Optimization of a Battery Manufacturing Line Using Computer Simulation

by

Gregory J. Mazenko

B.S. in Systems Engineering
United States Naval Academy, Annapolis, MD (1986)

Submitted to the Departments of Electrical Engineering and Computer Science and the Sloan School of Management in Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Electrical Engineering and Computer Science
and
Master of Science in Management

in conjunction with the
Leaders for Manufacturing Program
at the Massachusetts Institute of Technology
June, 1995

Signature of Author ______________________
Sloan School of Management
Department of Electrical Engineering and Computer Science
May 12, 1995

Certified by ______________________
David E. Hardt
Professor of Mechanical Engineering

Certified by ______________________
Roy E. Welsch
Professor of Statistics and Management Sciences

Accepted by ______________________
Frederic R. Morgenthaler, Chairman
Committee on Graduate Students

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

Optimization of a Battery Manufacturing Line Using Computer Simulation

by
Gregory J. Mazenko

B.S. in Systems Engineering
United States Naval Academy, Annapolis, MD (1986)

Submitted to the Departments of Electrical Engineering and Computer Science and the Sloan School of Management in Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Electrical Engineering and Computer Science
and
Master of Science in Management

Abstract

The importance of major capacity expansion projects in a manufacturing company cannot be overstated. The successes or failures of the expansion projects have tremendous influence on the company's ability to serve its markets, whether the strategic goal is to grow, to maintain market share, or to gain a foothold in new and developing markets. Innumerable decisions must be made throughout the life of an expansion project. Sometimes project teams have reliable, historical data at their disposal to help in the decision-making process. Other times, important decisions must be made on the basis of educated conjecture. Generally, the more relevant historical data and useful analysis tools that a project team has at its disposal, the better its decisions will be.

The goal of this thesis was to examine the effects that computer simulation could have on a large capacity expansion project in a major global manufacturing firm. The project studied was in its design and development phase, so any insight gained from the simulation work could assist the project team's decision-making process before actual production operations began.

This thesis examines the relevant issues surrounding the use of the computer simulations, the various statistical techniques used in their development, and the insights gained from their use. It describes the specific effects that simulation had on the project team's problem-solving process. It illustrates how computer simulation was used in conjunction with the theory of constraints to develop possible improvement strategies. Finally, the thesis examines how simulation was used to help make recommendations for future expansion projects.

Thesis Advisors:

David E. Hardt, Professor of Mechanical Engineering
Roy E. Welsch, Professor Statistics and Management Sciences

Acknowledgment

I wish to acknowledge Polaroid Corporation's IBAM team and the Leaders for Manufacturing Program for their support of this work.

Table of Contents

List of Figures

List Of Tables

Chapter 1 - Introduction

1.1 Computer Simulation

Computer simulation is receiving recognition as an excellent way to gain insight into complex problems that cannot be easily solved using other methods of analysis. Simulation is particularly well suited to solving problems in systems that have many different randomly occurring phenomenon. Manufacturing operations that have variability in raw materials, randomly occurring machine failures, variability of human performance, and many sources of variation, fit nicely into the category of systems that are well suited for simulation analysis.

This thesis addresses the use of computer simulation in the design, development, and construction of a new manufacturing line at Polaroid Corporation's Battery Manufacturing Division.

1.2 Corporate Structure

Polaroid Corporation is divided into three major business units entitled High Resolution Imaging, Core Photography, and Electronic Imaging Systems. Battery Division, which is located at the R-5 plant in Waltham, Massachusetts is part of the Integral Film Assembly Division in the Core Photography Business Unit. A tree diagram of the pertinent parts of the corporation follows:

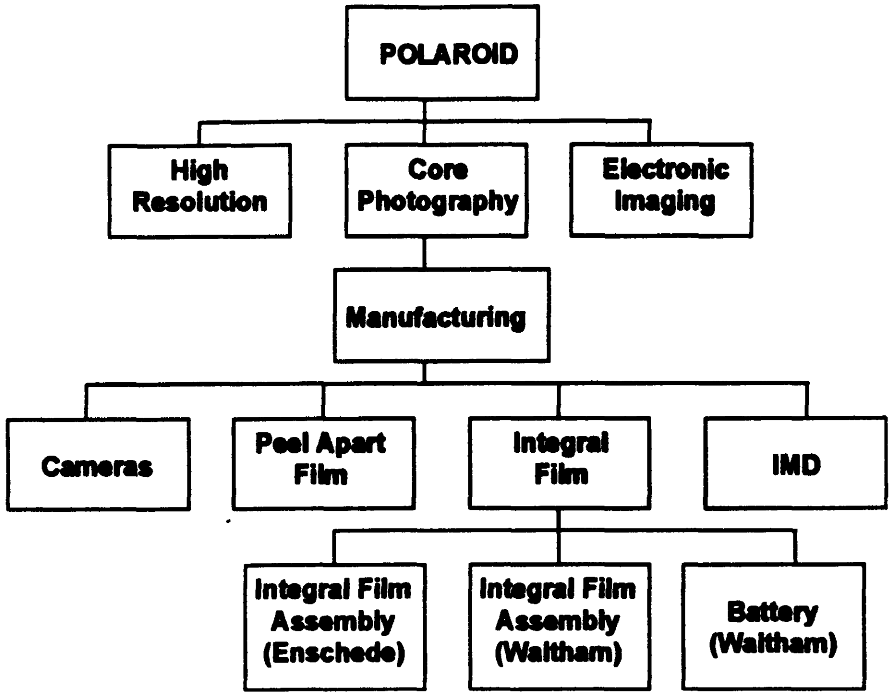

Figure 1: Polaroid Corporation Structure

Though Battery Division (R-5) is shown on the same tree level as Integral Film Assembly, all of the batteries which R-5 produces are subsequently used in the Integral Film Assembly operation at either the Waltham site or at the Enschede site in the Netherlands.

1.3 Battery Division (R-5)

R-5 is a relatively small plant, with roughly 250 permanent employees. They are the only battery manufacturing plant in the corporation, so they must conduct extremely high-volume operations year-round. A simple block diagram of operations at R-5 follows:

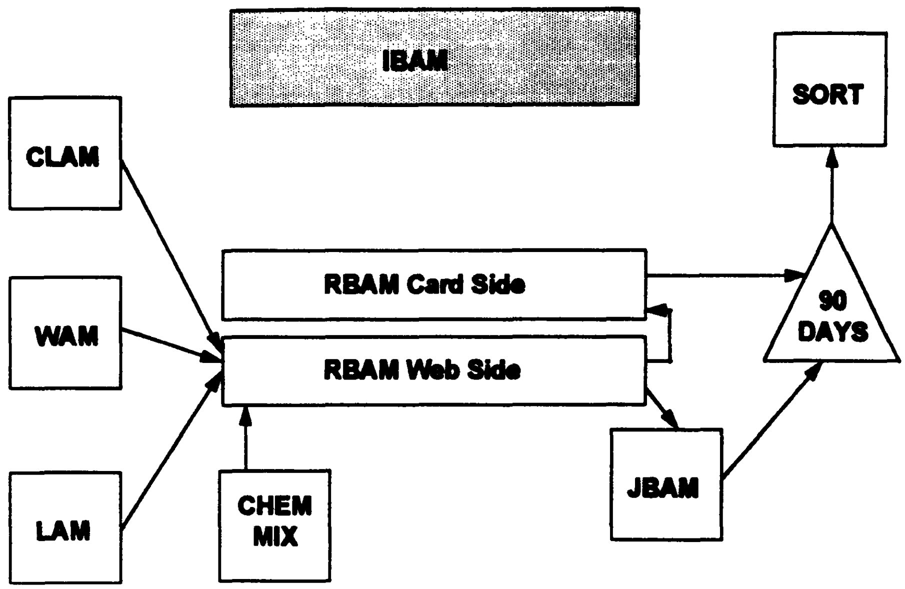

Figure 2: R-5 Process Flow

In order to promote an understanding of plant operations, a description of each acronym and operation shown in figure 2 follows.

CLAM Cathode laminator assembly machine. One of three pre-assembly machines in the plant. The CLAM produces rolls of cathode top plates for the batteries. The cathode is the positive battery terminal.

WAM Web assembly machine. The second pre-assembly machine. The WAM produces rolls of composite material that form the three inner layers of the battery between the cathode top plate and the anode bottom plate.

LAM Anode Laminator Assembly Machine. The third pre-assembly machine. The LAM produces the anode bottom plates for the batteries. The anode is the negative battery terminal.

CHEM MIX In the chem mix area, the slurry which is the electrolyte for the battery is produced and stored.

<u>RBAM</u> Rotary Battery Assembly Machine. The high-speed assembly line that takes materials produced in pre-assembly and chem mix and assembles them into finished batteries. The RBAM is so named because of its extensive use of rotary technology. It is called a machine, but is actually a manufacturing line made up of a series of fully-coupled machines. On the web side of the RBAM, batteries are constructed on a paper carrier web. The card side of the RBAM is a finishing operation in which the batteries are attached to thin cardboard squares to give the battery a more rigid structure. A series of voltage checks are also performed.

<u>JBAM</u> Joshua Battery Assembly Machine. Performs the same function as the card side of the RBAM. The card stock used on the JBAM is slightly smaller than that used on the RBAM card side. Batteries finished on the JBAM are used in film packs for the Captiva camera. Polaroid's internal project name for the Captiva is Joshua.

<u>90 Day Buffer</u> 90 days worth of finished goods are held at the plant. These seemingly high levels of inventory are necessary due to the nature of some types of battery defects that have thus far proven undetectable. These latent defects sometimes cause tiny shorts in the batteries that lead to very gradual losses of charge. These defects tend to be fully manifested after 90 days, so voltage tests are again performed after the 90-day incubation period has passed.

<u>Sort</u> Final voltage checks and packaging of the batteries for shipment are accomplished at Sort.

<u>IBAM</u> Incremental Battery Assembly Machine. The capital project that is the subject of this thesis. The IBAM is so named because it will give the plant an incremental increase in battery manufacturing capacity. It will be a battery manufacturing line that performs the same functions as the RBAM web side, but at much lower volumes. The IBAM is scheduled to produce batteries by January 1, 1996. At the time this thesis was

written, different portions of the IBAM were in varying stages of design, development, and construction.

1.4 Philosophy of IBAM

Sales projections for instant film in the mid to late 1990s dictated the need for substantially higher output from R-5 and the rest of the Integral Film Assembly Division. R-5 is currently the only location within Polaroid that manufactures batteries. As is shown in the block diagram of operations in figure 2, the RBAM web side is the only machine at R-5 which performs the assembly functions between pre-assembly and finishing operations performed on the RBAM card side or on the JBAM. The RBAM web side is currently stretched as far as it realistically can be. Short-term solutions like cutting maintenance and test time are impractical, since they could have devastating long-term effects.

A major project of some sort was needed in order to give R-5 the extra capacity required. Whatever the nature of this project, it was determined that its global objectives were the following:

- A step increase in annual battery manufacturing capacity.
- A step increase in yield.
- Continual improvement in customer satisfaction by the elimination of manufacturing defects.

The project would be constrained in that no modifications to current pre-assembly or chem mix procedures could be introduced, and the batteries themselves would have to be identical to those currently produced.

With these broad objectives in mind, several different approaches were examined. The first was to continue some on-going efforts to improve the existing equipment. The second was to make major modifications to the existing equipment by providing redundancy in some of the critical areas. The third approach was to build a new line which utilized existing rotary technology. The final possibility was a new machine which used intermittent motions in many of the areas that currently employed rotary technology.

For a variety of strategic reasons which are not the subject of this thesis, the high-yield intermittent motion, new line concept was chosen. A cross-functional team including R-5 members and members of the Integral Film Assembly's Machine Engineering Division (MED) was formed. Their task was to design, develop, and construct a high yield intermittent motion battery assembly machine which met the previously stated program objectives. The IBAM team began their work in late 1992.

1.5 My Expected Contribution

My internship at R-5 began on June 1, 1994, about 1 1/2 years through the 3 year project. At this point, different portions of the IBAM were in various stages of design, prototyping, and construction. Following a brief indoctrination period, the IBAM program manager and I decided that I could make a worthwhile contribution to the project by taking on the responsibilities associated with the development of computer models of the IBAM. The IBAM program had purchased a modeling package entitled Pro-Model for Windows version 1.10. Up to that point, the use of the tool had been limited to a single model of the IBAM's anode cell which had been developed by an outside simulation development firm.

Although the programmers from the outside firm were skilled modelers with extensive experience in modeling of manufacturing systems, it was difficult for them to develop models with the flexibility necessary to test the wide variety of issues which emerge in a rapidly changing project such as the IBAM project. It would be much more beneficial to the project if there were a team member who could develop many different flexible models in response to rapidly changing project decisions and constraints.

My charter was to develop models of the IBAM as a whole and of important components in order to examine the effects that different factors had on overall IBAM performance. By doing this, we hoped to optimize the performance of the line, and to perhaps uncover some hidden problems with the IBAM design.

1.6 Thesis Structure

- This thesis is divided into eight chapters.
- Chapter 2 is a brief chapter which describes the batteries themselves and some broad issues of concern in battery manufacturing.
- Chapter 3 takes a close look at the methods used to quantify the profiles of the behavior of the RBAM equipment.
- Chapter 4 examines some issues surrounding the development of the computer models of the IBAM, and introduces a problem involving slurry dry-out which became the focus of the IBAM team in the second half of the internship.
- Chapter 5 describes in detail the efforts of the IBAM team to solve the dry slurry issue.
- Chapter 6 focuses on the theory of constraints and how it was used to form recommendations for improved throughput of the IBAM.
- Chapter 7 explores some of the possibilities that would come into being if some of the IBAM's constraints were removed.
- Chapter 8 explains some overall conclusions and lessons from the internship.

Chapter 2 - Batteries

2.1 Overview

This chapter is intended to give a brief overview of the batteries that are produced at R-5. First, the battery and its components are discussed. Then, a very general description of the manufacturing methods currently employed on the RBAM web side and some of the surrounding issues is given. Finally, the methods used on the proposed IBAM along with the major areas of expected improvement are examined.

2.2 Battery Structure

An understanding of the structure of the battery is an essential prerequisite for an understanding of the issues which are addressed in this thesis. Below is an exploded view of a completed battery:

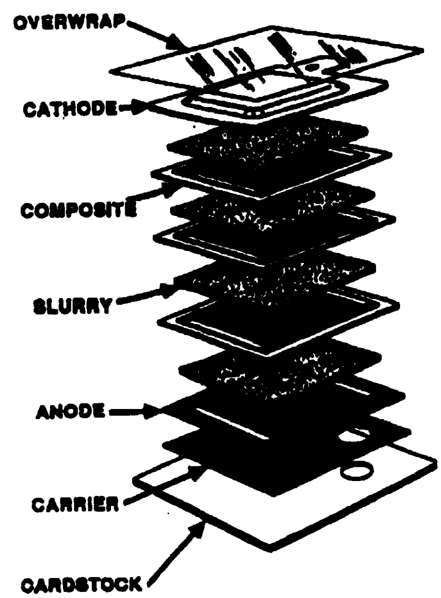

Figure 3: Battery Structure

The cardstock on the very bottom and the overwrap on the very top of the battery are added in the finishing stage of production. Recall from Chapter 1 that finishing operations are accomplished on either the RBAM card side or on the JBAM. All of the layers in between are assembled on the RBAM web side. The new IBAM will assemble the same components as the RBAM web side.

The carrier, also referred to as the carrier web, is nothing more than brown paper. The anode, or negative terminal, is aluminum based. It is the first layer applied to the carrier. Next is the first of four layers of slurry. The slurry is the electrolyte for the battery. Physically, it is a black paste which requires very special handling. One of the reasons it requires special handling is because it loses its desirable properties as an electrolyte if exposed to open air for an excessive amount of time. For this reason, if the RBAM web side stops, and the layers of slurry in the unfinished batteries have been exposed to open air for 15 minutes or more, the batteries must be discarded. This will prove to be a critical issue on the IBAM, and will be examined in detail in later chapters.

The next six layers in the battery are alternating layers of composite material and slurry. These interior layers give the battery greater voltage capacity. The final component of concern is the cathode top plate. The cathode is also aluminum based, and is the battery's positive terminal.

When fully assembled, the battery is less than one centimeter tall. This dimension is very important, since the battery is an integral component of an instant film package. A bulkier battery design would create a bulkier film package, which is undesirable for many reasons including cost and customer satisfaction.

2.3 Overview of Production on the RBAM Web Side

The RBAM web side is a series of machines which are fully coupled to one another through the carrier web. All of the pre-assembly machines, the LAM, the WAM, and the CLAM, produce rolls of their respective materials for use on the RBAM web side. Components on the pre-assembly rolls are laid out four across. Each 1x4 section of the roll is referred to as a row. The following diagram shows a pre-assembly roll. This diagram could represent an anode, composite, or cathode roll, since all are configured similarly:

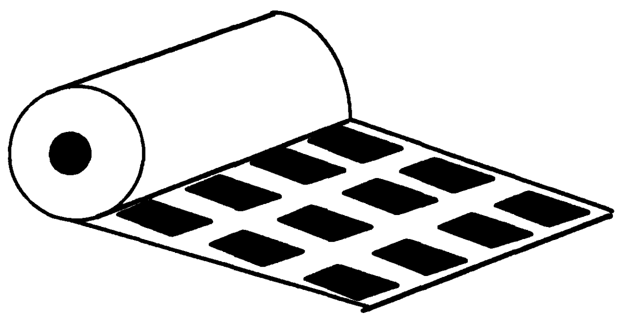

Figure 4: Roll From Pre-Assembly

The backbone of the RBAM web side is the carrier web. The carrier web is a large roll of brown paper. The carrier web roll is placed on an unwind stand at the beginning of the line and is fed continuously through the series of machines which make up the RBAM web side. As a section of the carrier web travels further down the line, components are placed on it one 1x4 row at a time. The rotary equipment which places the components is situated in-line with the carrier web, overhead. The slurry dispensing mechanisms are also in-line and overhead.

The rolls of pre-assembled components are placed on unwind stands located off-line. (In the context of discussions of the RBAM and IBAM, off-line means nothing more than being physically situated somewhere other than directly over or under the carrier web. Off-line does not imply that the machines are decoupled.) The pre-assembled rolls are fed continuously from their unwind stands to the in-line rotary equipment which cuts rows off of the rolls and places them row by row onto the carrier web.

The last in-line machine that the carrier web is fed through is a cut-off mechanism simply called the final cut-off (FCO). The FCO waits until 18 rows pass by before making its cut. The batteries are now configured in an 18x4 array as shown in the following diagram:

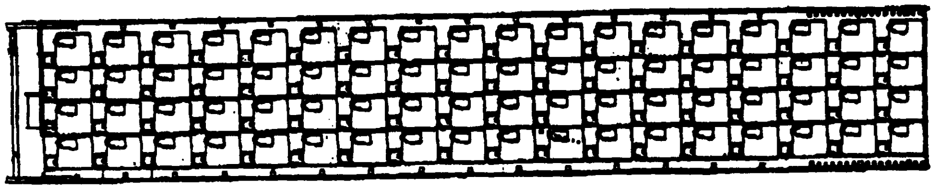

Figure 5: Battery Strip

These 18x4 battery arrays are referred to as *strips*. The strips proceed from the FCO to a metal vacuum table where they wait to be fed into the vacuum sealer. During steady state operations, there will be three strips waiting to be vacuum sealed.

The vacuum sealer can handle a single strip of 72 batteries at a time. The sealing operation marks the point where the slurry is no longer exposed to air. Following the sealing operation, battery strips are loaded into tubs where they wait for further processing on the finishing equipment.

2.3.1 Important Issues on the RBAM Web Side

The most pressing problem that the RBAM web side faces is low run time. This is because all of the machines which make up the RBAM web side are completely coupled to one another by virtue of the fact that all components are placed onto a continuous paper carrier web. If any machine on the RBAM web side goes down, the web must stop and production stops. The only buffering occurs on the vacuum sealer table, and in festoons on the unwind stands. Recall that the vacuum sealer table holds three strips. Each strip only takes about eight seconds to process, though, so this buffering capability is virtually negligible. The festoons on the unwind stands are simply layered rollers through which the roll of material is fed in a serpentine fashion. If the main roll runs out or stops feeding, the festoon rollers can move closer together and provide a short period of uninterrupted material flow upstream. A diagram illustrating the concept is shown in figure 6:

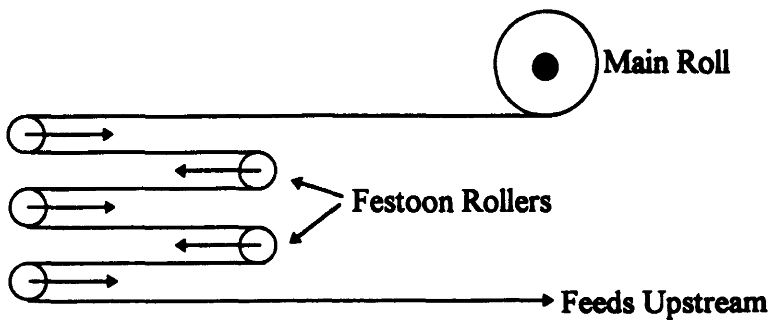

Figure 6: Festoon System

Most of the festoons have only enough material to feed the upstream processes for a few seconds. The real purpose of the festoons is to allow operators to splice in new rolls of material at the ends of old rolls without shutting the whole RBAM web side down.

Despite these two minor exceptions, the RBAM web side is fully coupled and non-buffered. This would not be a problem if all of the machines which comprise the RBAM web side were very reliable, but this is not the case. In the period that will be described in chapter 3 from which machine reliability data was taken, the RBAM web side was down due to individual machine failures for 33% of the time during which it was scheduled to run. Improvements in run time would dramatically reduce the cost of manufacturing each battery.

In addition to run time problems, yield loss problems also afflict the RBAM web side. The specific causes of yield loss are constantly evaluated, and possible solutions to yield loss problems are regularly pursued. Yield loss varies from period to period, but is generally on the order of 10% or so. Improvements in yield loss would also dramatically reduce the manufacturing cost of each battery.

It is worth explaining in very broad terms some of the causes of high levels of yield loss on the RBAM web side. The first issue is the serial correlation of many of the defects. Since the carrier web flows continuously and at a fairly high speed, many problems which require that material be discarded are not detected until quite a few battery rows have been affected. An example of such a defect would be skewed composite rows. Sometimes the rotary equipment which places the composite rows will place them with poor alignment. If one row is skewed, chances are that the next row will be too, and so on. Also, if problems occur with material before it is placed on the web, there is no way to discard the material. Bad material gets placed on the carrier web more often than is desirable.

Another general source of high yield loss is the finishing equipment's sensitivity to misaligned (misregistered) material. If one of the rows on an 18x4 battery strip appears to be misregistered, the whole strip is discarded. This may seem very wasteful, and it is, but it helps prevent the finishing machines from getting jammed and going down.

