DOE/ID/13614

Re-Engineering Casting Production Systems

Final Report – 03/02/1998 – 03/01/2001

**F. Peters
T. Van Voorhis**

June 2001

Work Performed Under Contract No. DE-FC07-98ID13614

**For
U.S. Department of Energy
Assistant Secretary for
Energy Efficiency and Renewable Energy
Washington, DC**

**By
Iowa State University
Ames, IA**

DOE/ID/13614

RE-ENGINEERING CASTING PRODUCTION SYSTEMS

FINAL REPORT
03/02/1998 – 03/01/2001

Frank Peters
Timothy Van Voorhis

June 2001

Work Performed Under Contract No. DE-FC07-98ID13614

Prepared for the
U.S. Department of Energy
Assistant Secretary for
Energy Efficiency and Renewable Energy
Washington, DC

Prepared by
Iowa State University
Ames, IA

Re-Engineering Casting Production Systems

-- FINAL REPORT --

**Frank Peters
Timothy Van Voorhis**

**Industrial & Manufacturing Systems Engineering Department
Iowa State University
Ames, IA 50011**

**515/294-3855 Voice
515/294-3524 Fax
fpeters@iastate.edu
tvanvoor@iastate.edu**

This project was funded by the U.S. Department of Energy (DOE) award No. DE-FC07-98ID13614. However, any opinion, findings, conclusion, or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the DOE.

<u>1</u>	<u>INTRODUCTION</u>	<u>3</u>
<u>2</u>	<u>SURVEY RESULTS</u>	<u>3</u>
2.1	General Demographics	4
2.2	Production Manager Responses	5
2.3	Differences In Perception Between Production And Sales	8
<u>3</u>	<u>ON-SITE ANALYSIS OF PRODUCTION SYSTEMS</u>	<u>10</u>
3.1	Development of Flowpaths	10
3.1.1	The Four Kinds of Steps	11
3.1.2	Branches	12
3.1.3	Why Flowpaths?	12
3.1.4	Limitations of Flowpaths	13
3.1.5	How the Data Was Collected	13
3.2	Results of On-Site Production System Analysis Studies	14
3.3	Definitions of Rework	20
3.4	General Recommendations	20
<u>4</u>	<u>LAYOUT IMPROVEMENTS</u>	<u>21</u>
<u>5</u>	<u>USE OF PRODUCTION SIMULATION</u>	<u>23</u>
5.1	Overview of Discrete Event Simulation	23
5.2	Steady-State Analysis of Simulation Output	24
5.3	Model	24
5.4	Output	25
5.5	Data Analysis	25
5.6	Continued Research	27
<u>6</u>	<u>SCHEDULING</u>	<u>27</u>
6.1	Introduction	27
6.2	Modeling Considerations	29
6.3	Data Requirements	29
6.4	A Mathematical Model for Pouring Scheduling	31
6.5	Solution Heuristic	32
6.6	Computational Experience	33
6.7	Future Work	34
<u>7</u>	<u>TECHNOLOGY TRANSFER</u>	<u>34</u>
<u>8</u>	<u>INDIVIDUAL FOUNDRY RESULTS AND IMPLEMENTATIONS</u>	<u>35</u>

1 INTRODUCTION

The goal of this three-year project was to improve the production systems in use by steel foundries in the United States. Improvements in the production systems result in less rework, less scrap, and less material handling, all of which would significantly reduce the energy demands of the process. Furthermore, these improvements would allow the companies to be more competitive, more responsive to customers' needs, deliver products with less lead time and require less capital. The ultimate result is a stronger domestic steel casting industry, which uses less energy.

A major portion of this research involved the deployment of student researchers at steel foundries, to study their production systems and collect data. Ten different students gained valuable industrial experience, as they studied the production systems of twenty-one different companies during the three-year study. Study periods ranged from 4-13 weeks. As a direct result of this project, three foundries have hired industrial/manufacturing engineers that have graduated from Iowa State University to help them continue making production system changes.

The methodology used to study the production systems will be outlined here. Results from the participating foundries will also be shown.

At the beginning of the project, a survey of production system practices and performances was conducted. The results of which are included here.

Layouts that were proposed to foundries to improve their production systems are included in this report. Some particular features of the layout and advantages are discussed.

A major problem that the project team discovered was the manner in which most foundries schedule their operations. The molding and pouring schedule is developed, with little or no consideration for down stream operations. Because of this, the research team developed a scheduling algorithm that is presented here, which was not originally included in the initial proposal.

Production simulation modeling was used to study the impact of casting variability has on the production system. As variability of the product increases, it becomes more difficult to develop a coherent schedule that can be followed.

The implementations and changes that resulted from this study are reviewed for those companies that have made or are making production system changes as a result of this project.

2 SURVEY RESULTS

Near the beginning of the project, a survey of production related issues was sent to all domestic steel foundries, and those international foundries that are members of Steel Founders' Society of America. There were two separate surveys, one for the production manager and another for the sales manager. The surveys were identical, except the production manager survey contained additional questions about capacities, equipment, and staffing. The purpose of the two surveys was to see if differences in perception among people within the same company existed across the industry. Both surveys were received from 21 foundries. Twenty-five additional foundries only sent the production or sales manager survey, or the same person

completed both surveys. For all of these cases, it was treated as if they only completed the production manager survey.

2.1 General Demographics

Figures 1-4 give profiles of the respondents, in terms of their molding type, casting weights produced, alloys produced, and number of employees. The average annual production of the respondents is 9224 tons. The average number of employees is 174.

Figure 1: Percentage Of Production Based On Molding Type

Figure 2: Percentage Of Production Based On Casting Weight

Figure 3: Percentage Of Production Based On Alloy.

Figure 4: Number Of Employees

2.2 Production Manager Responses

The survey asked several questions regarding the processing times for castings. Figure 5 shows the average number of processing days reported to complete the castings, starting from when the order was released to production. Figure 5 also shows the processing days required after shakeout. Forty-five percent of the production managers said their company collected data on processing times, from which their answers were based. It is assumed that the others based their answers on observations.

Figure 5: Days Consumed From Release Of Order Until Shipping, And Shakeout Until Shipping

The considerations when developing the weekly or daily production schedule were asked. Of the 46 foundries, 7 reported that the schedule was based solely on due date. Fifteen people said that capacity was a consideration, without specifying the particular areas whose capacities were taken into account when developing the schedule. Other people specifically state those areas whose capacities are considered when developing the schedule: molding (12), melting (8), core making (4), each department (3), cleaning operations (2), outsource (1), [unspecified] bottlenecks (1). If a person stated more than one area, they all were included. It is noteworthy that only three companies consider all areas when developing a schedule. Some other considerations when scheduling include: backlog, inventory, need, dollar value of products, order quantity, customer requirements, overdue orders, rush orders, and monthly goals. On a related topic, 72% of the respondents reported using software to assist in scheduling of production. The types of software in use varied from in-house products (6), spreadsheets (7), to commercial packages the most common of which was B&L (5).

The survey also asked several questions about expediting orders (not including sample castings.) Figure 6 shows the breakdown of order expediting. Twenty-one of the 46 companies reported expediting over 15% of their orders. The average percentage of orders expedited was 17.8 %. The reasons that require orders to be expedited were also asked, and summarized in

Table 1. The ranking was accomplished by assigning 7 points for a 1st place vote, 6 for a 2nd place vote, and so forth. For each reason, the number of first place votes is in parenthesis.

