimentation. It is important to note that the Metropolis Criterion shown in (18) is generically intended for minimization problems. With this research it is desired to minimize the sum of pairwise correlations, which is

Table 7
Details of problems used for experimentation

Problem set	Wear areas (m)	Product-mix	Range of tool wear	Feasible solutions	SA cooling rate
1 ^a	3	A(1), B(1), C(1), D(1), E(1), F(1), G(1), H(1)	1–3	40,320	0.8191
2 ^a	3	A(1), B(1), C(1), D(1), E(1), F(1), G(1), H(1), I(1)	1–3	362,880	0.9780
3 ^a	3	A(1), B(1), C(1), D(1), E(1), F(1), G(1), H(1), I(1), J(1)	1–7	3,628,800	0.9970
4 ^a	4	A(4), B(3), C(2), D(1), E(1), F(1)	1–8	1,663,200	0.9936
5 ^a	4	A(2), B(2), C(2), D(2), E(2)	1–10	113,400	0.9315
6 ^a	3	A(3), B(3), C(2), D(2), E(2)	1–25	1,663,200	0.9936
7	5	A(4), B(4), C(3), D(3), E(2), F(2), G(1), H(1)	1–9	2.993(10) ¹³	0.9999
8	3	A(12), B(3), C(2), D(2), E(1)	1–25	219,629,600	0.9999
9	3	A(12), B(3), C(2), D(2), E(1)	1–15	219,629,600	0.9999
10	3	A(4), B(4), C(4), D(4), E(4)	1–15	3.0554(10) ¹¹	0.9999
11	3	A(4), B(4), C(4), D(4), E(4)	1–30	3.0554(10) ¹¹	0.9999

^aDenotes problems where complete enumeration was performed to find optimal values of the three objective functions.

consistent with (18). For Wilks' Lambda and the gap measure, however, it is desired to achieve maximization. As a result, (18) undergoes a sign change when the objectives of a maximizing nature are addressed. The simulated annealing initial temperature (T_1) for all problems is 25 and the final temperature (T_F) is 1. It is also important to emphasize that for each problem set, the three objective function-based approaches via simulated annealing will process and evaluate the same number of solutions, so that fair comparisons can be made.

5.2. Tool-wear simulation

Each sequence obtained via this methodology is simulated in an environment where replacement of tooling is studied—this is done to obtain information on how much tooling replacement is required for the sequence. The simulation is actually quite straightforward. For each sequence, tool wear is simulated across the m areas for tool wear. The tooling available for area j is represented by $Tool_j$. The tool-wear matrix, again, is the “unlagged” part of the \mathbf{A} matrix. After product-set i is processed on tool area j , the remaining tooling is adjusted via the following relationship:

$$Tool_j = Tool_j - (a_{i,j} + (\hat{\sigma}_{a_{i,j}} \cdot z)), \quad (20)$$

where the value of $\hat{\sigma}_{a_{i,j}}$ is estimated by the following:

$$\hat{\sigma}_{a_{i,j}} = a_{i,j} \cdot cv. \quad (21)$$

For this research, the coefficient of variation (cv) is always set to 0.15. Also, the value of z is a random normal deviate sampled from a population having a mean of zero and a standard deviation of one [20]. After each product-set is processed by the m tooling areas, the depletion of tooling is accounted for via (20) and (21). If the level of tooling is less than the lower threshold value, the tooling is replaced, and the level of tooling is initialized to its upper threshold value. A count of tool replacements is also kept. For each simulation run, each sequence is repeated 2,000 times. Each simulation run is repeated 25 times for each problem and for each approach, so that a reasonable estimate of required tooling replacements is obtained. It is important to note that these simulation assumptions (coefficient of variation, tool replacement, etc.) have been adopted in accordance with “valid and credible” simulation models [21]. The “best” sequences will result in the smallest number of tooling replacements. It is also worth noting that previous research [7] has validated the simulated annealing parameters, as well as finding that the problem sets addressed here are of “industrial-size”.

