

ORDER COMBINATION METHODOLOGY FOR
SHORT-TERM LOT PLANNING AT AN
ALUMINUM ROLLING FACILITY

by
DAVID P. VENTOLA

B.S. Chemical Engineering, Princeton University (1985)

Submitted to the
Department of Chemical Engineering
and the
Sloan School of Management
in Partial Fulfillment of the
Requirements for the Degrees of

MASTER OF SCIENCE IN CHEMICAL ENGINEERING

and

MASTER OF SCIENCE IN MANAGEMENT

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June, 1991

© Massachusetts Institute of Technology 1991 / All rights reserved

Signature of Author _____

Department of Chemical Engineering
June 3, 1991

Certified by _____

Mark A. Kramer
Associate Professor, Chemical Engineering
Thesis Supervisor

Certified by _____

Anantaram Balakrishnan
Associate Professor, Management Science
Thesis Supervisor

Accepted by _____

William M. Deen
Graduate Officer, Department of Chemical Engineering

Accepted by _____

Jeffrey A. Barks
Associate Dean, Sloan Master's and Bachelor's Programs

ARCHIVES
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUN 21 1991

ORDER COMBINATION METHODOLOGY FOR SHORT-TERM LOT PLANNING AT AN ALUMINUM ROLLING FACILITY

by
DAVID P. VENTOLA

Submitted to the Department of Chemical Engineering
and the Sloan School of Management in partial
fulfillment of the requirements for the Degrees of
Master of Science in Chemical Engineering
and Master of Science in Management.

ABSTRACT

This thesis addresses a production and process planning problem in an aluminum rolling facility. The facility manufactures a wide range of sheet and plate products to customer specifications; these specifications include the alloy, temper, order quantity, physical dimensions, and tolerances for each order. The manufacturing process used to produce this variety of orders is a job shop, where discrete batches (in the form of ingots) of aluminum are processed. Traditionally, the facility has planned each order independently and assigned a dedicated ingot with a tailored process plan. The order sizes often matched the available ingot weights, hence the scrap due to excess production from each ingot relative to the order quantity was limited.

With the growing emphasis on just-in-time production, the customers are demanding smaller, but more frequent deliveries of sheet and plate products. In contrast, the company has recently expanded ingot casting capabilities so that large size ingots can be cast more economically. The manufacturing process exhibits economies of scale when processing the large size ingots, due to setup costs and scrap requirements at the processing centers; therefore the facility prefers to process large size ingots over small size ingots. Clearly, using a 20,000 pound ingot for a 6,000 pound order is uneconomical, since a large portion of the ingot is not assigned to an order and may be scrapped.

This thesis focuses on the following problem: given the current set of confirmed orders for various products, how best to combine the orders and assign them to ingots so that all orders are completed by their respective due dates. This problem raises many issues: 1) which orders can be combined? 2) what criteria to use for choosing combinations of orders? 3) how to identify the "optimal" combination?

The study identified order combination possibilities for sheet products of two different alloys. First, we defined methods of specification compatibility. A partial solution is to combine distinct customer orders that call for the same specification. In addition, orders with different tempers, gauges, and widths can also be combined within the process and metallurgical constraints.

The second question investigated is what criteria to use for grouping the orders from the list of compatible orders. The analysis of a hierarchy of scenarios (with increasing combination flexibility) to group orders shows improvement in some operational and financial measures; processing a given set of orders requires the use of fewer ingots which decreases inventory and the total cost of processing.

Thesis Supervisors: Professor Mark Kramer, MIT Department of Chemical Engineering
Professor Anantaram Balakrishnan, Sloan School of Management

Acknowledgements

The author gratefully acknowledges the support and resources made available to him through the Leaders for Manufacturing Program, a partnership between MIT and major U.S. manufacturing companies. The author is also grateful for the support and guidance he received from several members of the host company.

In addition, the author would also like to acknowledge the constant encouragement and push from his thesis advisors. They were instrumental in helping the document make sense.

Table of Contents

List of Figures	6
List of Tables.....	7
Chapter 1 -- Introduction.....	8
1.1 Highlights of Thesis	8
1.1.1 Organization of the Thesis	10
1.2 -- Manufacturing Environment	11
1.2.1 Products and Markets	11
1.2.2 Alloys and Tempers.....	12
1.2.3 Process Flow	14
Chapter 2 -- Problem Definition and Project Focus	17
2.1 Motivation for Research	19
2.1.1 Short-term Planning function.....	19
2.1.2 Impact of Ingot Casting Capabilities.....	20
2.1.3 Customer Order Quantities	22
2.1.4 "Uncommitted" Metal.....	22
2.1.5 Problem Summary.....	25
2.1.6 Combination Example	26
2.1.7 Focus on sheet products	28
2.2 Plant Measures.....	30
2.2.1 Operational Parameters	30
2.2.2 Financial Parameters	31
2.3 Operational Decisions for Order Combinations	35
2.4 Related Literature.....	35
2.4.1 Cutting Stock Problem.....	35
2.4.2 Bin Packing Problem.....	36
2.4.3 Corrugator Trim Problem.....	37
Chapter 3 -- Technological Considerations	39
3.1 Combination Methods.....	39
3.1.1 Alloy 1 Heat-Treatable and T tempers.....	40
3.1.2 Alloy 2 Non-Heat-Treatable and H tempers	47
3.1.3 Order Combination experience.....	47
3.2 Practical Considerations.....	50
Chapter 4 -- Expert Systems Application.....	51
4.1 A General Overview of Expert Systems.....	51
4.2 Implementation of Order Combination Rules.....	55
4.2.1 LEVEL5 OBJECT Environment.....	56
4.2.2 Alloy 1 Model.....	57
4.2.3 Alloy 2 Model.....	59
4.2.4 Future Possibilities	61

Chapter 5 -- Evaluation of Scenarios.....	62
5.1 Methodology	62
5.1.1 Scenarios.....	63
5.1.2 Heuristics.....	64
5.2 Alloy 1 and Alloy 2 Results.....	66
5.2.1 Operational Analysis.....	66
5.2.2 Economic Analysis	68
5.3 Custom Ingot Scenario	73
5.4 Steady State Analysis.....	77
5.5 Optimization Perspective on Order Combinations.....	79
5.5.1 Model Description.....	80
5.5.2 Model Formulation.....	81
Chapter 6 -- Conclusions & Recommendations.....	83
6.1 Conclusions	83
6.2 Recommendations for Future Work	85
List of References	87

List of Figures

Figure 1: Process flow diagram for sheet products	16
Figure 2: Information flow for short-term planning	19
Figure 3: Savings in head and tail scrap.....	21
Figure 4: Order size distribution for heat-treated sheet	23
Figure 5: Order size distribution for non-heat-treated sheet.....	23
Figure 6: Mismatch in size of order and ingot.....	24
Figure 7: Order combination tradeoffs	25
Figure 8: Processing path for order combination example.....	27
Figure 9: Frequency distribution of order sizes for alloy 1	29
Figure 10: Frequency distribution of order sizes for alloy 2.....	29
Figure 11: Schematic of breakdown in total variable cost	34
Figure 12: Relationship among physical properties, manufacturing process, and microstructure for aluminum alloys	40
Figure 13: Effect of cold work prior to solution heat treatment on the properties of 7475-T6 sheet.....	42
Figure 14: Generic processing paths for tempers of alloy 1	43
Figure 15: Schematic of cold rolling process	44
Figure 16: Three main components of an expert system	52
Figure 17: Representation of automotive trouble-shooting rules.....	53
Figure 18: Definition of an object within LEVEL5	54
Figure 19: Total variable cost for processing set of alloy 1 orders	70
Figure 20: Total variable cost for processing set of alloy 2 orders	70
Figure 21: Revenue from processing alloy 1 set of orders.....	71
Figure 22: Revenue from processing alloy 2 set of orders.....	71
Figure 23: Total contribution from alloy 1 orders	72
Figure 24: Total contribution from alloy 2 orders	72
Figure 25: Breakeven analysis comparing custom ingots vs gauge combination scenario for alloy 1	75
Figure 26: Finished goods and excess inventory profile for alloy 1 gauge combination scenario.....	78

List of Tables

Table 1: Aluminum alloys and tempers	13
Table 2: Sample customer order	18
Table 3: Data for combination example.....	27
Table 4: Gauge combination table for alloy 1 O/T4/T6 tempers.....	46
Table 5: Order data for alloy 1 example combination case	48
Table 6: Effect of order combination example case	48
Table 7: Structure of alloy 1 knowledge base.....	58
Table 8: Structure of alloy 2 knowledge base.....	60
Table 9: Summary of operational impact for alloy 1.....	67
Table 10: Summary of operational impact for alloy 2	67
Table 11: Custom ingots operational analysis.....	75

Chapter 1 -- Introduction

1.1 Highlights of Thesis

This thesis addresses a production and process planning problem in an aluminum rolling facility. The facility manufactures a wide range of sheet and plate products to customer specifications; these specifications include the alloy, temper, order quantity, physical dimensions, and tolerances for each order. The facility also manufactures a few standard products to stock. Broadly the manufacturing process consists of five stages: ingot casting, hot rolling, cold rolling, heat treatment, and finishing operations. These stages process discrete batches of aluminum (in the form of ingots). Traditionally, the facility has planned each order independently and assigned a dedicated ingot with a process plan tailored to the given order. The order sizes (measured in pounds of metal) often matched the ingot weights that could be cast; hence the scrap due to excess production from each ingot relative to the order quantity was limited.

With the growing emphasis on just-in-time production, the customers are demanding smaller, but more frequent deliveries of sheet and plate products. In contrast, the company has recently expanded ingot casting capabilities so that large size ingots can be cast more economically. The manufacturing process exhibits economies of scale when processing the large size ingots, due to setup costs and scrap requirements at the processing centers; therefore the facility prefers to process large size ingots over small size ingots. Clearly, using a 20,000 pound ingot for a 6,000 pound order is uneconomical, since a large portion of the ingot is not assigned to an order and may be scrapped.

One possible strategy to reduce scrap and increase recovery is to stock products in a semi-finished form, e.g., stock 14,000 pounds as an intermediate

sheet or plate product for subsequent use to produce another order, and complete processing the 6,000 pounds to satisfy the current order. However, at this time, the facility chooses not to proliferate semi-finished stocks; instead, it wishes to assign (or commit) all the production to confirmed orders.

Furthermore, processing unordered metal uses scarce production resources. Given this objective, one strategy to increase recovery (percent of ingot used for committed orders) is to assign multiple orders to a single ingot during the planning stage. The production and process planners in the facility, then, face the following problem:

Given the current set of confirmed orders for various products, how best to combine the orders and assign them to ingots so that all orders are completed by their respective due dates and all weight recovered from the ingot is as fully utilized as possible.

Although the problem is easy to state, this strategy of combining multiple orders raises a host of issues:

- Which orders can be combined (i.e. which orders have common operations along their respective processing paths), given the machine capabilities and other operational restrictions? Clearly, orders that require different alloys cannot be combined. In addition, numerous process and metallurgical constraints restrict the possible combinations of widths, gauges, etc.
- What criteria is appropriate for choosing combinations of orders? Combining orders has several performance implications (e.g. possibly increased production effort, impact on quality and recovery), as well as

economic implications (e.g. ingot cost, processing cost, scrap and holding costs).

- How to identify the “optimal” combination? Given the combinatorial structure of the problem, finding the optimal solution via enumeration of all combinations is impractical.

This thesis addresses some of these order combination issues by focusing on various specifications of sheet products for two different alloys. For these two alloys, we first probed the question of which orders can be combined. This analysis yielded some methods to determine the compatibility of different product specifications. One solution is to combine distinct customer orders that require the same specification. In addition to these simple combinations, certain orders with different tempers, gauges, and widths can also be combined within the process and metallurgical constraints. These order combination rules define a list of compatible orders.

The second question investigated is what criteria to use for grouping the orders from the list. Total variable cost and total revenue are used in grouping the orders based on a hierarchy of scenarios with increased combination flexibility. These scenarios show that processing a given set of orders requires fewer ingots which decreases processing time and the total cost of processing. We formulate a simple optimization model to find the optimal grouping of orders; however, we have not developed solutions methods or solved the problem.

1.1.1 Organization of the Thesis

The remainder of Chapter 1 provides a description of the manufacturing setting for the research. Chapter 2 outlines the motivation for this research into

order combinations and presents the costs and benefits of order combination. This chapter also discusses some related problems in production and process planning (e.g. the cutting stock problem). In Chapter 3 order combination methods are presented in the context of constraints that limit feasible combinations. Chapter 4 describes the implementation of the order combination feasibility rules within an expert systems shell. Chapter 5 provides an evaluation of the order combination methods developed. The operational and economic impact are presented for different groupings of orders. This chapter also models the tradeoffs involved with order combination as an optimization problem. Finally Chapter 6 provides some conclusions from the research in the context of the plant's manufacturing strategy.

1.2 -- Manufacturing Environment

1.2.1 Products and Markets

The plant at which the research was conducted has a strategy to be a high quality and low cost producer of differentiated, "hard-to-produce" products, on a global basis. The product mix consists of a wide variety of Plate and Sheet aluminum products (e.g. lithographic sheet stock and automotive trim). The distinction between Plate and Sheet is based on the thickness of the finished metal. Plate aluminum products have thicknesses between 0.250" and 6.00" and are sold either in rough form or as finished parts (reducing scrap at the customers operations). Sheet aluminum is sold in either coiled sheet or flat sheet format and has thicknesses in the range 0.006" - 0.249".

In the aerospace market the plant produces components such as: fuselage skins (highly polished sheet that does not require painting), structural

components of fuselages, wing skins, etc. Commercial aircraft, high performance military aircraft, and space vehicles use these components. For the automobile industry the plant provides engine components, structural components, exterior trim and bumper stock. Another big product in the transportation industry is truck and trailer sheet for roofs and exterior panels. Some other markets in which the plant competes are: lithographic plate stock, venetian blind stock. This wide range of products consists of 50-60 alloys in several tempers.

1.2.2 Alloys and Tempers

The main reason that the plant can participate in an array of different markets is the inherent versatility of aluminum. The properties of aluminum can be dramatically altered by adding a small percentage of an alloying element. The range of properties covers the full spectrum from hard, high yield strength, low deformation, high melting point alloys to soft, low yield strength, low melting point alloys. Table 1 provides a summary of the various alloy and temper designations for aluminum which includes a general description of the types of alloys. The aluminum industry uses a four digit number (i.e. 1100) to designate alloys. The first digit defines the general category (primary alloying element) of alloy and the remaining three digits define the amount of the primary alloying element that is added.

There are two broad categories of aluminum alloys: heat-treatable and non-heat-treatable. Heat-treatable alloys undergo specific thermal operations, typically after the rolling operations. These thermal operations alter the microstructure of the aluminum which changes the metallurgical characteristics of the metal. Table 1 describes the different types of thermal treatment applied

Table 1: Aluminum alloys and tempers

A. Non-Heat-Treatable Alloys & Tempers

1. ALLOYS

<i>name of alloy class</i>	1XXX	3XXX	5XXX
<i>primary element added</i>	Al	Mn	Mg
<i>general class</i>	soft	soft	inter/hard
<i># different alloys</i>	5	8	18

2. TEMPERS

- F -- as fabricated
 - O -- annealed (heat)
 - H -- strain hardened
 - H1 - strain hardened only (cold worked)
 - H2 - strain hardened, then partially annealed (heat)
 - H3 - strain hardened, then stabilized (low temp heat)
- Final Degree of Strain Hardening
- HX2 - 1/4 hard
 - HX4 - 1/2 hard
 - HX6 - 3/4 hard
 - HX8 - Full hard
 - HX9 - Extra full hard

B. Heat-Treatable Alloys & Tempers

1. ALLOYS

<i>name of alloy class</i>	2XXX	6XXX	7XXX
<i>primary element added</i>	Cu	Mg/Si	Zn
<i>general class</i>	hard	inter	hard
<i># different alloys</i>	13	6	7

2. TEMPERS

- W -- solution heat treated, unstable temper
 - T -- thermally treated to stable temper other than F, O and H
 - T3 - solution heat treated, then cold worked
 - T4 - solution heat treated, then naturally aged
 - T6 - solution heat treated, then artificially aged
 - T7 - solution heat treated, then stabilized
 - T8 - solution heat treated, cold worked, then artificially aged
- Form of Stress Relieving
- TX51 - by stretching
 - (TX52 - by compression)
 - (TX53 - by thermal treatment)

to aluminum. For the non-heat-treatable alloys, the rolling operations determine the metallurgical properties of the finished product.

1.2.3 Process Flow

Since a wide variety of products and specifications are produced across numerous processing centers, the plant functions as a make-to-order job shop. The manufacturing process consists of five basic stages: ingot casting, hot rolling, cold rolling, heat treatment, and finishing operations. The "upstream" process (ingot casting through hot rolling) is the same for both Plate and Sheet products, but the downstream processes are different. Since this thesis deals with sheet products, we focus on the processing paths for sheet products.