Of the respondents, 74% reported having regular meetings in which the expediting of orders was discussed. Of the 34 that had meetings, 21 were held weekly, 10 were held daily, and the remainder were held at other intervals. The average number of personnel hours consumed per week at these meetings was 11.3. However, it is not known how much of this time was spent discussing expediting orders.

Figure 6: The Percentage Of Orders Expedited

Table 1: Ranking of the Reasons That Require Parts To Be Expedited. (Number Of First Place Votes In Parenthesis.)

Reason		Pts.
customer needs sooner	(18)	227
customer late to order	(14)	209
quality problems	(1)	146
tool/design changes	(3)	120
production delays	(3)	115
equipment breakdowns	(2)	111
lack of personnel	(1)	70
other	(3)	18

Any attempt to improve overall productivity needs to be aware of which areas tend to be production bottlenecks. Thus, respondents were asked to rank the five areas with the smallest capacities. The survey also asked to again rank the areas on the basis of their tendency to cause production delays. The rankings (which were scored 5 points for first place, etc.) are given in Table 2.

The survey also asked the major capital investment projects completed in the last 5 years, and those projects that have the highest priority for completion in the next 5 years. The answers were grouped into categories and reported in Table 3. Since the total number of completed projects was nearly the same as the planned projects, the difference between the two could be compared directly. Those projects for which the planned instances outnumbered those

completed in the past are towards the top of the list. Molding and coremaking equipment, which accounts for the most projects, continues to be important. The planned number of cleaning room equipment upgrades nearly doubled. The planned purchases of machining equipment and melting furnaces drop off significantly from the past 5 years. The intended use of the additional buildings or additions was usually not specified.

Forty-eight percent of the companies reported that they batch castings for cleaning operations. For those who do batch castings, the smallest, average and largest batch size was queried. The average responses for these three questions, respectively, are: 5.7, 167, and 422. The most common method of determining batch size (7 responses) was the size of the shot blast unit or container used in the cleaning area. Each of the following three methods was used by five foundries: a) casting size, b) size of the production run, lot or heat, c) order size. No companies determined batch size on the basis of minimizing the WIP, or optimizing the cleaning operations.

Of the 46 foundries, 28 (61%) track castings through the plant and 12 (26%) reported the use of bar coding. However, only 3 companies use bar coding to track the castings through the plant. Other uses of bar coding (with number of responses in parenthesis) include: shipping in general (2), shipments to a particular customer (2), labor reporting (1), inventory (1), final product ID (1), work order entry (1), and scheduling (1).

There were a few other simple answer questions included in the survey. The average percentage of production employees able to work in 2 or more areas was 47%. The average number of people who process incoming orders is 3.9.

Table 2: Ranking Of Areas In Terms Of Being Most Under Capacity, And Their Tendency To Cause Production Delays. (First place votes in parenthesis.)

Under Capacity		Reason for Delay	
Area	Pts	Area	Pts
Molding (8)	82	Heat Treatment (5)	55
Heat Treat (5)	65	Welding (5)	44
Core making (8)	61	Molding (3)	41
Welding (5)	50	Coremaking (2)	35
Grinding	43	Grinding	35
Melting (3)	43	Machining (3)	31
Blast Clean	32	Inspect (1)	22
Inspect (1)	31	Melting (1)	19
Machining (4)	28	Pouring (1)	18
Riser Removal (1)	26	Arc Air	13
Arc Air	25	Blast Clean	12
Pouring (1)	18	Shakeout	10
Shakeout	15	Riser Removal	10

Table 3: Capital Projects Completed During The Past Five Years, And Planned For The Next Five Years

Capital Project	Completed	Planned	Difference
molding equipment	14	20	6
cleaning room equipment	5	10	5
building/addition	5	9	4
office	0	3	3
heat treating equipment	7	9	2
other sand equipment **	8	10	2
computer software/hardware	5	6	1
shakeout	2	2	0
environmental compliance equipment *	6	5	-1
investment molding equipment	4	3	-1
pouring	2	1	-1
blast cleaners	7	5	-2
other equipment	4	1	-3
material handling	10	5	-5
specialized equipment	10	5	-5
melting furnaces	12	6	-6
machining equipment	15	8	-7

* includes air handling

** other than molding equipment, includes sand reclamation

2.3 Differences In Perception Between Production And Sales

For 21 foundries, we received two different responses, one from the production manager and one from the sales manager. In general, the responses were quite similar, but there were a couple of noteworthy exceptions. Production managers report about a one week difference in process time (from order release to shipping) for the top four customers; sales managers do not report a significant difference, Table 4. Note that the production manager data included in Table 4 is only for those companies that also provided a sales survey, so these results will not be identical to those results reported earlier. Sales managers paint a significantly poorer picture of on-time performance than production managers, as can be seen in Figure 7.

Both groups of managers tended to agree on the reason why orders require expediting, Table 5.

Table 4: Processing Time Required, As Reported By Production And Sales

	Production Average	Sales Average
All Customers		
Days from Release to Shipping	36.4	36.4
Days from Shakeout to Shipping	19.5	22.1
Top Four Customers		
Days from Release to Shipping	29.9	35.6
Days from Shakeout to Shipping	17.8	17.3

Figure 7: Different Perceptions Of On-Time Performance

Table 5: Reasons That Orders Are Expedited

Production Manager			Sales Manager		
Reason		Pts	Reason		Pts.
customer needs sooner	(7)	98	customer needs sooner	(5)	102
customer late to order	(6)	90	customer late to order	(8)	99
production delays	(3)	75	production delays	(6)	75
tool/design changes	(2)	64	quality problems		65
quality problems		64	tool/design changes	(1)	57
equipment breakdown		48	equipment breakdown	(1)	57
lack of personnel	(1)	42	lack of personnel		28

3 ON-SITE ANALYSIS OF PRODUCTION SYSTEMS

In addition to the survey, more detailed information of production systems at several metalcasting facilities was collected by deploying student researchers for 4-12 weeks.

A major portion of the data collected was information on the flow of product through the plant. The product flow was based on observations by the student researchers and assisted by the plant personnel. At each company, the flow of product was summarized with flowpaths, of which a portion of an example is exhibited in Figure 8. The flowpaths include all of the process, inspection, material handling, and queuing steps.

Figure 8: A portion of a flowpath used to describe the steps within the production system.

3.1 Development of Flowpaths

The value of the flowpaths in understanding a plant's operations increase when you consider that at each plant, only 2 to 8 flowpaths account for at least 75% of the production.

Each of the companies makes a wide variety of castings; however, many castings have nearly the same sequence of operations, and could be summarized in the same flowpath.

These flowpaths varied greatly from the process plan sheets that many of the companies had on file for their products. The company's process plan sheets only included the scheduled operations and inspections. First, the flowpaths differed because they included material handling, which often is 80% of the total number of steps in the flowpath. Deviations from the company's process plan also occurred frequently, due to rework, attempt to keep all workers busy, equipment breakdowns, and supervisor preferences. We feel that the flowpaths are important, because it is these deviations from the expected process that adds unnecessary complexity to the system. This complexity results in increased work in process, lead times, and material handling.

3.1.1 The Four Kinds of Steps

Each step of the flowpath is put into one of the four subsequently defined categories:

- 1) process steps
- 2) inspection steps
- 3) material handling steps
- 4) queues.

Process Steps: A process step, denoted on the flowpath by a green circle, is any step in which some operation is performed which either creates or changes an object. For example, making a mold and pouring a casting are processes since they create new objects. Welding a casting and preheating a ladle are processes since they change an object. On the other hand, "cooling down" is not a process step because, even though a casting undergoes change, no operation is required in order to affect that change. Inspecting or changing the location of an object is also not a process step, since the object itself remains unaffected.