5.3. Research questions

The following research questions are constructed to assess the performance of the objective functions and the simulated annealing methodology as stated above:

1. Is the level of required tooling replacement sequence dependent? In other words, does the approach used to obtain production sequences have an effect on the required number of tooling replacements?
2. Do the objective function-based approaches provide production sequences (including sequences obtained via complete enumeration) with fewer required tooling replacements as compared to the approaches using minimized makespan and usage rates?

3. Do the objective function-based approaches provide production sequences (excluding sequences obtained via complete enumeration) with fewer required tooling replacements as compared to the approaches using minimized makespan and usage rates?
4. Which of the objective function-based approaches provides the fewest number of tooling replacements?

These research questions are addressed via univariate ANOVA and comparison of mean replacements.

Since there are 11 different problem sets explored, vastly differing numbers of required tooling replacements occur across the different problems—this complicates comparison across problem sets. To combat this, ratios are constructed for each problem set. For each problem set, the required number of tooling replacements (obtained via simulation) is divided by the minimum number of required tooling replacements for that specific problem set. For example, if a problem set is simulated and the following number of required tooling replacements are obtained: 1005, 1,000 and 1,010, the ratios would be as follows: 1.005 (1,005/1,000), 1.000 (1,000/1,000) and 1.010 (1,010/1,000). The value 1,000 is the denominator because it is the minimum value. The ratio of 1.005 could also be interpreted as the tooling replacement requirement being 0.5% inferior to the minimum for the problem, the ratio of 1.0 could be interpreted as minimum for the specific problem, while the ratio of 1.010 could be interpreted as 1.0% inferior to the minimum for the specific problem. When this ratio is constructed for the 11 problem sets, standardization allows meaningful comparisons of required tooling replacements to be made *across* all problem sets.

5.4. Computational experience

The objective function approaches to obtaining sequences via simulated annealing and complete enumeration were performed using C/C++ on a Windows platform with a Pentium II processor at 300 MHz. The tool-wear simulation was also performed with the same resources. While the results section briefly discusses CPU time, it is noted here that efforts were made to minimize CPU requirements in all circumstances.

6. Experimental results

With regard to the first research question, ANOVA shows that tool wear is sequence dependent—the sequencing approach has a statistically significant effect on the number of replacements ($F=20.02$, $p < 0.001$). This analysis considers all 11 problem sets, and the ratio metric described above is used as the response variable. Table 8 shows means and standard deviations for all approaches (both objective function based, as well as the minimal makespan and minimal usage approaches).

Table 8 details where differences in mean replacements exist across problem sets—univariate F -statistics are in the first column, listed underneath the Problem set number. Note that Problem Sets 1 and 2 are not sequence dependent. The nine remaining problem sets are sequence dependent, and all of these F -statistics have associated p -values < 0.001 .

Table 9 shows material usage rates and required setups (a surrogate measure for makespan) for each approach organized by problem set. It is obvious from inspection of Table 9 that the minimal usage approach does in fact result in minimal usage at the expense of several setups, while the

Table 8

Means (and standard deviations) of replacements by problem set and approach—minimal replacements for each problem set are italicized