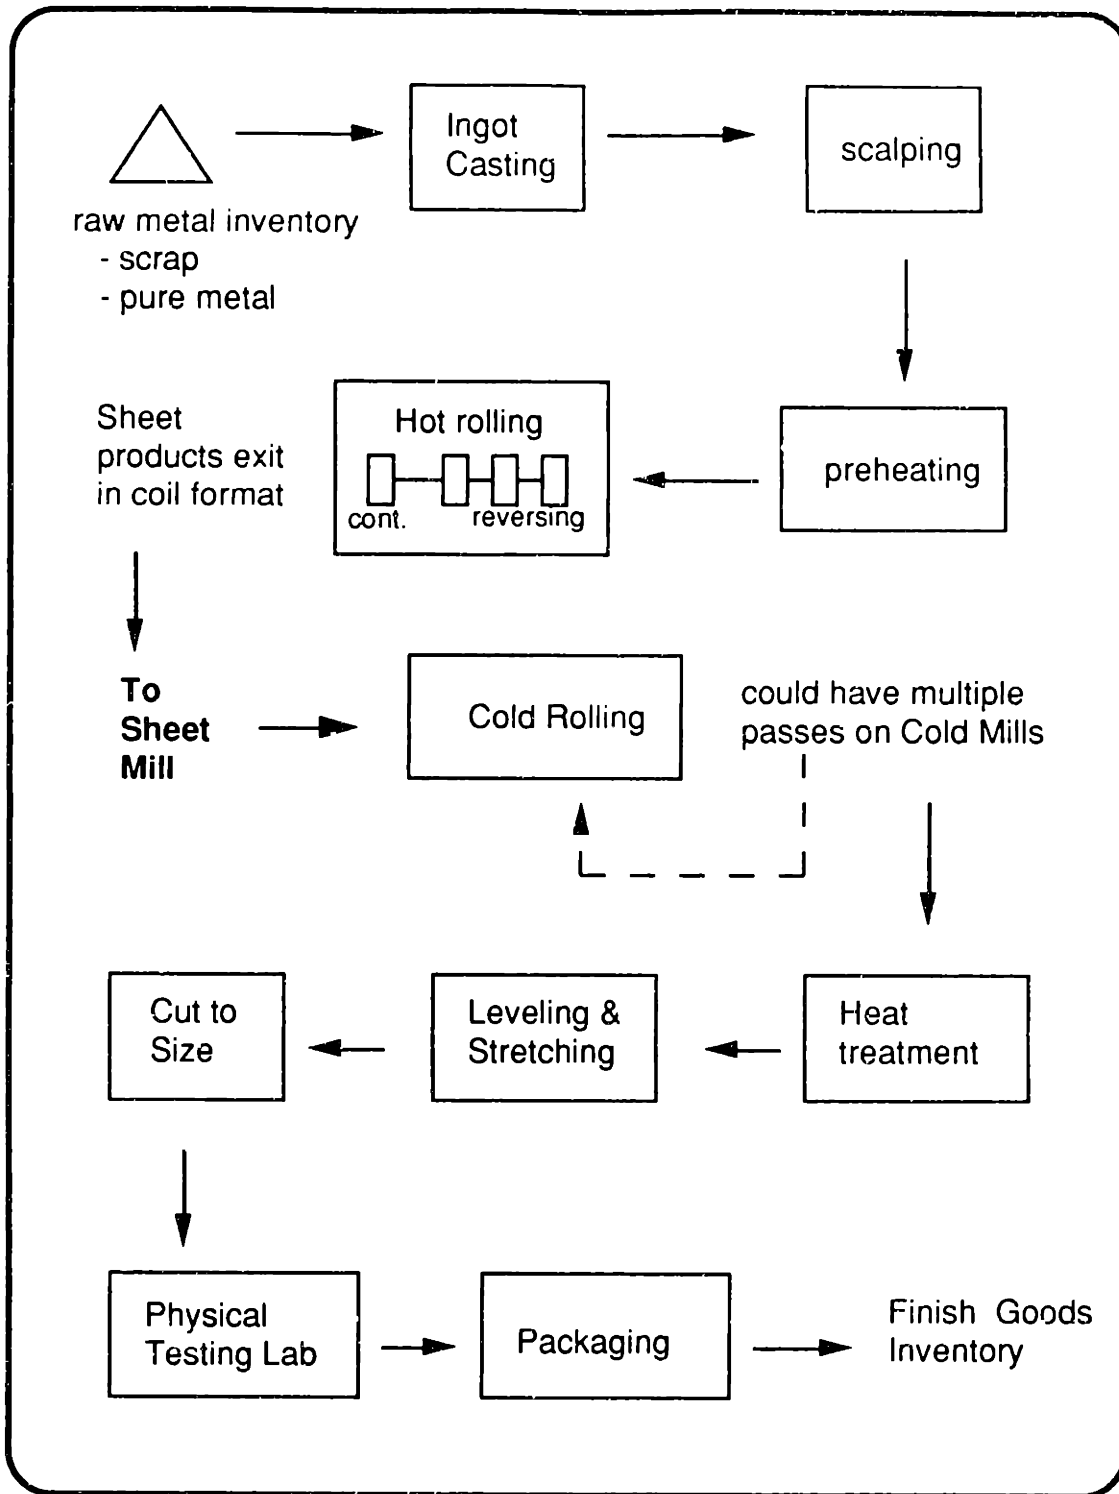
Figure 1 depicts the process flow for sheet products. Ingots (large, solid, rectangular blocks of aluminum) are cast from pure aluminum and scrap metal. For some alloys, the plant can cast ingots that will produce up to 30,000 pounds of finished metal. A milling machine then "scalps" the ingots. This operation removes the oxidized layer from the top and bottom surfaces of the ingot and provides a smooth, uniform surface prior to the rolling operation. The final step in preparing the ingot for hot rolling is preheating. The ingots are placed in furnaces and heated to the required temperature for the rolling operation.

The ingots are removed from the preheat operation and placed directly on the hot rolling station. This station has four rolling mills connected in series. The first three mills are reversing mills that reduce thickness and increase length by multiple passes through the rolling station. The fourth mill is a continuous one-pass mill with multiple rolling stations, each of which successively reduces the thickness. Overall, the hot rolling station can achieve large decreases in the thickness of the ingot. The sheet products exit the hot

rolling operation wrapped in coil format. In a "typical" processing path for a sheet product, the next step is cold rolling, which further reduces the thickness of the ingot, but has better tolerance control than the hot rolling operation. The plant has several cold rolling mills that operate in parallel, in contrast to the hot rolling operation which is sequential. Some orders require multiple passes on the cold mills to achieve the ordered thickness. For some tempers the cold rolling operation is a critical step in determining the final metallurgical properties, e.g., a certain minimum amount of thickness reduction must be achieved during cold rolling in order to meet temper specifications.

After cold rolling the finished coils go through a variety of heat treatment operations. For heat-treatable tempers, the metal will be processed at one of the temper furnaces or annealing furnaces. For non-heat-treatable alloys, the metal may pass through a stabilize furnace. Next (depending on the alloy), the metal may need to be leveled and/or stretched, to remove internal stresses. The finishing operations cut the metal to the final dimensions (length and width) specified by the customer. Finally, the physical testing lab samples and tests the alloy for appropriate metallurgical characteristics.

Figure 1: Process flow diagram for sheet products



Chapter 2 -- Problem Definition and Project Focus

The production and process planning functions consist of various stages: long-term, medium-term and short-term planning. Long-term planning addresses strategic issues; such as developing new products, new processes, new markets, and distribution channels. Medium-term process planning focuses on tactical decisions within the organization; demand forecast and available capacity determine operational policy decisions (e.g. inventory policy). Short-term process planning focuses on the routing (assignment to specific machines) and scheduling of specific orders, given the policies developed from the medium-term planning decisions. [1]

We address the short-term process and production planning decisions, which occur after the process planner receives a specific customer order. The facility wishes to assign all its production to confirmed customer orders, and simultaneously utilize large size ingots to realize benefits of economies of scale in the rolling process. When small customer order sizes (say 5000 pounds) are assigned large ingots (25,000 pounds), the process planner must assign the uncommitted production either to another customer order, excess inventory, or perhaps scrap.

The first part of this chapter provides background for the problem which motivates the study. We detail the mismatch between ingot weights and customer order quantities. The second part of the chapter outlines plant metrics relevant to the problem analysis. The third section presents the operational decisions which result from order combinations. In the final section of the chapter, we relate the order combination problem to other problems found in literature.

Table 2: Sample customer order

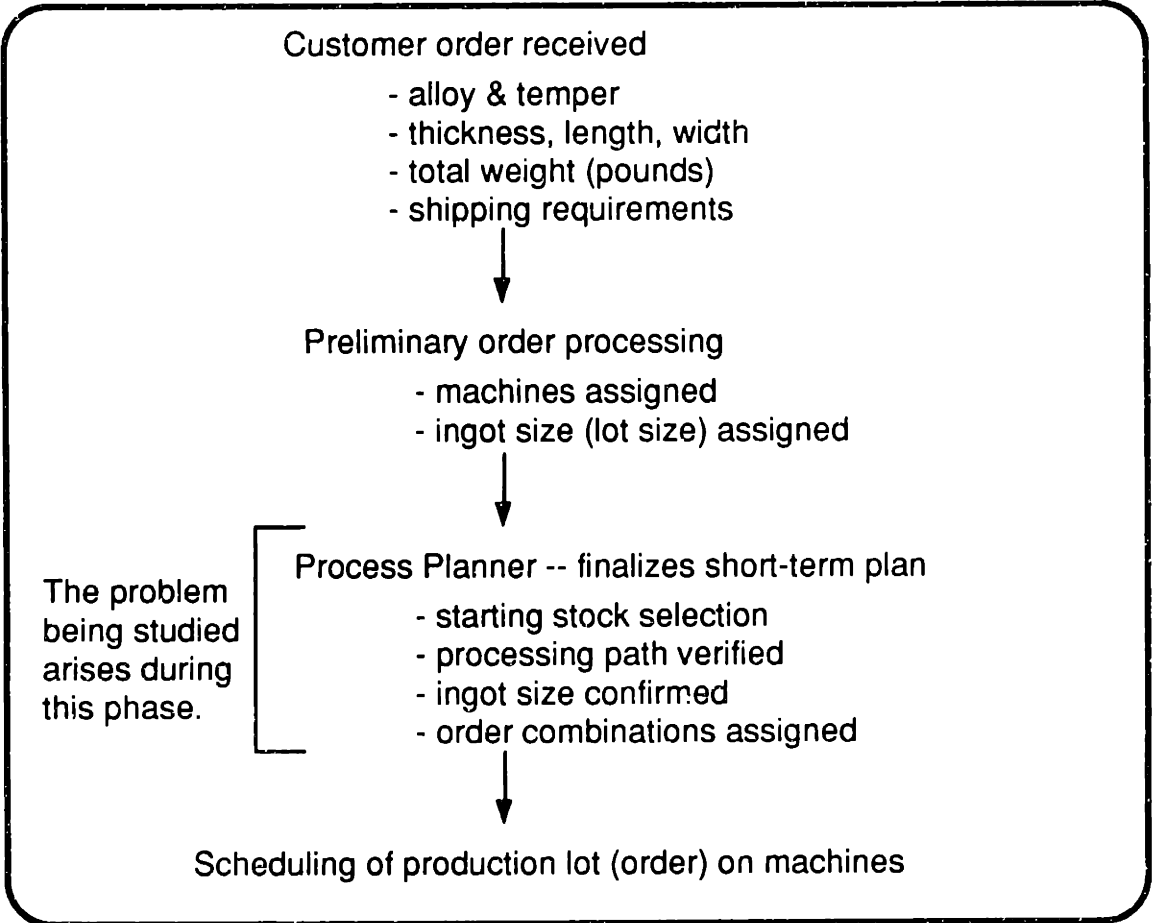
Product: Coil Sheet Mill Finish	
Specification Category	Value Ordered
Dimensions	
- gauge	0.050 in.
- width	48 in.
- length	
Tolerance	
- gauge	+/- 0.003
- width	std.
- length	
Alloy	Alloy 1
Temper	T6
Finish	mill finish
Edge Condition	trimmed
Physical Properties	
Mechanical Properties	
Quantity (item size)	5000 lb
- Quantity Tolerance	+/- 10%
Coil Size	
- Max O.D.	
- Max coil weight	5000 lb
- Min OD	
- Min Weight	2000 lb
Packing	
Reports/Samples	
Oiling/Interleaving	
Marking	
Transportation	

2.1 Motivation for Research

2.1.1 Short-term Planning function

Currently the short-term planning function, which is performed by production control at this facility, typically considers one customer order at a time and assigns the order to a dedicated ingot. Table 2 presents the characteristics (physical dimensions, tolerances, due date, etc.) of a sample customer order. Under the current system, when the process planner receives an order, she/he uses the automated plant order processing system to select an ingot and develop a preliminary processing path. Figure 2 summarizes the flow of information that impacts the short term planning decisions.

Figure 2: Information flow for short-term planning



Developing a process plan for an order involves several steps. The planner must first decide whether the order will be satisfied from finished goods stock, intermediate stock (an ingot that has been partially processed), or starting stock (ingot). The planner must also select the processing path, confirm the ingot size, and assign order combinations. Order combinations occur where the order quantity for a specific customer order is less than the recovery weight (starting ingot weight minus scrap loss in processing) of the chosen ingot. This is the specific type of situation that we will explore in greater detail.

2.1.2 Impact of Ingot Casting Capabilities

The plant casts ingots in several distinct sizes for each alloy. Over the past several years the plant has upgraded its ingot casting facilities to cast larger size ingots. The plant now has the capability to cast ingots that will yield (or recover) up to 30,000 pounds of finished metal. The yield (or recovery) is the ingot cast weight minus the scrap incurred during processing. Processing larger ingots in place of smaller ingots produces benefits in the form of economies of scale which are summarized as follows:

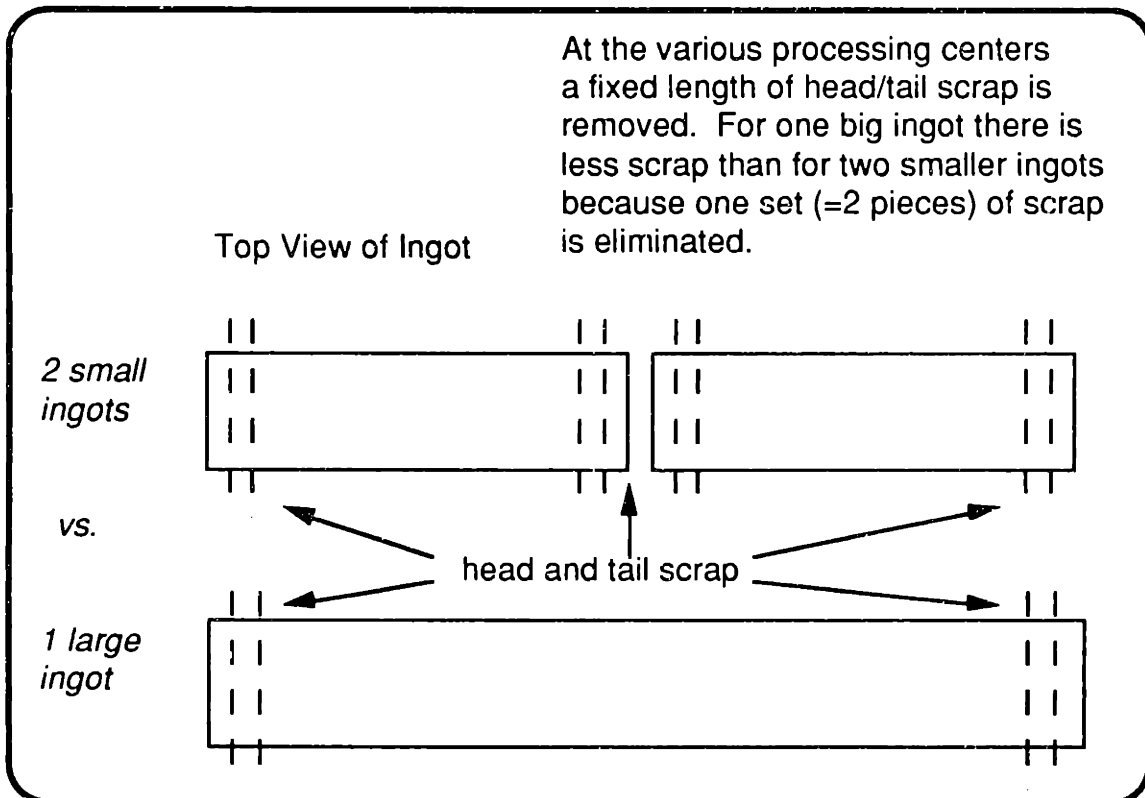
- Reduced scrap -- beginning and end of coils at each production center, scalping loss
- Reduced set up time -- because of fewer coils, and hence less between coil setup time at each production center
- Increased flow rate of good metal

The scrap level (as a percent of ingot weight) decreases for large ingots because many processing centers scrap a fixed length of metal at the beginning (head) and end (tail) of each ingot in the standard operation of the machine.

Since the scrap is a fixed distance, using fewer ingots reduces the overall percent scrap (as a percent of ingot weight, see Figure 3) due to lesser head and tail scrap. The reduced set up time results from the fact that each processing center has a fixed set-up or between-coil time. As with the scrap, if fewer ingots are being processed, the amount of set-up time decreases because fewer coils are processed. The higher flow rate of good metal results directly from less between-coil time (setup time) and lower scrap percentage. If big ingots are used in place of small ingots, then fewer ingots are required to produce a given weight of metal. As the throughput or flow rate of good metal increases, the lead time for any given product may decrease.