Material Handling Steps: A material handling step, denoted on the flowpath by a blue arrow, includes any operation that moves an object from one location to another. Perhaps the greatest contribution of the flowpaths is to quantify the details of each of these movements. For each material handling step, the mode of transportation (e.g. "fork truck" or "conveyor"), the horizontal distance, and the time taken by the move is included on the sheet. For manual moves, the vertical distance is also recorded. Note that only moves between, into, or out of stations were recorded. Moves that occurred within a station, such as repositioning or rolling over were not included.

Inspection Steps: An inspection is denoted by a purple/brown square; this is different from a process in that inspection is evaluating a castings quality, rather than changing it (other than possibly to temporarily mark defects). Inspection steps include visual inspection, hardness testing, magnetic particle, or other methods of assessing casting or mold quality. Many times, inspection stations are distinguished by the fact that the path forks in the two different directions associated with passing or failing inspection. The case of divergent paths is described below.

Queues: Queues are only indicated on the flowpath when one of the following three criteria are met: 1) castings must wait for an extended amount of time; 2) a significant amount of Work-in-Process inventory is accumulated in a particular area; or 3) castings enter a buffer zone between operations from which castings are not selected according to the FIFO ("first in, first out") rule. For example, if a batch of castings typically must wait for several hours or more before heat treatment, then, by criteria one, a queue exists. If waiting time is short (say a few hours), but there are a dozen bins of castings sitting in front of an area, then we have a queue by criteria 2. Finally, even if waiting time is short and the amount of WIP is small, if castings go into a waiting area where the order of processing is not simply a function of which castings arrived first (i.e. some castings can "pass" others), then we have a queue according to the third criteria.

Looking at the data as a whole, there are fewer queues compared to the total number of other steps (on average, a flowpath contains 55 material handling steps, 20 process steps, 11 queues and 3 inspections). This indicates that queues are not inserted any time some small amount of time elapses between two adjacent steps. The goal was to define a queue narrowly enough so that each identifiable queue is a worthy subject for further analysis. In general, if the area before an operation added complexity to the system, it was labeled as a queue.

3.1.2 Branches

Flowpaths often branch based on the results of inspections, and for a variety of other reasons (equipment availability, customer-specific requirements, or different supervisors changing the routing). In this case, our researchers attempted to estimate the proportion of castings that would follow either of these two (or more) branches. Generally, no data was taken to support these estimates. If we had no reason to believe that one path was more common than the other, each would be assigned a likelihood of 50 percent. In many cases, these loops represented rework cycles and thus one casting could go through the same part of the flowpath multiple times. In this case, the percentages often would vary depending on how many times the casting had already gone through the loop. The flowpaths reflect this by showing different estimated percentages for castings that have reached a branch for the first time, second time, or third time.

3.1.3 Why Flowpaths?

The flowpaths serve several purposes. For one thing, they are a valuable way to better understand the production process of each foundry. While foundries are already well aware of their own processes, we believe that the flowpaths nevertheless provide a valuable tool simply by filling in details that can easily be overlooked. Seeing the sequence of steps can sometimes be a basis of improvement suggestions, or at least identify parts of the flowpaths that are far from ideal.

Flowpaths facilitate analysis of material handling and of queues. First, they can be used to identify unusually time-consuming or expensive movements. While these are often a consequence of the foundry's layout (and thus not easily subject to change), they can also point out opportunities for reducing material handling expenses with minimal effort (e.g. changing the location of an inspection station, or shifting a grinder to do touch-up work).

After shakeout, the sum of all of the process times is generally much less than the time it will take that casting to get from shakeout to shipping. What accounts for this difference? Lead

time is largely a function of how much time a casting spends sitting in queues. Reducing lead time is very nearly equivalent in most cases to reducing the size of these queues. Thus, a strong incentive exists for further investigation of each queue identified on the flow chart.

Queues are not bad; rather, they perform several useful functions: they serve as a capacity buffers to ensure high utilization of people and equipment; they protect against unforeseen disruptions such as equipment failures. Queues allow certain operations (especially heat treatment) to be performed on large batches of castings, thereby lowering unit costs. A wise approach to queues is not to indiscriminately eliminate or reduce them, but to ask where queues are necessary and how big these queues need to be to fulfill their purpose.

Beyond the story told by the sequence of steps, the flowpaths are valuable because they assist in quantifying the costs of material handling for different types of castings. Since data was collected at twenty-one foundries, it also serves as a tool for measuring the performance of participating foundries relative to each other. Initially, we planned to simply record all of the material handling steps that occurred within a given sector of the foundry during a particular period of time. While this certainly afforded plenty of opportunities to observe the daily workings of each foundry, the data generated was not notably useful. The flowpaths brought a measure of organization to the collection of material handling data. They provide answers to relevant questions: How far does a casting travel? How many times is it moved by overhead cranes? How much personnel time is required to move a large casting from shakeout to shipping? While aggregate measures of material handling expense are readily available (e.g. number of employees engaged in material handling activity, annual maintenance expense for overhead cranes, etc.), the flowpaths have the added benefit of isolating costs for different types of castings and showing which moves are relatively expensive.

3.1.4 Limitations of Flowpaths

Due to some conscious decisions made at the start of the project, the flowpaths do not provide all the essential data of the casting production process. First of all, the decision was made not to time production workers, so no process times are available. Secondly, as was mentioned previously, very little data was kept with respect to the percentages of castings which passed/failed at various inspection points. Thus, the flowpaths do not give a comprehensive perspective on rework, although they do give some indication of its importance.

If a system was in place to collect this data, it could be easily added to the flowpaths, which would increase their value as a system improvement tool.

3.1.5 How the Data Was Collected

Horizontal distances between two locations were measured from the midpoint of each location. All times for manual material handling moves are averages based on direct observation. Times for mechanized material handling moves were developed in one of two ways. First, extensive data collection efforts were made to time various material handling steps; many of the times come directly from these measurements. Due to time limitations and the number of material handling steps on the flowpaths, some times were derived by developing equations which estimated the time required to move a casting any given distance via a particular material handling device (e.g. fork truck or crane). These equations were always based on data that had been collected at that particular foundry, and thus can be seen as a fairly straightforward extension of the recorded data. The calculated time would include a constant value (hook up, or fork truck loading) and a variable amount based on distance traveled.

There are some categories for which no data was collected. For example, no data was used to determine the "average casting weight" along the various flowpaths. These are estimates based on the observations of our researchers and their conversations with their contacts.

The material handling steps are described by the distance that the part traveled, the time consumed, the mode of transport, and the vertical distance for manual moves. The vertical distance is included for manual operations because large values could indicate potential ergonomic problems. From this collection of data, the material handling estimates can be made.

At certain points in the flowpath, castings could proceed in any of several possible options. These options are captured as branches in the flowpath, and the probability of each option is included, both on the flowpath and in the spreadsheet. In addition to calculating the material handling costs, this is being used to calculate the number of steps that a casting is expected to follow towards completion.

3.2 Results of On-Site Production System Analysis Studies

While much of the data requires little explanation, the "costs" at which are built upon several assumptions. Material handling costs include labor hours, as well as maintenance and depreciation costs of material handling equipment. Figures provided by the Steel Founders' Society of America Operations Committee Meeting to derive these costs included: a labor cost of \$25 per hour (including benefits); a fork lift cost of \$7,000 per shift year; and an overhead crane cost of \$36,000 per shift year. The value of a casting was assumed to be \$1.30 per pound.