Problem set	Makespan	Usage	SA Wilks'	Enum. Wilks'	SA corr	Enum. corr	SA gap	Enum. gap
1	1794.2	1795.2	1794.9	1794.4	1794.9	1795.3	1794.5	1795.1
($F = 1.4$)	(1.8)	(2.0)	(1.5)	(1.6)	(1.5)	(1.5)	(1.2)	(1.5)
2	2145.2	2144.5	2145.3	2145.8	2145.1	2145.0	2144.9	2145.6
($F = 1.0$)	(2.1)	(1.6)	(1.4)	(2.0)	(1.4)	(1.6)	(1.8)	(1.7)
3	4759.2	4759.7	4752.4	4763.4	4762.8	4762.0	4752.2	4748.8
($F = 47.4$)	(3.7)	(4.5)	(4.8)	(3.8)	(4.2)	(3.3)	(4.5)	(3.6)
4	7856.8	7845.9	7842.2	7845.1	7844.2	7842.0	7848.3	7848.8
($F = 23.4$)	(4.4)	(5.9)	(4.7)	(5.0)	(5.4)	(5.2)	(4.2)	(4.7)
5	8635.3	8645.0	8636.4	8642.8	8630.0	8630.7	8634.4	8622.9
($F = 31.1$)	(6.4)	(5.4)	(6.3)	(6.2)	(7.2)	(7.1)	(6.6)	(5.5)
6	16935.6	17263.8	16940.1	16945.1	16965.2	16968.3	17190.5	16965.0
($F > 1000$)	(22.5)	(20.1)	(19.8)	(19.3)	(16.9)	(22.8)	(12.8)	(23.4)
7	11876.6	11849.4	11851.2	N/A	11865.8	N/A	11868.7	N/A
($F = 108.4$)	(4.4)	(5.6)	(5.7)		(6.1)		(6.1)	
8	16580.6	16192.7	16705.9	N/A	16303.2	N/A	16503.8	N/A
($F > 1000$)	(21.7)	(13.2)	(13.9)		(15.1)		(16.8)	
9	23242.8	23145.8	23296.9	N/A	23224.7	N/A	23245.0	N/A
($F = 337.8$)	(12.4)	(14.5)	(15.4)		(14.6)		(17.1)	
10	18510.9	18491.2	18471.2	N/A	18550.6	N/A	18536.1	N/A
($F = 160.0$)	(16.1)	(11.0)	(15.0)		(9.7)		(10.7)	
11	39299.0	39419.8	38309.6	N/A	38211.8	N/A	38505.2	N/A
($F > 1000$)	(27.0)	(28.0)	(29.7)		(28.7)		(28.7)	

minimal makespan approach provides just the opposite result—minimal makespan (setups) at the expense of high usage rates.

Regarding the second research question, ANOVA shows that the objective function-based sequencing approaches (which include the optimal objective function approach sequences obtained via enumeration) result in fewer replacements as compared to the sequences obtained via minimal setups and minimal makespan approaches ($F = 66.46$, $p < 0.001$). The objective function approaches provide a mean replacement ratio of 0.44% inferior to the minimum (standard deviation of 0.55%), while the makespan and usage approaches provide a mean replacement ratio of 0.74% inferior to the minimum (standard deviation of 0.99%).

For the third research question, ANOVA shows that the objective function-based sequencing approaches (after the removal of the optimal sequences obtained via enumeration) still performs better than the minimal makespan and minimal usage approaches ($F = 20.33$, $p < 0.001$). The mean replacement ratio for the objective function-based approach is 0.54% inferior to the minimum (standard deviation of 0.66%). The mean replacement ratio for the sequences obtained via the minimal makespan and minimal usage approaches is, of course, the same as that mentioned above (0.74%).

Since the second and third research questions involve a comparison between sequences obtained via objective function versus makespan and usage approaches, and provide evidence that the objective function approaches provide generally superior results (respective F -statistics of 66.46 and 20.53 both