Given the potential benefits listed above, the facility prefers to process a large size ingot. The "ideal" ingot is a large size ingot that can produce a wide

Figure 3: Savings in head and tail scrap



variety of dimensional specifications for the alloy. If large size ingots are not utilized properly, the plant incurs costs, not benefits. When outlining the benefits, we made a key underlying assumption namely, all the metal produced will be shipped against customer orders in a timely fashion. If this assumption holds true then clearly the plant realizes the benefit from economies of scale in the process. For the orders where the ingot size is greater than the order quantity, many drawbacks related to excess metal exist (see section 2.1.4).

2.1.3 Customer Order Quantities

During the same time period that the plant increased its ingot casting capabilities, customers continued to order small quantities. Just-in-time production, which promotes reduced levels of inventory, became popular. This philosophy reinforced the small customer order quantities. In many cases the large ingot sizes yield more metal than is required for a particular customer order.

Figures 4 and 5 display the current frequency distribution of customer order sizes for heat-treated and non-heat-treated sheet products. The x-axis is normalized by the ideal ingot recovery weight (1 = ideal ingot recovery weight). The graphs show a large percentage of the orders that are much smaller than the desired ingot. These small order weights create a size mismatch between the large ingot size and the ordered quantity.

2.1.4 "Uncommitted" Metal

Figure 6 summarizes the size mismatch issue. This mismatch creates a quantity of "uncommitted" metal, i.e., production that is not assigned to confirmed customer orders. The uncommitted metal is the biggest drawback of using large ingots. When planning the order, the process planner tries to apply

Figure 4: Order size distribution for heat-treated sheet


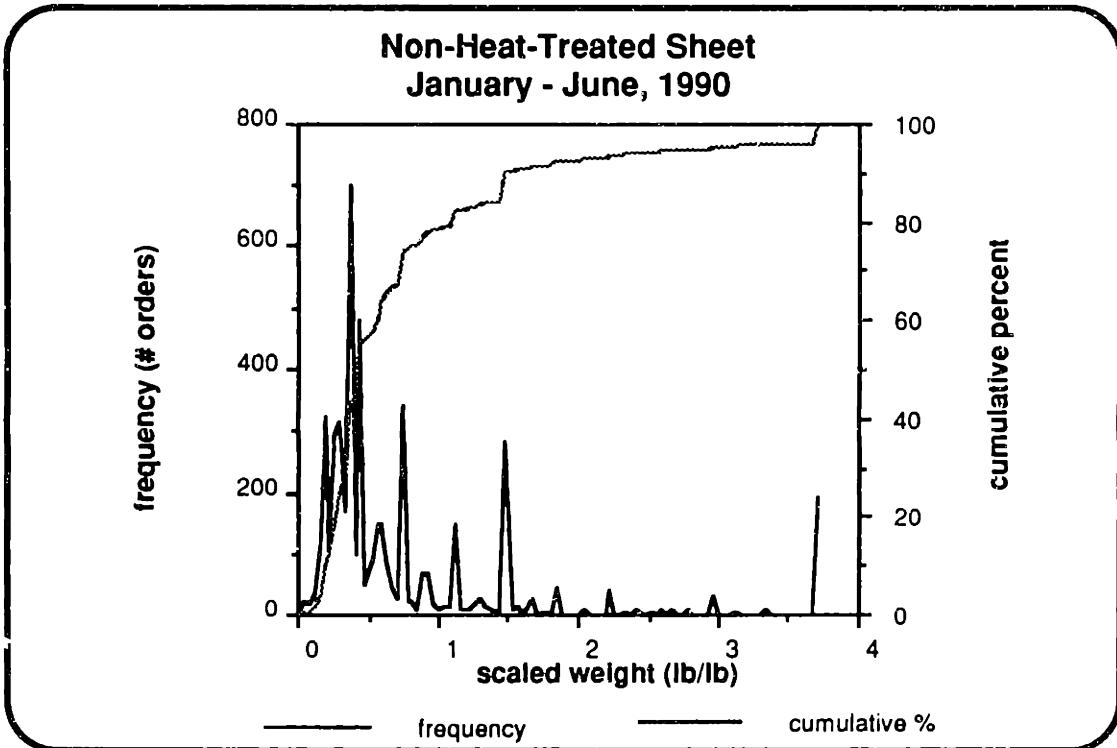


Figure 5: Order size distribution for non-heat-treated sheet



the uncommitted metal to another current order, a future order (i.e. combine the orders), or as a last resort to finished goods inventory. Metal placed in finished goods inventory may eventually be scrapped. The uncommitted metal results in:

- increased work-in-process inventory
- increased finished goods inventory
- opportunity cost of lost process time
- increased scrap

As a result, reducing the uncommitted metal in orders of this type has become the motivation for the research into order combinations.

While developing the short-term plan for an order, the planners search the order book for other orders with the same specifications, so that they can apply the uncommitted metal to another similar order and create a fully committed ingot. This step does not always eliminate the uncommitted metal. This paper explores taking this order combination concept a step further by defining other specifications that are compatible with the given specification, so that the planner can consider applying the uncommitted metal to an order of a different specification. This will allow the planner a wider range of order

Figure 6: Mismatch in size of order and ingot




Figure 7: Order combination tradeoffs

<u>Pros</u>	<u>Cons</u>
- benefits associated with big ingots: reduced scrap, higher throughput, etc.	- increase scrap within ingot when splitting orders, increase reprocessing cost
- limit processing of uncommitted/unnecessary metal	- decrease process consistency
- decrease total number of ingots that are required to process a given set of orders	- increase process complexity

combination possibilities. Figure 7 displays the pros and cons for using order combination as the strategy to achieve fully committed ingots. By assigning more than one order to an ingot, the plant can utilize a higher percentage of large ingots and realize the associated benefits described earlier. However, order combinations create added process complexity which results in increased scrap within a coil. This extra scrap occurs when two gauges are rolled within the same coil or two tempers are assigned to the same ingot.

2.1.5 Problem Summary

As outlined in the introduction, the problem is to take a current set of confirmed orders for various products, to combine and assign these orders to (large) ingots so that all orders are completed by their respective due dates and the ingots are fully committed. The first issue to consider is which orders can be combined. In addition to the requirement that all combined orders must be for

the same alloy, what are the constraints that limit the combination of orders for different specifications? Once a list of compatible orders is developed, what criteria to use for choosing combinations of orders? To answer these questions, we need to study the engineering and the performance implications of combining orders. Furthermore can we to identify an "optimal" set of combinations? This question in not addressed in this thesis. We next discuss an illustrative example to highlight the order combination tradeoffs.

2.1.6 Combination Example

We must develop the short-term process plan for two customer orders for alloy 1. Historically, these orders have been produced separately from two different ingots, but we can combine the orders and produce them from a single ingot, even though the specifications require different tempers and gauges. Figure 8 displays the processing path for the ingot which produces both orders. The combined processing path incorporates the commonalities from the two individual processing paths; the processing paths for the individual orders have common operations though the cold rolling step. The cold mill rolls the ingot to two distinct finish gauges (0.050" and 0.063"), splits the ingot and sends the orders to their remaining process operations.

Table 3 shows that when the orders are processed separately, the total uncommitted or excess metal is 39,000 pounds. However, a single ingot can satisfy both orders with 10,500 pounds of excess metal. The order combination reduces inventory, but creates additional scrap loss within the ingot, due to added process complexity at the cold mill. This is one of the tradeoffs which Chapter 5 studies in more depth.

Figure 8: Processing path for order combination example

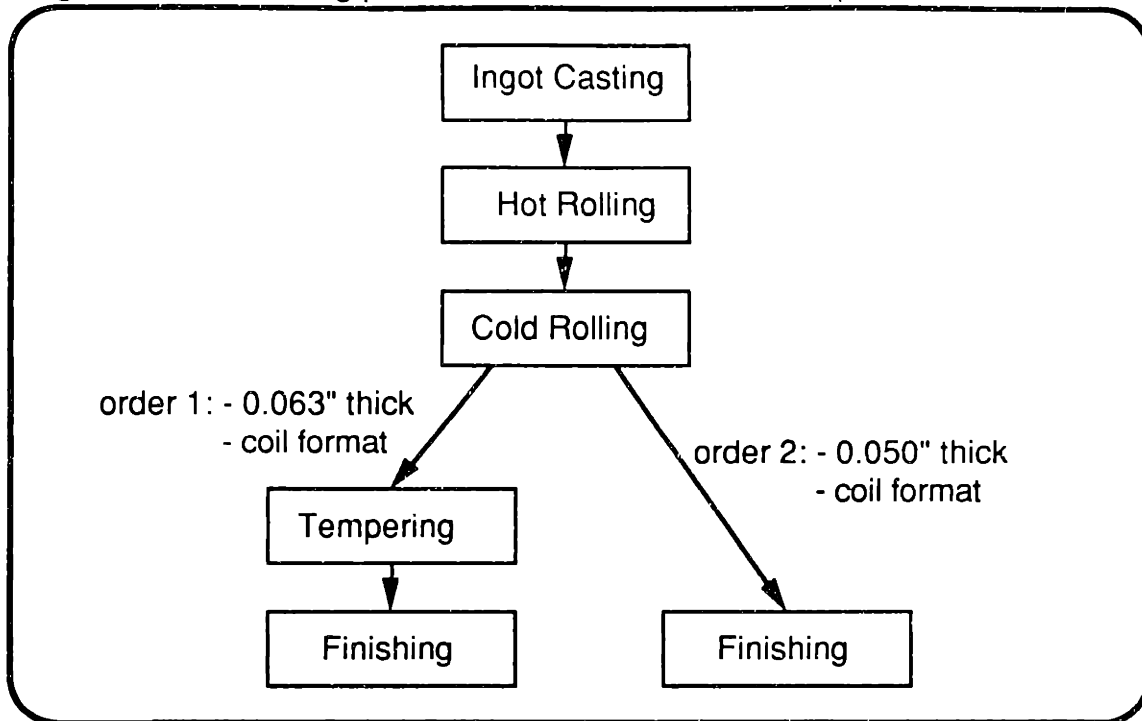


Table 3: Data for combination example

	<u>Order #1</u>	<u>Order #2</u>	<u>Combined</u>
<i>temper</i>	T4	O	O/T4
<i>thickness (inch)</i>	0.063	0.050	0.050/0.063
<i>width (inch)</i>	48	48	48
<i>length (inch)</i>	144	144	144
<i>weight (lbs)</i>	10,000	5,000	15,000
<i>Due Date</i>	90-08-25	90-08-02	90-08-02
<i>est. recovery of ingot (lbs)</i>	27,000	27,000	25,500
<i>uncommitted metal (lbs)</i>	17,000	22,000	10,500

2.1.7 Focus on sheet products

To study order combinations, we decided to focus on cold rolled sheet products for two specific alloys. There were several reasons for this decision. We decided to focus our efforts on a limited piece of the size mismatch issue. The data shows that sheet products experience a higher percentage of size mismatch than plate products. We restricted our study to orders that require cold rolling because the economies of scale are more prevalent along the processing paths for cold-rolled sheet products.

Since one goal of this effort was to increase utilization of large ingots, it was important to study alloys that the ingot plant can cast as large ingots. Finally, we wanted to choose those alloys whose demand was moderate (about 15 - 30 orders per month for cold rolled sheet), so that we would have an adequate sample size to study.

Using these criteria, we choose alloy 1 (heat-treatable) and alloy 2 (non-heat-treatable) for study. The general concepts or rules developed for alloy 1 would be associated with 'T' and 'O' tempers and thus should apply to other heat-treatable alloys. Similarly for alloy 2, the general concepts could be defined for 'H' tempers and thus be applied to other non-heat-treatable alloys.

Figures 9 and 10 show the order size distributions for alloy 1 and alloy 2. As in Figures 4 and 5 the x-axis is normalized by the ideal ingot recovery weight (1 = ideal ingot recovery weight). For both alloys, the graphs show that a high percentage of orders are much smaller than the ideal ingot size.

Figure 9: Frequency distribution of order sizes for alloy 1

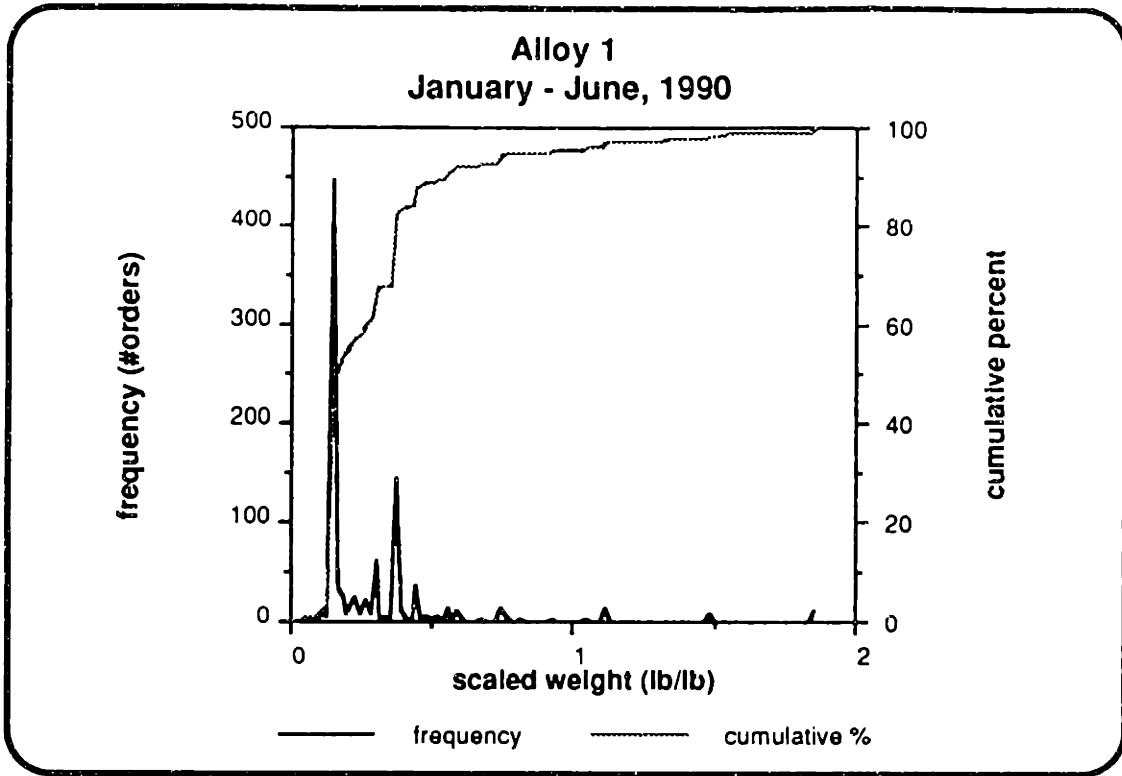



Figure 10: Frequency distribution of order sizes for alloy 2



2.2 Plant Measures

This section outlines the various operational and financial parameters that the plant uses to measure the production operations.

2.2.1 Operational Parameters

An order (Table 2 displays a sample customer order) is defined by the (i) alloy, (ii) temper, (iii) gauge, (iv) width, (v) length (vi) nominal order weight with quantity tolerance, and (vii) due date. One important parameter for this study is the order weight plus the quantity tolerance. When an order is being planned for production, the goal is to produce a quantity of metal that falls within the shipping tolerance. Depending on the available metal in the selected ingot, the planner, may assign production either close to the upper tolerance or the lower tolerance. Consider, for example, an order for 10,000 lbs of metal with a +/- 20% shipping tolerance. The order is satisfied when at least 8000 lbs, but up to 12,000 lbs, of finished metal is shipped to the customer. Since the quantity tolerance influences the actual amount of metal shipped to the customer, it also directly impacts the revenue from this order.

In addition to the shipped weight for each order there are several other operational measures used to evaluate an impact on performance:

- the number of ingots used to process the group of orders
- average size of ingot processed
- total recovered weight from the set of orders
- percent recovery from started weight
- actual quantity shipped for each order
- number of order combinations achieved for the group of orders.

We use these measures for the scenario analysis in Chapter 5.

2.2.2 Financial Parameters

From an economic perspective, for processing a given set of confirmed customer orders, the objective is to maximize contribution to profit. Equations (1) and (2) define the profit contribution.

$$\text{Profit} = (\text{UR} - \text{UVC}) * \text{Q} - \text{TFC} \quad (1)$$

$$= (\text{TR} - \text{TVC}) - \text{TFC} \quad (2)$$

where,

UR = unit revenue (\$/lb)

TR = total revenue (\$)

UVC = unit variable cost (\$/lb)

TVC = total variable cost (\$)

Q = total number of units sold (lbs)

TFC = total fixed cost (\$)

and,

(UR - UVC) is defined as the *unit contribution* (\$/lb)

(TR - TVC) is defined as the *total contribution* (\$) [2]

This study focuses on total contribution. The total quantity sold multiplied by the unit price determines the total revenue generated. Figure 11 and equation (3) outline the various components of the total variable cost.

$$C_{\text{total}} = C_{\text{metal}} + C_{\text{ingot}} + C_{\text{process}} + C_{\text{scrap}} + C_{\text{inventory}} + \text{OC} \quad (3)$$

where,

C_{total} = total variable cost

C_{metal} = metal value

C_{ingot} = variable ingot stock cost

C_{process} = variable processing cost

C_{scrap} = variable scrap reprocessing cost

$C_{\text{inventory}}$ = inventory holding cost

OC = opportunity cost

Metal Value

The metal value, C_{metal} , is the actual cost of the aluminum in a molten state, i.e. the costs involved in the mining, refining and smelting operations. These operations deliver molten aluminum to the ingot casting facilities. Thus, the metal value is the raw material cost of aluminum.