If the finishing equipment were able to process strips other than 18x4, it would be possible to decrease yield loss, because the bad parts of 18x4 strips could be discarded and the rest could be saved. Unfortunately, the finishing equipment can only process 18x4 strips, and cannot be reconfigured without large capital expenditures and large amounts of downtime. (Keep in mind that R-5 is the only plant in the corporation which manufactures batteries. Major equipment configurations which could keep the plant down for weeks are generally considered impractical.)

Run time losses and yield losses drive the manufacturing cost of each battery produced on the RBAM up substantially, and are worthy of the full-time attention of the entire staff at R-5. Specific initiatives for improvements to the RBAM equipment are an appropriate area for further work, but were not the focus of my internship. My efforts focused on the IBAM project and ways in which we could optimize the performance of the new line. Of course, an appreciation of the issues surrounding production on the RBAM was essential to the IBAM team's ability to design a better line. A basic understanding of these RBAM issues will contribute greatly to the reader's understanding of the IBAM issues introduced later in this thesis. The next section explains in broad terms some of the differences between the RBAM and the IBAM.

2.4 Production on the IBAM

The IBAM team tried to identify major problem areas on the RBAM which could be improved upon. In broad terms, I believe that the fundamental improvements designed into the IBAM can be summed up as follows:

- Ability to discard bad material before it is placed onto the carrier web.
- Buffering of certain machines so that failures will not cause the entire line to go down.
- Improvement in the ability to precisely place components due to intermittent motion instead of continuous motion.

The following diagram shows a sketch of the IBAM as of mid September, 1994.

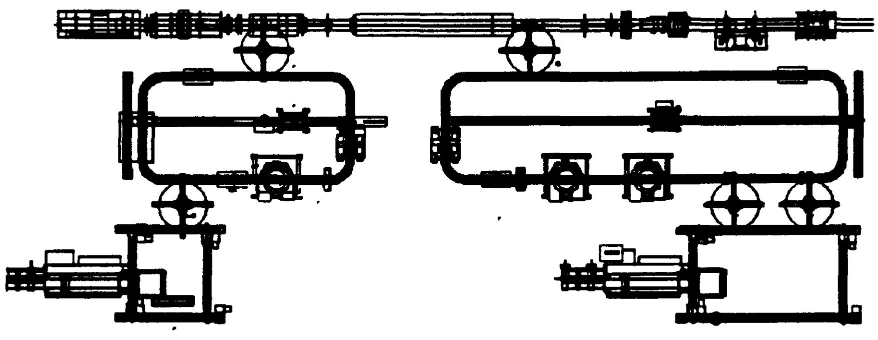

Figure 7: IBAM as of 9/16/94

The line can really be thought of in terms of the following five areas:

1. Web Transport System
2. Anode Cell
3. Composite Cell
4. Cathode Machine
5. Vacuum Sealer

A broad overview of each area follows.

Web Transport System - The same paper carrier web used on the RBAM will be used here. The major difference is that the web motion will be intermittent. The web will be stationary when components are placed on it. Following placement of the components, the web will advance an appropriate distance for the next row(s). Slack at certain points along the web will allow different sections to advance separately from other sections, as long as all sections advance at the same average rate over time.

Anode Cell - The unwind stand is buffered by the cut loop. Anode rows will be cut and placed into canisters. The canisters will move around the cut loop's conveyor system, where they will be picked up by a cambot (a pick & place device.) The anode rows will be placed onto pallets which travel counter-clockwise around the main conveyor. The pallet will travel under the slurry dispenser, where slurry patches will be applied, and then under the vision check system. If the vision system detects any problems, the pallet will go to the reject conveyor. If the pallet is not rejected, it will continue on to the next cambot, which will place the anode row onto the carrier web.

Composite Cell - Very similar to the anode cell, except that pallets will travel clockwise and will each carry two composite rows. The cambots between the cut loop and main loop will each place a single composite row onto separate positions on the pallets. The cambot between the main loop and the carrier web has a double head, and will pick and place both composite rows simultaneously.

Cathode Machine - In the configuration depicted, the entire cathode machine is in-line, overhead of the carrier web. The cathode rows are pocketed to give them their shape, and then placed on the top of the battery stacks.

Vacuum Sealer - Unlike the RBAM, the IBAM's vacuum sealer will be in-line, and will only seal three rows at a time instead of eighteen. There will be no vacuum table. Final cutoff is after the sealer instead of before, as on the RBAM.

Overall machine yield and run time should be better than that of the RBAM web side. The IBAM will produce about 1/6 the volume of the RBAM web side, so it will run much more slowly.

The machine configuration described above was the one which I began to develop computer models of. Details of the development of the models and of the problems that the model uncovered are given in the remaining chapters.

Chapter 3 -- Data-Based Characterization of Existing Process

3.1 Chapter Overview

This chapter is dedicated to a discussion of the data used to develop the various machine downtime distributions which were inputs into the various computer simulations developed during the internship. It begins with a discussion of why a detailed data analysis was necessary. Next, data collection procedures and issues surrounding the validity of the data are addressed. Finally, the chapter describes the specific statistical tools and techniques that added value and created confidence in the distributions themselves.

3.2 Importance of Data Analysis

A model of a system is only as good as the assumptions upon which the model is based. This certainly applies to computer simulations. Many questions regarding the performance and behavior of the IBAM can be easily answered using straight-forward analytical techniques. It is the questions that require insight into randomly occurring phenomenon such as machine failures that can most easily be answered using more sophisticated techniques such as computer simulation. The accuracy with which these randomly occurring phenomenon are characterized will determine the usefulness of the model.

In the case of the IBAM, the shapes of the machines' expected downtime distributions are critical. I will explain in detail why this is so in subsequent chapters. In addition to the shapes of the distributions, a detailed understanding of the specific types of failures which occur on the RBAM is essential when predicting performance of the IBAM. These performance predictions have profound effects upon design decisions of the IBAM machinery itself, and upon managerial decisions regarding such issues as scheduling, budgeting, and maintenance plans.

3.3 Procedures Currently in Effect for Data Collection

The plant currently has two separate processes in place for collecting machine downtime data for the various components of the RBAM. The first process is one in which the operators are responsible for entering appropriate data at a station containing a keyboard and a monitor which is tied into the plant's VAX computer. The second process is an automatic one which makes use of some data acquisition equipment. Data obtained automatically is also processed and stored by the VAX computer.

3.3.1 Manual Data Collection

The manual system is the plant's most reliable source of downtime information. When a machine on the RBAM goes down, operators are supposed to note the time of day to the nearest minute, and start a stopwatch in order to measure the duration of the downtime. After troubleshooting is complete, and the RBAM is up and running, the operator enters the time that the downtime began, the shift, the particular machine that went down, the sub-system of the machine, the nature of the problem, the corrective action taken, the duration of the downtime, and the number of consecutive occurrences of the problem. The number of consecutive occurrences is recorded in order to relieve the operators of the burden of having to make several identical entries which indicate short downtime durations. If the machine fails for the same reason five times within a two-minute period, the operators can simply call it a two-minute downtime and indicate five consecutive occurrences rather than attempting to break the occurrences into five separate entries. Little in the way of accuracy is lost by doing this, but much burden is taken off of the operators.

This manual data collection system has its shortcomings. First of all, many downtimes which last much less than one minute are simply not recorded. Secondly, since the plant has only one VAX computer, the system responds very slowly if other users are performing computationally intensive functions such as generation of yield reports and runtime reports. Frustration with the slow response time will sometimes deter operators from making proper entries. Great progress was made in addressing this problem when a new and faster VAX computer replaced the older one in use just prior to the beginning of the internship. But even with the newer computer, sluggish response can still be a problem. Thirdly, there is a tendency on the part of the operators to not use the stopwatch and to estimate the downtime duration by simply looking at the clock. This becomes evident when one looks at histograms of the downtime durations for the various machines. The number of eight, nine, eleven, and twelve-minute downtimes are generally noticeably lower than the number of ten-minute downtimes. It would appear that if a downtime is close to ten minutes, there is a tendency to simply call it a ten-minute downtime and not bother distinguishing it from a nine-minute downtime. This phenomenon was common around the five and ten-minute marks, but much less common as downtime length increased. Since the difference between a seven-minute downtime and an eight-minute downtime was of notable significance for reasons which I will explain later, it was necessary to smooth these areas out a bit. The details of how this was done are explained in section 3.5.4.

Overall, the quality of the data available for use in developing the computer models of the IBAM was high. In speaking with many of my colleagues about their experiences trying to obtain good data for various analyses of manufacturing operations in other plants, I realized that although the data available at R-5 is imperfect and could be greatly improved, it is basically sound data. The operators are usually diligent in recording the data, and the data is sufficiently detailed to be of great utility in many different types of analyses.

3.3.2 Automatic Data Collection

The automatic data collection process in place on the RBAM was still in the debug phase for the duration of the internship. Data which was collected by the automatic system was available, but the causes that the system assigned to downtime events were considered very unreliable. There is great potential to have data that does not have the inherent inaccuracies described in section 3.3.1 if the automatic data collection system is finally debugged and deemed reliable. But for now, the manually entered data is still the most reliable data source in the plant.

The IBAM will have state-of-the-art data acquisition equipment, so the potential to have an extremely high quality database of downtime, quality, and yield related data certainly exists. Proper utilization of this resource will greatly assist future analysis of the IBAM, and will help management greatly in their decision-making processes.

3.4 The Data Itself

The format in which the data was stored on the VAX was such that some up-front manipulation was necessary in order to put the data into a meaningful and useable format. A simple routine was available which downloaded the data in text format from the VAX computer to the user's personal computer hard drive via the plant's Ethernet-based local area network. Once the data was on the user's hard drive, whatever manipulation was necessary could be performed easily on a spreadsheet. A sample of the downtime data available follows. Note that each row identifies a single downtime event. The columns contain the following information:

1. Time of downtime event
2. Shift
3. Machine, or area which went down
4. Sub-system or sub-area
5. The problem
6. Operator initial
7. Action taken
8. Length of downtime in minutes
9. Number of consecutive occurrences

1	2	3	4	5	6	7	8	9
1/17/94 13:00	A	CATHODE	LAMINATOR	LAMINATION	O	UNKNOWN	10	1
1/19/94 4:30	C	CATHODE	DRUM	WEB BREAK	O	RESET	5	1
1/20/94 4:05	C	CATHODE	PIN BELTS	BROKEN BELT	M	REPLACED	20	1

The simulation package used could accept downtimes based upon clock times or times that the machine was actually in use. It seemed sensible to transform the data from its current clock-time form into a form which reflected machine usage time. If the downtime distributions were developed based strictly upon clock times, then inaccuracies resulting from not accounting for brief periods of non-scheduled time, maintenance time, and test time would be introduced into the model. Additionally, if machine usage times are the basis for the distributions, then the distributions for different machines will be independent of one another and exclusive of one another, and will more accurately describe the actual behavior of the individual machines.

This point requires some explanation: Since all machines on the RBAM are completely coupled to one another, if any one goes down, the entire line goes down. If for example, the cathode machine goes down at 13:00 for five minutes, and then runs for five minutes, and then goes down again at 13:10, the actual machine usage time between failures is only five minutes. The database indicates that ten minutes of clock time have expired since the last failure of the cathode machine. If the downtime distributions were constructed based upon clock times, they would be skewed in favor of longer times between failures. Although the simulation can use these less accurate distributions, the insight gained from studying the more accurate usage-based distributions is greater. The most problematic machines are more easily identified in the absence of the skewing that comes from using clock times.

If clock times are used, then all machine distributions are dependent upon one another in the sense that actual downtimes for each machine are incorporated into the times between failures of all other machines. But more important are the effects of the exclusive nature of usage-based downtimes on the simulation. If clock-based downtimes are used, then two or more downtimes can occur simultaneously during the simulation. This may adversely effect the simulation output. If usage-based downtimes are used, then it is impossible for two downtimes to occur simultaneously since once one machine goes down, the others are no longer in use. This more accurately reflects the behavior of the machinery itself.

One may argue that the effects of the skewing of the clock-based distributions in favor of long times between failures will be mostly offset by the occasional simultaneous occurrence of two or more downtimes. I argue that it is much better to attempt to characterize the machine behavior as accurately as possible rather than rely on two different inaccuracies to offset one another.

3.4.1 Period Studied

I initially elected to use data from the period from January 1, 1994 through June 30, 1994 to develop the machines' time-to-failure (ttf) and time-to-repair (ttr) distributions. I asked the engineer responsible for the data collection process and for generation of associated performance reports what he thought an appropriate time period was. He indicated that for the most part, he uses a three-month period when he is asked to generate reports which give insight into long-term trends. His philosophy was that any cyclical phenomenon related to machine performance usually repeat themselves several times each quarter. Increasing the timeframe studied would increase the workload without increasing the insight gained from the analysis.

Despite this advice, I elected to use a six-month period. I reasoned that the data was readily available, and a longer timeframe would increase the statistical significance of the distributions developed. It seemed reasonable to use the most recent six-month period in order to ensure that any recently developed trends were captured.

Once I began analyzing the data, I noticed that there were unusually large gaps between documented machine failures from January 1 through January 16. In asking around, I could not come up with any explanation for this such as the VAX computer being down, or unusually high non-scheduled time due to holidays, testing, or maintenance. The documented failures were so thinly spread during this period that there is no chance that it was due to particularly good machine performance. The machines never run as well as data from this period would indicate. In the end, it seemed best to simply discard the period from January 1 to January 16 as unreliable and do the study with data from January 17 on.

3.4.2 Converting to Machine-Usage Time

The process used to convert the clock times recorded for the downtime events into machine-usage times was easily accomplished on a spreadsheet. All times were compared to 12:00 midnight on January 17. The number of minutes which elapsed since this reference point were recorded in a column of the spreadsheet entitled *cumulative minutes*. For each downtime event after the first one, the durations of all previous downtime events were subtracted from the cumulative minutes in order to create a new column entitled *machine minutes*. Next, it was necessary to subtract all of the RBAM's non-scheduled time from the machine minutes.

Non-scheduled time at the plant consisted of many things. If the plant were not running due to a one-day holiday, there would be 24 hours of non-scheduled time. If scheduled preventive maintenance were performed, it would count as non-scheduled time. Shift meetings, machine test time, plant power failures, and full upstream buffers (recorded as 'no tubs') are several more examples of periods which would be defined as non-scheduled.

There were no entries in the database which indicated which periods of time were non-scheduled time. Fortunately, the appropriate six-months worth of hand-written time sheets were available in the operations office. On these time sheets, supervisors recorded the total amount of non-scheduled time for their eight-hour shift. The specific times were not recorded, only the shift totals. It was necessary to compare each shift total to the spreadsheet containing individual downtime events in order to determine where in the shift the non-scheduled time occurred. If, for example, the spreadsheet indicated 280 minutes between successive failures, and the time sheets indicated that there were four hours of non-scheduled time during that shift, it was easy to conclude that the 240 non-scheduled minutes were a part of the apparent 280 minutes between failures. All such easily identifiable periods of non-scheduled time were subtracted out and the actual machine minutes between failures estimates were improved.

Test periods were not always as easy to identify. This is because the RBAM machinery runs during test time, and operators still make some downtime entries. In all but a few cases, the reasons for downtimes did indicate that the machine was in test mode, and the period of test time could be identified and subtracted out of the data. In the few cases where there was ambiguity over when in the shift the actual non-scheduled time occurred, the appropriate length of time was subtracted from the period during the shift where there were the fewest recorded downtime events. There were only a few of these ambiguous situations in the entire period studied, so any inaccuracies in assigning an hour or two's worth of test time are certainly negligible.

After all of these subtractions were made, each downtime event had a time in actual machine minutes since 12:00 midnight January 17, 1994 associated with it. It was now very easy to determine times between failures (ttf) for any single machine, combination of machines, sub-system, or failure type by simply determining the differences in machine minutes between successive events of interest.

Duration of downtimes recorded in the database (ttr) did not require manipulation like the ttf data did. But both the ttf and ttr data still required smoothing, which was accomplished in the simulation software itself. Smoothing is explained in section 3.5.4.

3.5 Statistical Techniques Used

Several statistical tools were used in an attempt to simplify the distributions and make them more accurate. The four worth examining in this section are analysis of variance, the chi-square goodness-of-fit-test, least squares regression through the use of non-linear programming, and a simple type of smoothing which was applied to the distributions.

3.5.1 Analysis of Variance (ANOVA)

ANOVA was useful when deciding whether or not similar types of equipment had statistically significant differences in their mean times to fail and mean times to repair. The best way to describe ANOVA's usefulness is with a discussion of how it was used to develop distributions for the different unwind stands on the RBAM.

Unwind stands are the machines on which the rolled up materials from pre-assembly are loaded and slowly fed into the various stages of the battery assembly process. The RBAM has six unwind stands as follows:

- anode unwind (1)
- composite unwinds (3)
- cathode unwind (1)
- carrier web unwind (1)

There are two main reasons that it would be desirable to use common distributions for equipment which behaves similarly. The first is that it saves some work in the development of the distributions and in the subsequent creation of the computer simulation itself. The second is that distributions developed using more data are more likely to be representative of future machine behavior than those developed with less data. If, for example, all six unwind stands behaved similarly, a common distribution which covered all of them would be based on six times as much data as individual distributions would be. With this in mind, I attempted to see how similar the performance of the unwinds was by formulating a null hypothesis (H_0)and testing its validity using ANOVA. The null hypothesis was:

H_0 All unwind stands have the same mean time between failures.

In order to test H_0, times between failures for the six unwind stands were lined up in columns and an ANOVA was performed. The confidence level was chosen as 90%. The following table resulted:

Anova: Single Factor Unwind TTF

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
ANODE	92	143380	1558.48	2.79E+06
CARRIER	206	142784	693.13	5.26E+05
CATHODE	383	143789	375.43	1.82E+05
COMP 1	59	137625	2332.63	7.60E+06
COMP 2	71	142871	2012.27	3.65E+06
COMP 3	39	137565	3527.31	1.00E+07

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	6.24E+08	5	1.25E+08	69.83	3.74E-61	2.22
Within Groups	1.51E+09	844	1.79E+06			
Total	2.13E+09	849				

Table 1: ANOVA Results

An underlying assumption of ANOVA is that the data sets being compared have equal variance. Table 1 indicates that for these sets of data, the variance increases as the mean increases. Under these circumstances, it is often useful to transform the data so that the ANOVA assumptions are met. One useful method for determining what type of data transformation to use is to plot the natural log of the standard deviation as a function of the natural log of the mean for the data sets being compared. Then, determine the slope (β) of a line which fits the data fairly well. The data should be transformed by raising it to the 1-β power. If β=1, then the data should be transformed by taking the natural logarithm.

The following chart shows the the log(standard deviation) plotted against the log(mean) for the six sets of unwind ttf data, along with lines with slopes equal to 0.5, 1.0, and 1.5, which are included for purposes of comparison:

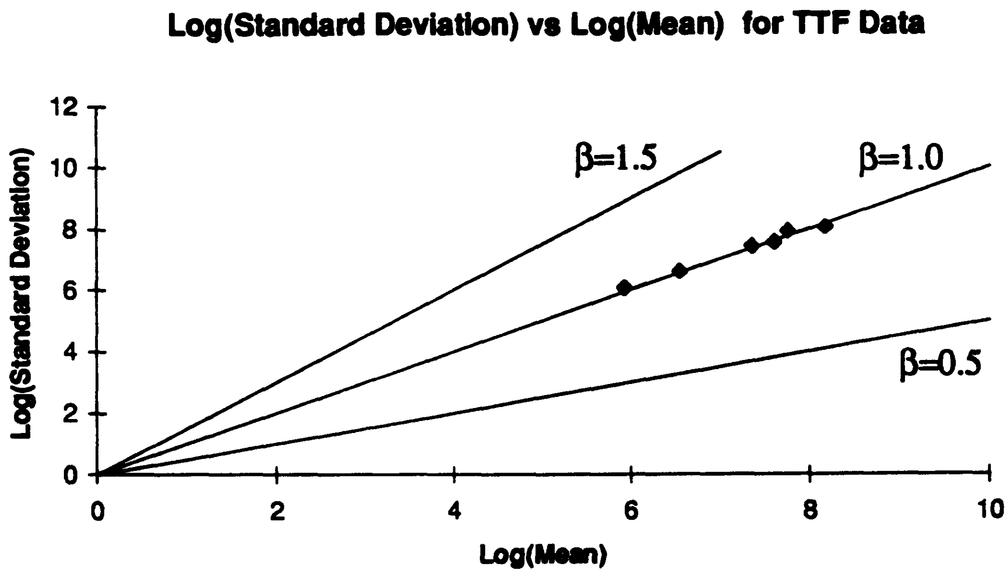

Figure 8: Log(Standard Deviation) vs. Log(Mean)

The β=1 line fits the data fairly well, so I took the natural log of the data and ran another ANOVA. The following table resulted:

Anova: Single Factor Log(TTF)

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
ANODE	92	614.64	6.68	1.98
CARRIER	205	1217.08	5.94	1.76
CATHODE	381	1992.61	5.23	1.88
COMP 1	59	422.04	7.15	1.58
COMP 2	71	500.09	7.04	1.82
COMP 3	38	289.18	7.61	1.67

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	521.76	5	104.35	57.22	2.90E-51	1.85
Within Groups	1532.00	840	1.82			
Total	2053.76	845				

Table 2: ANOVA Results

It is easy to see that based upon the 'F' statistic of 57.22 as compared to the critical 'F' value of 1.85, H_0 is clearly rejected at the 90% confidence level. This means that we can conclude that all of the unwinds do not have the same mean time to failure, and a single distribution for failure time would be inappropriate.

The results are not surprising. The unwind stands, while similar, have different materials loaded onto them. Different failure performance is expected.

The next step was to test whether or not the three composite unwinds were the same. The null hypothesis is as follows:

H_0 All composite unwind stands have the same mean time between failures.

The resulting ANOVA table follows:

Anova: Single Factor Log(TTF)

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
COMP 1	59	422.04	7.15	1.58
COMP 2	71	500.09	7.04	1.82
COMP 3	38	289.18	7.61	1.67

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	8.24	2	4.12	2.42	0.09	2.34
Within Groups	280.76	165	1.70			
Total	289.00	167				

Table 3: ANOVA Results

Since $F > F_{critical}$, we reject H_0. This means that the three composite unwinds apparently do not behave similarly. One further test compared composite unwind #1 to composite unwind #2. The following table resulted:

Anova: Single Factor

SUMMARY

Groups	Count	Sum	Average	Variance
COMP 1	59	422.04	7.15	1.58
COMP 2	71	500.09	7.04	1.82

ANOVA

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	0.39	1	0.39	0.23	0.64	2.75
Within Groups	219.10	128	1.71			
Total	219.49	129				

Table 4: ANOVA Results

This ANOVA tells us that at the 90% confidence level, we cannot reject H_0. A reasonable assumption is that composite unwinds 1 and 2 behave very similarly and have the same mean ttf. Further ANOVAs not documented here were performed comparing the various unwind stands to one another in varying combinations. The only case in which H_0 was not rejected was in the case of composite unwinds 1 and 2. Why composite unwind 3 was significantly different from 1 and 2 is unknown.