The material handling costs presented understate the true costs because they only include the time of the actual moves. For example, the time for the fork truck to move to the next pick up, or sitting idle, is not included. Our analysis assumes 100% utilization of all the material handling resources. A conservative estimate would be to double the material handling costs that we report. Sample studies in some production areas indicate that a more reasonable factor would be 3 or 4 times the costs reported.

The following figures represent data from the twenty-one foundries. From each foundry, 2-8 flowpaths were used to capture the production system. There are seventy-four data points, each representing one flowpath.

Figure 9 through Figure 14 include data from the entire production system, while Figure 15 through Figure 18 only include data after shakeout. Figure 9 plots the average casting weight versus the number of material handling steps that the casting encounters. Note that the average casting weight is the average weight of the castings represented by that particular flowpath. Figure 10 shows the horizontal distance traveled by the castings. Figure 11 shows the expected cost of the material handling, plotted versus the average weight of the castings represented by the particular flowpath. Figure 12 highlights the material handling costs for the castings with weights under 500 pounds. The percentage of the casting value consumed by the material handling steps is shown in Figure 13, and again in Figure 14 for those castings less than 500 pounds.

Since the majority of the material handling occurs after shakeout, the following graphs only consider this portion of the production process. Figure 15 and Figure 16 show the material handling costs. The percentage of casting value consumed by material handling of the castings after shakeout is plotted in Figure 17 and Figure 18.

Figure 9: Average casting weight versus the number of moves for the total production process

Figure 10: Average casting weight versus the horizontal distance traveled, for the total production process

Figure 11: Average casting weight versus the material handling cost, for the total production process.

Figure 12: Average casting weight versus the material handling cost, for the total production process

Figure 13: Average casting weight versus the percent of casting value consumed by material handling costs, for the total production process

Figure 14: Average casting weight versus the percent of casting value consumed by material handling costs, for the total production process

Figure 17: Average casting weight versus the material handling cost, for the post shakeout portion of the production process.

Figure 18: Average casting weight versus the material handling cost, for the post shakeout portion of the production process.

3.3 Definitions of Rework

Clearly, the post shakeout steps are responsible for a large portion of the material handling costs. A major reason behind this is the rework of the cleaning operations that occurs. The cause for the rework may be due to problems prior to shakeout (namely molding, coremaking, or metal quality in the mold), or the problems may occur during the post shakeout operations (such as arc air, riser removal, grinding, and welding.) Regardless of the source of the problem, the cleaning operations typically have the responsibility of resolving the issue.

At the project review meetings held in conjunction with this project, there were several discussions on what was meant by 'rework,' for this industry. Three definitions have prevailed during these discussions:

1. Any work that is in addition to what was planned for a particular castings
2. Any work in addition to the bare minimum that is needed to remove the gates and risers
3. Any work that is in addition to that required by the best X% (e.g. 5%) for that particular casting.

The advantage of the first definition is that it takes into account that it may be more beneficial to expend additional resources on cleaning, instead of changing the pattern. This would be particularly appropriate for short run castings. Definition 2 leaves no ambiguity on what is rework. By measuring rework as any effort expended past the bare necessity would drive for an overall reduction of finishing time. The third definition takes into account that some castings turn out better than others. This definition would encourage a foundry to search out the variables that caused other castings to come out worse, and eliminate the problem. The third definition does not take into account, for instance, how much time a particular casting would take at another foundry. The absolute measurement attained by Definition 2 does consider this, so that a company would not develop a false sense of confidence in their performance. One could argue that these definitions are somewhat academic, and not practical on the shop floor. However, a change of philosophy needs to occur within the industry to overcome the problem of excessive rework.

Rework, by any definition, is obviously causing throughput times, WIP, manpower requirements, and material handling needs to increase. However, without a method to track the history of a casting, the magnitude of the problem will go largely unnoticed. Systems need to be implemented, such as a barcoding system, which enables this critical information to be collected and utilized. With this knowledge, problem castings can be identified, and specific problems can be eliminated.

3.4 General Recommendations

The data collected has not yielded a "silver bullet": a simple fix that could be applied to any complex system to dramatically reduce costs. However, the flow path data, along with observations made by the field researchers, did lead to some recommendations:

- Material handling costs can be reduced by moving people to castings rather than moving castings to people whenever possible. This is especially relevant for inspections.
- Typically, castings must be sorted at least once after pouring (often at shakeout). Additional sorting operations increase material handling costs unnecessarily.

- Whenever possible, a good goal for each department is to move castings directly to the next functional container so that they are properly organized for their next process. For example, smaller castings should be placed directly on a heat treat rack after the process immediately preceding heat treatment. This eliminates several moves, and shortens the lead-time.
- Besides increasing lead times, re-work substantially adds to material handling distance and cost. Some foundries track the percentage of castings that fails an inspection and thus requires additional work; but most do not. The impact of re-work justifies making an effort to reduce it. For the success of such efforts to be measurable, a system must exist for quantifying re-work.
- Queues serve a valuable purpose as buffers; they help to ensure high utilization of key equipment and personnel. On the other hand, for most castings, the majority the time between shakeout and shipping is spent in queues. Reducing lead time and WIP requires working to reduce queue size by analyzing where strategic buffers are needed and how large these buffers should be. The size of the queue should only be large enough to keep key personnel and equipment busy most of the time.
- The survey of steel foundries revealed that many companies were planning to construct new buildings. At many of the companies that we worked with, there were many opportunities for ‘finding’ space by eliminating unnecessary queues. In addition, layouts that we have developed to reduce material handling also reduce the amount of floor space required.
- Since material handling often consumes more than 10% of the casting value, it deserves to have personnel responsible for reducing it. The detailed flow paths have proven to be an invaluable tool for identifying areas for improvement.
- These personnel who are dedicated to improving the production system can also work on other system issues. Examples include, determining the optimal size of the buffers, and developing mechanisms to determine and reduce the amount of rework and scrap.

4 LAYOUT IMPROVEMENTS

Facility layouts can play an important role in having an efficient production system. However, some of the companies that were studied had very poor facility layouts, but their material handling costs compared to companies producing similar castings were low. Nevertheless, it takes much greater production system discipline to manage a facility with a poor layout. Even for these companies with a relatively good production system, additional savings could be realized with a better layout.

An ideal time to consider making layout improvements is when large new pieces of equipment are being considered. The tendency has been to locate these pieces of equipment where the old one was located, or where the other similar pieces of equipment are located. One company used the flowpaths developed from this project to determine the optimal location for a new blast cleaner. The best feasible location was 200’ from the original intended location, which was where the machine being replaced was located. This change saved the company an estimated 1500 miles of fork truck moves annually.

Several facility layouts were developed for participating companies. Figure 19 displays one of these layouts, which represents many of the plant layout improvements that would be

advantageous. In this layout, castings originate from the blast cleaners on the right side and proceed on 2' wide roller conveyors. The castings would be placed on shallow bins or transport boards. These bins or boards can be rolled directly into the process areas, and can be easily accessed by the operator, minimizing ergonomic problems. The stations are organized in a 'cellular' format. This means the different cells contain all of the equipment necessary to complete the casting. For instance, the cell could include the different types of grinding, arc air, and welding. Traditionally, many companies have organized their plants with all of stand grinding in one area, hand grinding in another, and welding in yet another area. Cellular layouts will minimize the material handling, because the casting never has to leave the cell to be finished. The central area would house all of the inspection operations. In this location, the inspectors can monitor the output of all of the cells, and work with the operators to make any changes in their procedures. This is in contrast with having inspection at a remote location, with little or no communication with the process operators. This layout also separates the castings being cleaned before heat treatment (top of picture) from those that are coming after heat treatment (bottom of the figure.)