Ingot Stock Costs

The ingot stock cost, C_{ingot} , is the variable cost incurred when casting molten metal into an ingot, scalping the ingot, and preheating the ingot. These operations prepare the ingot for hot rolling. For this study, we assign a per pound variable cost ("ingot stock cost") to ingots delivered to the hot rolling operation. The ingot variable cost factors include direct labor and materials, heat to maintain molten metal, and water to cool the metal during casting.

Processing Costs

Processing cost, $C_{process}$, is the variable cost associated with the process operations from hot rolling through packaging the metal for shipment. These operations (see process flow diagram in Figure 1) include hot rolling, cold rolling, heat treating, and finishing, which represent the metal forming and shaping operations that take the cast ingot and generate the finished metal product. The variable expense at these processing centers is directly related to the processing time at each center. Once we calculate the processing time, a variable charge per unit time is applied for each processing center. The plant calculates the charge per unit time using the variable costs to run the processing center.

A key step is to calculate the process time at each machine. An average processing rate (in feet per minute of metal), based on historical data, is available. For each lot being processed, we calculate the total linear feet of metal based on the coil weight plus the width and thickness of the metal. Using

the coil length as the base, we calculate the process time from the average processing rate. If the operation involves setup time (i.e. the time to remove lot A and load lot B on the equipment), the set up time is added to the processing time to give the total process time at the equipment. As described above the cost factor and total process time determine the variable processing cost for the mill operations.

Scrap Reprocessing Costs

Scrap reprocessing costs, C_{scrap} , is the variable cost required to remelt scrap generated in the various process operations. The process of forming ingots into finished metal products inherently involves generation of process scrap. The scrap results from side trim when reducing the width to the specified width or scrap at the beginning or end of processing a coil (head and tail scrap) due to machine limitations. The ingot casting facilities reprocess this scrap metal to ingot form; the scrap reprocessing cost accounts for this expense.

Inventory Holding Costs

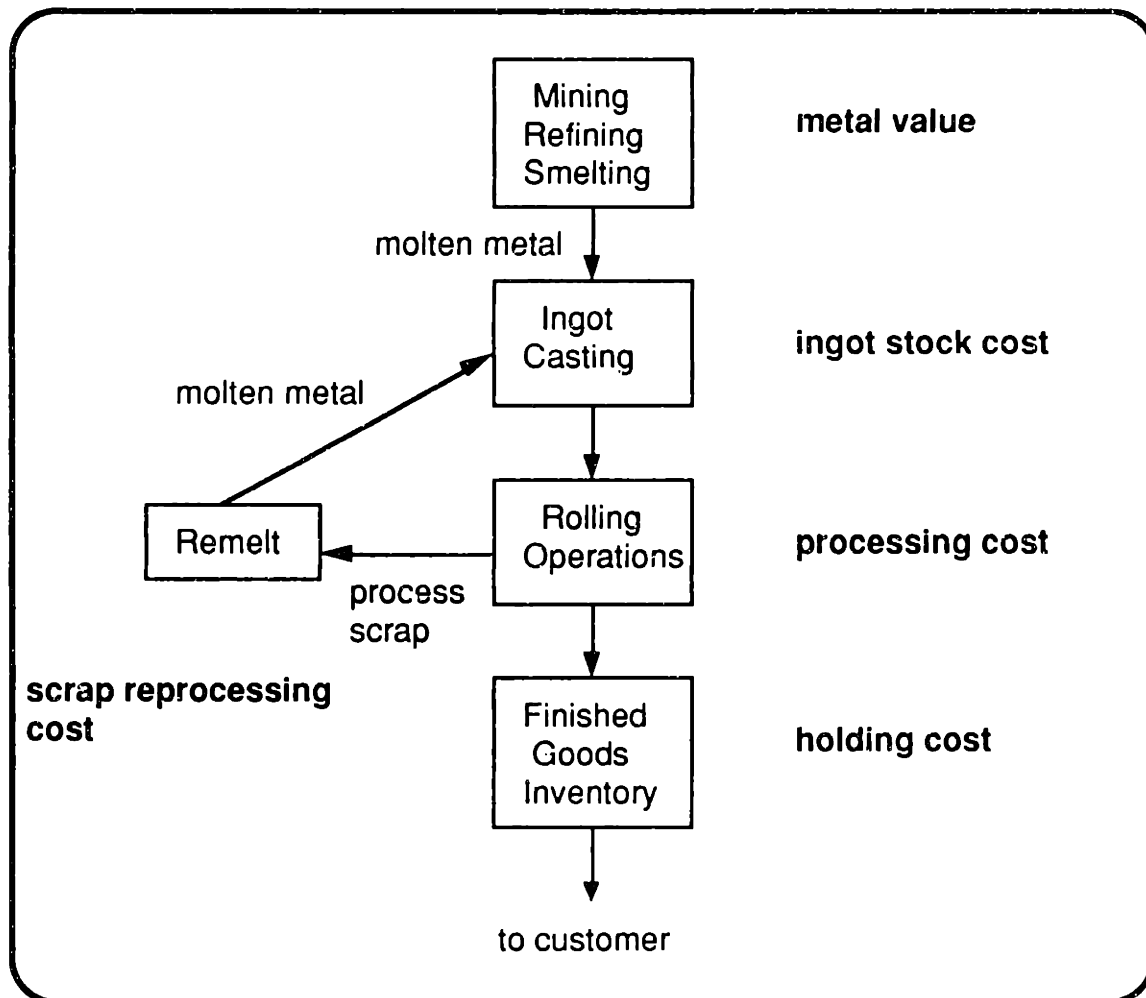
Inventory holding cost, $C_{inventory}$, is the inventory holding cost associated with carrying finished goods inventory. The total value of metal in inventory is the sum of the value added in processing plus the metal cost itself. The dollar value of material in inventory is assessed at a standard rate of interest per year.

This simple evaluation does not consider any of the other tangible (i.e. material damage) and intangible costs (i.e. hiding quality process problems) of carrying inventory. Nor does it consider the opportunity cost of processing metal before it is needed to be shipped. (This is only an issue for capacity constrained production centers.)

Opportunity cost

The opportunity cost, *OC*, is the cost associated with processing metal that is not assigned to a current order (excess metal or future order). For processing centers at full capacity, the opportunity cost is important because processing uncommitted metal leads to tardy orders. The opportunity cost is difficult to quantify, but we must recognize that it exists. This research does not attempt to quantify the opportunity cost, but reduction of uncommitted metal tends to reduce to opportunity cost.

Figure 11: Schematic of breakdown in total variable cost



2.3 Operational Decisions for Order Combinations

Implementing order combinations involves three decisions: choosing which orders to combine, choosing an ingot for the grouped orders, and specifying the planned shipping weight. The order combination rules define a list of compatible orders. From this list, how should the orders be grouped? We use net contribution as the basis for grouping orders. Should the grouping be done manually or by some optimization routine?

Related to the order grouping decision is ingot selection. Since there are fixed size ingots, the sum of the combined order weights cannot exceed the ingot weight. Recall that the quantity tolerance provides flexibility in the order weight. Thus, the shipping weight decision relates to ingot selection. When we choose the grouped orders and assign them to a specific ingot, we have implicitly identified a target production and shipping weight for each of the grouped orders. We must make all three decisions simultaneously.

The evaluation in Chapter 5 uses a hierarchy of scenarios with increasing combination flexibility to answer these questions for a group of orders. This evaluation is one method which provides answers to these questions. We could also use a more sophisticated method, such as an optimization routine. Section 5.5 explores one possible optimization method.

2.4 Related Literature

2.4.1 Cutting Stock Problem

Cutting stock problems, which are common in paper, adhesive tape, steel bar and plate, and glass manufacturing, have some similarities to order combination. A typical cutting stock problem arises in a situation where large

master rolls or sheets are produced to a semi-finished form and then cut to size to satisfy particular customer orders. These industries sell "standard product types" which are standard in their thickness and physical properties. Customers order these standard products by specifying the desired length and width. The company receives many different orders (with differing lengths and widths) for the same product type. The problem is to determine the best finishing sequence to produce the distinct orders from the master rolls or sheets of standard product. The order combination problem also considers grouping orders, but for order combinations we allow different gauges and tempers, as well different lengths and widths.

One common objective function for cutting stock problems is to minimize the total trim loss for processing a given set of orders. Another way to state the objective function is to minimize the amount of input material. Others have taken the traditional models and extended them to incorporate additional factors in the cutting process: set-ups, inventory, etc. The associated objective function is to minimize the total cost of processing the set of orders. Linear programming is the most common solution technique. Other utilized methods are dynamic programming, integer programming, and heuristic algorithms. ([3], [4], [5], [6])

2.4.2 Bin Packing Problem

In the classic bin-packing problem one seeks to minimize the number of equal capacity bins needed for the packing a given collection of pieces. There are many applications of the bin-packing problem in industry. In computer science important storage allocation problems appear as bin-packing problems; these include packing records into auxiliary storage and word lay-out problems. The bin packing problem is also applicable to some of the stock cutting situations described above. The bin-packing problem is similar to order

combination, if we think of the cast ingot as the bin and the customer orders as the items being packed.

The bin-packing problem is NP-complete; therefore it is difficult to find an optimal solution to this problem. A variety of approximation algorithms have been developed to provide solutions to bin-packing problems. ([7], [8], [9])

2.4.3 Corrugator Trim Problem

Schedulers in corrugator box plants must determine the least-cost method of combining customer orders on the corrugator, where one of the major costs is to avoid waste or excess trim lost from the materials used. This problem is also similar to the cutting stock problem described above. A corrugator processes stock rolls of liner and medium into corrugated material from which shipping containers are made. While on the corrugator, the corrugated material must be cut into rectangular blanks of the sizes needed to fill customer orders. The solution to the corrugator trim problem is simply a specification of what must be done in order to obtain the needed quantities of rectangular blanks. This involves specifying the stock roll sizes to be used, the lineal feet of corrugated material to be produced from each stock size, and the cutting instructions for each stock size which produce the required order sizes and quantities. The corrugator works by first slitting the material to width and then chopping it to length. To obtain a high degree of corrugator utilization, orders for various widths are combined to utilize the full width of the corrugator. The strips can be chopped to a limited number of lengths (2 or 3); this restricts the orders than can be combined.

The objective is to minimize the total cost of processing a given set of orders. The costs include the production and inventory of the various input stocks, the corrugation process and trim loss. Linear programming solutions

are common; a set-partitioning algorithm is one solution procedure employed.
[10]

These three problems are similar to our order combination problem because they satisfy multiple customer orders with different specifications from a single master roll of material. The customer orders for the three problems above differ only in the length and width of the finished product; these problems study a geometric grouping of orders. The order combination problem also attempts to satisfy multiple customer orders from a single master roll (which in our case is the ingot). In addition to allowing different widths and lengths among the grouped orders, the order combination problem allows different thicknesses and physical properties (tempers) to be produced from a single ingot (master roll). The inclusion of a wider range of specifications adds complexity to the order combination problem. Also, the three problems described above typically focus on a single production center for grouping the customer orders, whereas the order combination problem involves multiple processing centers because the product dimensions of temper and gauge impact many production operations.

Chapter 3 -- Technological Considerations

The simplest form of order combination is to assign two distinct customer orders of the same specification to a single ingot. However, what happens when two orders for the same specification do not exist? This question leads us to consider combining orders of different specifications. If produced separately, the different orders would have distinct processing paths. The concept is to choose orders that share a large number of the processing operations, and separate the orders downstream so that the economies of scale for the large ingots are realized upstream of the separation point. Furthermore, we could assign more than two orders to each ingot until the ingot becomes fully committed.

Each order requires specific properties: physical dimensions and physical properties of the metal (elongation, yield strength, and tensile strength). The order specifications and machine capabilities determine the processing path in the plant. In order to identify orders that are compatible with a given order, we must consider what other specifications are achievable along the processing path or similar processing paths. The first part of this chapter discusses combination rules based on property compatibility. This discussion incorporates the physical limitations of the equipment. The final section highlights practical considerations that are important for implementation of order combinations.

3.1 Combination Methods

A finished product has many dimensions which characterize the product; Table 2 displays these dimensions in a sample customer order. The first

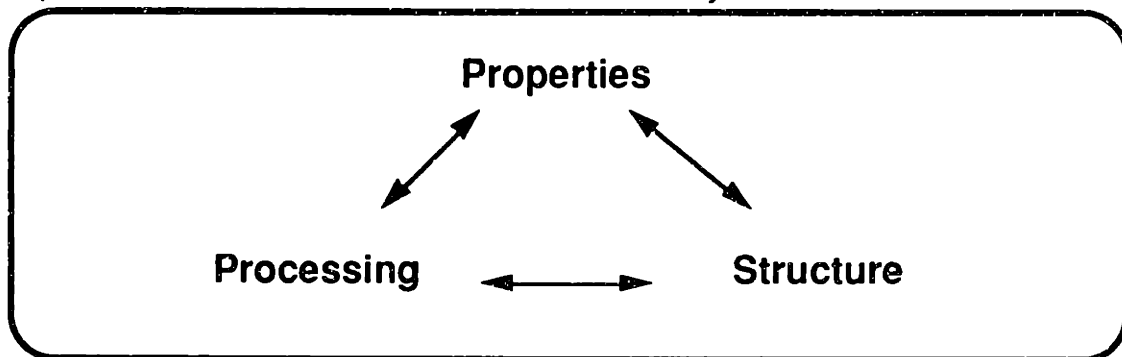
restriction for order compatibility is that the orders request the same alloy. Other characteristics of a finished product are the temper, the gauge (or thickness of metal), width, and length for flat sheet, or coil size for coil sheet. Along each of these dimensions, other compatibility restrictions exist. Since length combinations are flexible, the process planners currently perform length combinations, as well as simple width combinations (similar widths). This study defined rules for temper and gauge compatibility. Using the tools described in this chapter and Chapter 4, the process planners have the capability to perform order combinations using several dimensions: gauge, temper, width, and length.

3.1.1 Alloy 1 Heat-Treatable and T tempers

Temper

We associate a general set of physical properties with each alloy. For example, one alloy (e.g. 1100) may be soft having low yield strength and high elongation, and another alloy (e.g. 2024) may be hard having high yield strength and low elongation. Within the family of properties for a given alloy, the temper designation defines the specific physical characteristics for the finished

Figure 12: Relationship among physical properties, manufacturing process, and microstructure for aluminum alloys.




metal. Figure 12 displays the relationships among microstructure, manufacturing process, and properties. Each parameter influences the other two. A given microstructure results from a specific manufacturing process which also defines the final properties of the metal. Microstructure is a term which includes several components: crystal structure, composition of the phases, grain shape, and grain boundaries.

Recall from Table 1 that T tempers are formed by solution heat treatment after cold work. The purpose of solution heat treatment is to put the maximum practical amount of hardening solutes such as copper, magnesium, silicon, or zinc into the metal matrix (i.e. no precipitates). The heat treatment process allows the diffusion of the alloying elements into the metal matrix to form a solid solution. When the metal is properly cooled, most of the alloying element remains within the matrix.

The grain size, a component of the microstructure, is a function of both the cold work and the subsequent heat treatment. For aluminum alloys, the grain size typically does not influence the physical properties. However, with small amounts of cold work, coarse grains may develop because there was not enough cold work to break up the grains, therefore heat treatment acts to grow the grains rather than form new ones. Coarse grains can have negative effects on the metal properties. Small grains increase the strength by limiting the dislocation movement across the numerous grain boundaries; large grains, having much fewer grain boundaries, have less resistance to dislocation movement and hence a lower strength.