The IBAM will have four unwind stands as opposed to the RBAM's six. This is because only one composite unwind will be used in the IBAM instead of three. ANOVA made it clear that it was necessary to use separate ttf distributions for the anode, carrier, and cathode unwinds in the IBAM simulation. The distribution for the single IBAM composite unwind was developed using data from the RBAM's composite unwinds 1 and 2. This decision was based on the premise that it was best to be conservative in the analyses which will be described in subsequent chapters. It would be very undesirable to underestimate the effects of a problem because of overly optimistic assumptions. The performance of composite unwinds 1 and 2 was consistently worse than that of composite unwind 3. The conservative approach was to use failure data from 1 and 2 instead of from 3.

The ttr distributions for the unwind stands yielded interesting results. The following null hypothesis was tested using ANOVA:

H_0 All unwind stands have the same mean time to repair.

The resulting table follows:

Anova: Single Factor Unwinds TTR

SUMMARY

Groups	Count	Sum	Average	Variance
ANODE	92	588	6.39	47.21
COMP 1	59	452	7.66	486.16
COMP 2	71	331	4.66	65.20
COMP 3	39	182	4.67	22.39
CARRIER	206	1446	7.02	81.03
CATHODE	383	2217	5.79	45.33
	850			

ANOVA

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	588.69	5	117.74	1.38	0.23	2.22
Within Groups	71835.48	844	85.11			
Total	72424.17	849				

Table 5: ANOVA Results

In this case, $F<F_{critical}$, so we cannot reject H_0 at the 90% confidence level.

Based upon table 5, I developed a single ttr distribution for all of the unwind stands. But later reflection showed that as in the ttf case, the variances seem to increase as the mean increases. A plot of log(standard deviation) vs. log(mean) follows:

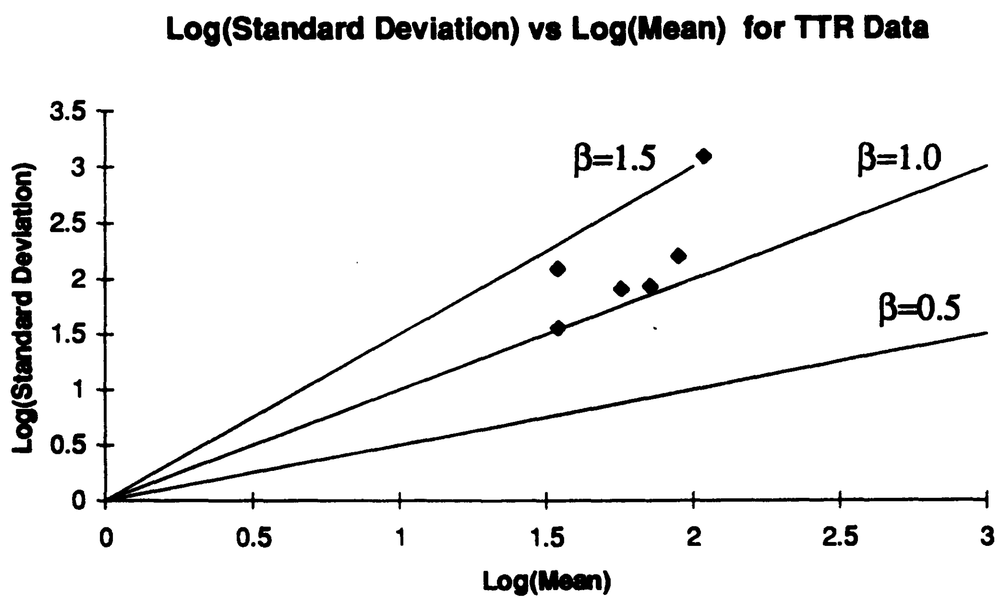

Figure 9: Log(Standard Deviation) vs. Log(Mean)

Although the point which corresponds to composite unwind stand #1 is an obvious outlier since its standard deviation is so large, a line with β=1 seems to fit the best, so I took the natural logarithm of the data and performed another ANOVA. The following table resulted:

Anova: Single Factor Log(TTR)

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
ANODE	92	149.18	1.62	0.37
CARRIER	206	338.09	1.64	0.49
CATHODE	383	602.61	1.57	0.29
COMP 1	58	83.30	1.44	0.61
COMP 2	71	89.37	1.26	0.32
COMP 3	39	50.94	1.31	0.36

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	11.48	5	2.30	6.12	1.39E-05	1.85
Within Groups	316.09	843	0.375			
Total	327.57	848				

Table 6: ANOVA Results

The transformed data adheres to the equal variance assumption much better than the raw data does. The ANOVA based upon the transformed data indicates that all unwind stands do not have the same mean ttr, since $F>F_{critical}$ and H_0 is therefore rejected. One more ANOVA is useful:

Anova: Single Factor Log(TTR)

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
ANODE	92	149.18	1.62	0.37
CARRIER	206	338.09	1.64	0.49
CATHODE	383	602.61	1.57	0.29

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	0.66	2	0.33	0.92	0.40	2.31
Within Groups	245.61	678	0.36			
Total	246.27	680				

Table 7: ANOVA Results

In this case, since $F<F_{critical}$, we cannot reject H_0. This means that the repair times of the anode, carrier, and cathode unwind stands are similar.

Tables 6 and 7 suggest that it would have been better to separate the composite unwind ttr data from the ttr data for the anode, carrier, and cathode unwind stands before developing ttr distributions. The ttr distribution which was developed for all of the unwind stands had a mean ttr of 6.14 minutes. The anode, carrier, and cathode unwind stands together have a mean ttr of 6.24 minutes, while the composite unwind stands have a mean ttr of 5.71 minutes. The effects that combining the unwinds' ttr data had on simulations which used the distributions were negligible. Total downtime for the unwind stands was still correct, but slightly more downtime was allocated to the composite unwind, and slightly less to the other three unwind stands.

3.5.2 Chi-Square Goodness-of-Fit Test

It was important to decide exactly how to represent the ttf and ttr distributions in the simulation itself. The modeling software offered a number of built-in distributions in which it was only necessary for the developer to identify the distributions' appropriate parameters. For example, the exponential distribution, written E(a), required a single parameter: the mean. The normal distribution, written N(a,b) required two parameters: a mean and a standard deviation. Distributions were assigned in the model's downtime editor. TTF was assigned in a column entitled *frequency*, while ttr was assigned in a column entitled *logic*.[3] In the following example, Machine A's failure and repair times are exponentially distributed with means of 284 minutes and 14.7 minutes, respectively. Machine B's failure times are normally distributed with a mean of 284 minutes and a standard deviation of 50 minutes, while its repair times are normally distributed with a mean of 14.7 minutes and a standard deviation of 3 minutes:

	Frequency	Logic
Machine A	E(284) min	E(14.7) min
Machine B	N(284,50) min	N(14.7,3) min

Table 8: Sample Programming Code

It is also possible to develop user-defined distributions if none of the built-in distributions fit adequately. In the following example, Machine A's ttf and ttr distributions are user-defined empirical distributions entitled machine_a_ttf and machine_a_ttr:

	Frequency	Logic
Machine A	machine_a_ttf() min	machine_a_ttr() min

Table 9: Sample Programming Code

Note that no parameters are needed in the parentheses, because the user must completely define the distributions in the software's table editor. These definitions will look something like the following:[3]

machine_a_ttf Percentage	Value	machine_a_ttr Percentage	Value
25	15	25	5
25	30	25	10
25	45	25	15
25	60	25	20

Table 10: Example Distribution

According to table 10, Machine A will go between 0 and 15 minutes between failures 25% of the time, between 15 and 30 minutes between failures another 25% of the time, and so on.

I wanted very much to be able to define the different machines' performance in terms of the convenient built-in distributions. This would allow the flexibility to define the distributions' various parameters with variables that could be easily modified from test to test. For example, suppose that Machine A's behavior was defined as follows:

	Frequency	Logic
Machine A	E(variable1) min	E(variable2) min

Table 11: Sample Programming Code

It would be extremely easy to test the effects that improved ttf or ttr performance of Machine A had on the overall system. In order to test these effects, it would only be necessary to change the values of variable1 and/or variable2 for different replications. The different values for variable1 and variable2 could be set up in an external file that the simulation received input from. Outputs from the different replications could then be compared easily, and the desired effects could be readily quantified.

With this in mind, I tried to determine if the machines' performance could be accurately defined by built-in distributions. I began by creating histograms of the ttf and ttr data for the different machines. All of the histograms which I created had exponential shapes to them. The following example is a histogram of the RBAM's cathode ttf data:

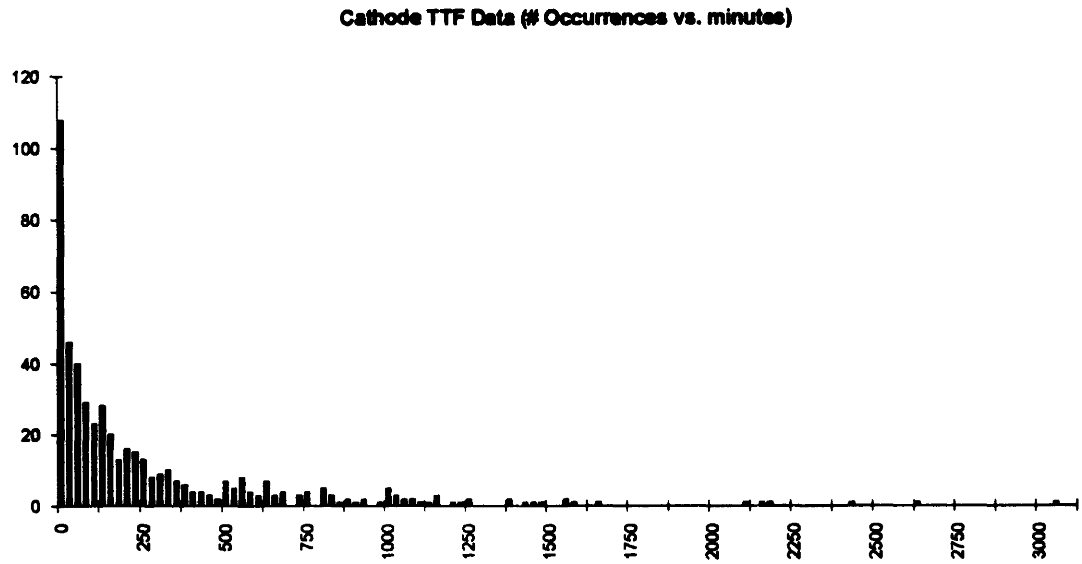

Figure 10: Cathode TTF Histogram

Since the distributions' shapes were exponential in nature, I attempted to define them in terms of the exponential distribution. The exponential probability density function (pdf) is given by:[1]

$$f(x) = \frac{1}{\beta} e^{-x/\beta} = \lambda e^{-\lambda x}$$

where:

$$\lambda = \frac{1}{\beta} = \text{failure rate of distribution}$$

In order to convert the histogram into a pdf, it was necessary to do a couple of things. First, the numbers of occurrences in each bin had to be converted into a percentage of the total occurrences. Then, the percentage associated with each bin had to be divided by the width of the bin. This was necessary to ensure that the area under the pdf curve would equal 1. Finally, the value of the (% of occurrences)/(bin width) was assigned to the midpoint of each bin. The resulting pdf is plotted below:

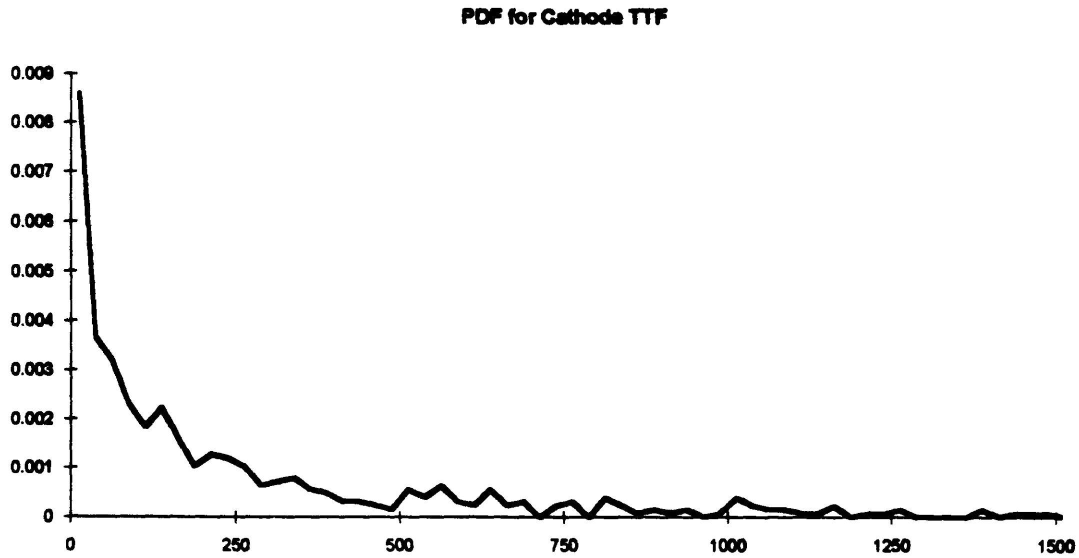

Figure 11: PDF for Cathode TTF

Note that the independent axis is only plotted out to 1500 minutes as opposed to 3000. This was done to better show the shape of the most important part of the curve.

The data used to create the above empirical pdf had a mean value of 284.7 minutes. In other words, the RBAM cathode machine has a mean time to failure of 284.7 minutes. For this reason, I first examined an exponential pdf with a mean of 284.7 minutes. The empirical pdf and the exponential pdf with a mean of 284.7 are plotted in figure 12:

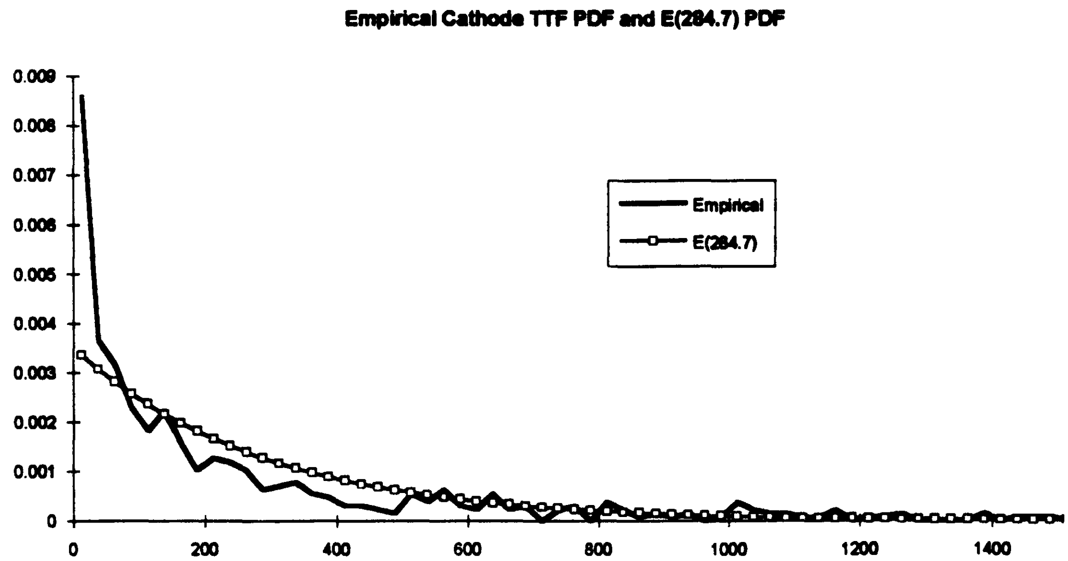

Figure 12: Distribution Comparison

I separated the data into 12 bins and performed the chi-square goodness-of-fit test. The test is used to see whether data fits a particular type of distribution. The test says that for n data points, if the data is broken into k mutually exclusive and exhaustive bins, the variable Y_i denotes the number of occurrences in each of the k bins, and p_i is the probability that the distribution that the data is being fitted to assigns to each of the k bins, then;

$$q_{k-1} = \sum_{i=1}^{k} \frac{(Y_i - np_i)^2}{np_i}$$

has an approximate chi-square distribution with k-1-h degrees of freedom, where h is the number of parameters estimated in order to define the distribution.[1] (In this case, h=1 since λ=1/284.7 was estimated from the observed data.) The null hypothesis H_0 that the data has the underlying distribution is questioned if the computed q_{k-1} is larger than:

$$\chi^2(\alpha; k\text{-}1\text{-}h)$$

The following table shows how the calculations were performed to see if the data passed the goodness-of-fit test when fitted to an exponential distribution with $1/\lambda=284.7$:

Bin	Expected Value (=np_i)	Observed Value (=Y_i)	$(Y_i-np_i)^2/np_i$	
100	113.51	115	0.02	
200	79.89	66	2.42	
300	56.23	45	2.24	
400	39.57	37	0.17	
500	27.85	20	2.21	
600	19.60	16	0.66	
700	13.80	16	0.35	
800	9.71	18	7.08	
900	6.83	10	1.47	
1000	4.81	8	2.12	
1200	5.77	11	4.75	
More	5.65	21	41.68	
			65.16	=sum
		Chi-test	23.21	
		α=.01		

Table 12 : Goodness-of-Fit-Calculations

Since q_{k-1} =65.16, and $\chi^2(.01;10)$ is only 23.21, H_0 is convincingly rejected, even at the α=.01 level. The exponential distribution with a mean ttf of 284.7 does not fit the data.

Verification tests which are described in section 3.6 also showed that the E(284.7) distribution did not fit the data adequately. From figure 12, it appears as if there are many more instances where the times between subsequent failures are on the order of 0 to 50 minutes than the E(284.7) distribution would indicate. Also, the E(284.7) curve is consistently above the empirical curve between about 100 and 500 minutes. It appears as if the outlying points in the 2000-3000 minute range (shown in figure 10) greatly increased the average value of the data, and an exponential distribution with the same mean value as that of the data is not quite right. It falls off too gradually in the 0-500 minute range.

3.5.3 Exponential Regression

I next attempted to fit an exponential curve to the data by using a least squares regression. I figured that a least squares regression would minimize the skewing effects of the data points representing extremely long times between failures. I set up a simple non-linear program to minimize the sum of the squares of the errors between the empirical curve and the fitted curve by changing the value of λ, the failure rate. The program structure is shown below:

			λ=	0.006965
			1/λ=	143.574
minutes (x)	empirical	$=\lambda e^{-\lambda x}$ (fitted)	square of error (fitted-empirical)^2	
12.5	0.008588	0.006384	4.86E-06	
37.5	0.003658	0.005364	2.91E-06	
62.5	0.003181	0.004507	1.76E-06	
•	•	•	•	
•	•	•	•	
3062.5	7.95 E-05	3.8 E-12	6.23E-09	
			1.68E-05	=sum of squared errors
Minimize the sum of the squared errors by changing λ				

Table 13: Non-linear Program

As shown in table 13, the value of 1/λ which provided the best fit curve was 143.6 minutes. A plot of all three curves is shown in figure 13:

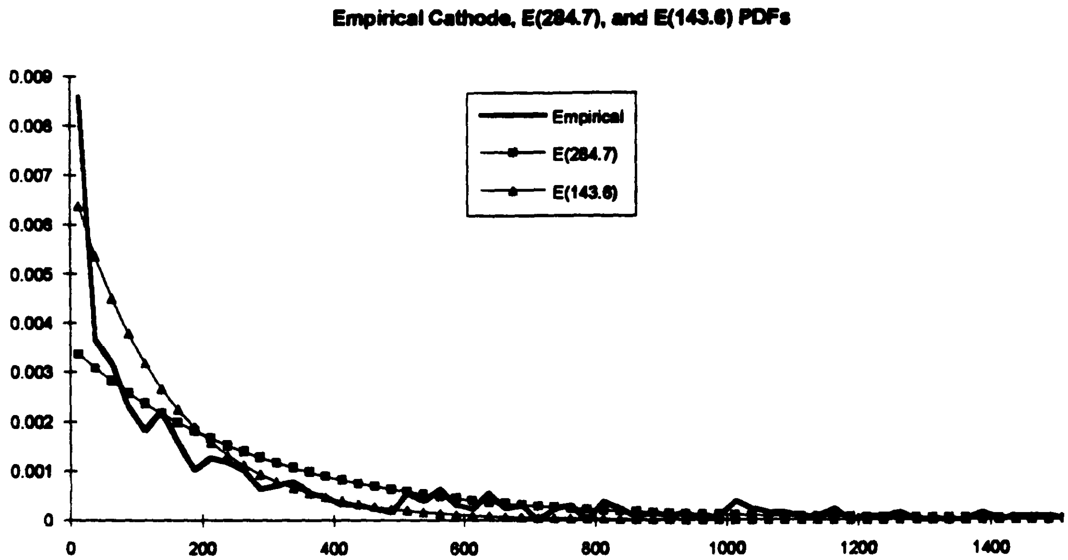

Figure 13: Distribution Comparison #2

The fitted curve did not accurately represent the machine behavior any better than the E(284.7) curve did. In fact, it was much worse. Although the calculations are not shown, a chi-square goodness-of-fit test using the same bins as before was performed to determine if the E(143.6) curve fit the data. The q_{k-1} value was 5,870, which is vastly greater than $\chi^2(.01;10)$, which is only 23.21. By minimizing the effects of the outlying points representing long times between failures, the distribution became too heavily weighted in favor of short times between failures. During validation tests, the machine failed too often when the E(143.6) distribution was used.

In order to best capture the true behavior of the cathode and all other machines, I decided to use empirical distributions rather than the convenient built-in distributions. As was described previously, a consequence of this decision was that it would not be as easy to test the effects of varying degrees of machine reliability from replication to replication of the computer simulation. By using a different approach, however, it was possible to partially offset this limitation. Rather than use variables to define the distributions' parameters, it seemed to work well enough to simply use the same distribution in every replication, and multiply the times returned by the sampling process by some constant. By changing the constant from replication to replication, the effects of different machine performances could be tested. An example of how this logic would look in the downtime editor follows:

	Frequency	Logic
Machine A	variable1*machine_a_ttf() min	variable2*machine_a_ttr() min

Table 14: Sample Programming Code

By using variables in this way, the ttf and ttr distributions would maintain their shapes from replication to replication, but their independent axes would expand or shrink according to the values of their corresponding variables. Despite the initial concern over the lack of flexibility that empirical distributions would introduce, this approach of sampling from the same distribution and then multiplying by a constant seemed to work quite well, and provided the required flexibility.