Figure 19: Proposed layout for the cleaning area of a metalcasting plant that produces castings in the approximate weight range of 5-150 pounds

Similar to the improvements recommended for a facility producing castings less than 150 pounds, cellular layouts were also proposed for areas cleaning larger castings that are moved via jib cranes. Figure 20 displays the traditional layout organized by the type of process. Figure 21 shows a cellular layout. The castings enter a cell, which includes all of the grinding, welding, arc air and inspection operations within it.

Figure 20: The current typical layout which has all of the same operations grouped together

Figure 21: Improved layout of a metalcasting cleaning area for larger castings that are moved by jib cranes.

5 USE OF PRODUCTION SIMULATION

The research team has used production simulation modeling to investigate the role that product variability plays in the steel casting process. Specifically we have concentrated on cleaning room operations and how the throughput time of a casting is affected by varying degrees and different types of variability. The goal of this part of our research is to define quantitatively the role of variability and how much it negatively affects the casting process and the final product of the process. We have worked closely with a large member foundry to collect data and other information about the process under investigation. The research is a case study of a specific foundry, but the methodology and results are applicable to other foundries and the steel casting process in general. We used Arena 3.0 discrete event simulation modeling software as the primary tool for our research.

5.1 Overview of Discrete Event Simulation

Many types of simulation are used in industry as tools to model the performance of systems or to test the effect that a variable has on a system. In our research we are using discrete event simulation, which is commonly used in such areas as manufacturing system design,

inventory systems, communications systems, and transportation applications. Discrete event simulation models a system as it changes over time. The goal of the model is to gain some understanding about how the corresponding system behaves. In some cases when the model is simple enough it can be modeled mathematically. In many cases however, the system of interest is extremely complex and the only realistic modeling technique is simulation.

Events within the system are defined as occurrences that may change the state of the system. In our model an event is usually the movement of a part from one place to the next in its flow path. The state of the system is defined as the collection of variables necessary to describe a system at a particular time. In our model we typically use for simplicity the variable 'total number in system' to define the state though there are several variables that describe the system at a point in time. For example, the number of parts at a certain station gives information about the utilization of that resource and if the parts are moving through the system in a balanced way. The discrete model includes a simulation clock that determines when various events occur in the system. At each interval of the simulation clock the state of the system changes based on the events that occur at that time. All of these events are tracked in detail by the model and are available as outputs for analysis.

5.2 Steady-State Analysis of Simulation Output

At the beginning of the model (as the simulation clock is just starting) the model is said to be in its transient-state. That means that the number in the system is very low and is not yet reflecting the level that is needed to reflect the actual system. The transient-state lasts until the model 'fills up' and enters into the state called steady state. Steady-state means that the steady-state variable (total number in system) will have approximately the same distribution at any time. In other words, the model is reflecting the actual system and the variable for number in system reflects this. It is important to differentiate between the transient-state and steady-state since the output would not be valid if sampled from data that is taken during any transient-state.

5.3 Model

We have designed a detailed model of all post-shakeout operations for a foundry and are using real data and production schedules as inputs to the simulation model. Following are the main features of the model. Note that casting refers to a particular model of a product, and a part is the individual pieces.

1. Each of the castings has its own unique flow path through the cleaning room.
2. Each path begins after shakeout and ends at shipping.
3. An actual schedule for one year of production from the member foundry is read into the simulation model. By using this historical information versus arrival distributions, the model is called trace-driven.
4. Included in the schedule for each casting are process time parameters for each station on its flow path. These process times reflect the combined process and loading time for a part at a station.
5. All of the parameters for process times in the schedule can be changed to reflect different values. In this way, the process time distributions can produce different levels of variability.
6. There are three work shifts at our model foundry and those are programmed as well so that stations are only in operation for a certain number of hours per 24-hour period.

There is also a model feature that implements overtime work at stations if/when they exceed a certain number of parts.

7. The queue discipline used throughout the model is earliest due date first (i.e. of the parts waiting for a certain operation, the part that is due to be shipped the earliest is always selected first for processing).
8. At the blast stations, the batches are selected as follows. The part with the earliest time due is selected along with all parts of the same casting type. There is a defined station capacity for weight in the model, and the total weight of the parts in the batch cannot exceed the defined capacity. These parts are processed as a batch and the batch process time is the sum of the blast process time for each part.
9. The heat treat car and temper stations also process parts in batches. The batches are selected somewhat like the blast stations, with one main exception. If the most urgent part and all like it do not fill the HT car to its weight capacity, then parts of similar size are selected to complete the batch. It is wasteful to process a HT car that is not at full capacity since almost the same amount of energy is used regardless of the batch size.

All of these features are based on observations and data received from the foundry and are representative of the actual system.

5.4 Output

Each part is tracked in detail as it proceeds through the simulation model. The following information is included in the simulation output:

1. Wait time at each station
2. Process time at each station
3. Total throughput time for each part
4. Station utilization
5. Average number of parts at a station
6. Batch sizes for Heat Treat and Blast operations
7. Total number in system

These data allow us to analyze the casting process in a number of different ways. Figure 22 represents the average time for each part in the system. The data was sampled after the model had reached its steady state. Note that this figure represents example data because the actual data represents proprietary. All times are represented in minutes.

Figure 23 shows another example of the valuable results that can be gained from a simulation model. This figure displays the utilization rate for a variety of key resources in a cleaning department. This data, for example, could be used to make decisions about future capital expenditures, staffing, and product schedules.

5.5 Data Analysis

Preliminary analysis of the data shows a high percentage of the total throughput time (TTPT) to be made up of the total non-process time (TNPT). Also, it seems that the castings with the higher total process time (TPT) do not necessarily have a higher TNPT and thus a higher TTPT. The TTPT seems to be most affected by two things, castings that have longer flow paths

and castings that visit stations that have high process times. The longer flow paths force a part to wait at more stations and the high process time stations force the parts to wait longer. .

Figure 22: Simulated Throughput Time shown by Part Number

Figure 23: Simulated utilization rates for various resources in a cleaning area.

5.6 Continued Research

This model will also be used to test different schedules, as described in the section on scheduling below and determine how various pouring schedules can lessen the negative impact of downstream operations and variability on the throughput of a casting.

The relationship between wait time and process time at a station give insights about the bottlenecks that exist in the process. Also of interest will be the affect on the above data as the level of variability is increased in the model. We are able to quantify the affect that variability has on the model and more generally on the casting process.

6 SCHEDULING

As a part of this project, we have studied the process by which the foundries develop their daily pouring schedule. Typically, the foundries are unable to consider the impact of their pouring schedule on downstream operations because developing a feasible schedule is a labor-intensive task involving numerous constraints. To address this, we are working with ZYX Co. to develop software that can both automate the current scheduling process and improve it by estimating the impact of the pouring schedule on downstream Work-in-Process (WIP) inventory levels. An integer-programming model minimizes a comprehensive cost function that includes the costs of pattern tooling set up, late delivery, WIP inventory, and under-utilization of assets. The software implements a heuristic that finds multiple solutions to this integer program, each of which corresponds to a feasible schedule. This software is capable of handling realistically sized scheduling problems in a reasonable amount of time. Upon its completion, this system could be extended for implementation at many other steel foundries.