Above a certain minimum percent reduction, the grain size is independent of the amount of cold work. Percent reduction (on the finish cold mill pass) in thickness of the metal measures the amount of cold work performed (see section on gauge below). To illustrate this phenomenon, Figure

Figure 13: Effect of cold work prior to solution heat treatment on the properties of 7475-T6 sheet (adapted from [11])



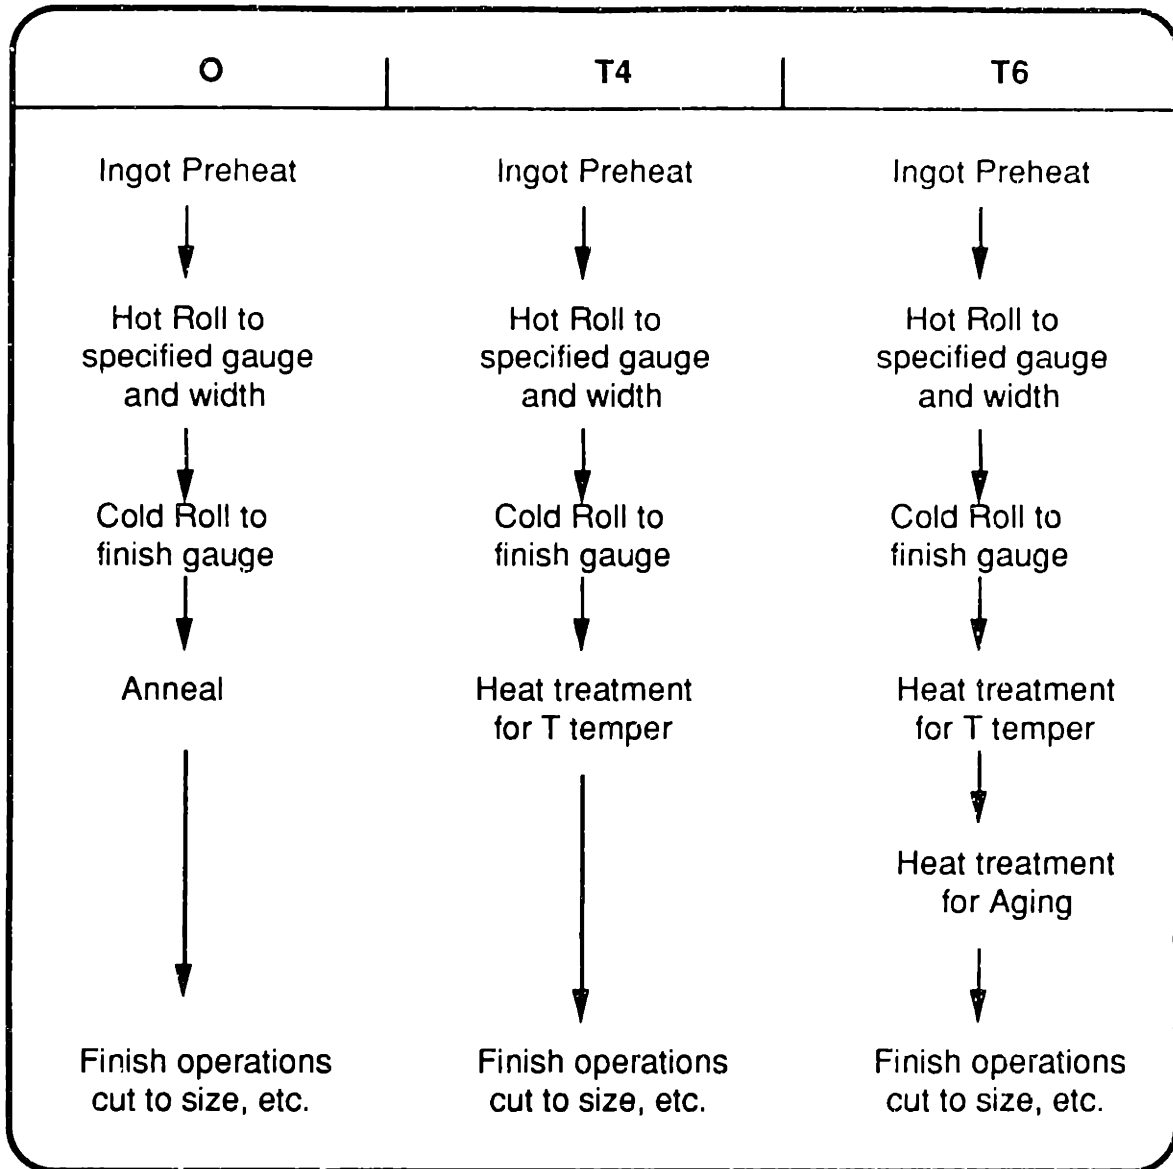
13 displays the relationship between tensile strength and percent reduction for 7475-T6 sheet. Above 30% percent reduction, the physical properties are independent of the amount of cold work; the solution heat treating process determines the physical properties and microstructure. To consistently achieve temper properties, a minimum percent reduction constraint results. [11]

For alloy 1 only three tempers are important: O, T4, and T6. We determined this by an order history analysis. These three tempers have common processing paths through the completion of cold rolling. After cold rolling the processing paths diverge because T temper requires solution heat treating, while O temper is annealed. Figure 14 displays generic processing paths for each of these tempers.

Gauge

The other main product dimension studied is the gauge or thickness of the metal. The capabilities of the rolling operations (both hot and cold rolling)

Figure 14: Generic processing paths for tempers of alloy 1



determine the specific processing path for the thickness reduction of each ingot. As discussed in Chapter 1, the hot rolling operation has four stations. The first three stations are reversing mills and the final station is a continuous (one pass) mill. Since there are four mills, the hot rolling operation has the capability to achieve large thickness reductions and produce a wide variety of intermediate gauges.


Finish gauges cannot be produced from any arbitrary intermediate gauge, however. As mentioned above the cold rolling operation is important for determining temper properties because certain tempers impose a minimum cold work constraint. For the purpose of gauge combination in alloy 1, we also define the maximum possible thickness reduction which is chiefly determined by the capabilities of the rolling mill. The reduction is computed by:

$$r\% = 100\left(1 - \frac{t_f}{t_o}\right) \quad (1)$$

where: $r\%$ = percent reduction
 t_f = exit thickness from cold mill
 t_o = entry thickness to cold mill

The maximum reduction can be expressed as a function of several parameters: back tension, front tension, friction, entry thickness (t_o), roll radius, mill spring (plastic distortion), roll flattening, roll bending, width of metal, and yield stress of the alloy. Figure 15 displays a schematic of the cold rolling process. For a

Figure 15: Schematic of cold rolling process



given cold mill entry thickness (t_0), we define a minimum percent reduction (e.g. 30%) based on temper requirements and a maximum percent reduction from this expression. Thus, for a specific cold mill and a given t_0 , these two constraints create a range of possible exit gauges (t_f) which can be produced in a single pass on the cold mill. Due to these two constraints, we cannot combine widely varying gauges. ([12], [13], [14])

For alloy 1 we constructed a list of these compatible gauge ranges based on empirical performance data (i.e., practical experience). Table 4 provides a summary of these ranges. The maximum percent reduction constraint for each cold mill corresponds to the '*min - order 1 finish gauge*' in the table; the minimum percent reduction constraint for temper corresponds to the '*order 2 - max combinable gauge*.' Practical considerations for the cold mill determine '*max - order 1 finish gauge*.' Each row represents one specific set of compatible gauges and an associated processing path. The reason for overlap of the various sets (rows) of compatible gauges is due to differences in the processing paths. For example, different processing paths may have two cold mill passes or differing cold mill entry thickness.

We locate ranges of potential gauge combinations by searching the table for the rows that contain the finish gauge of the current order. First we find the appropriate min/max range for order 1 finish gauge; this row highlights one set of possible gauge combinations. In the second step we identify the rows where the given finish gauge is greater than the '*max - order 1 finish gauge*' and less than the '*order 2 - max combinable gauge*.' These rows highlight additional gauge combinations. At each step we also verify that the width of the order is less than '*width 1*' or '*width 2*' listed in Table 4 under the heading '*max CM entry width*.' These widths describe physical limits of the cold mills.

Once the list of compatible gauges is defined, customer orders that fall in this gauge range can be combined with the given order. For example, if two orders are being considered for combination with specified gauges of 0.050 inches and 0.075 inches and the width of each is narrower than 'width 2,' they are compatible based on the sixth row of the Table 4 (they fall between 'min - order 1 finish gauge' and 'order 2 - max combinable gauge' of row 6). However two orders with specified gauges of 0.050 inches and 0.085 inches are not compatible since none of the rows define a processing path that can achieve these two gauges from a single ingot.

Table 4: Gauge combination table for alloy 1 O/T4/T6 temper

<i>process path</i>	<i>Order 1 Finish Gauge Min</i> \geq	<i>Order 1 Finish Gauge Max</i> $<$	<i>max CM entry width</i>	<i>Order 2 - Max Combinable Gauge</i>
1	0.011	0.020	width 1	0.064
2	0.020	0.023	width 1	0.064
3	0.023	0.064	width 1	0.064
4	0.020	0.032	width 2	0.072
5	0.020	0.044	width 2	0.060
6	0.044	0.060	width 2	0.080
7	0.060	0.080	width 2	0.100
8	0.080	0.100	width 2	0.126
9	0.100	0.140	width 2	0.140

3.1.2 Alloy 2 Non-Heat-Treatable and H tempers

Non-heat-treatable alloys have various temper designations (see Table 1). We studied only H32 temper for alloy 2. As the description "strain hardened" implies, the amount of cold work performed on the metal determines the final metallurgical properties. As mentioned above, the percent reduction on the finish cold mill pass measures the amount of cold work.

In the strain hardened condition, the H3X tempered alloys tend to age soften at room temperature. Therefore, these alloys are usually heated at low temperature to complete the age-softening process and to provide stable mechanical properties and improved working characteristics. H32 temper is strain hardened to 1/4 hard (much less cold work than the full hard condition) and then stabilized at low temperature. [11] Based on the historical process data, we found some flexibility in the processing paths used to achieve the final properties for alloy 2-H32. Specifically, the data shows ranges of percent reduction for the finish cold rolling pass with a different stabilize operation associated with each range. The combination of the degree of cold work and the thermal operation achieve the desired temper properties. Thus the historical data provided the gauge combination rules for alloy 2-H32.

3.1.3 Order Combination experience

During this study, we planned and implemented an order combination experiment that included four orders combined in a single ingot. The purpose of this study was: 1) to investigate the flexibility of some processing centers with regard to order combination, 2) to define practical considerations for order combinations, 3) to identify implementation issues for order combinations.

Table 5 displays the relevant information for the four orders. Note that all four orders required a finished quantity that was significantly smaller than the

Table 5: Order data for alloy 1 example combination case

	<u>Order #1</u>	<u>Order #2</u>	<u>Order #3</u>	<u>Order #4</u>
type	coil sheet	flat sheet	flat sheet	flat sheet
temper	O	T6	T6	T4
thickness (inch)	0.018	0.020	0.025	0.032
width (inch)	48	48	48	48
length (inch)	-	144	144	144
weight (lbs)	7,000	5,000	4,000	4,000
wt. tolerance	40%	40%	10%	10%
est. recovery of ingot (lbs)	27,886	18,948	18,695	28,987
Due Date	90-08-25	90-09-08	90-09-08	90-08-25

Table 6: Effect of order combination example case

	<u>Separate orders</u> <u>4 dedicated ingots</u>	<u>Combined orders</u> <u>1 ingot w/ 4 orders</u>
total ingot yield (lbs)	90,000	27,000
ordered weight (lbs)		
- lower	14,400	14,400
- nominal	20,000	20,000
- upper	25,600	25,600
shipped weight (lbs)	25,600	24,000
excess metal (lbs) (inventory)	64,400	3,000
total processing time	3.7Y minutes	Y minutes
total processing cost	\$3.4X	\$X

yield from the ingot initially chosen to process the order. Also note the quantity tolerance for each order allows some flexibility in the quantity of metal shipped to satisfy the order.

From the alloy 1 rules defined above, we see that all four orders are compatible through the final cold rolling step, and all the gauges are compatible on processing path 1 of Table 4. After cold rolling, order 1 is separated because the O temper processing path diverges at this point. Based on Figure 14 which describes temper processing path compatibility, orders 2 - 4 are compatible for the T tempering step and then order 4 is separated because T4 and T6 tempers diverge at this point.

Table 6 summarizes the differences between processing the four orders with dedicated ingots versus processing them combined in a single ingot. By combining the four orders into a single ingot, this experiment saved over 60,000 pounds of inventory or recycled excess metal. For the combined orders, the processing time was one-third of the time required for processing the four orders separately and similarly the processing cost was one-third of the of the separate cost. The ratios are not one-fourth because combining orders with different specifications into a single ingot adds some processing and set up time which also increases the processing cost. In addition splitting the orders at various points in the processing path generates extra scrap. Overall, the combination of the four orders benefits the plant by reducing inventory, total processing time and costs, while still satisfying the four customer orders.

3.2 Practical Considerations

During the order combination experiment described above, many practical issues surfaced. The current administrative system is geared to assign one order to one ingot. The definition of process routing becomes more cumbersome when more than one order is assigned to an ingot. A related issue is that when more than one order (with different specifications) is assigned to an ingot the production operations have difficulty maintaining the identity of the different orders. If the cold rolling operation reduces a coil to two different thicknesses, they have no easy method to mark the gauge change so that subsequent operations can distinguish the separate orders.

When orders of different specifications are assigned to the same ingot, extra process scrap results. For instance, if a cold mill rolls two different gauges within the coil, extra scrap is generated within the coil at the transition point between gauges. Also if two orders in the same ingot have different processing paths, the processing center which separates the orders incurs extra head and tail scrap because it is handling two coils at its exit rather than one.

In the experiment performed, the four orders had four different gauges. We concluded from this experience that assigning more than two gauges to a single ingot is not practical at this time. One gauge per coil has been the standard operating procedure; an incremental change to two gauges in a coil seems more practical at this point.

Chapter 4 -- Expert Systems Application

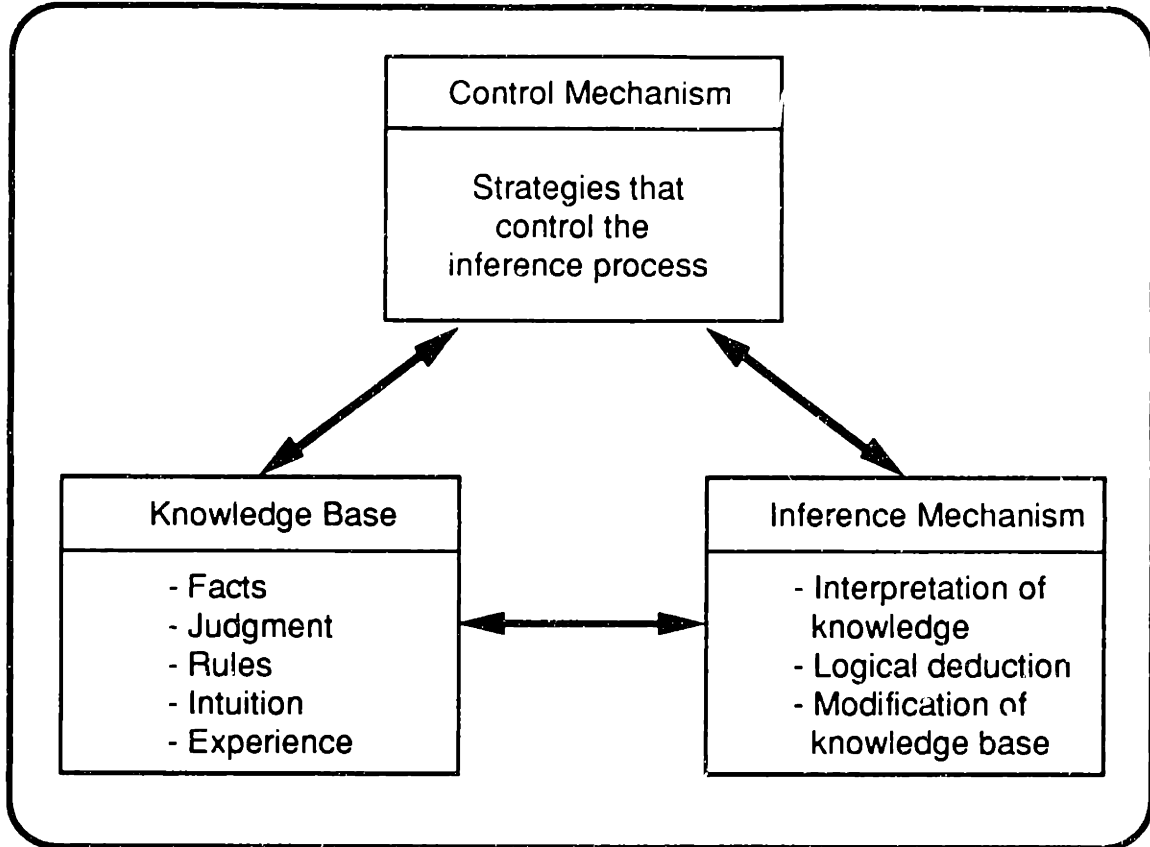
We automated the order combination methods, described in Chapter 3, in the form of a computer program to simplify the use of these methods. For programming, we choose an expert systems shell, instead of a traditional platform, to exploit some of the advantages associated with AI and expert systems. Expert systems are programming languages that can process more than just numeric data and numeric problems; they can process symbolic problems. This feature expands the application of expert systems to a wider range of problems than traditional programming. [15]

The first section of this chapter provides a general description of expert systems. The second section of the chapter describes LEVEL5 OBJECT, the expert systems shell we used for programming. This section also provides a summary of the knowledge bases developed for the alloy 1 and alloy 2 order combination methods.

4.1 A General Overview of Expert Systems

There were several reasons for choosing a commercially available expert systems application over a conventional programming approach: "human-like reasoning mechanism" that uses certain data and has a modular structure, representational clarity, and user-friendly interface. All of these features, which will be expanded upon in this section, offer advantages over the conventional programming approach. One drawback is that the programmer must work within the limits of the chosen expert system shell. Fortunately a wide range of products exist on the market to suit most needs. The evolution of these expert systems programming tools has brought the concepts of AI to the

Figure 16: Three main components of an expert system [15]



novice because they are easy to use and have a wide range of applicability.