3.5.4 Smoothing

Smoothing of the data was accomplished as a natural consequence of the way in which empirical distributions had to be defined in the simulation software's table editor. Recall the example of the user-defined empirical distribution entitled machine_a_ttf given in the previous section. Consider the 2nd quartile of that distribution, which indicates that 25% of the time, the machine would go between 15 and 30 minutes between subsequent failures. The software assumes that values drawn from this quartile are uniformly distributed between 15 and 30 minutes. This assumption of uniformly distributed values smoothes out irregularities in the data between these points. Recall from section 3.3.1 that the operators sometimes had a tendency to round downtime durations to the nearest 5-minute interval. If the empirical distribution were developed so that the endpoints of one interval were 8 and 12 minutes, then the assumption of uniformity between the intervals' endpoints would smooth out the irregularities noticed in the data. The likelihood that the distributions would accurately represent machine behavior would increase. If certain apparent irregularities in the data are deemed to be important to the definition of the machine behavior, then it would be important to choose the endpoints of distribution intervals wisely so that the irregularities were not lost in the smoothing process.

3.6 Testing the Validity of the Empirical Distributions

After developing empirical ttf and ttr distributions for the various RBAM machinery, it was important to perform a reality check to see how well the simulation software interacted with the distributions and how closely the simulation results reflected the documented machine behavior. In order to accomplish this, a very simple simulation named test1 was developed in which there was a single machine which had an entity pass through it every 2.7 seconds, the overall rate at which the proposed IBAM would produce battery rows. The actual time chosen was unimportant, since the percentage of downtime caused by the distributions was the value of interest.

Initial ttf and ttr distributions were developed for the following RBAM areas:

- Cathode machine
- Unwind stands (all unwinds combined)
- Slurry delivery
- Vacuum sealer
- Carrier web

For each replication, the test1's single machine was programmed to fail according to one of the above machines' empirical ttf and ttr distributions. Ten replications of 200 hours duration were run on test1 for each of the above machines with the exception of the unwinds. They were tested over the weekend, so 20 replications of 300 hours each were run. Each 300 hours of simulation time took about an hour of actual time. Percentages of downtime were recorded for each replication and compared to expected percentages.

The expected downtime percentages were determined by comparing the total number of downtime minutes caused by each machine to the total number of minutes in which the RBAM actually ran during the period studied. For example, the cathode machine was the cause of 7,531 minutes of downtime during the period studied. During the same period, the RBAM actually ran for roughly 143,800 minutes. The expected percentage of downtime from test1 when sampling from the cathode's ttf and ttr distributions would then be (7,531)/(7,531+143,800), or 5.0%. The following table shows the results of the test runs. The 90% confidence intervals are based upon the student's T distribution since the number of replications was probably too low to warrant the use of the normal distribution:

Machine	Replications	Expected %	Mean of results	90% Conf Int
Cathode	10	5.0 %	5.77%	4.37%-7.16%
Unwinds	20	3.5%	3.41%	3.05%-3.77%
Slurry System	10	5.4%	5.70%	4.34%-7.07%
Vacuum Sealer	10	8.8%	7.33%	5.83%-8.84%
Carrier Web	10	12.4%	12.62%	11.2%-14.02%

Table 15 : Verification Test Results

The results were very encouraging. The mean values are all fairly close to the expected values. The furthest off is the vacuum sealer, but the expected result still lies within a 90% confidence interval of the mean of the replications. The real aim is to have all of the distributions, when taken together in a large computer simulation, accurately represent the overall manufacturing line behavior. The above results suggest that the methods used to quantify the machines' behavior profiles were sound and analysis based upon these methods should be reliable.

The ttf and ttr distributions for the cathode, the vacuum sealer, and the carrier web were ready for use in the computer models, but the unwinds and slurry system required further breakdown into more specific sub-system distributions to ensure that they were properly applied to the simulations. Complete listings of the distributions used in the various models are contained in Appendices B and C.

Chapter 4 - Development of the Computer Models

4.1 Initial Modeling Approach

The software used to develop the computer simulations was ProModel for Windows version 1.10. ProModel is a package best suited for the modeling of discrete part manufacturing systems, although continuous flow or batch processes such as chemical mixing operations can be modeled as long as the system can be described by discrete entities like gallons of fluid. The anode and composite cells lent themselves nicely to modeling in ProModel, since lots of discrete entities like battery components and pallets were the objects of interest. The web transport system and cathode machine required more creativity to model since they involved continuous webs of material. The vacuum sealer was simply treated as a part of the web transport system.

When the internship began, one model of the IBAM's anode cell which was developed in ProModel by a third party was available for me to study. This model, in conjunction with ProModel's documentation, offered enough insight into modeling of manufacturing systems for me to get started. My objective was to improve the model of the anode cell as necessary, and develop additional models of the composite cell, the cathode machine, and the web transport system. After these four modules were completed, I would then merge them together into a single model of the entire IBAM. My intention was to use this very detailed IBAM model as a platform to perform designed experiments and sensitivity analysis on so that the performance of the IBAM could be optimized.

This initial modeling task was quite substantial. I wanted to capture as fine a level of detail as I could. For example, I broke motions of the cambot heads down into distinct events which took only fractions of a second. I tried to make the graphics sufficiently detailed so that a user could tell the status of any pallet, canister, battery component, or machine in the system at any time. This level of detail took its toll on the modeling process in three ways. First, it turned the development of the model into a process which took about 8 weeks' worth of attention. Second, it slowed the performance of the model itself down. The model was developed and run on a personal computer with the following pertinent performance parameters:

- 486DX2/66 processor.
- 8 MB RAM
- 19 MB virtual memory swap file established on the hard drive.

Despite the highly capable configuration of the computer, the IBAM simulation would only run at about 1/3 the speed of real time. In other words, it would take the computer 3 seconds to simulate 1 second of machine run time. When the model was tested on a computer with a 90 MHz Pentium processor and 16 MB of RAM, the speed roughly doubled, but was still aggravatingly slow. With the graphics disabled so that machine motions were simulated but not displayed on the monitor, the simulation would run nearly 10 times as fast as real time. This is much faster than with the graphics enabled, but even at this speed an attempt to simulate 6 months worth of production would take 18 days.

The third way that the high level of detail adversely affected the modeling effort was by making it very difficult to keep pace with the rapidly changing nature of the IBAM project. Any time configuration changes occurred, the model would have to be changed, sometimes dramatically. I learned to make my work as modular as possible in order to protect myself against changes in the machinery configurations, but these changes were always setbacks despite my efforts to modularize. The changes also underscored the need to have an actual team member doing the modeling rather than trying to contract the work out to third parties.

In this thesis, the data analysis was introduced before the modeling was discussed. During the internship itself, however, the modeling effort was begun well before the data analysis was. By the half-way point of the internship, the detailed IBAM model was 99% complete. All that remained in order to make the model functionally complete was some debugging of a problem that surrounded the interface between the web transport system and the anode cell. It was at this point, just after the midstream internship review, that I began the data analysis described in detail in chapter 3.

In retrospect, I believe that in general, it would be more valuable to perform a detailed statistical analysis and develop machine behavior profiles before beginning a substantial modeling effort. Since I knew nothing about the process of manufacturing batteries when I began the internship, the modeling effort was an excellent way for me to learn about the process and gain exposure to important issues surrounding production.

However, if a modeler already knows a good deal about the process that he intends to model, it would be more beneficial to perform some statistical analysis first. This is because the statistical analysis is of great benefit by itself. It is very likely that insight gained from the statistical analysis will help to streamline the modeling process by ensuring that it is focused on answering important questions. A statistical analysis which is not followed by a modeling effort will very likely be of great benefit to a project. A model that is not based upon a sound statistical analysis is not very likely to be of great value and will probably only help to answer the simplest of questions.

My initial perspective was that the statistical analysis would help me develop a better model. Now that the internship is complete, my perspective has changed. I now believe that modeling is an excellent way to supplement a good statistical analysis and to further explore questions and insights which developed from the statistical analysis. My approach of developing a detailed model, supplementing it with a statistical analysis, and then seeing what I could learn was not the optimum approach. The following section explains how the statistical analysis greatly narrowed the focus of the modeling efforts and pointed them in an extremely productive direction.

4.2 Discovery of the Dry Slurry Yield Loss Issue

Following the development of the machine ttf and ttr profiles described in chapter 3, I began to wonder if yield loss due to dry slurry would be a substantial problem on the IBAM. On the RBAM, it takes about 1 minute from the time that the first slurry patch is applied until the battery strip passes through the final cutoff mechanism. From there, the battery strip usually waits a little over 30 seconds until it is vacuum sealed. Due to the extra length and slower operating speed of the IBAM as it is shown in figure 7, in steady state, the first slurry patch would be exposed for about 1 minute before it was even placed onto the carrier web. Then, it would take about 7 minutes for the battery row to make its way down to the vacuum sealer. All the while, of course, the slurry would be exposed and subject to dry-out. (Recall that the plant's policy is to discard any batteries whose slurry has been exposed for 15 minutes or more.) The following plot shows intuitively why one might expect the dry slurry yield loss issue to be a pressing one on the IBAM:

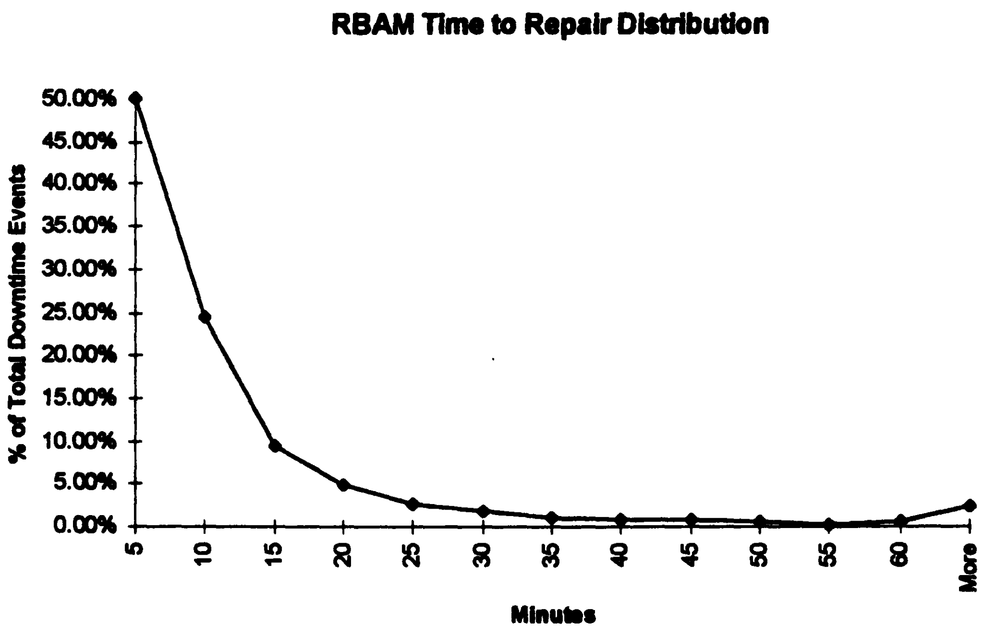

Figure 14: RBAM TTR

On the RBAM, all downtime occurrences in excess of 15 minutes cause the line's shift registers to fill with a signal indicating that the contents of the web must be discarded. Intuitively, if the area under the above curve represents all downtime events, then the area under the curve to the right of the 15 minute mark on the x axis would represent the portion of downtime events which lead to dry slurry yield loss. On the IBAM, the area under the curve to the right of the 8 minute mark on the x axis would represent the portion of downtime events which lead to dry slurry yield loss. The 8 minute mark is the relevant mark for the IBAM since it takes 7 minutes for a row to get from one end of the machine to the other, leaving only 8 minutes of slack time. Note that the current policy of discarding material on the RBAM following downtimes in excess of 15 minutes really amounts to having a 16 minute threshold, since it takes about 1 minute for batteries on the RBAM to travel form the point of initial slurry application to final cut-off. Despite this small inconsistency, I assumed that on the IBAM, a 15-minute threshold would mean that the total exposure time could not exceed 15 minutes.

With the specific intention of testing the impact of the 15-minute slurry threshold on IBAM production, I developed a new simulation of the IBAM. Unlike the detailed IBAM model which had taken me weeks to develop, this new model was much simpler and took only a day and a half to develop. The only graphics in the new model were a few counters which showed the number of battery strips in a few different places throughout the line. In discussions, I referred to this model as my block diagram model of the IBAM. In my initial tests, I used the machine profiles of the RBAM to describe the ttf and ttr performance of the corresponding IBAM machines. The assumption was that the IBAM cathode machine would fail according to the same distributions as the RBAM cathode machine, and so on. The simulation time-stamped the application of the first slurry patch, and then tabulated the total number of strips over a period of time which were sealed after 15 minutes had passed. This number was compared to the number of strips which were sealed before 15 minutes had passed, and an estimated dry slurry yield loss was calculated. My estimate based on this first set of assumptions was that IBAM dry slurry yield loss would average about 8% over time.

Two items are worthy of note at this point. The first is that since the block-diagram simulation was much simpler than the detailed simulation, it ran at a rate about 100 times faster than real time. This enabled me to start a simulation at about 4:00 P.M., run it for 15 hours or so overnight, and have the results waiting when I arrived at 7:30 A.M. the next morning. By doing this, I could readily simulate 1500 hours, or roughly two months worth of production each night.

The second noteworthy item is that the simulation did not *reveal* the problem. It was very useful in testing the effects of the problem, but it was the statistical analysis and the distribution shapes which provoked the thoughts which led to the analysis of the dry slurry issue.

I informed the IBAM program manager of my preliminary findings, and we discussed the impact that the dry slurry issue would have on the project.

4.3 Establishment of Need to Reduce Exposure Time

4.3.1 Refinement of Assumptions

Recall from chapter 1 that the objectives of the IBAM are the following:

- A step increase in annual battery manufacturing capacity.
- A step increase in yield.
- Continual improvement in customer satisfaction by the elimination of manufacturing defects.

The program manager was clearly concerned that the fundamental objectives of the program would not be met if the dry slurry issue was not addressed. He first requested that I refine my assumptions a bit and perform further analyses. The ttf and ttr distributions used in the block diagram simulations were based upon data that was manually entered into the system by the RBAM operators. While the operators were generally faithful to their responsibility to enter the downtime data, sometimes the entries were not made. This meant that the machines actually failed at a rate that was somewhat higher than the ttf distributions indicated. In fact, the initial simulation runs indicated that overall IBAM run time was on the order of 75% of scheduled time. The RBAM run time during the period that the distributions were based upon was actually about 67% of scheduled time. Only two differences should have existed between the RBAM run time performance and the simulated IBAM run time performance. The first was due to the fact that the IBAM had only three baby hoppers in its slurry delivery system instead of four. The second was that the IBAM had fewer unwind stands, and that there were buffers between the unwind stands and the rest of the line. Since neither the baby hoppers nor the anode and composite unwind stands were major sources of downtime, these differences should have been minimal, and certainly should not have increased expected run time from 67% to 75%. The difference was mostly due to the bias in the ttf distributions that resulted from operators sometimes neglecting to make the proper entries.

The program manager wanted to know what dry slurry yield loss would result if the IBAM only achieved 67% run time. This gave rise to the question of how the distributions should be modified in order to have 67% run time from the simulation. We decided that the most reasonable and workable assumption to make was that the downtime events which were not entered were evenly distributed among all machines for which distributions were developed. We had no reason to believe otherwise. This assumption was implemented by using the multiplicative factor described in section 3.5.3 on all of the ttf distributions. By trial and error, I determined that a factor of 0.7 was appropriate. In effect, the simulation would determine, via sampling from the appropriate ttf distribution, a time until the next failure for each machine, and then multiply the time by 0.7 in order to make it fail sooner. This compensated for the entries that the operators did not make. The simulation consistently averaged 67% or 68% run time when the factor of 0.7 was used. The estimate of dry slurry yield loss increased to about 10.4% based upon these assumptions, and went all the way up to 12.3% in the end when 1 final refinement was made.

4.3.2 Initial Ideas

My initial feeling was that the length of the carrier web would have to be shortened. The design of the IBAM at this point had a carrier web with 154 1x4 battery rows in process with exposed slurry. At 2.7 seconds per battery row, this equates to about 7 minutes of work in process (WIP) with exposed slurry. Cutting rows out of the carrier web would reduce exposure time and improve the dry slurry yield loss issue.

The IBAM program manager had many years experience at the plant, and wondered whether the 15-minute slurry threshold was really a valid threshold, or if perhaps using a higher threshold would still allow for batteries of sufficiently high quality. The threshold of 15 minutes had been established in tests conducted in the early 1980s, and had been the plant's policy ever since. But none of the people who were involved in the test were still employed at the plant and no adequate documentation of the insight gained from the tests could be found. Since the dry slurry issue would be very prominent on the IBAM, upper management decided to completely revisit it and begin a thorough round of new tests to determine if the slurry threshold could be increased, and to develop a better understanding of the issue in general. The tests would take in excess of 270 days, since 270-day incubation periods were sometimes necessary for minor defects to fully manifest themselves. Management would be unwilling to make policy changes until the 270-day test results were thoroughly analyzed.

At the time the relevant decisions were being made on the IBAM, and later when this thesis was written, the slurry threshold tests were incomplete and whether or not the slurry threshold could be increased remained unknown.

Simulation was the perfect tool for determining the effects of various carrier web lengths and possible slurry threshold levels. With all of these possibilities in mind, the IBAM program manager wanted to see a family of curves that showed dry slurry yield loss on the dependent axis and slurry threshold time on the independent axis, with a curve for each possible length of carrier web. Then, he wished to bring the entire IBAM team into the process of solving the problem. The following set of curves was developed and presented to the IBAM team:

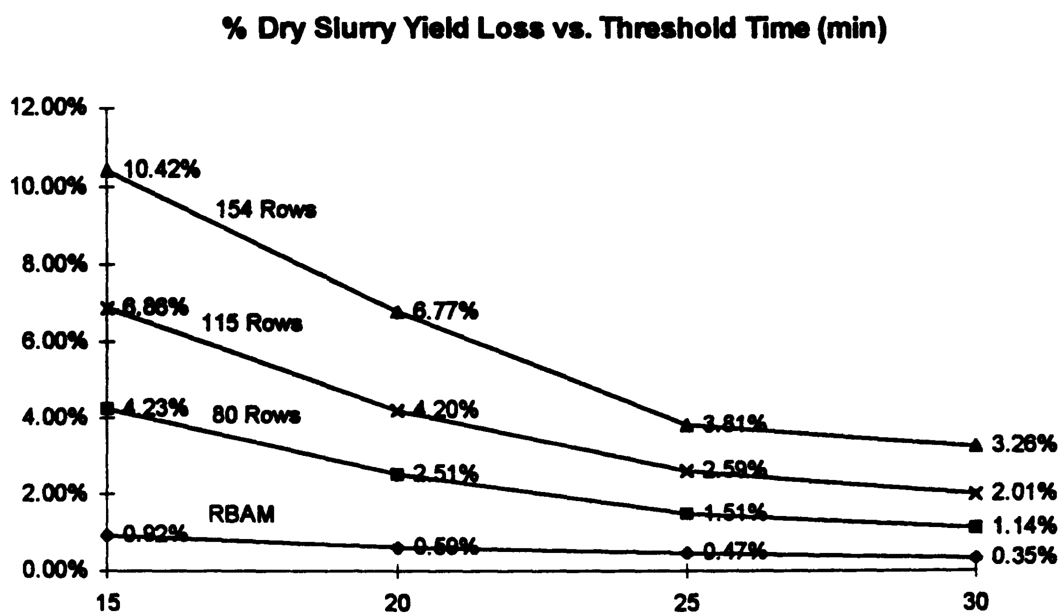

Figure 15: Initial Simulations

The then-current IBAM design had 154 rows, and was considered the baseline. The team engineer with primary responsibility for the web transport system and all associated CAD drawings estimated that the web length could be reduced to 115 rows with moderate layout changes such as taking the cathode unwind stand off-line. More radical changes such as moving the composite cell to the opposite side of the carrier web could further reduce web length to about 80 rows. Curves are plotted above for the baseline 154 row web, the 115 row web resulting from moderate changes, and the 80 row web resulting from radical changes. Also plotted is my estimate of the dry slurry yield loss that the RBAM would currently experience at each slurry threshold level. Even the 15-minute point, which should represent the current mode of RBAM production, is a simulation-based estimate. Overall yield loss is currently tracked on the RBAM, but individual types of yield loss are not. By comparison, the dry slurry yield loss is very low on the RBAM. This is because the RBAM runs at about six times the speed of the IBAM, so the slurry is exposed for a much shorter time.

Chapter 5 - The IBAM Team's Solution of the Dry Slurry Problem

5.1 Chapter Overview

This chapter explains the decision-making process that the IBAM team used in order to arrive at a proper course of action in response to the dry slurry yield loss problem. The first two sections explain the information that the team felt it should have in order to weigh all options adequately. The remaining sections examine the results of further simulations which were run in support of the team's efforts, the implications of these results, and the decision-making methodology used to arrive at a final solution.

5.2 Initial Input from the IBAM Team

At the time that the entire IBAM team was introduced to the dry slurry yield loss problem, ideas for solutions were limited. Only the IBAM program manager and I had given the problem any thought, and we were primarily considering three options: either shortening the carrier web, increasing the slurry threshold, or both. We also considered the possibility of speeding up the machinery. Although speeding up the machine would decrease slurry exposure time and actually increase throughput, it was clear that this approach alone would not be sufficient to bring dry slurry yield loss down to an acceptable level. The speed of the IBAM might be able to increase by 25% or so, but it could not triple or quadruple. Other measures would be required.

The creativity of the IBAM team members was noteworthy. Possibilities ranging from re-moistening the slurry with a mist of water to inverting the battery stack so that the slurry was applied to less absorbent material were considered. Each idea was evaluated on its merits, but in the end, fundamental problems with each of these ideas precluded them from receiving serious consideration for implementation. The team decided that some combination of the following three measures would be the answer to the problem:

- Increase the speed of the IBAM.
- Reduce the length of the carrier web.
- Increase the slurry threshold level.

Since tests to determine if the slurry threshold could be increased beyond 15 minutes were already being conducted by appropriate personnel, the IBAM team focused on determining ways to increase the speed and reduce the web length. The team felt that more refined information than was available from the first set of simulations would be required. The analyses that the team requested me to perform are described in the next section.

5.3 Refinement of Assumptions and Scenarios

Most of the creative energy of the IBAM team was spent figuring out ways in which the carrier web length could be reduced. Fortunately, the web transport system was not yet constructed, and was conceptually less mature than the other major subsystems. This meant that changes to the web transport would be less expensive than changes to other more fully developed areas. The team gave serious consideration to roughly a dozen new IBAM configurations. The concepts considered had web lengths ranging from 154 rows all the way down to 70 rows. Each row took 2.7 seconds to process, so on a 70 row web it would take just over 3 minutes to get from the point where the anode layer was applied to the web to the vacuum sealer. In steady state, the slurry patch on the anode row would be about 1 minute old by the time it was applied to the web, so the total exposure time with a 70 row web would be roughly 4 minutes. A 154 row web would, on the other hand, take about 8 minutes in steady state. The different web lengths were for the most part a function of whether or not the anode cell and composite cell were on the same side of the carrier web, and of how much machinery was placed off-line. Recall that in the context of the IBAM, a machine is considered off-line if it is not physically located above or below the carrier web. Being off-line does not mean that the machine is decoupled from the line.