6.1 Introduction

The authors observed and chronicled scheduling practices at many steel foundries that have participated in this project. The majority of these foundries schedule molding, melting, and pouring operations manually. In response to the survey, several foundries acknowledged that

they used software to assist in scheduling, but the nature of the software varied significantly. Some foundries had developed their own software, while others listed a spreadsheet as their scheduling tool. The company that produces the most commonly mentioned commercial package informed us that their software does not determine the initial sequence of jobs. Instead, after receiving this initial sequence as input, it helps coordinate subsequent operations. These responses seem to indicate that there is no widely accepted software system that has essentially solved the pouring scheduling problem across the industry. Similarly, textbook scheduling approaches are poorly suited for generating valid pouring schedules. For example, foundries are related to the traditional job shop in that they produce multiple products, which each require several sequential operations. However, the standard job shop model cannot incorporate realistic requirements such as flask availability, melting constraints and the coordination of molding and pouring. The lack of prior success in scheduling steel foundries is shown in the fact that we were unable to find any such references in the literature.

Approaching foundry-wide scheduling from the perspective of the daily molding and pouring schedule seems to be an attractive alternative for three reasons. The first reason is that, according to the survey, the majority of foundries do not consider downstream operations. This practice can lead to temporary bottlenecks, even when upstream operations have smaller capacities on average than downstream ones. For example, let's consider a foundry with the capacity to melt 20,000 tons per month and the capacity to heat treat 25,000 tons per month. If this foundry only produced a single type of casting, then heat treatment would never be a bottleneck (barring breakdowns, or other contingencies). However, if this foundry made a variety of different castings each week, and one week's pouring schedule includes an unusually high number of castings which have especially high heat treatment requirements, then this would potentially cause a bottleneck in front of the heat treatment operation in succeeding days [Johnson, 1998]. Since, as in the example, capital-intensive upstream operations will tend to have smaller capacities than downstream operations, such bottlenecks are often difficult to anticipate. Unanticipated bottlenecks are especially disruptive since no plans can be made ahead of time to deal with the upcoming problem.

While the WIP inventory described above should be transient, it can be expensive nonetheless. According to the survey, the average number of days between the time a part was separated from its mold until the time the part was shipped was 18. The average part would require no more than three days to complete, were it being worked on continuously. Thus, parts are spending, on average, approximately two weeks waiting in queues. Since WIP and throughput times can be a result of unsophisticated techniques for developing a pouring schedule (such as the example given above), it makes sense to address this by proposing an improved pouring schedule methodology.

For many foundries, which manufacture a variety of castings, developing the daily molding and pouring schedules is a labor-intensive process. It is not unusual to have an employee whose primary (perhaps only) task is to create these schedules. This illustrates that this task is certainly not trivial and underscores the benefits that could be derived from automating this process.

Finally, most textbook approaches to scheduling (e.g. the job shop model) require a much greater availability of production data than would be attainable at most of our partner foundries. Although one might argue that a foundry needs to concentrate on collecting such data before attempting to upgrade the scheduling process, our goal was to create a scheduling tool that can provide assistance in the near future, even in the absence of detailed data. Although the quality

of the solutions generated is dependent on the quality of the supporting data, the aggregate nature of the data requirements greatly reduces the data collection burden. This is particularly important in this case where the castings produced are constantly changing.

6.2 Modeling Considerations

Developing a feasible pouring schedule is difficult because several constraints must be satisfied. The research team worked very closely with a partner foundry, to be identified in this report as ZYX Co. This company's scheduling system and constraints was pretty typical of the foundries we worked with, and they were interested in improving their scheduling method. Each casting produced by ZYX Co. requires a flask to contain the mold. This flask is required until the part has solidified and cooled, which can be up to two days. For ZYX Co., the most difficult scheduling complication is the availability of copes and drags with the required dimensions. An additional complication is introduced by metallurgical requirements. Some alloys cannot follow other alloys in the furnace, because of contamination. Thus, the scheduler must consider the sequence of heats associated with a particular furnace. Since molding and pouring schedules must be closely coordinated, a third consideration is molding capacity. Furnace capacity is also a consideration; the pouring weight of each heat must fall between a minimum and maximum value for the heat to be feasible.

A successful scheduling model must have an objective function that reflects the goals and values of the foundries that use it. This model's objective balances three competing criteria. First of all, it assumes that there is an incentive to complete orders on time, and assesses a penalty for late orders. Second, it incorporates opportunity costs, which can be assessed when a foundry does not fully utilize an important asset. For example, a foundry might have two melting furnaces: a more efficient one that they desire to fully utilize and another that is used only when needed. In this case, a cost would be assessed if the efficient furnace is not fully utilized, but no cost would be assessed if the other one were idle. Similar costs can be assigned to molding, heat treatment and cleaning lines. The third type of cost is the holding cost associated with projected WIP at downstream operations. For example, WIP costs would discourage scheduling several parts with high cleaning requirements and with the same cleaning line, since this would lead to a bottleneck and a high level of WIP.

6.3 Data Requirements

Since our goal was to develop a scheduling tool that does not require detailed production data, the data required is relatively minimal, and essentially consists of information which is either readily available or can be estimated without extensive time studies. The program reads data from four files; the first holds a list of current orders, including the type of casting, the quantity to be delivered and the date on which the parts are due for delivery. Orders can be assigned a priority value of 1 for "hot" orders, or -1 for orders that should not be poured.

The second file is a master data file for every casting, which the foundry has produced (or, at least, has on record). Although this file could potentially be extremely large, the program only stores the production data for castings for which there is a current order. In this way, storage concerns are largely eliminated. Most items on this file are standard production information:

- the molding line on which the casting is produced, and the time required to make its mold
- the flask models (cope and drag) used to hold the mold (if a flask is required)
- the pouring weight of the casting

- the alloy of the casting
- the cleaning line to which the casting proceeds for finishing operations (essentially, this includes all operations after the casting is removed from its mold with the exception of heat treatment)
- the number of parts produced per mold

Each of the above items are relatively straightforward, although in some cases (flasks, for example) there could be more than one possible alternative. Also included in the second file is information on downstream operations, where information is less precise. Detailed data is generally not available, and rework and the inherent variability of the casting process make it difficult to accurately predict the times required for individual operations. Thus, rather than requesting specific times for each operation, the model seeks to get an accurate sense of the expected aggregate resources required to finish a particular part. Thus, only a single number is expected for the labor hours required to complete all cleaning operations.

One limitation of the model is that it completely ignores variability in cleaning processes. If the time required to complete each process were independent of each other, then the relative variability (measured by coefficient of variation) of the entire cleaning line would be less than that of the individual processes. However, if process times are not independent (which is more likely since cleaning times are all related to the initial quality of the part), then variability is more problematic. Additional research into process variability could be beneficial and might significantly impact scheduling. At this point, no attempt is made to include variability for two reasons: first, none of the foundries we have worked with have data to support this type of detail; and second, methods to incorporate the variability of downstream operations into a deterministic pouring model require additional research and analysis.

The second file also includes the heat treatment requirements, which are dealt with separately in the program, since this resource frequently is a bottleneck. The reason for this area being a problem is not due to lack of total capacity, but because of uneven levels of work in this area. The relatively long processing times (ranging from several hours to a day) also causes scheduling problems. Heat treatment requirements are straightforward in the absence of rework. However, unlike cleaning requirements, heat treatment requirements cannot be simply measured in units of time. The size of the casting is also an important factor in determining the heat treatment resources it requires. For example, assume that casting A is a medium-sized casting that requires 8 hours in the heat treatment furnace; ten of these parts can undergo heat treatment simultaneously. Casting B also requires 8 hours in the furnace, but, due its large size, it takes up the entire furnace. In terms of the resources consumed, the heat treatment requirement of casting B is equivalent to ten of casting A. The dimensions of the molding flasks are used to estimate the space consumed by the part.