Table 9
Material usage rates (and required setups) for problem sets

Problem set	Makespan	Usage	SA Wilks'	Enum. Wilks'	SA corr	Enum. corr	SA gap	Enum. gap
1	10.5 (8)	10.5 (8)	10.5 (8)	10.5 (8)	10.5 (8)	10.5 (8)	10.5 (8)	10.5 (8)
2	13.33 (9)	13.33 (9)	13.33 (9)	13.33 (9)	13.33 (9)	13.33 (9)	13.33 (9)	13.33 (9)
3	16.5 (10)	16.5 (10)	16.5 (10)	16.5 (10)	16.5 (10)	16.5 (10)	16.5 (10)	16.5 (10)
4	47.11 (6)	7.77 (12)	8.77 (12)	13.98 (12)	19.94 (12)	23.68 (12)	19.44 (12)	23.68 (12)
5	28.00 (5)	8.00 (10)	8.00 (10)	16.00 (9)	8.00 (10)	8.00 (10)	8.00 (10)	16.00 (10)
6	48.58 (5)	7.08 (12)	16.25 (11)	19.25 (10)	24.92 (12)	20.42 (11)	33.58 (12)	29.83 (11)
7	162.30 (8)	16.10 (20)	48.10 (17)	N/A	106.5 (20)	N/A	82.80 (20)	N/A
8	208.75 (5)	10.45 (17)	35.05 (13)	N/A	38.35 (19)	N/A	22.35 (17)	N/A
9	208.75 (5)	10.45 (17)	73.75 (11)	N/A	18.65 (15)	N/A	12.45 (17)	N/A
10	216.00 (5)	17.80 (20)	42.00 (17)	N/A	176.50 (17)	N/A	56.00 (20)	N/A
11	216.00 (5)	17.80 (20)	29.20 (16)	N/A	104.00 (17)	N/A	34.00 (20)	N/A

with p -values < 0.001), additional comments are made regarding their differences. In general, the minimal makespan approach performs quite poorly regarding mean replacements. The only time this approach provides a desirable result is for Problem set 6. It should be intuitive that minimization of makespan actually encourages non-uniformity of tool wear, without any intermixing of the same product-set, individual tooling areas are subjected to the same wear over and over again. The minimal usage approach is a bit of an enigma. With this approach, there are three occasions where it actually performs well with regard to tool wear—Problem Sets 7, 8, and 9. There are also, however, two occasions where it actually performs quite poorly—Problem Sets 6 and 11. The minimal usage rate, as previously described, maximizes the intermixing of the unique product-sets requiring sequencing. This will, to some effect, assist in distributing some uniformity among the areas for tool wear (presumably, this is what happened for Problems Sets 7, 8 and 9). Unfortunately, the minimal usage rate approach never actually considers the values of the tool-wear matrix, which can yield sequences resulting in large numbers of replacements, despite the fact that product inter-mixing is optimized (presumably, this is what occurred for Problem Sets 6 and 11). Due to these issues, one cannot rely on the non-objective function approaches of minimal makespan and minimal usage for any consistently desirable results. Regarding this minimal usage approach, one could reason that tool replacement performance is “at the mercy” of the tool-wear matrix, which is not considered by the heuristic, which can provide some potentially disappointing results—an intuitive argument against the use of the minimal usage approach.

Table 10

Mean replacement inferiority ratio (standard deviation) for objective function-based approaches

Approach	Mean replacement inferiority ratio (standard deviation)
Simulated annealing Wilks' Lambda	0.575% (0.91%)
Simulated annealing correlation sum	0.375% (0.22%)
Simulated annealing gap measure	0.668% (0.64%)
Enumerated Wilks' Lambda	0.289% (0.14%)
Enumerated correlation sum	0.281% (0.14%)
Enumerated gap measure	0.234% (0.11%)

For the fourth and final research question, Table 10 shows mean replacement inferiority ratios for all of the objective function approaches. Recall that the mean replacement inferiority ratio is used so that comparisons and analyses can be made across *all* problem sets.

From inspection of Table 10, it is immediately obvious that the objective function-based sequences obtained via complete enumeration provide desirable results for all three objective functions. The mean replacement inferiority ratio for these enumerated solutions is 0.268% (standard deviation of 0.134%), compared to 0.539% (standard deviation of 0.662%) for the sequences obtained via simulated annealing. These differences are significant ($F = 73.96$, $p < 0.001$). While it is unfortunate that it is not practical to obtain optimal solutions to all of these problems, the strong performance of the enumerated approaches does suggest the general effectiveness of these objective functions to obtain sequences assisting to provide uniformity in tool wear.