The dozen or so concepts that the team developed fit into several categories with regard to web length. The baseline was 154 rows, several were around 120 rows, several others were around 100 rows, and the shortest was 70 rows. The team felt that by giving up rows, operator accessibility and machine maintainability may suffer because the machinery would be located closer together.

There were also questions regarding the expandability of some of the options. One firm criterion for any concept was that, if implemented, it must allow sufficient factory floor space for construction of another identical line should expansion become necessary in the future. The options with the shortest web lengths were not necessarily the easiest to expand. This is because if an option was shorter, it was probably wider, and the factory floor space dedicated to the IBAM and its possible expansion was rectangular, not square.

In order to evaluate the different options thoroughly, the team wanted to know how big a difference in estimated dry slurry yield loss there was between say, a 100-row option and a 70-row option. The accessibility issues might very well outweigh the yield loss improvements if the yield loss improvements were small. For this reason, they asked me to run simulations on web lengths of 100 rows and on 70 rows in addition to the ones which were already complete. The 120-row options were close enough to the 115-row scenario already tested that further tests of this web length were deemed unnecessary given the time constraints we were facing.

5.3.1 Best Case/Worst Case Approach

The team also felt that it was very important to see analysis which was based on an alternative set of assumptions. The assumptions that the first simulations were based on were as follows:

- All major IBAM systems will fail according to the same ttf and ttr distributions as the corresponding major RBAM systems.
- Since the data on the VAX computer is based upon manually entered data, all ttf distributions will be adjusted consistently so that overall simulated line run time is about 67%. (A multiplicative factor of 0.7 was determined through trial and error.)

- All areas of new technology on the IBAM for which there is no data available will operate at 100% efficiency.

The most questionable of these assumptions is that new technologies will operate at 100% efficiency, but the rationale for this will be explained later. The above set of assumptions was considered a worst case scenario by the IBAM team. It was considered worst case because the team collectively believed that they had designed a line which would be much more reliable than the current RBAM. Nobody believed that the IBAM would be less reliable overall once the plant had proceeded down the learning curve which would certainly exist.

In order to incorporate into the simulations the belief that the IBAM would be more reliable than the RBAM, the team analyzed the RBAM's major sources of downtime events and predicted how the IBAM would be different with regard to each type of failure. The following spreadsheet shows the top 70% of RBAM failures and the team's collective estimate of how the IBAM performance would change in each case. The total downtime for the period studied was 55,641 minutes. Major RBAM systems and their corresponding total numbers are in bold type. Individual failure types are in regular type. The columns should be interpreted as follows:

Avg. TTF - The average time between failures in minutes.
Avg. TTR - Avg. time to repair, or, avg. duration of a downtime event, in minutes.
Total Mins - Total down minutes for that system or failure type in the period studied.
% of Total - Total minutes for that system or failure type/55,641 minutes.

55641			Avg. TTF	Avg. TTR	Total Mins	% of Total	Modification
CATHODE			**284.70**	**14.71**	**7531**	**13.53%**	
Cathode	Creaser	Web Break	2038.94	5.77	404	0.73%	Reduce by 67%
	Drum	Web Break	1347.08	6.17	654	1.18%	Reduce by 50%
	Drum	Web Jump Pins	2098.74	15.23	990	1.78%	Reduce by 50%
	Laminator	Lamination	4266.58	30.64	1011	1.82%	Reduce by 75%
	Pin Belts	Web Jump Pins	1714.23	13.07	1072	1.93%	Reduce by 25%
	Pocketer	Drive Fault	4600.52	8.37	226	0.41%	Reduce by 75%
	Pocketer	Web Break	4161.35	4.39	136	0.24%	Reduce by 75%
					4493		
			Avg. TTF	Avg. TTR	Total Mins	% of Total	
SEALER			**125.88**	**11.79**	**13871**	**24.93%**	
Sealer	FCO	Jams	2558.82	6.49	357	0.64%	Reduce by 25%
		Knife not Cutting	5028.39	9.50	266	0.48%	No Reduction
		Poor Cut	6598.55	20.70	414	0.74%	No Reduction
	Pick&Place	Not completing cycle	815.86	8.68	1527	2.74%	Reduce by 100%
	Positioner	Seal to Carr	2933.88	16.65	799	1.44%	Reduce by 100%

	Table	ALL	285.55	9.47	4761	8.56%	Reduce by 100%
	VAC Sealer	Chamber not coming up	8431.94	22.44	359	0.65%	No Reduction
		Jams	3835.53	4.89	176	0.32%	Reduce by 67%
		Losing/No Vacuum	6210.30	11.83	272	0.49%	Reduce by 50%
		Open Seals	4271.55	18.09	597	1.07%	No Reduction
		Unknown	6319.43	29.19	613	1.10%	No Reduction
		Water Leak	8170.76	24.24	412	0.74%	Reduce by 100%
		Weak Bonds	5408.54	17.96	467	0.84%	No Reduction
					11020		
			Avg. TTF	Avg. TTR	Total Mins	% of Total	
SLURRY			**299.37**	**16.25**	**8155**	**14.66%**	
Slurry	Applicators	Dirty Tray	6828.65	3.94	303	0.54%	No Reduction
		Low Weights	6536.14	20.73	1658	2.98%	No Reduction
		Poor Spread	11238.77	7.40	355	0.64%	Reduce by 33%
		Short Patch	11281.96	22.15	1041	1.87%	No Reduction
		Slurry in Seal	15566.67	5.73	172	0.31%	No Reduction
		Slurry Leak	17754.22	6.26	169	0.30%	No Reduction
	Hoppers	Mother & Babies	1960.79	30.75	2245	4.03%	No Reduction
					5943		
			Avg. TTF	Avg. TTR	Total Mins	% of Total	
UNWINDS			**171.02**	**6.12**	**5236**	**9.41%**	
Unwind	Anode	Web Break	2022.17	4.70	334	0.60%	Reduce by 60%
	Carrier	Jumped Pins	3898.64	7.39	266	0.48%	Reduce by 80%
		Web Break	1044.38	5.35	733	1.32%	Reduce by 80%
	Cathode	Festoon not Holding	3726.33	6.83	123	0.22%	Reduce by 100%
		Web Break	445.31	5.14	1661	2.99%	Reduce by 60%
	Comp 1-3	Web Break	3114.46	3.33	450	0.81%	Reduce by 60%
					3567		
			Avg. TTF	Avg. TTR	Total Mins	% of Total	
WEB			**107.73**	**14.67**	**20400**	**36.66%**	
Web	Anode	Lamination	11816.17	10.58	127	0.23%	Reduce by 50%
		Mistracking	4358.00	12.12	400	0.72%	Reduce by 70%
		Registration	3600.69	14.28	514	0.92%	Reduce by 70%
		Web Break/Jump Pins	5713.48	15.16	379	0.68%	Reduce by 70%
	Carrier	Mistracking	5059.71	24.93	698	1.25%	Reduce by 40%
		Splice Mistracked	7504.11	9.26	176	0.32%	Reduce by 60%
		Square Hole NP	3673.66	17.59	686	1.23%	No Reduction
		Web Break	5700.92	13.80	345	0.62%	Reduce by 40%
		Web Jump Pins	856.01	11.66	1924	3.46%	Reduce by 40%
	Comp 1-3	Mistracking	5640.80	6.96	195	0.35%	Reduce by 70%
		Poor Cut	6019.03	8.50	578	1.04%	Reduce by 100%
		Registration	4873.92	18.80	1617	2.91%	Reduce by 70%
		Skewed	13973.29	5.71	120	0.22%	Reduce by 40%
	General	PCM Off-line	4229.26	40.06	1362	2.45%	Reduce by 60%
		Unknown	5840.91	17.77	391	0.70%	No Reduction
		Web Jump Pins	1022.47	14.50	2030	3.65%	Reduce by 40%
	Surge	Out of Sequence	3318.39	6.44	264	0.47%	Reduce by 100%
	Tamp 1-3	Web Jump Pins	3399.09	12.39	1946	3.50%	Reduce by 40%
					13752		

Table 16: Detailed Failure Information

In order to incorporate these reductions into the simulations, I developed new ttf distributions. I calculated the time between consecutive failures of each type listed above. Then, I multiplied each time between consecutive failures by a factor equal to:

$$\frac{1}{1-\%\text{Reduction}}$$

For example, if a failure type had a 25% reduction, I multiplied all times between failures by 1/(1-0.25), or, 1.333. All failures were then sequentially listed and assigned a new cumulative number of machine minutes which would have passed since the beginning of the test period had the reductions actually taken place. Failures which went beyond the end of the test period after the times had been increased were eliminated. If a failure type was reduced by 25%, on the average, 25% of the events should have fallen out. I then developed new user-defined distributions in the same way that the old ones had been developed, and maintained the multiplicative factor of 0.7 for the sake of consistency and ease of comparison to the worst case simulations. This gave the following set of assumptions for another set of simulation runs:

- All major IBAM systems will fail according to the distributions which incorporate improvements expected by the IBAM team.
- For the sake of consistency, all ttf distributions will be adjusted by the same multiplicative factor of 0.7 that was used in the worst-case simulations.
- All areas of new technology on the IBAM for which there is no data available will operate at 100% efficiency.

Note that the first two assumptions changed from the worst case scenario. The simulations based upon this new set of assumptions were considered best case scenarios by the IBAM team.

The assumption that all new technologies not present on the RBAM will perform at 100% efficiency warrants explanation. The IBAM team certainly did not believe that new technologies such as the cambot pick and place heads would be perfect. We gave consideration to trying to estimate the performance of all new technologies by doing benchmarking around the corporation with other plants which might be using similar technologies. Given more time, this might have been the optimum approach. But the time pressure associated with our decision-making precluded us from pursuing such a time-consuming strategy.

We considered using manufacturers estimates of reliability of the machines, but determined that manufacturers estimates were so overly optimistic that they were of little value. The team firmly believed that the overall IBAM run time performance would exceed that of the RBAM. The first set of assumptions which forced the simulation to have run time equal to about 67% of scheduled time, the RBAM's performance for the period studied, seemed to the team to represent a worst-case scenario. The thinking was that at worst, the IBAM would have run time performance as good as the RBAM. When it came to incorporating all of the expected improvements into the distributions for the best case scenario, the team members realized that the true performance of the IBAM would most likely be worse than the best case estimate, because new technologies would introduce some additional downtime which had not been accounted for. We figured that the true run time performance of the IBAM would probably lie between the worst case (lower bound) and best case (upper bound) which we had established.

Remember that the purpose behind the simulations was to develop insight into the effects that different conditions had on dry slurry yield loss. The assumptions that the IBAM team made were perfectly adequate for the purposes of developing this insight and using it to assist the decision-making process. The future may show that the IBAM performs better than our best case estimates or worse than our worst case estimates. But this is irrelevant. The assumptions and resulting analysis were adequate to the task given the pressures of program cost and schedule.

The results of the simulations run in support of the IBAM team's efforts to solve the dry slurry yield loss problem are shown in the following several charts. These are the charts that the team used in their decision-making efforts. The first 2 charts show worst case and best case curves for 154; 100; and 70-row webs. The estimated current RBAM performance is also plotted for reference. Note that the scales of the dependent axes vary.

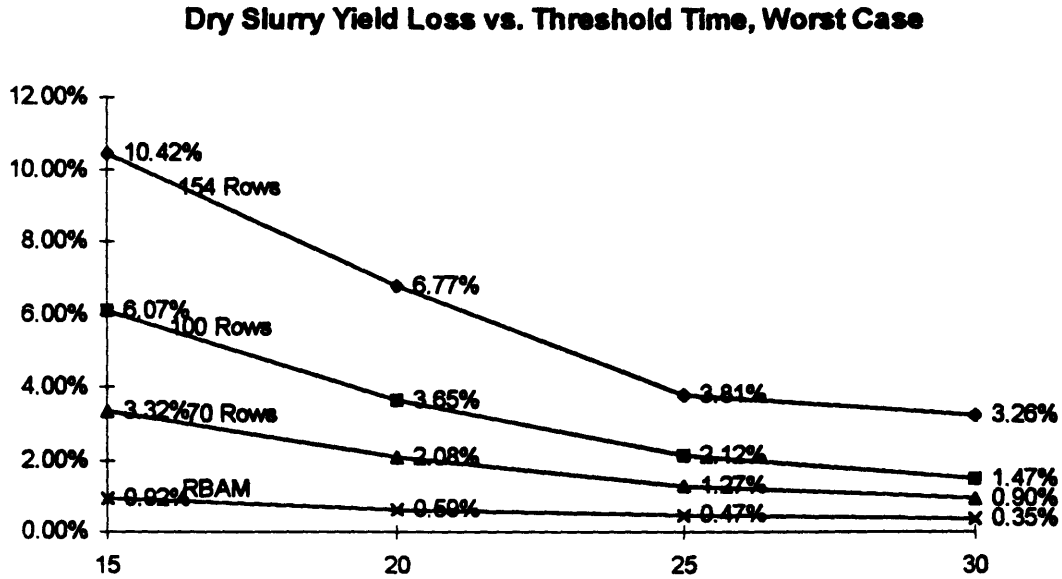

Figure 16: Worst Case

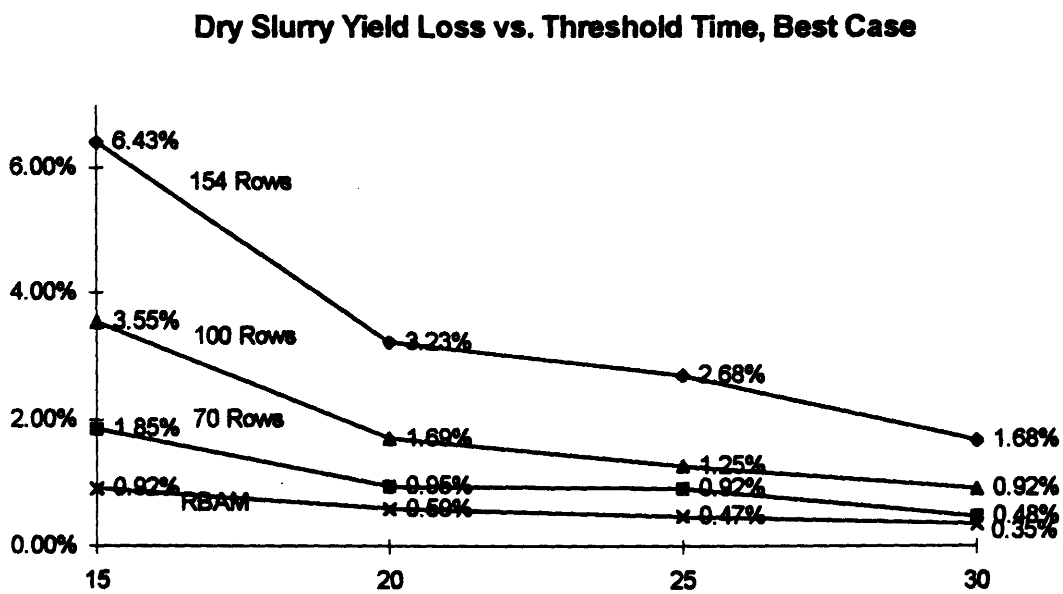

Figure 17: Best Case

The next 3 charts show the worst case, the best case, and the arithmetic average of the worst and best case for the different web lengths individually.

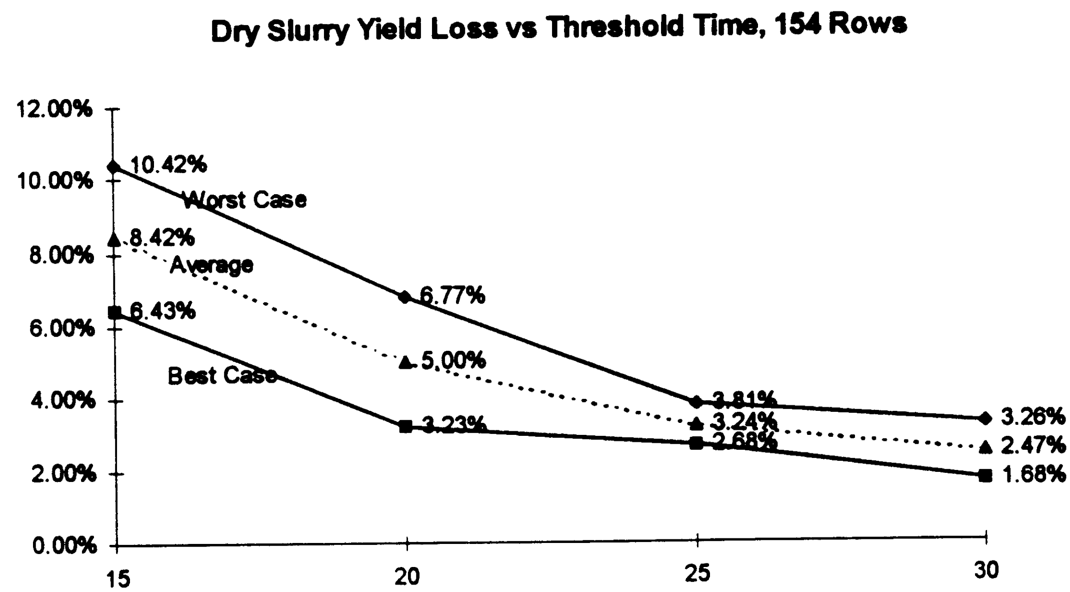

Figure 18: 154 Rows

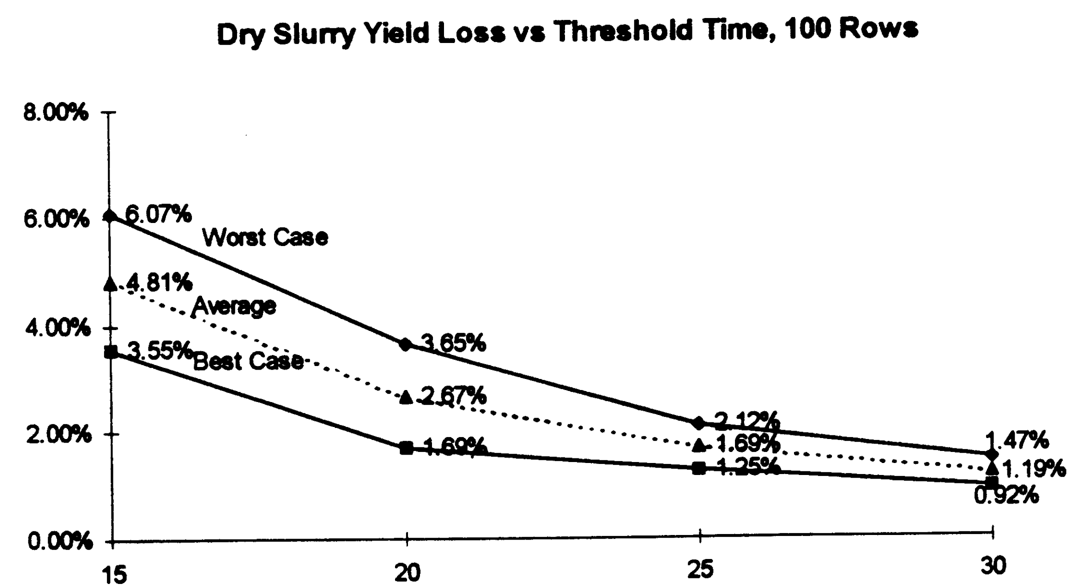

Figure 19: 100 Rows

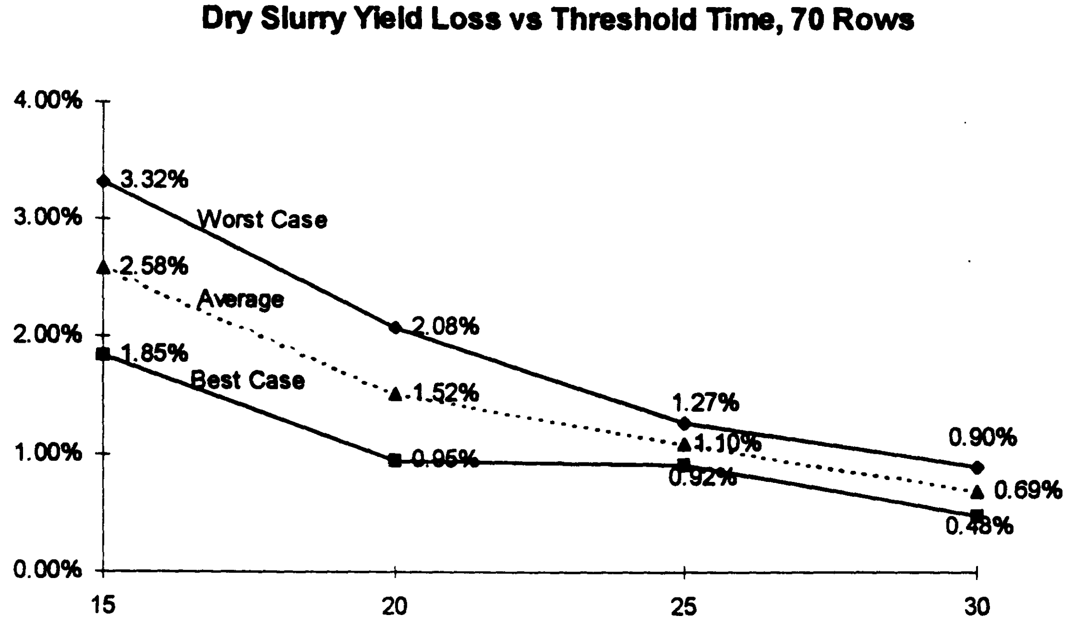

Figure 20: 70 Rows

Each point was obtained by averaging the yield losses obtained from five replications with 100 hours of simulated production time per replication. Some apparent irregularities exist in the shapes of some of the curves, but this can be attributed to the statistical variation from replication to replication.

Since each point was only based upon five replications, there was some concern about the statistical significance of the differences between points which were close to each other. For example, consider the points at the 25 minute threshold for the 100 row and 70 row curves from figure 16. The 90% confidence intervals derived from the student's t-distribution are [0.80%, 1.74%] for the 70 row, 25 minute threshold point, and [1.46%, 2.78%] for the 100 row, 25 minute threshold point. This is somewhat troublesome, since the confidence intervals overlap. We knew that there would be some measurable difference between these points, but we were not confident in the magnitude of the difference. Although it would have been ideal to use more replications per point, time pressure made this approach impractical.

After the internship was complete, further simulations were run in order to narrow the confidence intervals around each data point. Plots of curves based upon 20 replications per point are included in Appendix A. The 90% confidence intervals around the corresponding points on the curves in figure 29 are [1.45%, 1.73%] for the 70 row curve at 25 minutes, and [2.62%, 3.1%] for the 100 row curve at 25 minutes. These confidence intervals do not come close to overlapping.

The curves based upon 20 replications per point did not really provide any new insight. The group was forced to rely on curves based upon five replications per point because of time pressure, but these curves were adequate to the task, and did not result in any sub-optimal decisions.