The third file provides information regarding the cost structure and the capacity of the foundry. Capacity information is relatively straightforward:

- maximum number of heats per day for each furnace
- minimum and maximum weight for each furnace
- maximum number of molds per day for each molding line
- the daily capacity of each cleaning line (in terms total labor hours available)
- the daily capacity of heat treatment (in ft³-hours)

Also, an estimate of the current WIP level at each cleaning line and in the heat treatment department is required (using the same units as their capacities).

Cost information is more difficult to accurately estimate (e.g. what is the true cost of failing to ship an order on time?). However, since these are among the factors that differentiate good schedules from poor schedules, an attempt at quantifying costs is necessary for implementing an optimization approach. Since costs can easily be adjusted, users have the opportunity to experiment with different values until finding ones that generate desirable results. Costs include opportunity costs for failing to fully utilize a resource (i.e. a particular furnace, molding line, cleaning line, or heat treatment resources), holding costs for WIP inventory, lateness costs, costs for postponing "hot" orders, and fixed setup costs associated with beginning production of a particular casting.

6.4 A Mathematical Model for Pouring Scheduling

Our model is an integer program whose structure facilitates a natural expression of the various constraints that complicate the scheduling process. Two types of integer variables are defined, corresponding to the two basic decisions made in developing a pouring schedule. The first decision is the sequence of alloys to pour from each furnace. The second decision is the particular castings to produce from each alloy assigned to a heat. The model's variables can be expressed as follows:

$$z_{i,j} = \begin{cases} 1, & \text{if alloy } i \text{ is poured during heat } j \\ 0, & \text{otherwise} \end{cases}$$

$$y_{k,j} = \begin{cases} 1, & \text{if casting } k \text{ is poured during heat } j \\ 0, & \text{otherwise} \end{cases}$$

While the software can generate schedules for several days, it generates pouring schedules one day at a time. Its objective function quantifies the trade-offs between high asset utilization, WIP inventory, on-time delivery, and prioritizing hot orders. All costs are user-assigned and (except holding costs) are allowed to be set equal to zero. Constraints ensure that

- each heat can include only castings from a single alloy
- all orders must be produced, either by placing them on the current day's schedule or postponing them
- the weight of each heat must be within a proper range
- the parts on each daily pouring schedule must not violate daily molding capacity
- sufficient cokes and drags are available to hold the molds needed for the daily schedule
- parts poured today will impact future WIP in downstream operations
- lateness is computed based on the level of WIP in downstream operations

A few additional points are worth noting. First, the model neither generates a molding schedule nor attempts to explicitly coordinate molding and pouring schedules. While it is obvious that the molding department must stay ahead of the pouring department, the timing of when the molds should be made (e.g. should molding be one day ahead of pouring? or a half-day?) is left to the discretion of the foundry. Instead, the model creates a framework within which effective coordination should be achievable. Since the program requires that the molds needed for the scheduled castings can also be produced in one day's time, it seems reasonable to assume that the day's pouring schedule does not put impossible demands on molding. As long as

a sufficient buffer period is maintained, molding should be able to complete the molds needed for the daily pouring schedule.

Secondly, in response to ZYX Co.'s comments, we are working to extend the model to reflect the fact that more than one day can elapse between pouring and shakeout, in which case a flask cannot be used in consecutive days.

Finally, the costs of postponing late orders are calculated by multiplying a penalty constant by the square of the order's lateness. This is done so that the marginal cost of failing to complete a late order continually increases, thereby ensuring that an increasingly late order will become an increasingly high priority.

6.5 Solution Heuristic

Due to the number of integer variables in this formulation, finding an exact "optimal" solution for real problems does not appear to be a practical goal. Although the scheduling program implements a branch-and-bound algorithm of sorts, it limits the number of open branches. The heuristic finds feasible pouring schedules in two steps. The first step generates multiple feasible alternatives for assigning alloys to heats (i.e. it finds multiple feasible integer values for the z_{ij} variables). The second step generates a pouring schedule (i.e. the castings produced during each heat) for each of the alternative assignments.

Each iteration t of the alloy assignment procedure does two things. First, it compares the costs of each alternative assignment of the first t heats. For each of these assignments, the heuristic solves a linear programming (LP) relaxation of the original integer program (IP). The optimal objective function values of these LPs provide a lower bound on the optimal value of the IP, given their partial assignment. The quality of these alternative assignments is determined by these lower bounds. The parameter γ prevents exponential growth in the number of alternatives. Only the γ partial assignments with the smallest lower bounds are explored further. High values of γ allow more alternatives to be explored while lower values lead to decreased CPU time.

Secondly, each of the γ best partial assignments generates multiple "children", each of which extends its "parent" partial assignment by assigning a different alloy to heat $t+1$. Rather than considering all possible alloys, the parameter β limits the number of alternative alloy assignments to each heat. The LP solution is used to predict the desirability of assigning any given alloy i to heat $t+1$. Alloys are ranked based on their LP solution values, and the β highest-ranking alloys are assigned to for heat $t+1$. As β increases, the procedure considers a greater variety of scheduling alternatives with the consequent increase in CPU time.

To illustrate how these parameters restrict the number of alternatives, consider how the algorithm proceeds with $\beta = 3$ and $\gamma = 5$. Three alternative assignments of the first heat would be evaluated. If each of these assignments was feasible, each could generate three alternative assignments of the second heat. Thus, iteration 2 could consider nine alternative assignments for the first two heats. Only the five assignments with the smallest bounds would be further explored. Each of these could generate three assignments of the third heat and iteration 3 (and every iteration thereafter) could explore 15 alternatives. Thus, the total number of LP relaxations solved during the course of the alloy assignment procedure is bounded by $\beta \times \gamma \times$ (the number of heats).

When the assignment algorithm is completed, it has identified several alternative assignments of all of the heats. The second phase of the heuristic focuses on exactly which castings to produce for each of the alloys that are assigned to at least one heat. For each assignment, an LP is solved to determine two things: which castings are candidates to appear on

its final pouring schedule and how many resources (e.g. molds, cleaning time, flasks) will be allocated to each alloy. After solving this LP, the second stage of the heuristic decomposes the scheduling problem by alloy and solves separate integer programs for each alloy. Its objective function minimizes the sum of setup costs, postponement costs, and the costs assigned to unutilized molding capacity. Constraints enforce the weight requirements for each heat and reflect the inter-dependence of the various alloys, particularly how they compete for scarce resources. These are handled in two different ways. Molding and flask constraints must be satisfied in order to generate a workable schedule and are therefore included. However, the right-hand side values of these constraints are reduced by subtracting the resources allocated to castings of other alloys. Constraints related to downstream operations do not impact a schedule's feasibility and therefore are deleted. Obviously, this represents a discontinuity between the solutions generated in the assignment LP and the solutions of the IPs for each alloy. However, to ensure some consistency between these solutions, castings whose LP solution values equal zero are not included in the IP.

The parameters β and γ restrict the number of candidate schedules to be no greater than $\beta \times \gamma$. Although several IPs must be solved for each alloy assignment, limiting the allowable castings to those with nonzero LP solution values allows these IPs to generally be solved quickly using a standard branch-and-bound procedure. Furthermore, feasible schedules are readily available since non-integral IP solutions can be made feasible by setting each fractional variable to 1. The solutions of each alloy IP are combined to generate feasible solutions to the original IP model. Cost information for the candidate schedule is displayed in an output file to facilitate easy analysis. Regardless of their costs, each candidate schedule is included in this output to give users the opportunity to evaluate each alternative. The candidate schedule with the lowest total cost is displayed at the end of the program as the recommended schedule.