[15]

Figure 16 displays the three basic components of an expert system. The knowledge base contains the data and information that the user has gathered to apply to the context. This data and information can take a variety of forms as seen in Figure 16. The inference mechanism interprets the knowledge in the knowledge base and performs its reasoning process based on this knowledge. The control mechanism organizes and controls the strategies taken to apply the inference process. [15]

Figure 17: Representation of automotive troubleshooting rules [16]

Rule 408:

C is a car.

IF: the pattern observed by attaching an oscilloscope to the charging circuit of the car is fluctuating.arches, and the alternator of the car C responds properly to different loads,

THEN: there is strong evidence <0.9> that the cause of the problem with the car C is voltage.regulator.bad.

Rule 428:

C is a car

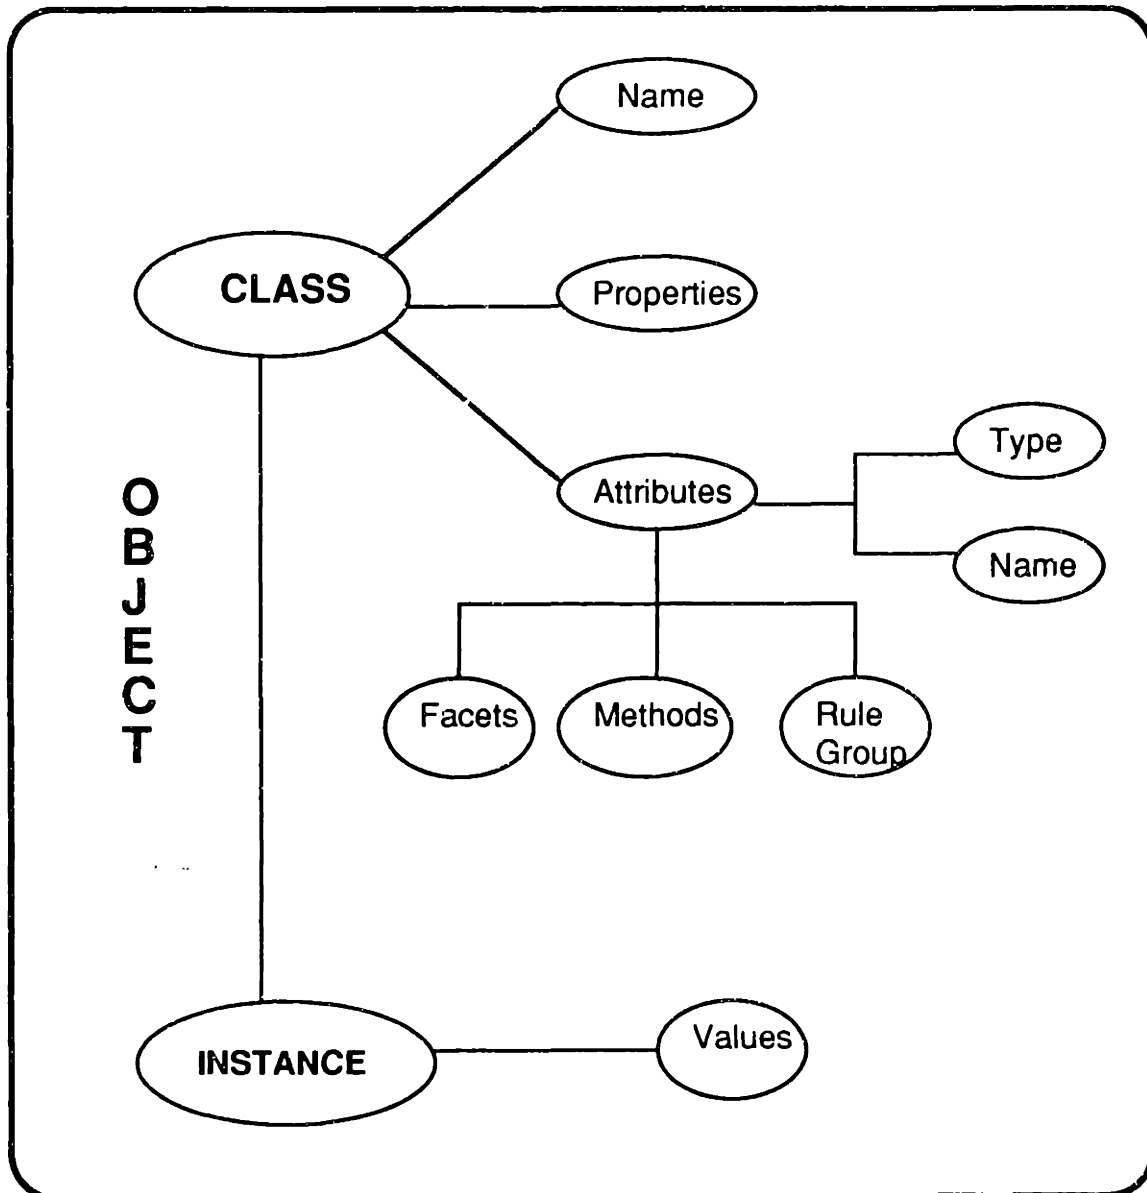
IF: the pattern obtained by attaching an oscilloscope to the charging circuit of the car C is straight.line, and the result of pulling out the field connector is no.flash, and the field connector does not have a voltage, and the dashboard lights do not glow when their ground circuit is completed, and the fusible link is getting voltage, and the fusible link is not conducting power,

THEN: it is definite <1.0> that the cause of the problem with the car C is fusible.link.bad.

One goal of an expert systems application is to solve the types of problems that experts solve. The programmer organizes the knowledge of the expert so that the program will provide the same answers to a problem as the expert would provide. "IF/THEN rules" are one way to code within an expert system application. This facility helps organize the expert's knowledge so that the program's reasoning mechanisms can effectively utilize the knowledge. Figure 17 displays an example of two IF/THEN rules used in troubleshooting automotive problems. The rules represent the encoded knowledge of someone who is an expert in troubleshooting automotive problems.

Many expert systems applications have "object-oriented programming" features. Objects are entities that combine the properties of procedures and data, since they can trigger computations as well as contain information. Messages sent between objects form the basis for the program's computations. Objects have properties of their own, but also can inherit properties from other objects. Figure 18 shows a pictorial representation of an object within the

Figure 18: Definition of an object within LEVEL5 [17]



LEVEL5 environment. In LEVEL5 OBJECT, the *class* forms the core of the object. The *attributes* within the *class* define specific features and data types of the object. *Methods* attached to the *class* determine the operations triggered by the object. The *properties* of the *class* further specify the structure of the object. However, the object remains a skeletal structure until the user declares *instances* of the *class*. The *instances* contain individual "data records" for specific occurrences of the *class*. ([17], [18])

Another attractive feature of expert systems is the user interface. Many shells run within an icon driven environment (such as *Microsoft Windows*). Also some development shells contain many built in "programming" features that simplify the programmer's task. This allows the developer to choose the desired option from the menu of choices, rather than having to create the code for this option. Running the coded program is also icon driven which makes it simple for the programmer to add instructions for the user to follow.

A final advantage of an expert systems application is the modular structure of the code. A program can be built in specific segments and additional features can then be added to the code without having to rewrite the existing pieces. It is similar to the use of subroutines in traditional programming, yet it is easier to incorporate new features without altering the existing lines of code. Once the foundation for a program is established, the program can be easily expanded to incorporate other problem solving abilities.

4.2 Implementation of Order Combination Rules

We implemented the order combination rules within the expert system shell LEVEL5 OBJECT from Information Builders, Inc. This package offers the variety of features described above. It also runs in the *Microsoft Windows 3.0*

environment. Two simple models were written using this shell: one each for alloy 1 and alloy 2. These models function as stand alone systems. (There were no links developed with the order database or any other outside databases or programs.)

4.2.1 LEVEL5 OBJECT Environment

In the LEVEL5 OBJECT development environment the main features used to code an application are the *Objects Editor*, the *Agenda Editor*, *Display Editor*, and the *Rule/Demon Editor*. The *Objects Editor* allows the programmer to build the structure of various objects. In addition the programmer creates *classes* and assigns *attributes* to these *classes*. A *class* with its specified *attributes* defines the skeleton of the object. By declaring *instances* of the *classes*, the user adds data to the object. The programmer attaches *facets* to the *attributes* which specify various properties of the *class*. Similarly the programmer attaches *when changed* or *when needed* methods to the *attribute*. These methods are procedures that execute on the value of the *attribute* or define the value of the *attribute* at runtime.

Within the *Rule/Demon Editor* the programmer can create either *rules* or *demons*. The format of *rules* and *demons* is similar, yet LEVEL5 OBJECT associates *rules* with the backward-chaining reasoning process and *demons* with the forward-chaining process. The *rules/demons* typically have the "IF.....THEN..." structure, but LEVEL5 also has some built-in commands that can be incorporated into the *rules/demons*. Both the antecedent and conclusion portions of the *rules/demons* contain the *attributes* defined through the *Objects Editor*. The *rules/demons* manipulate the values of these attributes. LEVEL5 has specific syntactical requirements for the format of *rules* and *demons*. The *rules* and *demons*, specified by the programmer, define the reasoning process

of the application. Within a single knowledge base application, mixing *demons* and *rules* provides the advantages of both forward-chaining and backward-chaining reasoning.

Backward-chaining systems require the use of the *Agenda Editor*. *Attributes*, posted in the *Agenda* as goals for the system, control the specific path followed by the backward-chaining reasoning process. The *Display Editor* allows the programmer to design displays for the user interface. The programmer can create multiple displays and link them together or simply have one display that contains the relevant information for running the knowledge base. The *Display Editor* contains various graphical tools that allow the programmer to create a unique environment for the knowledge base interface. The various tools can link to *attributes* and display the value of the *attribute*, update the value of the *attribute*, etc. In addition LEVEL5 offers capabilities to track the reasoning process while running a knowledge base. These features are extremely helpful when trying to debug an application.

4.2.2 Alloy 1 Model

The application, written as a backward-chaining model, requires several pieces of information about the order that the user must supply (because the system is a stand alone tool). The goal (or *Agenda*) of the knowledge base is to provide suggestions of combination possibilities for the specified order. The *Agenda*, which controls the inference mechanism, determines the reasoning process for the alloy 1 knowledge base. The knowledge base consists of several objects. The *current order* object is the "data input field". The *gauge combo table* object contains the core knowledge about specification compatibility. The *possible combos* object is an empty structure; during run time the program creates *instances* of *possible combos*. These instances

Table 7: Structure of knowledge base for alloy 1

ALLOY 1 - KNOWLEDGE BASE

Agenda

- controls the knowledge base reasoning session.
- the goal is to provide the user suggestions about possible order combinations
- the attributes have the 'exhaustive' facet active so that all possibilities are found.

Objects

- *'current order'*
 - class takes input data from the title display
 - uses 'when changed methods' to check the accuracy of the information that is input
- *'gauge combo table'*
 - holds the "data" for the model to use
 - this is the core of the knowledge in the system
- *'possible combos'*
 - is an empty framework for "windows" of order combinations
 - specific instances get created during runtime by the MAKE command.
 - these attributes are linked to 'listboxes' in the conclusion display to output all possible combos to the user.
- *'temper notes'*
 - instances contain the "knowledge" about the temper combinations that are possible
- *'temper output'*
 - during runtime gets assigned an appropriate value from 'temper notes'

Rules

- rules are used to search for gauge combinations that meet the compatibility criteria
- rule for temper assignment
 - searches 'temper notes' for information about temper compatibility
 - assigns the appropriate information to 'temper output'

Input Display

- contains promptboxes for the values of the attributes for 'current order' class.
- displays error messages if illegal values are input for the attributes.

Output Display

- displays all windows of combination opportunities with the associated flow paths necessary to process the orders.

define the various windows of specification compatibility. There are also several rules which operate on the "data" specified by the user and the data contained in the *gauge combo table* object. These rules search the *gauge combo table* to find the appropriate, compatible specifications. Two displays, an input display and an output display, define the user interface.

4.2.3 Alloy 2 Model

Similar to the alloy 1 system, the application was written as a backward chaining model that requires information about the order which the user must supply. For alloy 1, the *gauge combo table* object represents the knowledge for specification compatibility, whereas for alloy 2 a series of rules best describes combination possibilities. Since the core knowledge in the application is represented differently for the two alloys, the alloy 2 knowledge base has a different structure than the alloy 1 knowledge base.

The goal for the alloy 2 knowledge base is the same as for alloy 1; to provide suggestions of combination possibilities for the specified order. As in the alloy 1 knowledge base, the *Agenda*, which controls the inference mechanism, determines the reasoning process for the alloy 2 knowledge base. There are several objects that make up the knowledge base. The *current order* object is the "data input field". The *possible fin from fin*, *fin from inter* and *inter for fin* objects are empty structures. Three empty structures are necessary for this knowledge base (whereas only one was needed for alloy 1) because there are three distinct reasoning paths that define compatible specifications. During run time the program creates instances of these objects. These instances define the various windows of compatibility. The rules in this knowledge base contain the information about how to determine compatible specifications. These rules, which operate on the data input into the current order object,

define the windows of compatibility. As in the alloy 1 knowledge base, two displays define the user interface.

Table 8: Structure of alloy 2 knowledge base

ALLOY 2 - KNOWLEDGE BASE
<p><u>Agenda</u></p> <ul style="list-style-type: none">- controls the knowledge base reasoning session.- goal is to provide suggestions to the user about possible order combinations- the attributes have the 'exhaustive' facet active so that all possibilities are found.
<p><u>Objects</u></p> <ul style="list-style-type: none">• <i>'current order'</i><ul style="list-style-type: none">- class takes input data from the title display- uses 'when changed methods' to check the accuracy of the information that is input• <i>'possible fin from fin'</i><ul style="list-style-type: none">- has structure for one type of combination possibility- instances are created during runtime• <i>'possible fin from inter'</i><ul style="list-style-type: none">- structure for a second type of combination possibility- instances are created during runtime• <i>'possible inter for fin'</i><ul style="list-style-type: none">- structure for third type of combination possibility- instances are created during runtime• <i>domain</i> (LEVEL5 default class)<ul style="list-style-type: none">- contains simple attributes that control the sequence of rule firing in the backchaining reasoning process
<p><u>Rules</u></p> <ul style="list-style-type: none">• multiple rules contain the information about specification compatibility
<p><u>Input display</u></p> <ul style="list-style-type: none">- data input regions for the information the knowledge base will need in order to determine the combination possibilities for the given order.
<p><u>Output display</u></p> <ul style="list-style-type: none">- contains listboxes of attributes that describe the different types of combinations that are possible (based on 'current order' being planned)

4.2.4 Future Possibilities

The order combination applications currently function as stand alone systems. The concept and applicability of AI and expert systems is much broader than simple applications like the ones described above. One possible improvement is to provide a link with the order database to automate the grouping of compatible orders. In the current system, the planners must search the order book for orders that meet the compatibility requirements.

If this link were made, the planner would be able to select a specific order from the order book (by highlighting it with the mouse); the program would then search for orders that match the compatibility criteria. A simulation of this approach was performed using LEVEL5 OBJECT which has the capability to interface with some types of outside databases. But in our case the link to the order database was not established. Instead a fixed set of orders was input into the LEVEL5 knowledge base. The evaluation discussed in Chapter 5 uses the order data and program as an aid in the analysis.

Given a group of orders that must be planned, one of the issues raised by order combinations is: how best to group the orders? Also, how to identify an optimal set of combinations? Potentially an expert system could be employed to help resolve these issues, but this approach was not addressed in this thesis.

Chapter 5 -- Evaluation of Scenarios

Chapter 3 described the rules which determine feasible combinations. For a given order, these rules provide us with a list of compatible orders. Now we must address the issue: how to choose a specific combination from the list of feasible ones? Chapter 2 indicated that combining orders has several performance implications. The plant measures described in section 2.2 were used as the criteria for evaluating the order combination decision.

This first section of this chapter describes a methodology which uses a hierarchy of scenarios with increasing combination flexibility to assign specific order combinations. This section also reviews the decision criteria employed in this analysis. Section 2 outlines the results for these various planning scenarios. The third section compares the order combination strategy to a custom ingot strategy. The fourth section proposes an analysis of steady state plant operations using the gauge combination strategy. Since we can consider this problem from an optimization perspective, the final section presents a simple optimization model for grouping the orders.

5.1 Methodology

This study considered only cold rolled sheet products. For this analysis we captured the order book which existed on November 21; we collected orders listed in the database for the months of December '90, January '91, and February '91. Appropriate parameters were obtained for each order:

- alloy and temper
- due date

- order weight with quantity tolerance
- ingot assigned by order processing system
- estimated ingot recovery weight
- physical dimensions (length, width, gauge).

We classified the orders into two groups. The first group consisted of “combination candidates,” i.e., orders that had the size mismatch problem described in section 2.1. For the second group of orders, the order size and estimated ingot recovery matched closely. This second group of orders was not considered for the analysis.

For alloy 1 and alloy 2, we developed operational plans for the candidate group of orders under various sets of criteria/assumptions (scenarios): *no combinations*, *due date combinations*, *temper combinations* and *gauge combinations*. For each of these scenarios, we performed an operational analysis and financial analysis to help quantify the impact of the proposed methods.

5.1.1 Scenarios

The *no combinations* scenario is a theoretical case because in today's operating environment the process planners perform some combinations. In this scenario, we assigned each order to a dedicated ingot. Also, for each order we assigned the smallest ingot that would satisfy the order weight. The estimated recovery weight of the assigned ingot is greater than the lower quantity tolerance for the order. In this scenario, excess metal was scrapped.