5.4 The IBAM Team's Problem-Solving Methodology

By the time that the curves in the previous section had been generated, the IBAM team still had 10 possible line configurations under consideration. The baseline 154-row machine was carried as one of the 10 options and was evaluated alongside the other new configurations, but the team felt very confident that the 154-row web was unacceptable. Estimated dry slurry yield loss was simply too high on a 154-row carrier web, even under the most optimistic assumptions. But at the same time, the team had worked for months to develop the baseline configuration. The dry slurry yield loss issue had been overlooked during the development of the baseline system, but the baseline had many very desirable features that were in jeopardy of being lost in varying degrees depending upon which new configuration the team chose. Operator accessibility is one such feature. Simplicity of the overhead slurry delivery system is another.

The team decided that it should attempt to preserve as many of the baseline system's desirable features as possible when choosing another configuration. In order to accomplish this, we decided to spend several hours brainstorming the features of the baseline system which were desirable and worth preserving if possible. Also solicited were ideas for any new properties which, though they may not be characteristic of the baseline system, were deemed important and worthy of inclusion in the newly configured system. The list of brainstormed ideas were then compiled and distributed to all of the team members in the form of a ranking sheet. Each team member ranked all of the criteria individually giving each an importance rating of 1, 3, or 5, with 5 being very important and 1 being minimally important. A list of the brainstormed criteria and their cumulative importance rankings is shown below:

Criteria List

	Total
Shortest web maximizes yield	45
Output configuration is same as current RBAM we side (18x4)	45
Designed for CBAM expansion	43
Took yield loss operations off the web (Slurry lay down/Buffering)	41
Automatic reject in cells results in less added value and a cleaner well-functioning web (Optimized)	41
Web sections separated by catenaries (Precise positioning at placement)	39
Intermittent motion (Takes advantage of dwell period)	39
The mandate to lower capital with lower capacity but allow for expansion (Made simple in-line seal feasible)	39
Current configuration uses a simple sealer (Minimizes capital)	37
System allowed us to use a proven battery assembly method (Less risk, conservative approach)	35
Continuous operator access along one side of web	35
Most efficient for unit operations (buffered), as well as operational and maintenance access	35
Simple IBAM half cluster expands easily to full cluster	35
Opens up critical maintenance areas (Sealer/Composite area)	35
3 pitch sealer results in shorter web length	33
System lends itself to raw material flow in , and finished goods flow out	33
System utilizes existing anode cell (Minimizes capital)	33
Best utilization of building area (Minimum impact on facility)	33
System configuration driven by R5 production schedule	33
Fits with vacuum sealer cycle time of 8.1 seconds	33
Few operations on web line (Simplicity of web)	31
Unobstructed line of sight is maximized with current system	31
Easy operator access (Possibility to run with two operators)	30
Three active slurry areas are in the same general vicinity (minimizes labor)	29
Manufacturing process approach is one component at a time	29
Single composite loop design (Results in a shorter web length than the dual loop)	27
Decision to go with in-line cathode (Less expensive)	27
Desire for straight line slurry delivery allowed us to use current system engineering (Minimizes capital and risk)	27
Double pallet resulted in a shorter web with one web apply station (Minimized capital)	27
Anode is driving the composite configuration (Capital, operation, maintenance)	27
IBAM Over head slurry delivery system (in-line and simple)	25
Cathode and sealer have a common operator side (Minimizes labor and allows for a double-duty platform)	25
Web utilizes existing die module design (Saves money)	25
Current system configuration facilitated safety zone and tension zone design	25
Has pocketing operation in line	25
Cathode over web (Best one-sided access to machine)	23
Existing configuration of cells is a standard approach (Special not yet explored)	23
Have not looked at non-standard approaches to positioning of commercial equipment	23
In-line cathode 4-up vs. 8-up (Resulted in a shorter web length)	19
In-line cathode uses less floor space (Less perimeter to guard)	19
In-line cathode results in shorter cathode web line, but also a longer main web line	19
Less automation in tub handling (Shortened overall machine length)	19
Sample and reject station in line	19

Allowed for acceptable electrical cabinet scheme	19
Cut loop width driven by rear enclosure and slurry elevator interface	19
Single mother hopper (Minimizes capital)	17
Moving cutoff eliminated a catenary (results in shorter web)	17
Material flow left to right (Gives us access)	13

Table 17: Criteria Scores

Since there were so many criteria, the engineer who took responsibility for compiling the responses decided to place the criteria into four general groups and then average the scores of each category. He then presented his work to the team and asked for comments. The team agreed that evaluating all 10 IBAM configurations on each of the criteria would be excessively time consuming, and welcomed the opportunity to consolidate the criteria into a manageable number of categories. The categorized criteria are listed below:

Criteria List

SHORTEST WEB LINE AND NUMBER OF COMPONENTS BUFFERED (WEIGHT=7)	
Shortest web maximizes yield	45
Took yield loss operations off the web (Slurry lay down/Buffering)	41
Automatic reject in cells results in less added value and a cleaner well-functioning web (Optimized)	41
Most efficient for unit operations (buffered), as well as operational and maintenance access	35
3 pitch sealer results in shorter web length	33
Few operations on web line (Simplicity of web)	31
In-line cathode results in shorter cathode web line, but also a longer main web line	19
AVERAGE SCORE	35
EXPANDABLE TO CLUSTER (WEIGHT=5)	
Designed for CBAM expansion	43
Simple IBAM half cluster expands easily to full cluster	35
In-line cathode uses less floor space (Less perimeter to guard)	19
AVERAGE SCORE	32.333333
OPERABILITY AND MAINTAINABILITY (WEIGHT=3)	
Web sections separated by catenaries (Precise positioning at placement)	39
Intermittent motion (Takes advantage of dwell period)	39
Continuous operator access along one side of web	35
Opens up critical maintenance areas (Sealer/Composite area)	35
System lends itself to raw material flow in, and finished goods flow out	33
Fits with vacuum sealer cycle time of 8.1 seconds	33
Unobstructed line of sight is maximized with current system	31
Easy operator access (Possibility to run with two operators)	30
Three active slurry areas are in the same general vicinity (minimizes labor)	29

Cathode and sealer have a common operator side (Minimizes labor and allows for a double-duty platform)	25
Current system configuration facilitated safety zone and tension zone design	25
Cathode over web (Best one-sided access to machine)	23
Sample and reject station in line	19
Allowed for acceptable electrical cabinet scheme	19
Cut loop width driven by rear enclosure and slurry elevator interface	19
Material flow left to right (Gives us access)	13
AVERAGE SCORE	27.9375
MINIMUM IMPACT AND RISK TO COST AND SCHEDULE (WEIGHT=1)	
Output configuration is same as current RBAM we side (18x4)	45
The mandate to lower capital with lower capacity but allow for expansion (Made simple in-line seal feasible)	39
Current configuration uses a simple sealer (Minimizes capital)	37
System allowed us to use a proven battery assembly method (Less risk, conservative approach)	35
System utilizes existing anode cell (Minimizes capital)	33
Best utilization of building area (Minimum impact on facility)	33
System configuration driven by R5 production schedule	33
Manufacturing process approach is one component at a time	29
Single composite loop design (Results in a shorter web length than the dual loop)	27
Decision to go with in-line cathode (Less expensive)	27
Desire for straight line slurry delivery allowed us to use current system engineering (Minimizes capital and risk)	27
Double pallet resulted in a shorter web with one web apply station (Minimized capital)	27
Anode is driving the composite configuration (Capital, operation, maintenance)	27
IBAM Over head slurry delivery system (in-line and simple)	25
Web utilizes existing die module design (Saves money)	25
Has pocketing operation in line	25
Existing configuration of cells is a standard approach (Special not yet explored)	23
Have not looked at non-standard approaches to positioning of commercial equipment	23
In-line cathode 4-up vs. 8-up (Resulted in a shorter web length)	19
Less automation in tub handling (Shortened overall machine length)	19
Single mother hopper (Minimizes capital)	17
Moving cutoff eliminated a catenary (results in shorter web)	17
AVERAGE SCORE	27.818182

Table 18: Categorized Scores

The four categories were then assigned importance weightings of 7, 5, 3, and 1, in order of descending importance. The plan was to rank each of the 10 possible IBAM configurations on the four broad criteria defined by the categories. It took the team the better part of a full day to evaluate all of the configurations. Each configuration was given a score of high, medium, or low for each criterion, depending upon how well the team collectively felt that each configuration addressed each particular criterion. A matrix showing how the team scored the different configurations is shown below. The 10 IBAM configurations evaluated are labeled 'A' through 'J'. Note that the third and fourth categories were broken down into a few subcategories each:

Hi, Medium, Lo Rankings	**OPTION**									
	A	**B**	**C**	**D**	**E**	**F**	**G**	**H**	**I**	**J**
7: Yield Shortest Web Line & # Components Buffered	Lo	Lo	M	M	Hi	M	M	Lo	M	M
5: Expandability Expandable to Cluster	Hi	Hi	M	M	M	M	Hi	Lo	Lo	Hi
3: Operability and Maintainability										
Accessibility to high tending areas	Hi	Hi	M	Hi	Lo	M	M	Lo	Lo	M
Good material flow	Hi	Hi	M	Hi	M	M	Lo	Lo	Lo	M
Minimum # of operators	Hi	Hi	M	Hi	M	M	M	Hi	Hi	M
1: Minimum Impact and Risk to Cost and Schedule										
Lo Cost	Hi	Hi	M	M	M	M	M	Lo	Lo	M
Lo Risk	Hi	Hi	M	Lo	Lo	M	M	Hi	Lo	M

Table 19: Concept Evaluations

The next step was to assign scores to the grades of high, medium, and low, and to numerically rank the 10 configurations. The rankings of 41, 21, and 1 were used for high, medium, and low, respectively. The average score for each configuration was calculated, along with a weighted average in which each score was weighted by the importance weighting of the criterion (7, 5, 3, or 1). The numbers 41, 21, and 1 were chosen by the engineer compiling the scores so that the minimum difference in average score was 1. This was somewhat arbitrary, but for the sake of clarity, it seemed better to compare whole numbers. The matrix of scores along with the straight averages and weighted averages is shown in table 20:

Hi, Medium, Lo Rankings	**OPTION**									
	A	**B**	**C**	**D**	**E**	**F**	**G**	**H**	**I**	**J**
7: Yield Shortest Web Line & # Components Buffered	*1*	*1*	*21*	*21*	*41*	*21*	*21*	*1*	*21*	*21*
5: Expandability Expandable to Cluster	*41*	*41*	*21*	*21*	*21*	*21*	*41*	*1*	*1*	*41*
3: Operability and Maintainability										
Accessibility to High tending areas	41	41	21	41	1	21	21	1	1	21
Good material flow	41	41	21	41	21	21	1	1	1	21
Minimum # of operators	41	41	21	41	21	21	21	41	41	21
AVERAGE FOR CATEGORY	*41*	*41*	*21*	*41*	*14.3*	*21*	*14.3*	*14.3*	*14.3*	*21*
1: Minimum Impact and Risk to Cost and Schedule										
Lo Cost	41	41	21	21	21	21	21	1	1	21
Lo Risk	41	41	21	1	1	21	21	41	1	21
AVERAGE FOR CATEGORY	*41*	*41*	*21*	*11*	*11*	*21*	*21*	*21*	*1*	*21*
STRAIGHT AVERAGE	**31**	**31**	**21**	**24**	**22**	**21**	**24**	**9**	**9**	**26**
WEIGHTED AVERAGE	**30**	**30**	**21**	**23**	**23**	**21**	**25**	**8**	**10**	**26**

Table 20: Concepts' Numerical Scores

One further matrix is of interest. In the following matrix, cumulative weighted scores are calculated rather than weighted averages. Each score of 41, 21, or 1 is multiplied by the importance factor of that criterion, either 7, 5, 3, or 1. The scores are entered in the matrix and summed to give a total score for each of the ten configurations. Note that this matrix was generated after the internship was complete:

Weighted Scores	**OPTION**									
	A	**B**	**C**	**D**	**E**	**F**	**G**	**H**	**I**	**J**
7: Yield Shortest Web Line & # Components Buffered	**7**	**7**	**147**	**147**	**287**	**147**	**147**	**7**	**147**	**147**
5: Expandability Expandable to Cluster	**205**	**205**	**105**	**105**	**105**	**105**	**205**	**5**	**5**	**205**
3: Operability and Maintainability	**123**	**123**	**63**	**123**	**43**	**63**	**43**	**43**	**43**	**63**
1: Minimum Impact and Risk to Cost and Schedule	**41**	**41**	**21**	**11**	**11**	**21**	**21**	**21**	**1**	**21**
CUMULATIVE SCORES	**376**	**376**	**336**	**386**	**446**	**336**	**416**	**76**	**196**	**436**

Table 21: Cumulative Scores

The IBAM team reviewed the matrix in table 20 the day after the team had completed the scoring. The conclusion was that the decision-making methodology which we had employed had failed. We concluded this because according to the weighted average scores, option A, which was the baseline 154-row machine, received the highest score. We knew that the baseline system was unacceptable. The fact that it received the highest score indicated that our numerical ranking scheme was faulty. However, we figured that the time spent on the decision matrix was not necessarily time wasted, because we all developed greater insight into a wide variety of issues which we may not have considered previously.

At this point, we spent another day or so discussing our options, and decided that we would go with option E, which was the shortest option. (Option E had only 70 battery rows on the carrier web.) We arrived at this decision based upon our belief that yield was more important than other considerations, and that operator accessibility was probably equal to that of option A, the baseline. The team collectively felt very good about option E. The slurry threshold tests would not be completed for several months, but we felt that we had done all we could to address the dry slurry problem given the program's cost and schedule constraints.

The interesting point is that based on table 21, our decision-making methodology was actually pretty sound. The weighted cumulative scores show that option E was the best, based upon our criteria. I only discovered this after the internship was complete. The major shortcoming of our process was the way in which weighted average was computed, which tended to minimize the importance rankings of the criteria. Fortunately, the IBAM team was wise enough to resist the temptation to make what we all knew was a bad decision by going with the option with the highest weighted average score.

Below is a crude diagram of option E. Unless further changes occur in the future, this is the way the IBAM will look when production starts in 1996:

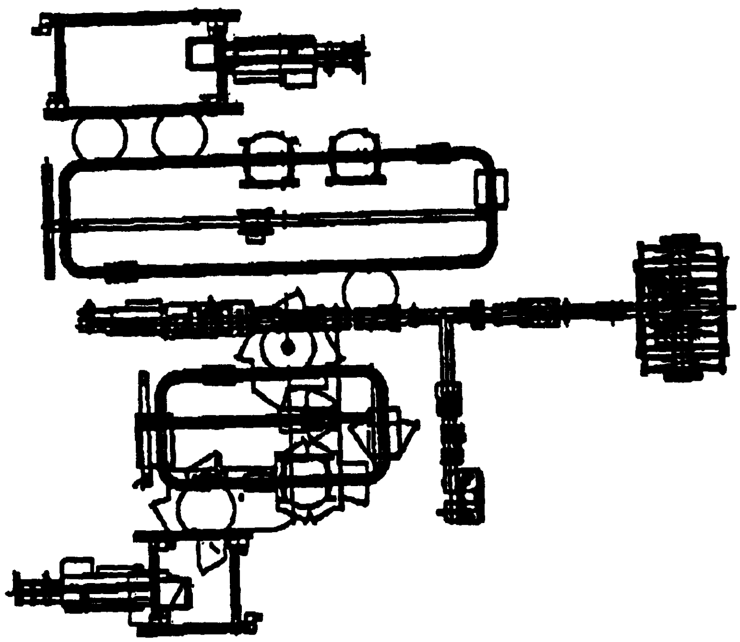

Figure 21: 70 Row IBAM Configuration

5.5 Yield Loss vs. Run Time Curves

Following the decision to go with the 70-row option, I decided to do some further analyses to determine the effects of different run times on dry slurry yield loss. The curves shown in this section were not used by the IBAM team during their decision making efforts, but were presented to the program manager for reference and to provide further insight into the dry slurry yield loss issue.

Figures 16 through 20 show dry slurry yield loss as a function of slurry threshold time for 2 specific IBAM run time scenarios, a best case and a worst case. Since the future run time performance of the IBAM is very uncertain, I thought it would be useful to run simulations which showed yield loss for a continuum of run time levels. It was necessary to hold slurry threshold constant for each family of curves. The multiplicative factor set at 0.7 on previous simulations was now varied in order to get different run time performance from replication to replication. Slurry threshold times of 15 minutes and 25 minutes were chosen. The families of curves are shown below. Note that the dependent axes are scaled differently:

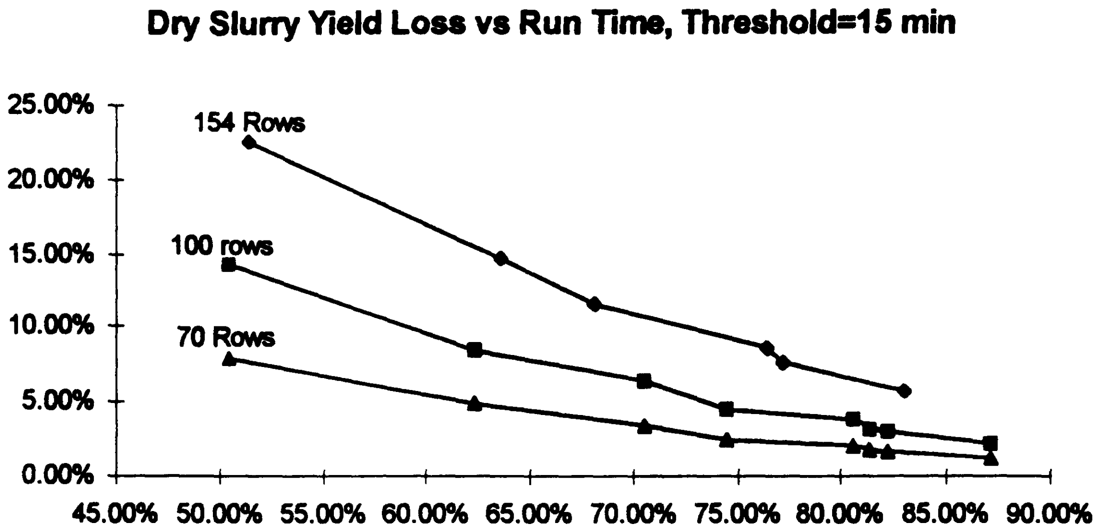

Figure 22: Yield Loss vs. Run Time

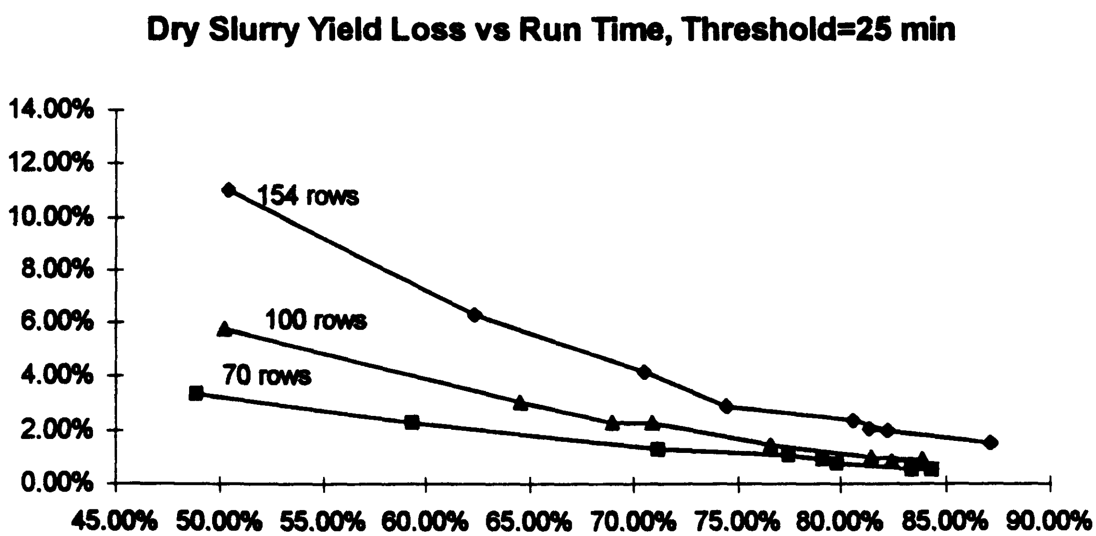

Figure 23: Yield Loss vs. Run Time

The insight one can gain from these curves is that in the event that run time performance of the IBAM is very bad, dry slurry yield loss is dramatically worse with longer carrier web lengths. This analysis further validates the decision to go with the 70-row carrier web. The 70-row curve is much flatter, meaning that the 70-row web is much more robust against dry slurry yield loss.

Chapter 6 - Application of Theory of Constraints

6.1 Theory of Constraints

The basic idea behind the theory of constraints is that in any process in which the final output is dependent upon a number of different sub-processes, the speed of the overall process is defined and limited by the slowest sub-process.[2] This applies to complex business systems involving suppliers, manufacturers, and distributors, and to smaller systems at individual plants. This chapter examines how the theory of constraints was applied to the IBAM in order to recommend improvements in the throughput of the line.

Recall that the IBAM team decided that in order to reduce the expected dry slurry yield loss, three basic approaches were appropriate. These were:

- Increase the speed of the IBAM.
- Reduce the length of the carrier web.
- Increase the slurry threshold level.

Chapter 5 described in detail the IBAM team's successful efforts to reduce carrier web length from 154 rows to roughly 70 rows. I say roughly 70 rows, because final design iterations may necessitate a small increase in length, or may allow a small reduction in length. Retesting of the slurry threshold time was being addressed by a team at the plant. The remaining step, to increase the speed of the IBAM, did not receive as much of the team's attention. This was because the team did not believe that dramatic speed increases were attainable.

In crude terms, the battery assembly process on the IBAM could be diagrammed as follows:

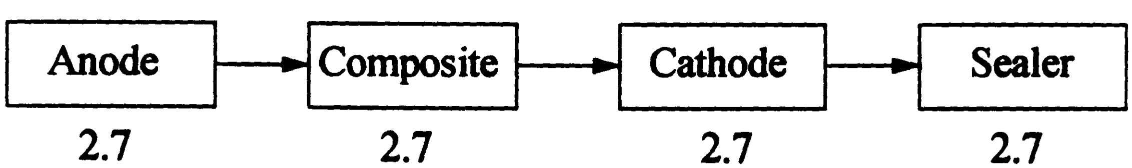

Figure 24: IBAM Process Flow

Essentially, there are four major processes which all process one battery row (4 batteries) in 2.7 seconds. In a layout such as this in which every process is coupled to the next, and there is no possibility of building inventory between processes, it is essential that each process be executed at the same speed as the next. The system as diagrammed above has no bottleneck. All processes would have to have the same increase in speed in order to achieve greater throughput.