The heuristic does not guarantee an optimal solution to our formulation of the pouring scheduling problem. However, considering the large degree of uncertainty with respect to the parameters that quantify downstream requirements, and the inherent difficulty in quantifying the costs associated with late orders or underutilized assets, the value of achieving "optimality" is questionable. Thus, generating and evaluating several candidate solutions is preferred to spending extensive computational effort in an attempt to find a single optimal solution to the given model.

6.6 Computational Experience

The goal of this project was to develop a procedure that could generate good pouring schedules for large-scale problems in a reasonable period of time. Table 6 shows how much CPU time is required to obtain a daily schedule for ZYX Co. for different values of β and γ . These problems scheduled 13 total heats, and used inputs that were based on historical data, including 290 orders for 88 distinct castings, segregated into 32 alloys. All computational tests were run on a Hewlett Packard Kayak XU with 64 megabytes of RAM.

Table 6: Daily Schedules Can Be Generated in a Reasonable Amount of Time

β	γ	CPU Time (minutes)
2	2	11.3
2	3	16.5
2	6	29.4
3	3	23.2
3	6	42.5

Since scheduling is done one day at a time, the time required to generate a weekly schedule would be approximately five times the time required to find a daily schedule. Thus, the program is able to find schedules for realistically sized problems in a timely fashion. The costs of the schedules generated by the software is not easily compared to those generated manually since manual scheduling makes no attempt to costs of any sort. Our hope is that each alternative would be a workable daily schedule, and that, by comparing relevant cost information, these alternatives could quickly analyzed. Furthermore, providing foundries with information on how today's pouring schedule will impact downstream operations could be beneficial in its own right.

6.7 Future Work

ZYX Co. evaluated the quality of the schedules generated by the original program running off of historical data. This review identified necessary extensions to the original mathematical model that are currently being addressed. Furthermore, in the interest of demonstrating the viability of this general approach, certain modeling issues have been avoided (e.g. core-making must be closely coordinated with molding and pouring). Thus, we are aware of opportunities for strengthening the model should it be adopted for daily use. Our goal is that these initial tests will lead directly to the completion of practical scheduling software that can be incorporated into their daily planning process.

Besides surmounting all the obstacles that must be overcome to bring the software into industrial use, there are significant optimization challenges involved in developing these schedules. Improved methods for formulating and solving the integer programs could perhaps substantially improve the speed of the program and the quality of its results. Finally, our initial work has focused almost exclusively on generating and evaluating pouring schedules. In the future, our desire is to further analyze the results of the linear relaxations and how they can be utilized to yield additional useful information regarding the value of scarce resources, and the costs imposed by current system constraints.

7 TECHNOLOGY TRANSFER

Several different methods of technology were used for the transfer of the information and findings from this project.

The most direct method of technology transfer was the actual deployment of student researchers at the participating foundries. The foundries that volunteered to host the students were the companies who were most likely to gain from the project. The researchers provided an outside and unbiased look at the operations. The researchers interacted with personnel at all levels of the company. In doing this, they were able to provide advice, identify strengths, and

offer suggestions. The senior personnel on the project also visited each of the host companies to observe and make recommendations.

Steering committee meetings were held which were opportunities for the researchers to present their findings. In addition, these meetings provided a forum for the team to offer suggestions, and for the committee to modify the ideas into workable solutions. Four such meetings were held during the project period, plus an additional meeting immediately before the project officially began.

In addition to the steering committee meetings, several other presentations were made to the industry. These included the annual meeting of the Steel Founders' Society of America, three presentations (plus another planned) to the Steel Founders' Society of America Technical & Operating Conference, the North Central Division of Steel Founders' Society of America meeting, the Eastern Division of the Steel Founders' Society of America, the Wisconsin Chapter of the American Foundry Society, the Northwest Chapter of the American Foundry Society, and the American Foundry Society Casting Congress. Articles for *Modern Casting* magazine are planned, with one in press.

A half-day hands on workshop was also presented to the member companies of the Steel Founders' Society of America. This workshop was used to demonstrate some of the production system principles that we were advocating for their facilities. Over eighty personnel participated in this event, which was very well received.

8 INDIVIDUAL FOUNDRY RESULTS AND IMPLEMENTATIONS

During the three-year project, twenty-one companies hosted student researchers that studied the company's production system for a period of 4 to 12 weeks. The following are descriptions of those companies that have changed or are changing their production system since their project participation. The company's names have been excluded for confidentiality reasons. Three companies have hired manufacturing engineers as a result of this project, and their influence on the company's production system are included in the following summaries.

- A large steel foundry in the Midwest producing carbon and low alloy castings has made significant changes to their production system as a result of the project. These changes have eliminated the need for a planned expansion; it has decreased the labor expended on products making the company more competitive, and reduced WIP and lead-time. The company created a new position of manufacturing engineer, and hired the student researcher that was deployed at their facility, who is behind many of the company's past and future production system changes.
- A medium sized steel foundry located in the southern tier of the United States that hosted a student researcher continues to benefit from the study. The company has also played an active role in the steering committee sessions. This involvement has begun to change some of the philosophical barriers to making production system improvements. The company has plans to implement a major change in their layout and production system, based on the recommendations of the project team.
- A medium sized Midwestern foundry, which hosted a student researcher, has purchased components to implement a new layout based on the project team recommendations. This new layout should realize significantly reduced material handling costs, and allow

the company to respond much more quickly to orders, reduce WIP, and increase the quality of the products.

- A large steel foundry in the Midwest, which produces a wide variety of casting alloys, hosted a student researcher during the project. As a result, they have hired an industrial engineering, the first for the company, from Iowa State University as a plant engineer. In this position, the engineer is implementing production system changes with respect to new equipment installations and improvements. They have utilized the results of the project to assist them in making changes.
- This large steel foundry, located in the east, has used the results of the student researcher on site to help focus on areas to improve. Since the project, the company has made a major purchase of automation equipment to eliminate a troublesome area of their production system.
- A medium sized steel foundry in the eastern United States that did not directly participate in the study has also gained benefits of the project. The company changed their hiring plans as a direct result of one of the technology transfer presentations that was made as part of the project. They decided to hire a manufacturing engineer, which worked on the project, to improve some of the production system problems they were having. The company decided that an industrial engineer, with the experience gained from the project, would be a good fit into their organization. This engineer has been very successful in developing and implementing many programs that has reduced manufacturing costs, lead-time, WIP inventories, and improved product quality.
- This small steel foundry in the South is planning a new addition at the time of writing this report. The company is continuing to work with the project team to develop and implement a new layout, which will improve their production system. This company hosted a student researcher.
- This small Midwestern foundry used the results of their project participation to help justify the need for an expansion of the metalcasting facility.
- This western, mid-sized foundry has used the results of the study to plan future changes to their manufacturing facility. At the time of this writing, they are looking to change an area of the plant, which the on-site researcher identified as being problematic. The project team continues to work with them on making these improvements.
- This mid-sized Midwestern foundry has been a very active participant in the project. They have sent several different managers and floor supervisors to the project review meetings, and hosted a student researcher. They have actively used the results and suggestions, and have seen dramatic improvements in their production systems. Improvements include reduced labor costs, significantly less WIP, less material handling, and less lead-time. Among other benefits realized from the project participation, the project helped them to focus on the need to reduce material handling in order to improve their system.
- Three mid-sized foundries have recently participated in the project, so we are unable to assess their changes as a result of hosting a researcher. However, the companies' interest in utilizing the results has been noted.