The *due date combinations* scenario simulates the current operating environment. In this scenario we performed combinations of orders with the same specifications (i.e. same gauge and temper), as well as simple length and

width combinations. We selected the smallest standard ingot that would satisfy the ordered weight(s). As we planned each order, we scanned the rest of the "order book" for a match. If we found a compatible order, the two orders were planned on the same ingot. In cases where we combined two orders, metal for one order may be placed in finish goods inventory for a period of time. All uncommitted metal was scrapped.

The *temper combinations* scenario allows temper combinations in addition to due date combinations (e.g. for alloy 1 T4 and T6 tempers are acceptable combinations). Due to the added complexity of implementing a temper combination, we added a 2% recovery loss to the ingot yield as a penalty. (From our order combination experiment described in Chapter 3, we estimated that a temper combination resulted in a yield loss of approximately 2%.) We did this to account for losses due to coil splits after cold rolling.

The *gauge combinations* scenario allows gauge combinations as well as temper combinations and due date combinations. Due to the added complexity of implementing a gauge combination, we applied a 5% recovery loss to the ingot yield as a penalty for a gauge combination. (From our order combination experiment described in Chapter 3, we estimated that a gauge combination resulted in a yield loss of approximately 5%.) We did this to account for losses due to changing gauges within a coil and for splitting the coil. In addition only two different gauges were allowed within the same ingot (e.g. one ingot could produce 0.050" and 0.063") because we wanted to limit the manufacturing complexity of the combination schemes.

5.1.2 Heuristics

For each of the above scenarios we planned the orders in ascending order of due dates, i.e., we planned the earliest due dates first and then

proceeded in chronological order through the list. When the order being planned had several feasible combinations, we executed the combination by 1) trying to commit the greatest percent of the ingot, and 2) choosing orders with the closest due dates. For example, consider a due date combination (all orders have the same specifications) where three orders are compatible. Order 1 requires 10,000 lbs (20% tolerance) and is due 90-08-02, order 2 requires 10,000 lbs (40% quantity tolerance) and is due 90-08-30, and order 3 requires 6,000 lbs (10% shipping tolerance) and is due 90-08-16. Based on the two rules listed above, we choose an ingot which yields 25,000 pounds of finished metal and assign orders 1 and 2 to this ingot (which fully commits the ingot).

The above scenarios represent different levels of combination flexibility used to plan the set of orders. While developing the process plans for each scenario, we did not have a quantitative measure of the economic impact of our order grouping decision. After we developed the plans for each scenario, we performed an economic analysis. Using the results of this analysis, we hope to determine the set of heuristics that allows us to maximize total contribution. We consider these economic analyses in the following perspective:

Objective Function:

$$\text{Max } \Pi \quad (\text{total contribution}) \quad (1)$$

where

$$\Pi = TR - C_{\text{total}} \quad (2)$$

$$TR = Q * p \quad (\text{total revenue}) \quad (3)$$

Q = quantity sold (pounds)

p = sale price (\$/lb)

$$C_{\text{total}} = C_{\text{metal}} + C_{\text{ingot}} + C_{\text{process}} + C_{\text{scrap}} + C_{\text{inventory}} \quad (4)$$

(total variable cost)

Section 2.2 describes the various components of this framework in detail.

5.2 Alloy 1 and Alloy 2 Results

The planning scenario evaluation consists of two parts. The first part is an analysis of the various operational parameters described in section 2.2. The second part of the analysis is a determination based on the above framework.

The operational analysis performed for each scenario included: the total number of orders planned, the total numbers of ingots used to process this set of orders, the size distribution of ingots, the number of combination sets developed, and the total pounds of finished metal produced. The financial analysis performed for each scenario consisted of calculations for total revenue, total variable costs and total contribution. In the analysis revenue differed from scenario to scenario because orders may have different shipping weights within the specified tolerance.

5.2.1 Operational Analysis

For December '90, January '91 and February '91, the plant received a total of 62 orders for alloy 1 in cold rolled sheet format. Of these 62 orders we identified only 45 as combination candidates; the remaining orders were large enough to justify dedicated ingots. We performed all analyses on the subset of 45 candidate orders. Table 9 summarizes alloy 1 operational parameters. For alloy 2, from a total of 95 orders, we identified 65 as combination candidates; we performed all analyses on this subset of 65 orders. Table 10 summarizes the operational parameters for alloy 2.

Table 9: Summary of operational impact for alloy 1

Alloy 1 - Planning Scenarios	<i>Types of Combinations</i>			
	NO COMBOS	DUE DATE	TEMPER COMBOS	GAUGE COMBOS
upper ship tolerance	597,879	597,879	597,879	597,879
nominal order weight	514,538	514,538	514,538	514,538
actual shipped weight	577,774	568,191	558,907	553,568
# ingots processed	48	40	37	33
est. recovered weight	856,137	801,397	747,759	700,679
avg recovery per ingot	17,836	20,035	20,210	21,232
# of combination sets	-	9 sets	12 sets	15 sets
shipped/recovered	0.675	0.709	0.747	0.790

Table 10: Summary of operational impact for alloy 2

Alloy 2 - Planning Scenarios	<i>Types of Combinations</i>		
	NO COMBOS	DUE DATE	GAUGE COMBOS
upper shipping tolerance	1,296,832	1,296,832	1,296,832
nominal order weight	1,090,892	1,090,892	1,090,892
actual shipped weight	1,290,133	1,223,792	1,174,352
# ingots processed	77	60	53
est. recovered weight	1,585,832	1,410,060	1,294,122
avg recovery per ingot	20,595	23,501	24,417
# of combination sets	-	18 sets	24 sets
shipped/recovered	0.814	0.868	0.907

The ratio of shipped weight to recovered weight increases about 10% from no combinations to gauge combinations. This indicates the degree to which the processed ingot is assigned firm orders. The number of combination sets increases and the total number of ingots required to process the order set decreases as we allowed more types of combinations (orders are being grouped together and assigned to an ingot). One result of this combination strategy is that the average ingot size increases, which means that a larger percentage of big ingots are utilized to process the orders. Section 2.1 describes the potential benefits of using big ingots to process orders. Some of these benefits are realized because the ingots are more fully committed (an increase in the ratio shipped/recovered), but we identified recovery penalties associated with combining orders that partially offset some of these benefits. The due date combination scenario portrays the current operating environment. If we adopt a gauge combination strategy, the order set is processed with over 10% fewer ingots. This result may have benefits to the plant such as freeing up some machine capacity. In general the operational parameters reflect positive trends as the combination flexibility increases.

5.2.2 Economic Analysis

Chapter 2 shows that the total variable cost consists of several components: metal value, ingot cost, processing cost, scrap reprocessing cost, and inventory cost. For this analysis, we determined the metal value based on the shipped weight for each scenario. We calculated ingot costs by assigning a cost per pound to the cast ingot. Processing costs were calculated for each ingot by production center, based on the method described in Chapter 2. We calculated inventory holding cost based on metal value plus the value added in processing at the standard interest rate. Finally the scrap reprocessing cost

contains two items: reprocessing cost for the standard process scrap during metal forming, as well as reprocessing cost for excess metal scrap.

Figures 19 and 20 display graphs of normalized total variable cost for each alloy. The x-axis represents the planning scenarios being considered. The flexibility of order combinations increases to the right. As the flexibility of combinations increases the total cost of processing the set of orders decreases. This decrease in cost is directly related to the decrease in the number of ingots being processed (see Tables 9 and 10). Figures 21 and 22 display the normalized total revenue which also decreases as the flexibility of combinations increases because the total shipped weight (see Tables 9 and 10) decreases. The decrease in revenue is not a necessary feature of order combinations. In our analysis the shipped weight decreases across the scenarios because in some of the combinations orders are produced to a lower level within the quantity tolerance.

Figure 23 and 24 show that the total contribution, which is total revenue minus total variable costs, increases slightly as the number of order combinations increases. Since both total revenue and total variable cost decrease by similar amounts, the total contribution does not significantly change from scenario to scenario. The static nature of this evaluation satisfies only the current confirmed set of orders; it does not allow new orders to be received. In practice however, some of the uncommitted metal would be applied to future orders, so that the true relationship for total contribution may be different than shown in Figures 23 and 24.

Increasing combination flexibility decreases the total variable cost for processing a given set of orders, but in our analysis total revenue decreases as well. What would happen if we required revenue to remain constant (i.e., we removed the quantity tolerance and aimed for a specific quantity for each

Figure 19: Total cost for processing set of alloy 1 orders

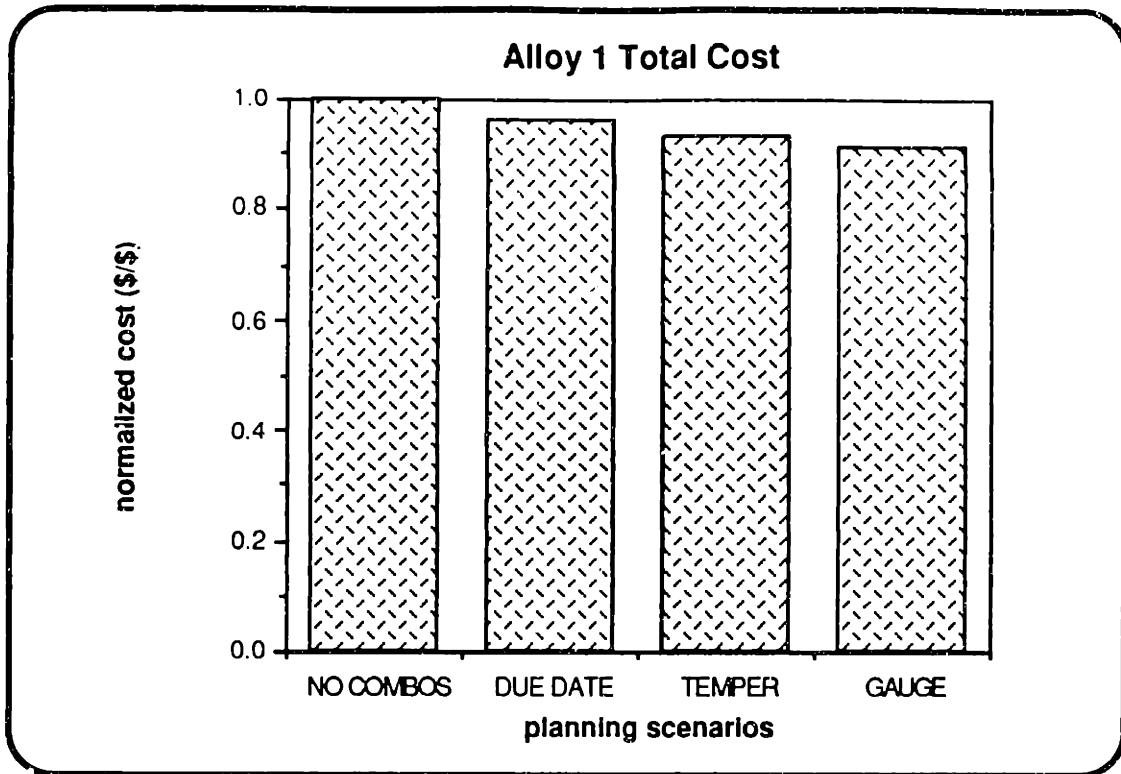


Figure 20: Total cost for processing set of alloy 2 orders




Figure 21: Revenue from processing set of alloy 1 orders




Figure 22: Revenue from processing set of alloy 2 orders




Figure 23: Total contribution from alloy 1 orders


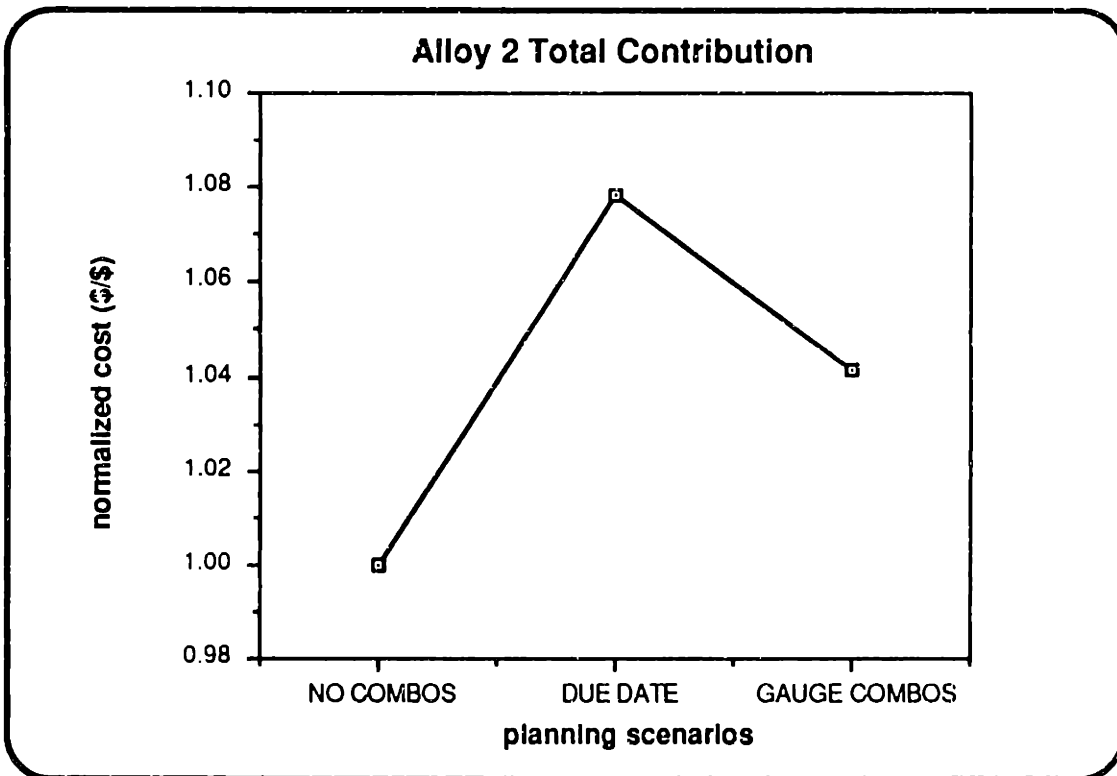


Figure 24: Total contribution from alloy 2 orders



order)? In this situation, it seems that total variable costs would also decrease. Since revenue is constant, this implies that profits would increase with increasing combination flexibility. An estimate for the magnitude of the increase in total contribution is the reciprocal of normalized total cost for the combination scenario (e.g. $1/0.9$ for gauge combinations of alloy 1). Perhaps we need to reconsider order combinations in light of its impact on the quantity of metal shipped for each order. What other implication does this have? If a customer receives less than the nominal order weight for today's order, will the customer order more metal in the next order?

This financial analysis indicates that order combinations do not generate additional profits for the plant, but order combinations decrease total variable costs. If the revenue were held constant across the scenarios, perhaps profits would increase with increasing combination flexibility. We also noted positive operational results above. A more detailed analysis which considers other products and alloys is necessary to determine the full impact of the financial and operational results. We can answer questions such as: does the decrease number of ingots and total recovered weight create more capacity for other products? Are there benefits beyond those noted for alloys 1 and 2?

5.3 Custom Ingot Scenario

Given the recent popularity to low inventory philosophies such as just-in-time, and the push in many industries to lot sizes of one, the large ingot philosophy of this plant seems unusual. Rather than process large ingots and use order combinations, why not cast custom ingots for each order? A custom ingot satisfies the quantity for a specific order without generating excess metal.

To provide some insight to this custom ingot question, we performed an analysis of the alloy 1 order set. We decided to compare a custom ingot scenario to the gauge combination scenario for alloy 1. To perform this analysis we used the following principles:

- applied same shipped weight for each order as in the gauge combination scenario
- used the width and thickness for each ingot from the gauge combinations scenario, and adjusted the length to fit the order weight
- assumed that the percent yield from the ingot was the same as the next smallest ingot for alloy 1. For ingots smaller than the smallest alloy 1 ingot, we used the percent recovery for the smallest standard ingot.
- used the same processing path as in the gauge combinations scenario
- treated the ingot stock cost factor as a parameter and performed a breakeven analysis (see Figure 25). Assume all custom ingots cost the same on a per pound basis.