The theory of constraints and the concept of bottlenecks applies when the question of how much speed increase is possible is posed. Although the maximum speed of each of the 4 sub-processes shown above is, strictly speaking, unknown, the team felt nearly certain that the composite process had the least opportunity for speed increases, and was therefore, the system's bottleneck. One engineer offered the idea of putting another double-headed cambot into the composite loop in order to give it potential for greater speed increases. The initial response to the idea was rejection. Other team members brought up concerns about extra conveyors which would be required to feed the second cambot, and about the capacity of the other equipment in the composite loop, so the idea temporarily died. The next section describes how a closer look along with refinements to the idea introduced a whole new option for the program manager to consider.

6.2 Effects of Adding a Second Double Cambot in Composite Loop

There is less opportunity to speed up the process of applying composite rows to battery stacks than there is in other areas, primarily because three composite rows have to be applied to the battery stack. The anode and cathode areas have only one component to apply, and the vacuum sealer seals three rows simultaneously, so it is believed that there is significant opportunity for speed increases in the sealing operation. Recall from chapter 2 that each pallet in the composite loop has 2 composite rows on it, and the cambot which takes the composite rows off of the pallet and places them onto the battery stack on the carrier web is a double-headed cambot. It processes the two composite rows simultaneously. After placing three pairs of rows, one pair on top of the next, the cambot waits while the carrier web indexes two positions so that the next pair of incomplete battery stacks is in position to receive composite rows. The following slide, which I presented to the IBAM team during one morning meeting, shows the timing of events in the composite loop:

$$\frac{1.4+1.4+1.4+1.2}{6\text{ comp rows}} = \frac{5.4\text{ sec}}{6\text{ comp rows}} = \frac{\mathbf{0.9\ sec}}{\text{comp row}} \text{ X3} = \frac{\mathbf{2.7\ sec}}{\text{battery row}}$$

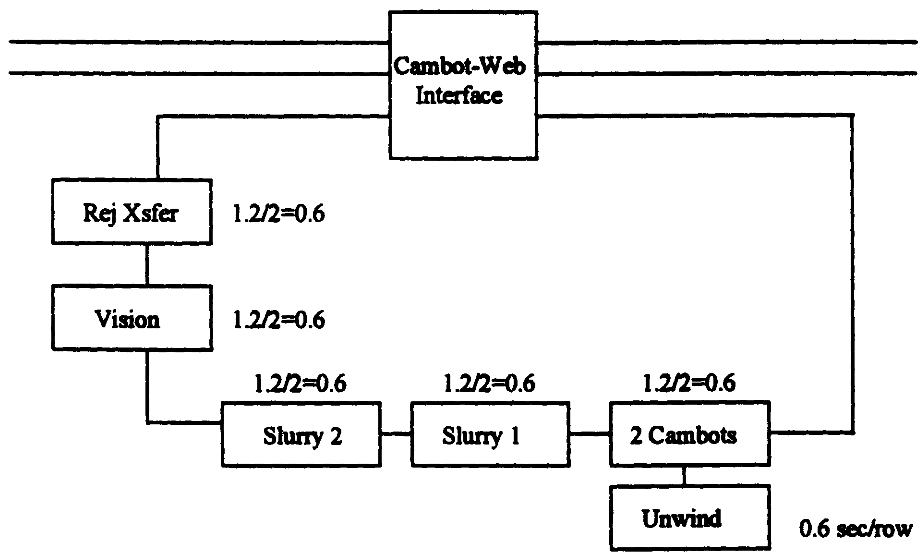

Figure 25: Current IBAM Timing

The diagram shows that the various processes in the composite conveyor loop all process two composite rows in 1.2 seconds, for a total of 0.6 seconds per row. The cambot-web interface encompasses all motions of the conveyor's walking beam, the double headed cambot, and the indexing motion of the carrier web. The relationship shown above the block diagram shows how for each pair of battery stacks on the carrier web, the cambot goes through three consecutive 1.4 second pick and place operations, followed by a 1.2 second double index of the carrier web. This means that it takes 5.4 seconds to process six composite rows, or conceptually, 0.9 seconds per composite row. The fact that the cambot-web interface requires 1 composite row every 0.9 seconds, coupled with the fact that the composite loop can provide 1 row every 0.6 seconds, shows that there is excess capacity in the composite loop.

To go after some of this excess capacity, it would be necessary to increase the speed of the cambot-web interface. There are a number of reasons why this is difficult. First is the fact that initial testing has shown that trying to speed up the actual 1.4 seconds worth of cambot motions is not practical. The arms are already swinging fast, and the heads are moving fast as well. Speed increases would force the cambots to operate beyond their intended parameters. The remaining opportunity lies in reducing the 1.2 second double index time of the carrier web. This could be accomplished by either speeding up the pin wheel drives, or overlapping the first 1.4 second cambot motion with the web index. Overlapping the two motions completely is impossible due to the requirement for dwell times at various points throughout the cambot's cycle. A simple speed increase, while feasible, may not be particularly desirable. The web is already indexing twice as far in less time than it is at both the anode and cathode application points. But certainly some overall speed improvements could be achieved by a combination of index time reduction and overlapping of motions. The actual speed increase possible will be determined by later tests, but the initial feeling is that a total of 2.4 seconds per battery row might be possible.

The next diagram is also a slide which I presented to the IBAM team to show how the timing would change if another double headed cambot was placed in the composite loop, as was once suggested and shortly thereafter dismissed as impractical. This slide shows a hypothetical case, and does not represent what is actually achievable:

$$\frac{1.4+(1.4+1.2)}{6\text{ comp rows}} = \frac{4.0\text{ sec}}{6\text{ comp rows}} = \frac{0.667\text{ sec}}{\text{comp row}}\ X3= \frac{2.0\text{ sec}}{\text{battery row}}$$

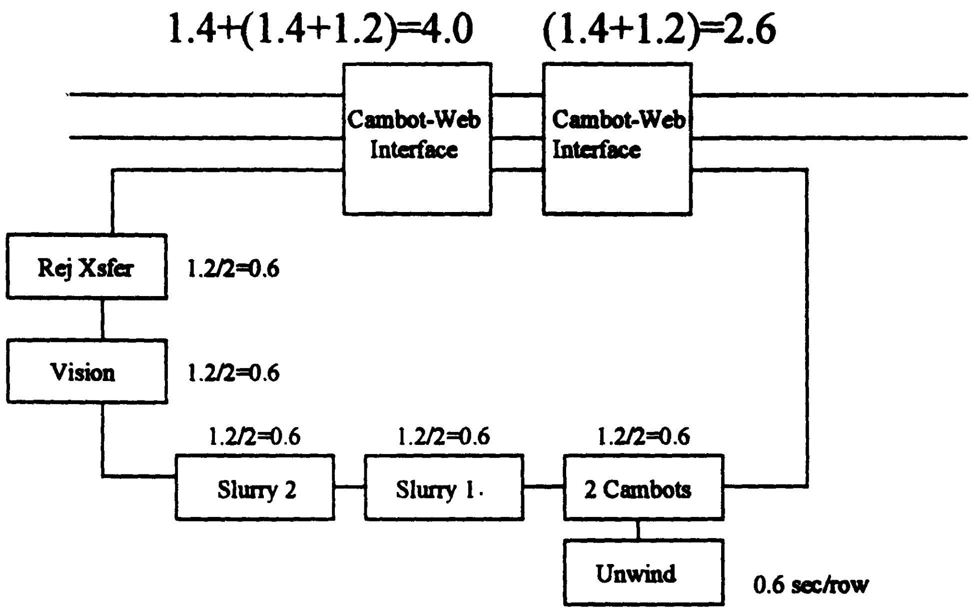

Figure 26: Hypothetical IBAM Timing

The idea is that the first cambot could place the first 2 pairs of composite rows onto the waiting battery stack. The carrier web would then complete its 1.2 second double index. The third pair of composite rows could then be placed onto the battery stack by the second cambot, a few feet downstream on the carrier web. According to the diagram above, this would mean that the overall time would be reduced to 2.0 seconds per battery row. While the first cambot placed its second composite pair, the second cambot would simultaneously place the third onto battery stacks downstream. The events which occur simultaneously are shown in parentheses. Note that the 1.2 second index of the continuous carrier web is actually a single event, even though the diagram makes it look like two separate indexing motions take place simultaneously. The relationship at the top of the diagram shows that existing capacity in the composite loop could support a second cambot. The cambots taken together would effectively require one composite row every 0.667 seconds, and the loop is capable of providing one every 0.6 seconds.

The final slide presented to the IBAM team was intended to show the practical effect of the addition of the second cambot. Again, the situation is hypothetical and does not represent what is actually attainable by the IBAM:

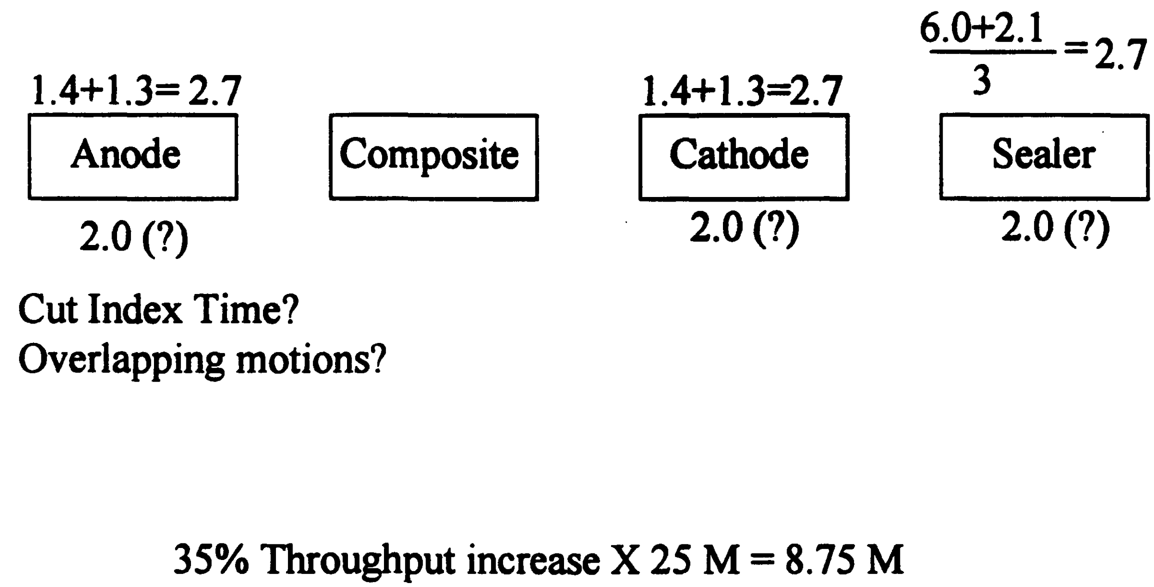

Figure 27: Hypothetical Throughput Increase

The slide shows that if 2.0 seconds/row were attainable, this would represent a 35% increase in throughput over the 2.7 seconds/row case. The diagram also highlights the need to increase the speed of the other major IBAM sub-systems to 2.0 seconds/row in order to actually achieve the throughput gains. The slide suggests that cutting index times and overlapping of some motions may be ways to achieve this.

The question of whether or not capacity existed in the composite loop to support a second cambot was answered by the slides, but the technical problems and increased complexity associated with the possible need for a conveyor system to support the second cambot had not yet been addressed. The next section describes how computer simulation was used to prove a concept, and what was actually attainable in the case of the IBAM.

6.3 The Role of Computer Simulation

I developed a crude simulation to test whether or not a simple sequencing scheme could be used to get pallets to the cambots on the existing conveyor system. The concept was simple, the first cambot would take the pair of composite rows off of the first two pallets which passed by on the conveyor system, but would allow the third pallet to pass. The second cambot would then see two empty pallets, which it would ignore, followed by a third pallet which carried a pair of composite strips. The second cambot would process the composite pair, let the next two empty pallets pass, process the next pair, and so on. The beauty of this concept was that all concerns about the requirement for additional conveyors and equipment to feed the second cambot were eliminated.

Believing that my crude simulation proved that the concept worked, I presented my simulation to the IBAM team. The team was not convinced. The simulation did not show in sufficient detail whether or not the actual motions of the cambot arms and heads would conflict with the pallets passing underneath. The program manager asked me to develop a more detailed model which would convince him and the rest of the team that the concept would really work. I did this. The model took perhaps three full days worth of work to finally debug. The program manager and the rest of the team now believed that the concept was worthy of consideration.

It is important to point out that simulation showed that the 2.0 seconds/row in the hypothetical case described in the previous section was not attainable. The lowest time/row possible turned out to be 2.1 seconds. This was because the walking beams which positioned the pallets under the cambot heads operated at 1.4 seconds/cycle. Three walking beam cycles would be required for each pair of battery rows, for a total of 4.2 sec/2 battery rows or 2.1 seconds/row. This would represent a 28.5% increase in throughput over the 2.7 seconds/row case, or a 14.3% increase in throughput when compared to the 2.4 seconds/row which, you may recall, was the most optimistic estimate of what the current composite system could achieve.

Note the very different role that computer simulation played in addressing this issue when compared to the dry slurry yield loss issue. The simulations written to test dry slurry yield loss were block-diagram type simulations, and were specifically designed without any graphics in order to improve computing efficiency. The double-cambot simulations, on the other hand, were used to give the team members a visual picture of how the pallets and cambot heads would interact under the proposed sequencing scheme. A flexible and versatile simulation package and someone with the ability to effectively use it provide a program manager with the ability to examine many different kinds of issues.

The next section explores the issues surrounding the decision of whether to include a second cambot or not.

6.4 What Was Actually Decided

I was very much in favor of adding the second cambot to the composite loop in order to achieve the gains in throughput. Before meeting with the program manager to discuss the issue, I compiled the following list of what I saw as the benefits and drawbacks of the idea:

<u>Benefits</u>

- 2.1 seconds/row vs. 2.7 seconds/row = 28.5% throughput increase.
  vs. 2.4 seconds/row = 14.3% throughput increase.
- Can start at 2.7 seconds/row until down the learning curve and comfortable with 2.1.
- Cambots can serve as backups for eachother during failures, so run time would improve.
- Less load on the cambots. (Still a 1.4 second cycle, but doing it fewer times.)
- Dry slurry yield loss will decrease with faster line speed.
- Even if the future shows that additional throughput is not needed, the increased IBAM throughput would allow the overworked RBAM to be offloaded of some production. It is better to make more batteries with the higher yield process. (Lower material and labor costs.)
- May allow for delay of the decision to expand to the CBAM[1] for a couple of years, or may even eliminate the need altogether. In either case, this is lots of $.

[1] CBAM refers to the possible cluster-style expansion of the IBAM to a second line. The decision to expand or not will come sometime in the future.

<u>Drawbacks</u>

- Additional $ required for extra equipment. ($500K estimated)
- Longer carrier web. (By about 6 feet or 18 rows.)
- May need to make the composite loop larger.
- All other equipment would be operating at a higher rate, and may get worn sooner.

It is important to remember that whether or not the other sub-systems of the IBAM could achieve 2.1 seconds/row was still an unknown. The engineers responsible for each sub-system predicted that 2.1 seconds/row was attainable, but their predictions, particularly in the cathode area, were less than certain.

As stated in the list of drawbacks, adding another cambot to the composite loop would necessitate increasing the carrier web length by an estimated 18 rows. If 2.1 seconds/row was possible, then the increased speed would more than offset the extra carrier web length, and dry slurry yield loss would improve. If, however, 2.1 was not attainable by the other IBAM subsystems, dry slurry yield loss might worsen, depending on what speed was actually attainable. Following the weeks' worth of efforts to reduce web length and minimize dry slurry yield loss, the program manager was disinclined to pursue a strategy that due to uncertainty, might actually increase dry slurry yield loss, albeit by a very small amount.

This reluctance, coupled with the fact that the program manager was uncertain whether he could get roughly $500K extra, caused him to decide against immediate implementation of the second cambot. Rather, he would keep the idea in his hip pocket as an option for future expansion. Once the IBAM produced a few million batteries, he would know for certain what speeds were attainable by the major IBAM sub-systems.

The major drawback associated with this decision is that it will cost much more to add the second cambot after production starts than it would have cost had it been implemented in the initial construction of the IBAM. Lost production during rework and actual rework costs will be much higher than $500K.

The decision to hold off on implementation of the second cambot is certainly a defensible decision. In fact, testing of the vacuum sealer which occurred just prior to the publication of this thesis indicated that the sealer may only be able to operate at 2.4 seconds/row.

The actual bottleneck of the system cannot be conclusively determined until thorough testing is done on the actual production equipment. The addition of a second cambot in the composite loop would reduce its cycle time to 2.1 seconds/row. If another system such as the sealer were not capable of achieving 2.1 seconds/row, it would be the new bottleneck, and should be the focus of extensive improvement efforts. The financial benefit associated with increases in IBAM throughput are tremendous, and should be weighed carefully against the cost of implementation.

Chapter 7 - Effects of Decoupling the Carrier Web

7.1 Chapter Overview

This chapter is dedicated to a discussion of some of the implications of the continuous paper carrier web. The finishing machines in use at R-5 require a paper-based carrier web. This chapter examines the effects that a non paper-based system could have on dry slurry yield loss. Due to cost and schedule constraints, the concepts presented in this chapter do not have direct applicability to the IBAM project. However, if the IBAM project is expanded in the future to include a second line and a new set of finishing machines, the ideas presented in this chapter should be considered during the design phase.

7.2 Decoupling the Carrier Web

Table 16 shows that 36.7% of all downtime on the RBAM web side is caused by problems with the carrier web. This fact suggests that the carrier web system (also referred to as the web transport system) is a reasonable area to focus improvement efforts on and, if necessary and practical, to change.

However, the web transport system developed for the IBAM is very similar to the web transport system used on the RBAM. Major exceptions are that the IBAM's web transport system will make use of intermittent motion instead of continuous motion, and that the IBAM will operate at a dramatically reduced speed. Additionally, the IBAM's web transport system incorporates many of the desirable features of the RBAM card side's drive system, which is less problematic than the corresponding web side system. It is hoped that these factors will contribute to improved IBAM run time performance.

Still, it is worthwhile to consider the effects that alternative web transport designs could have on dry slurry yield loss and overall IBAM efficiency. Consider a system in which the long, continuous paper carrier web is replaced by a conveyor system carrying pallets, each with a single 1x4 battery row. Just like on the current IBAM, components would be placed on top of one another until a complete battery stack was built. Then, the battery stack would be vacuum sealed. The line would be nearly the same as the current IBAM, with the obvious exception that the backbone of the line would be the conveyor system carrying pallets instead of the current paper carrier web transport system. This concept immediately gives rise to several problems:

- The output would be incompatible with the existing finishing equipment, which requires 18x4 battery strips.
- It is unclear how functions such as laminating and vacuum sealing could be accomplished when the battery stack rests on a pallet.
- The requirement for a return conveyor could introduce accessibility problems.

These and many other problems would need to be resolved. IBAM cost and schedule constraints currently preclude this concept from receiving any consideration, but the analysis which follows suggests that concepts such as this one are worthy of consideration during any future battery line projects.

The decoupled conveyor system described above would have some buffering between the different points at which components were placed onto the pallets, but more importantly, it would allow downstream processes to continue to run during upstream failures. The continuous paper carrier web prevents this from occurring on the current IBAM. Intuitively, one would expect dry slurry yield loss to be reduced, because only a portion of the work in process would be discarded during long downtimes, instead of all of the work in process. Of course, if the most downstream process is the cause of the long downtime, then all of the work in process would still have to be discarded. The vacuum sealing operation is the furthest downstream operation of interest, since slurry is no longer exposed after being sealed. On the RBAM, the sealer accounts for about 25% of the downtime. Intuitively, one would expect some dry slurry yield loss improvements during the other 75% of the downtime.

The chart that follows shows the results of simulations run using worst case ttf and ttr profiles on a system with the decoupled pallet conveyor system described above.

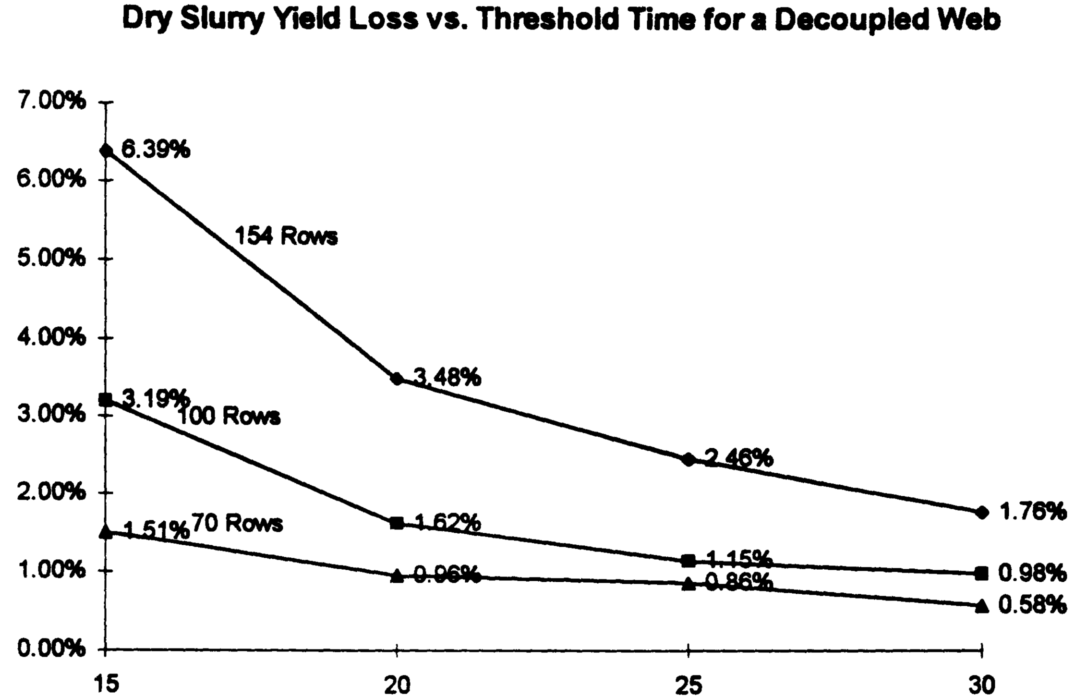

Figure 28: Decoupled Carrier Web

In order to isolate the effects of decoupling, several important assumptions were made before developing these curves:

1. The decoupled conveyor system fails according to the same profiles as the current RBAM carrier web.
2. Failures are evenly distributed across the length of the conveyor system. (Some failures occur upstream, and some occur downstream.)
3. The anode apply, composite apply, cathode apply, and vacuum sealer have speeds of 2.1 sec/row, 2.3 sec/row, 2.5 sec/row, and 2.7 sec/row, respectively, in order to ensure that the buffers fill up after failures. Note that the overall speed of the system is still 2.7 sec/row.

Figure 29 shows the curves for a coupled paper carrier web which serve as a basis for comparison to figure 28. Note that dry slurry yield loss is cut roughly in half when the decoupled system is used.

Figure 28 was generated in order to demonstrate that decoupling alone improves dry slurry yield loss. But one very important point has thus far been ignored. In the above scenario, times were adjusted to ensure that buffers filled up after failures. Actually, this is not desirable. By simply controlling the number of pallets between the point at which slurry is initially applied until it is vacuum sealed, the amount of work in process can be dramatically reduced.