Of the three simulated annealing approaches, the sum of pairwise correlation measure clearly outperforms the Wilks' Lambda and gap measure approaches. The reason that the sum of pairwise correlation outperforms the Wilks' Lambda approach is because the sum of pairwise correlation approach only considers correlation between the *same* wear area, while the Wilks' Lambda approach does consider correlation between all wear areas, including *differing* wear areas. The gap measure approach via simulated annealing provides results which are disappointing as compared with the other two objective approaches. While the approach is well intentioned because it only considers similarity between the *same* wear area (as with the sum of correlation approach), the results are nonetheless a disappointment. The authors are at a loss to explain the relatively inferior performance of the simulated annealing gap measures approach, especially considering the fact that the gap measure approach via complete enumeration provides desirable results.

It is also worth mentioning that the Wilks' Lambda approach is about 20 times more CPU intensive than the gap measure approach. This is due to the fact that the Wilks' Lambda approach involves two matrix decompositions—once to obtain the eigensystem of values needed to find the $\rho_{11}^{-1/2}$ matrix, and again to find the eigenvalues of the characteristic matrix required for the ultimate Wilks' Lambda calculation. Numerical iteration via the Jacobi method is used here to capture the eigensystems of the appropriate matrices [22]. The sum of pairwise correlation measure is about four times more CPU intensive than the gap measure approach—correlations are calculated, but there are no matrix manipulations required. The gap measure approach is the least CPU intensive of the three approaches, because the objective function calculation is so straightforward.

7. Concluding comments

The presented research provides methodology to obtain production sequences requiring minimal tooling replacements via uniformity of tool wear, as compared to more traditional sequencing approaches. The skeptical reader may consider that the presented methodology does provide sequences that provide only *slight* improvement over the more traditional sequencing approaches, in terms of required tooling replacements. The authors have a twofold response to this concern. First, the superiority of the presented methodology as compared to the more traditional sequencing approaches of minimal makespan and usage rates is statistically significant, with p -values < 0.001 . Second, the superiority offered by the presented approaches, albeit arguably slight, could result in large annualized savings for manufacturers using very expensive tooling, such as carbide tooling, or tooling comprised of industrial-grade cutting diamonds. For example, if a firm were to implement the presented methodology, and they found that because of the implementation, they could reduce replacements by 25 tooling sets per year, and the tooling was expensive (such as industrial-grade cutting diamonds)—costing \$8,000/set, an annual savings of \$200,000 could be realized.

Simulated annealing is the search heuristic of choice for this research. The authors are aware of the fact that there are many formidable search heuristics to choose from, such as tabu search, genetic algorithms, artificial neural networks and ant-colony optimization approaches. Simulated annealing was chosen here because of its well-regarded general performance and its relative simplicity. Comparing the various search heuristics with respect to uniformity of tool wear is beyond the scope of this research, but is an opportunity for future research. Another, perhaps more substantial, opportunity for future research is addressing the same problem but with other objective function measures similarly dedicated to evenly distributing tool wear.

This research is concerned with the minimization of tooling replacements, by finding sequences offering uniformity of tool wear. Although this has not been a heavily researched issue, the authors consider it important nevertheless. Throughout this paper, comparison has been made between the presented approaches and the more traditional approaches which minimize makespan and usage rate. The criticism of these two more traditional approaches is made only in association with their relative lack of tool-wear uniformity. The authors are fully aware of the importance of minimal makespans in production runs, as well as the flexibility offered when the stability of the material usage rate is optimized, especially with JIT systems. Furthermore, the authors consider it critically important to find production sequences that offer desirable levels of makespan, stability of material usage rates *and* uniformity of tool wear. Further research on production sequencing should consider all three factors in tandem.