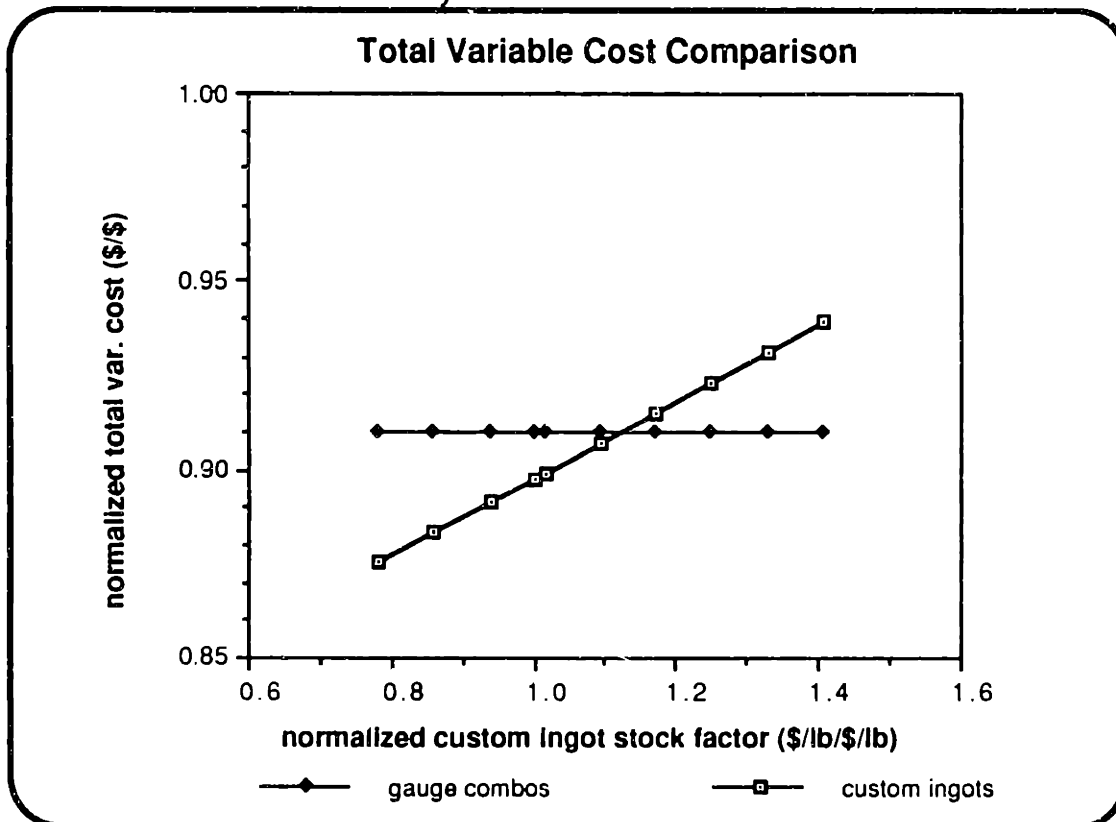
For the custom ingot analysis we performed the same operational and financial calculations described in the scenario evaluation in this chapter. Table 11 summarizes the results for the operational evaluation. Excess metal is eliminated, but more ingots are needed to satisfy the set of orders.

Since total revenue remains constant, Figure 25 displays a breakeven analysis which compares the total variable cost for processing the order set by 1) using the gauge combination scenario and 2) using a custom ingot scenario. The total variable cost for custom ingot scenario is expressed as a function of the ingot stock cost factor (cost per pound of a cast custom ingot - see section 2.2). The x-axis of the graph is normalized by the ingot stock cost factor for the

Table 11: Custom ingots operational analysis

	GAUGE COMBOS	CUSTOM INGOTS
upper ship tolerance	597,879	597,879
nominal order weight	514,538	514,538
actual shipped weight	553,568	553,568
# ingots processed	33	45
est. recovered weight	700,679	553,568
avg recovery per ingot	21,232	12,302
# of combination sets	15 sets	-
shipped/recovered	0.79	1

Figure 25: Breakeven analysis comparing custom ingots vs gauge combination scenario for alloy 1



most expensive standard ingot used in the gauge combination scenario. Based on this analysis, the breakeven point for total cost occurs at a value 10% higher than the most expensive current ingot. For processing the given order set, this is the point of indifference between a custom ingot scenario and gauge combination scenario. Below the breakeven value, the total cost for processing the order set using custom ingots is less than using gauge combinations.

This custom ingot analysis raises several issues. Is a custom ingot strategy feasible? Can the ingot casting facilities cast custom ingots for 50 different alloys? What impact would custom ingots have on the other production operations? Should the plant consider adopting a custom ingot strategy? What is a reasonable cost per pound for custom ingots? What would be the actual impact on total processing costs?

Custom ingots could provide several benefits to the plant. Reduced inventory levels is one benefit. Custom ingots would eliminate ingot stock inventory, excess metal inventory, and reduce the work-in-process inventory. Custom ingots would also allow more flexibility and responsiveness to customer needs. Some drawbacks exist for custom ingot, however. One problem is that order lead time would increase. Currently ingots are cast to stock. When the plant receives an order, a stock ingot satisfies the order. Custom ingots, however, are cast after the plant receives the order. Using a custom ingot strategy, the order lead time increases by the time to cast an ingot for the order. Custom ingots would also impact the manufacturing processes in the plant. Some of the production centers are not designed to process custom ingots; processing custom ingots would be a difficult transition for these operations.


5.4 Steady State Analysis

The planning scenario evaluation discussed above analyzes order combinations on a one time, static basis. A more realistic analysis would portray ongoing plant operations, and allow new orders to enter the order set. For alloy 1, this section presents one estimate of the ongoing plant performance assuming that the gauge combination scenario is standard operating procedure. For this analysis we estimate the steady state levels of two parameters for alloy 1:

- finished goods inventory - quantity of metal (per time period) that is produced for a confirmed order earlier than its due date (as a result of an order combination)
- excess metal inventory - quantity of metal that is produced, but not assigned to a confirmed order. This metal is held in inventory and applied against future orders received for this specification. If the metal is not utilized within six months, it is scrapped.

Figure 26 displays the performance of these two parameters for the gauge combination scenario for alloy 1. (The excess metal is accumulated over the length of the analysis, rather than being scrapped.) Since the finished goods inventory consists of confirmed orders within the time period, the level returns to zero at the end of the period. Assuming that the given order set is representative of future periods, we approximate the steady state average finished goods inventory level as the average level during our scenario analysis. We find that the average level of finished goods inventory for alloy 1 using a gauge combination strategy is 29,800 pounds.

Figure 26: Finished goods and excess inventory profile for alloy 1 gauge combination scenario



Excess metal inventory has three components which determine its level. Equation (5) describes the parameters affecting excess metal.

$$Ex(t + \Delta t) = Ex(t) + \Delta t * \{P(t) - F(t) - S(t)\} \quad (5)$$

where

$Ex(t)$ = level of excess metal inventory at time t (lbs)

$P(t)$ = production rate of excess metal at time t (lbs/week)

$F(t)$ = use rate of excess metal inventory for future orders at time t (lbs/week)

$S(t)$ = scrap rate of excess metal inventory at time t (lbs/week)

Δt = time increment used for analysis (weeks)

In the scenario analysis, excess metal inventory steadily grew because we assumed $F(t) = S(t) = 0$, while $P(t) > 0$. For ongoing plant operations, we recognize that $F(t) > 0$ and $S(t) > 0$. Over a long time horizon, the net accumulation of excess metal inventory is zero; therefore over a long time period (large Δt), $\{P(t) - F(t) - S(t)\} = 0$, which implies that $Ex(t + \Delta t) = Ex(t)$. We cannot determine the steady state level of excess metal inventory from equation (5), but we can establish an upper bound.

We assume that excess metal will not remain in inventory for more than six months (it is scrapped after this length of time). If we also assume that the order set analyzed is representative of future orders, then the maximum level of excess metal inventory is twice the ending level for the scenario analysis because the scenario analysis lasted three months and excess metal is scrapped after six months. The upper bound for alloy 1 excess metal inventory, $Max[Ex(t)]$, is 350,000 pounds.

5.5 Optimization Perspective on Order Combinations

Sections 5.1 and 5.2 describe an analysis of order combinations using heuristic procedures to group the orders. We employ simple criteria to determine the order groups. But, the nature of the order combination problem leads to the question; how to identify optimal combinations? The order combination methods generate a set of all compatible orders (from the given list of orders). Order combinations involve an economic tradeoff: incremental inventory holding and scrap costs for combining orders versus opportunity cost and excess metal scrap cost for not combining orders. This section pursues a more rigorous treatment of the combination problem by presenting one possible optimization model.

5.5.1 Model Description

We describe the model framework in five segments. First we consider the types of decision the model will support. These statements define the questions the model attempts to answer. The second segment highlights the tradeoffs which the model considers. In the third segment we define the inputs necessary to run the model. Next we list the assumptions which make the problem tractable. Finally, the last section specifies the constraints governing the problem solution.

1. Decisions

- which orders to group
- which ingot to select for each combination
- what weight to ship for each order

2. Tradeoffs

- inventory holding and scrap reprocessing costs for combining orders versus excess metal reprocessing cost for not combining orders

3. Inputs

- data from orders (alloy, temper, gauge, weight, due date)
- list of feasible combinations
- cost of processing a specific combination set
- revenue per pound by product type

4. Assumptions

- only confirmed orders generate revenue
- uncommitted metal is scrapped

- a specific combination set is produced just-in-time to satisfy the earliest due date within the set
- discrete ingot sizes apply and the estimated recovery was used as basis for yield weight

5. Constraints

- all orders must be processed by due dates
- machine availability constraints
- ingot size constraints (implicit -- see last assumption)

5.5.2 Model Formulation

The model requires the prior enumeration of all feasible order combinations and associated binary decision variable X_j with each feasible combination. The model also requires prior calculation of the total variable cost for processing each order-combination and the total revenue associated with each combination. The calculation of total variable cost includes the metal value, ingot stock cost, the processing cost, scrap reprocessing cost and inventory holding cost. Variable A_{jk} refers to resources required to process orders. These resources include manufacturing capacity at the relevant production centers (e.g. time available for hot rolling or cold rolling), ingot inventory of standard sizes, as well as other manufacturing resources.

Definition of Variables

J = set of all possible order combinations

$I(j)$ = set of orders i included in order combination j

C_j = total cost of combination j (includes ingot, processing, holding and scrap costs)

R_j = total revenue generated by order combination j

A_{jk} = amount of resource k used for combination j

B_k = amount of resource k available in planning horizon

$X_j = 1$ if combination j is selected /

= 0 otherwise

Objective and Constraints

$$\text{Maximize } \sum_j (R_j - C_j)X_j \quad (6)$$

subject to:

$$\sum_{j \in J(i)} X_j = 1 \text{ for all orders } i \quad (7)$$

($J(i)$ = set of combinations containing order i)

$$\sum_j A_{jk} X_j \leq B_k \text{ for all resources } k \quad (8)$$

$$X_j = 0 \text{ or } 1 \quad (9)$$

The model represents a set partitioning formulation. Equation (7) specifies that each order is contained in only one selected combination set. ([19], [20]) The model ignores the multi-period nature of the production operation.

The output of the model defines an optimal grouping of the order set and supports decisions in the short-term planning function. In theory an automated system could be developed which takes a given set of confirmed orders, then generates a list of feasible combinations, performs the cost and revenue calculations and then determines an optimal grouping of these orders.

Chapter 6 -- Conclusions & Recommendations

This thesis attempted to answer several issues raised by the concept of order combination: 1) what orders can be combined? 2) what criteria should be used to group orders? 3) can an optimal set of order combinations be determined? In our study we focused on the first two questions. Also as the study proceeded, we asked the question: what impact will an order combination strategy have on the plant operations?

The first section of this chapter reviews the conclusions for each of these questions. The second section provides recommendations for future work in order combinations.

6.1 Conclusions

Based on our study of alloy 1 and alloy 2, we present the conclusions and findings for each question listed above.

What orders can be combined?

We found the production process has the flexibility to produce orders with different specifications from the same ingot. Customer orders, for a single alloy, with different tempers, gauges, lengths and widths can be combined. Chapter 3 outlined the following combination possibilities:

- Different gauges can be produced from a single ingot in a single processing path. The machine and metallurgical constraints determine compatible gauges (for a single cold rolling pass).
- Different tempers can be produced from a single ingot. Commonality of production operations determines compatible tempers.

- Combinations of different widths and lengths were not studied, but are feasible, as well

What criteria should be used to group orders?

The selection of specific combinations from the list of feasible combinations is based on maximizing the net contribution for the combination. In picking specific combinations from a list of feasible combinations, we used two criteria: 1) minimize uncommitted metal and 2) minimize the difference in order due dates. If we group two orders in an ingot and their combined order quantity fully commits the ingot and their due dates are within two weeks, this is considered a "good" combination.

A secondary criterion is to minimize the processing complexity. Consider an order that has two feasible combinations (with same total weight): one is a due date combination (same specifications with due date three weeks in future) and the other is a gauge combination (with due date two weeks in future). Since the gauge combination adds processing complexity, we choose the due date combination in this case.

What impact will an order combination strategy have on the plant operations?

Based on the scenario analysis in Chapter 5, total variable cost for processing an order set decreases as the combination flexibility increases. The impact on plant profitability is not clear because in the scenario analysis the revenue also decreased as combination flexibility increased. (Figures 19 - 24 display these results.) This was due to the way the we performed the analysis, but the revenue decrease is not necessary. If we performed this analysis with the added restriction that revenue (shipped weight) remain constant for all the scenarios, the total variable cost for processing the order set would still

decrease. Since the revenue is constant, the total contribution would then increase with increasing combination flexibility.

Order combinations have several operational impacts. The number of ingots needed to process a given set of orders decreases as the combination flexibility increases. Also the quantity of excess metal decreases. Order combinations have negative impacts on the operation, as well. Grouping orders with different gauges and tempers into a single ingot increases the processing complexity and decreases the overall yield for the ingot.

6.2 Recommendations for Future Work

In the near term the plant can make several enhancements to the simple order combination framework presented in this thesis. First the plant needs to more fully understand the manufacturing process implications of combining different specifications. What is the true recovery impact of a gauge combination? What is the recovery loss due to a temper combination? Can the plant improve the administrative and tracking systems to better handle different specifications assigned to a single ingot (see section 3.2)?

To expand the scope of combination rules, the plant should study other heat-treatable alloys. The framework presented in Chapter 3 could apply to other heat-treatable alloys. Can similar gauge combination tables be developed for the other alloys? Is there a better format to represent this knowledge? Do the other T tempers have common processing operations?

The alloy 2 study focused on H32 temper. Do the other H tempers have similar processing paths? Can different H tempers be produced from the same ingot? Are there opportunities for gauge combinations in other non-heat-

treatable alloys? Or does the cold work for H tempers preclude gauge combinations for these alloys?

If the plant adopts a broader order combination strategy, then the decision process for order combinations must be automated. It is possible to develop a program that can identify feasible combinations of specific (confirmed) orders, choose an optimal set of combinations to execute and assign the combinations to ingots.

List of References

- [1] Nahmias, S. *Production and Operations Analysis*. Homewood, IL: Richard D. Irwin, Inc., 1989
- [2] Anthony, R. N., and J.S. Reece. *Accounting Text and Cases*. 8th ed. Homewood, IL: Richard D. Irwin, Inc., 1989
- [3] Wascher, G. "An LP-based Approach to Cutting Stock Problems with Multiple Objectives." *European Journal of Operations Research* 44, 1990, pp.175-184.
- [4] Goulimis, C. "Optimal Solutions for the Cutting Stock Problem." *European Journal of Operations Research* 44, 1990, pp.197-204.
- [5] Stadtler, H. "A One-dimensional Cutting Stock Problem in the Aluminum Industry and its Solution." *European Journal of Operations Research* 44, 1990, pp.209-223.
- [6] Tokuyama, H., and N. Ueno. "The Cutting Stock Problem for Large Sections in the Iron and Steel Industries." *European Journal of Operations Research* 22, 1985, pp280-292.
- [7] Rhee, W. T. "Optimal Bin Packing with Items of Random Sizes." *Mathematics of Operations Research* 13, 1988, pp. 140-150.
- [8] Frederickson, G. N. "Probabilistic Analysis for Simple One- and Two-dimensional Bin Packing Algorithms." *Information Processing Letters*, 11, 1980, pp. 156-161.
- [9] Coffman, E.G., K. So, M. Hofri, and A.C. Yao. "A Stochastic Model of Bin Packing." *Information and Control* 44, 1980, pp. 105-115.
- [10] Haessler, R.W., and F.B. Talbot. "A 0-1 Model for Solving the Corrugator Trim Problem." *Management Science* 29, 1983, pp. 200-209.
- [11] Hatch, J.E., ed. *Aluminum Properties and Physical Metallurgy*. Metals Park, Ohio: American Society for Metals, 1984.
- [12] Hosford, W.F., and R.M. Caddell *Metal Forming: Mechanics and Metallurgy*. Englewood Cliffs, NJ: Prentice Hall, 1983.
- [13] Harris, J.N. *Mechanical Working of Metals*. Oxford: Pergamon Press, 1983.
- [14] Avitzur, B. *Metal Forming: Processes and Analysis*. New York: McGraw Hill Book Company, 1968.

- [15] Rauch-Hinden, W.B. *A Guide to Commercial Artificial Intelligence*. Englewood Cliffs, NJ: Prentice Hall, 1988.
- [16] Hayes-Roth, F. "Rule-Based Systems." *Communications of the ACM* 28, 1985, pp. 921-932
- [17] *LEVEL5 OBJECT, User's Guide*. Information Builders Inc. New York, 1990.
- [18] Stefik, M., and D. Bobrow. "Object-Oriented Programming: Themes and Variations." *The AI Magazine* 6, 1986, pp. 40-62.
- [19] Vasko, F.J., F.E. Wolf., and K.L. Stott. "A Set Covering Approach to Metallurgical Grade Assignment." *European Journal of Operations Research* 38, 1989, pp.27-34.
- [20] Fisher, M.L., and P. Kedia. "Optimal Solution of Set Covering/Partitioning Problems Using Dual Heuristics." *Management Science* 36, 1990, pp. 674-686.