References

- [1] Kee PK. Development of computer-aided machining optimization for multi-pass rough turning operations. *International Journal of Production Economics* 1994;37:215–27.
- [2] Yeo SH. A multipass optimization strategy for CNC lathe operations. *International Journal of Production Economics* 1995;40:209–18.
- [3] Leung LC, Tanchoco JMA. Multiple machine replacement analysis. *Engineering Costs and Production Economics* 1990;20:265–80.
- [4] Luxhoj JT. Replacement analysis for components of large scale production systems. *International Journal of Production Economics* 1992;27:97–110.

- [5] Barad M, Hoang K. Cutting tools versus machine tools investments in FMS design. *International Journal of Production Economics* 1995;42:29–40.
- [6] Johnson RA, Wichern DW. *Applied multivariate statistical analysis*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 1992.
- [7] McMullen PR, Tarasewich P, Frazier GV. Using genetic algorithms to solve the multi-product JIT sequencing problem with setups. *International Journal of Production Research* 2000;38:2653–70.
- [8] Goldberg DE. *Genetic algorithms in search, optimization & machine learning*. Reading, MA: Addison-Wesley, 1989.
- [9] Michalewicz Z. *Genetic Algorithms + Data Structures = Evolution Programs*. New York: Springer, 1996.
- [10] Glover F. Tabu search: a tutorial. *Interfaces* 1990;20:74–94.
- [11] Glover F. A user's guide to tabu search. *Annals of Operations Research* 1993;41:3–28.
- [12] Bland JA, Dawson GP. Tabu search and design optimization. *Computer Aided Design* 1991;23:195–201.
- [13] Bonabeau E, Dorigo M, Theraulaz G. *Swarm intelligence: from natural to artificial systems*. New York: Oxford University Press, 1999.
- [14] Bonabeau E, Theraulaz G. Swarm smarts. *Scientific American*, 2000;282:72–9.
- [15] Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by simulated annealing. *Science* 1983;220:671–9.
- [16] Eglese RW. Simulated annealing: a tool for operational research. *European Journal of Operational Research* 1990;46:271–81.
- [17] Metropolis N, Rosenbluth A, Rosenbluth N, Teller A, Teller E. Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 1953;21:1087–92.
- [18] Miltenburg J. Level schedules for mixed-model assembly lines in just-in-time production systems. *Management Science* 1989;35:192–207.
- [19] Ding F, Cheng L. An effective mixed-model assembly line sequencing heuristic for just-in-time production systems. *Journal of Operations Management* 1993;11:45–50.
- [20] Box GEP, Muller ME. A note on the generation of random normal deviates. *Annals of Mathematical Statistics* 1953;29:610–1.
- [21] Law AM, Kelton WD. *Simulation modeling & analysis*, 2nd ed. New York: McGraw-Hill, 1992.
- [22] Matthews JH, Fink K. *Numerical methods using Matlab*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall, 1999.

Patrick R. McMullen is an Associate Professor of Management at Auburn University's College of Business. He received his Ph.D. in Decision Sciences from the University of Oregon. He has held teaching positions at the University of Oregon, the University of Maine and Harvard University. His interests lie in the area of combinatorial optimization and production scheduling.

Mark M. Clark received his MS and Ph.D. in Industrial Engineering from Auburn University. He is currently Product Manager for the Alabama Technology Network at Auburn University and also serves as an Adjunct Assistant Professor in the Department of Management.

David Albritton is a Ph.D. candidate in Management at Auburn University. He holds a master's degree in Business Administration from Auburn and a bachelor's degree from the University of Florida. Mr. Albritton spent the majority of his professional career in advertising and strategic marketing. His research interests include operations management and strategic uses of technology.

Major John E. Bell earned a Bachelor of Science Degree and a Masters of Science Degree in Logistics Management from the Air Force Institute of Technology. He is currently pursuing his Ph.D. in Management at Auburn University. His research and academic interests include logistics management, production and operations management, transportation management, and heuristic search methods.