Say, for example, that the conveyor ran at 18 inches/second, and that dimensionally, it was the same as the 154-row paper carrier web shown in figure 7. Smooth steady-state operation could be attained by using fewer than 20 pallets. The number of pallets on the return conveyor would be irrelevant. Only the number of pallets between the anode apply station and the vacuum sealer would need to be controlled.

By controlling the number of pallets, it would be possible to effectively have a decoupled 20-row system. A simulation was run to determine a single data point for a 20-row decoupled system with a dry slurry threshold of 15 minutes. The expected dry slurry yield loss was only 0.5%. This is even better than the corresponding RBAM benchmark of 0.92%.

7.3 Comments

This chapter was included in order to provide some useful ideas for expansion projects in the future. The most important issue throughout the examination of the dry slurry yield loss problem has been the length of the carrier web, which corresponds to the amount of work in process. By changing the carrier web concept, the amount of work in process could be reduced to levels unattainable with the current concept. It is also worth repeating that the paper carrier web was the cause of 36.7% of the RBAM downtime during the period tested. It may very well be that a new concept would be much more reliable and reduce downtime substantially. The future will tell whether the IBAM web transport system is problematic or not. If it is problematic, then concepts such as the one presented in this chapter may be useful.

Chapter 8 - Final Conclusions

8.1 Chapter Overview

This chapter is dedicated to a discussion of some overall thoughts and conclusions. Conclusions which pertain to the IBAM project itself were discussed in sufficient detail in the preceding chapters. This chapter describes some thoughts on the use of simulation in manufacturing projects general.

8.2 Conclusions

1. Simulation should be done as early as possible. The earlier in the course of a project that the dynamics of a system are understood, the better. Computer simulation is an excellent way to develop insight into these dynamics. Simulation should be used early to test the validity of concepts themselves, rather than just to try to optimize the performance of a particular concept once it has been chosen.

2. A statistical analysis should be done first. This point was mentioned previously, but is worth repeating. A detailed statistical analysis of the dynamics of the system should precede any large modeling effort. The insight gained from the statistical analysis will help to streamline the modeling efforts and ensure that they are pointed in a productive direction. A statistical analysis which is not followed by a modeling effort will probably be very beneficial by itself. A model which is not based upon a sound statistical analysis will probably not be of very great benefit.

3. It is better to have an internal capability to develop simulations. When third parties are contracted to develop computer simulations, it is unlikely that they will be able to develop models with sufficient flexibility to respond to rapidly changing project conditions. Third party simulations will probably be fine for testing different system variables like conveyor speeds once the design concept is frozen, but before design freeze, concepts themselves change, and multiple models will usually be required. Internal

modeling capability gives a development team the ability to test greatly varying scenarios as they come up throughout the course of the project.

4. Data acquisition resources should be given a high priority. Statistical analyses, computer models, and other analytical tools will only be as good as the data upon which they are based. State-of-the-art data acquisition systems can automatically collect a wealth of information on a manufacturing process. This information will promote an understanding of the process dynamics, and will be very valuable during any future improvement projects. Project managers should resist the temptation to cut data acquisition resources when facing budgetary pressures.

In summary, computer simulation is an extremely powerful tool when used effectively, and should not be ignored by project managers. The IBAM is scheduled to begin production on January 1, 1996. The machine which will go into production is much superior to the one which would have gone into production had computer simulation not been utilized.

References

[1] Hogg, Robert V. and Ledolter, Johannes. Applied Statistics for Engineers and Physical Scientists. New York, NY: Macmillan Publishing Company, 1992.

[2] Goldratt, Eliyahu M. and Cox, Jeff. The Goal. Croton-on-Hudson, NY: North River Press, 1992.

[3] ProModel version 1.10 Users' Manual.

Appendix A: Charts Based Upon 20 Replications Per Point

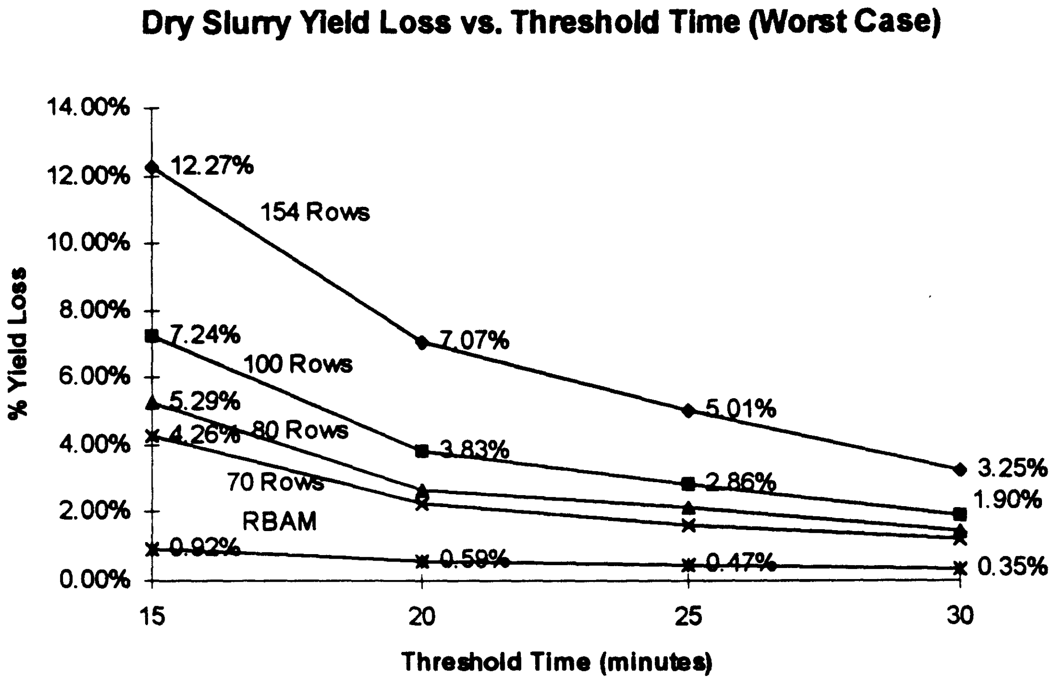

Figure 29: IBAM Worst Case

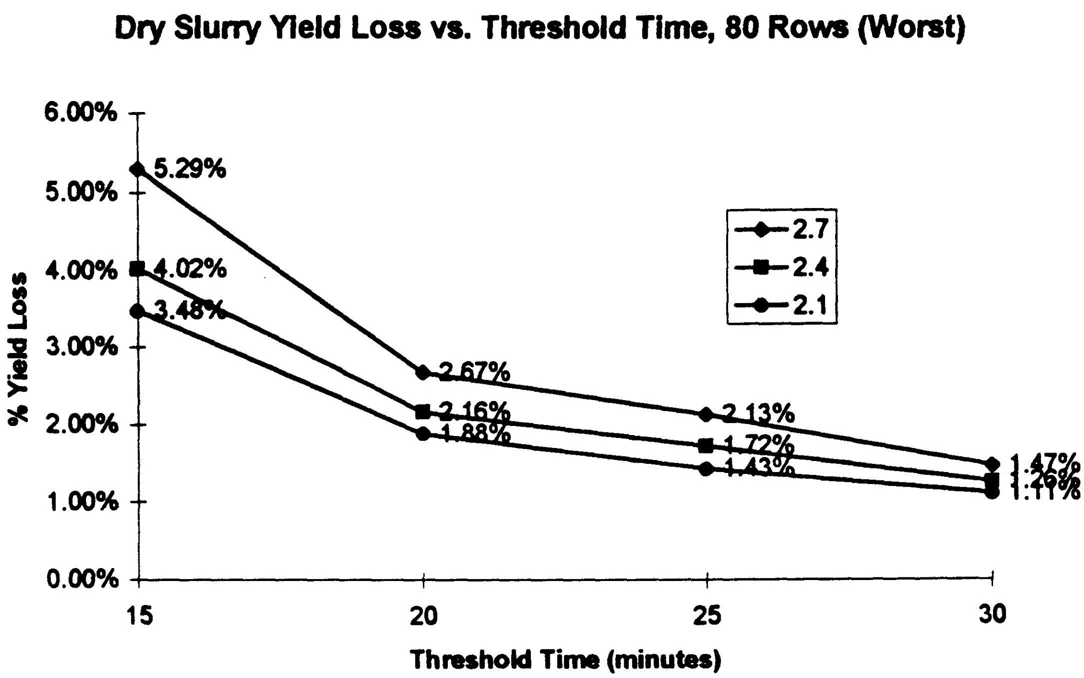

Figure 30: Varying Row Times

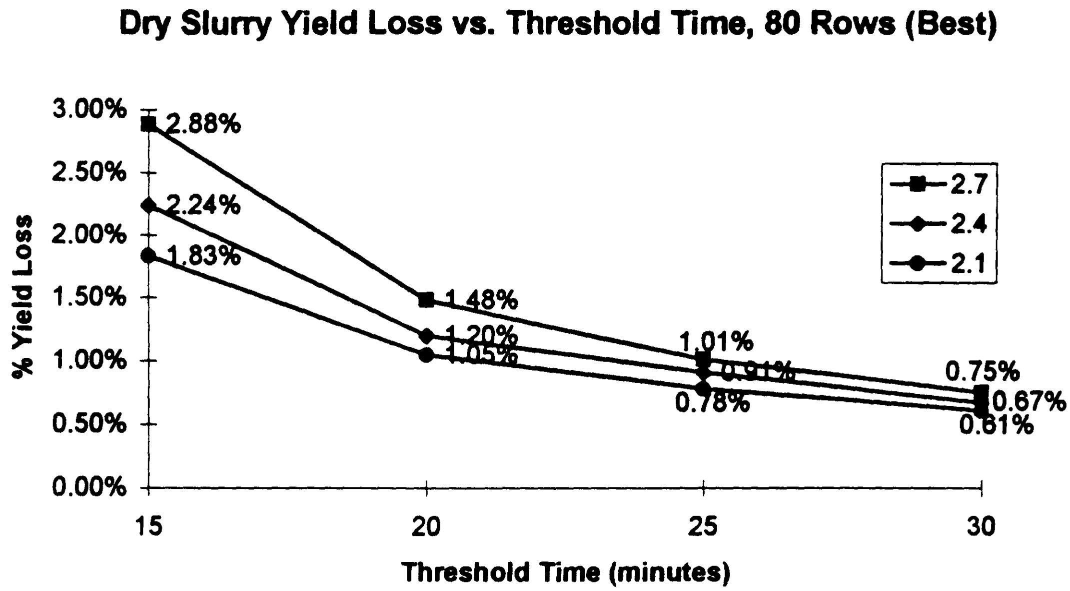

Figure 31: Varying Row Times

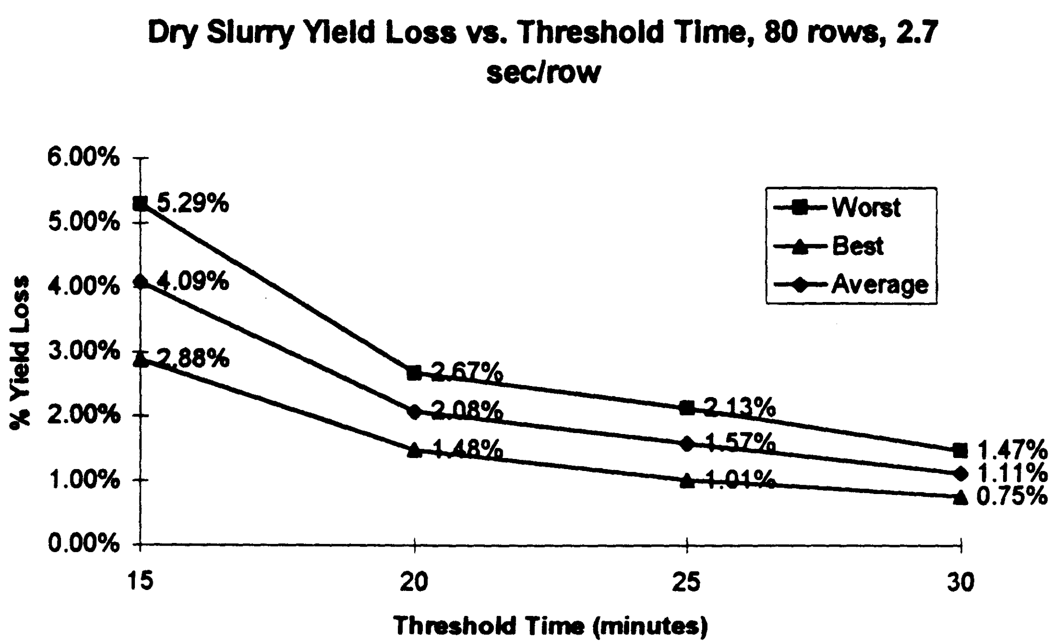

Figure 32: 80 Rows, 2.7 sec/row

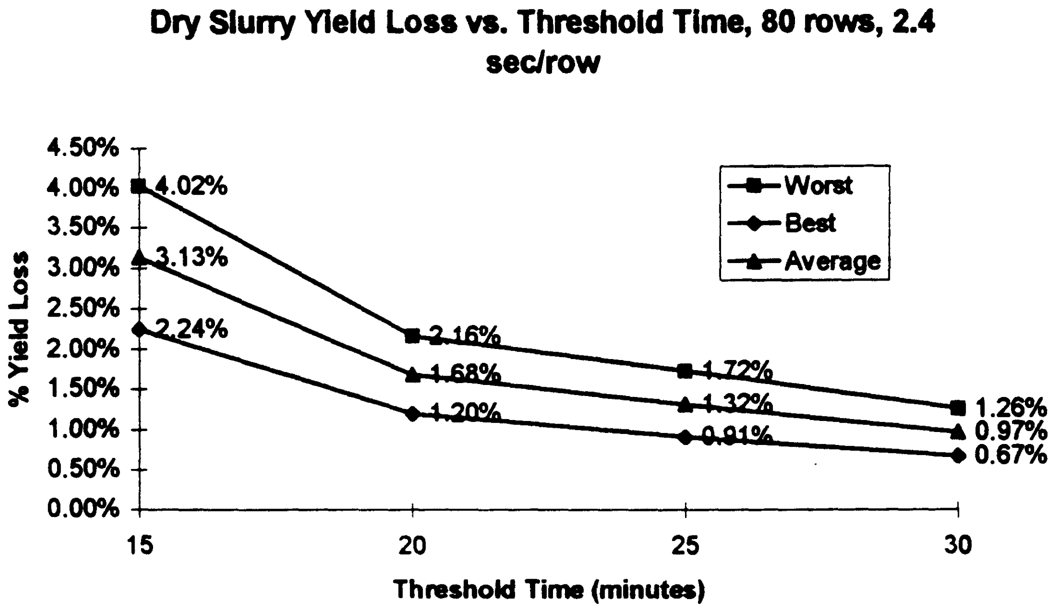

Figure 33: 80 Rows, 2.4 sec/row

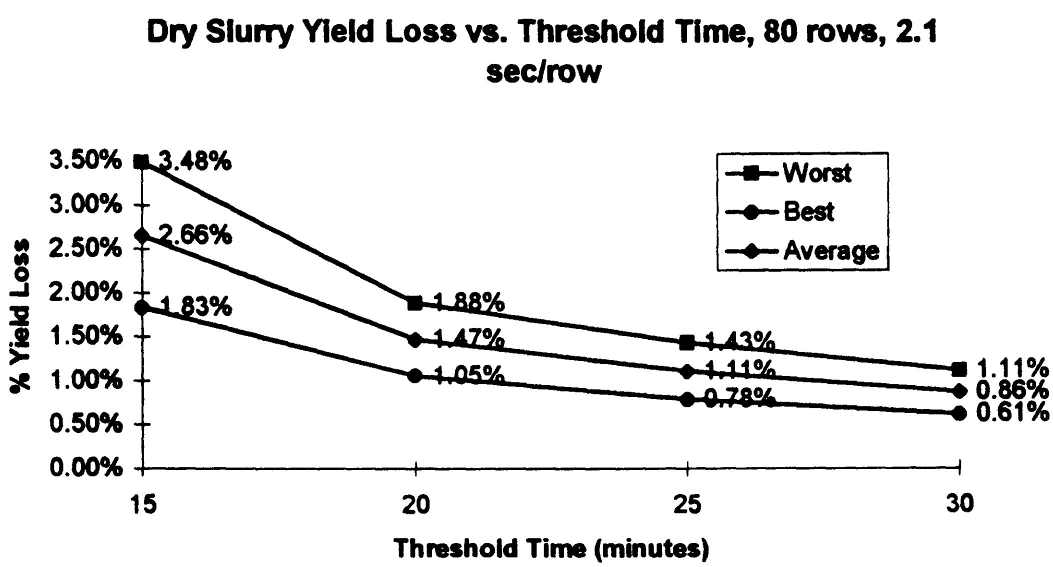

Figure 34: 80 Rows, 2.1 sec/row

Appendix B: Worst-Case Distributions

	%	Minutes
cathode_ttf	6.76	5
	10.74	10
	15.31	15
	20.48	22
	25.05	33
	30.02	45
	35.19	64
	39.96	80
	44.73	102
	49.7	128
	54.67	152
	59.64	180
	64.61	227
	69.58	270
	74.75	343
	79.52	461
	84.49	578
	89.46	757
	94.43	1036
	99.4	2194
	100	3050

	%	Minutes
cathode_ttr	2.93	1
	11.52	2
	21.68	3
	24.22	4
	53.13	5
	55.47	6
	58.2	7
	61.13	8
	62.11	9
	76.56	10
	78.52	12
	84.96	15
	88.28	20
	92.97	35
	97.66	120
	100	350

	%	Minutes
sealer_ttf	5.61	4
	10.08	5
	15.86	10
	21.3	15
	25.68	20
	30.06	27
	35.23	35
	39.96	48
	45.05	58
	50.04	70
	56.09	85
	60.3	100
	64.94	116
	69.94	137
	74.93	164
	79.93	199
	84.93	242
	89.92	316
	94.92	445
	95.97	500
	98.42	700
	99.47	900
	100	1208

	%	Minutes
sealer_ttr	0.93	1
	5.69	2
	14.53	3
	18.35	4
	46.64	5
	49.53	6
	53.44	7
	56.92	8
	58.03	9
	73.58	10
	76.13	12
	84.45	15
	85.39	18
	90.23	22
	95.58	40
	96.69	50
	97.79	60
	98.22	70
	98.73	80
	99.24	90
	99.75	120
	99.92	150
	100	185

	%	Minutes
web_ttf	8.76	5
	11.01	8
	15.51	11
	20.37	15
	25.24	22
	30.11	27
	35.36	37
	40.3	45
	46.37	55
	50.19	63
	55.21	73
	60.37	83
	65.47	97
	70.64	115
	75.36	146
	80.3	179
	85.32	214
	90.34	266
	95.36	380
	98.65	600
	99.7	900
	99.93	1200
	100	1602

	%	Minutes
web_ttr	0.93	1
	5.68	2
	11.5	3
	14.09	4
	39.11	5
	41.62	6
	46.08	7
	49.68	8
	50.83	9
	69.3	10
	71.82	12
	81.67	15
	86.7	20
	91.95	30
	95.97	50
	98.63	100
	99.35	150
	99.57	200
	99.64	250
	99.78	300
	99.93	350
	100	390

	%	Minutes
unwind_ttr	1.29	1
	13.45	2
	32.28	3
	40.47	4
	75.56	5
	77.89	6
	81.75	7
	85.5	8
	88.32	9
	93.57	10
	96.73	15
	98.95	30
	99.3	45
	99.53	60
	99.65	70
	99.77	100
	99.88	120
	100	170

	%	Minutes
carrier_unwind_ttf	4.88	49
	9.76	72
	15.12	87
	19.51	146
	24.39	158
	29.27	205
	34.15	250
	39.02	313
	43.9	393
	48.78	433
	53.66	494
	58.54	559
	63.41	680
	68.29	829
	73.17	923
	78.05	1093
	82.93	1311
	87.8	1539
	92.68	1725
	97.56	2353
	100	4949

	%	Minutes
anode_unwind_ttf	10.87	136
	21.74	377
	32.61	481
	43.48	751
	54.35	977
	65.22	1502
	76.09	2174
	86.96	3431
	100	7123

	%	Minutes
composite_unwind_ttf		
	10	219
	20	490
	30	771
	40	1090
	50	1402
	60	1934
	70	2442
	80	3321
	90	4531
	100	14967

	%	Minutes
cathode_unwind_ttf		
	4.99	19
	9.97	36
	15.49	47
	19.95	61
	25.2	73
	29.92	103
	34.91	117
	39.9	144
	44.88	179
	49.87	217
	54.86	252
	59.84	320
	64.83	373
	69.82	426
	74.8	522
	79.79	644
	84.78	750
	89.76	902
	94.75	1208
	99.74	2125
	100	2187

	%	Minutes
mother_hopper_ttf		
	19.44	17
	38.89	601
	58.33	1535
	77.78	4950
	100	34519

	%	Minutes
mother_hopper_ttr		
	22.22	10
	52.78	20
	58.33	25
	77.78	50
	100	150

	%	Minutes
baby_hopper_ttf		
	7.27	5
	12.27	20
	17.5	45
	22.95	86
	28.18	133
	33.41	188
	38.64	247
	43.86	313
	49.09	445
	54.32	585
	59.55	674
	64.77	856
	70	1105
	75.23	1386
	80.45	1717
	85.68	2282
	90.91	2858
	96.14	3759
	99.77	5189
	100	11225

	%	Minutes
baby_hopper_ttr		
	2.2	1
	8.79	2
	18.68	3
	20.6	4
	36.54	5
	37.36	6
	39.84	7
	42.58	8
	43.41	9
	57.42	10
	68.41	15
	81.04	20
	86.26	25
	90.38	30
	92.58	35
	93.96	40
	96.98	45
	98.63	60
	99.18	75
	99.73	120
	100	164

Appendix C: Best-Case Distributions

cathode_ttf	%	Minutes
	5.15	5
	9.62	12
	14.78	20
	19.24	35
	24.05	55
	28.87	90
	33.68	118
	38.49	142
	43.3	170
	48.11	233
	52.92	304
	57.73	365
	62.54	441
	67.35	518
	72.16	646
	76.98	796
	81.79	964
	86.6	1108
	91.41	1450
	96.22	1894
	100	3572

cathode_ttr	%	Minutes
	2.93	1
	11.52	2
	21.68	3
	24.22	4
	53.13	5
	55.47	6
	58.2	7
	61.13	8
	62.11	9
	76.56	10
	78.52	12
	84.96	15
	88.28	20
	92.97	35
	97.66	120
	100	350

sealer_ttf	%	Minutes
	6.33	5
	9.49	12
	14.87	25
	19.3	45
	23.73	65
	28.8	95
	33.54	129
	37.97	154
	42.72	198
	47.47	252
	52.22	299
	56.96	355
	61.71	392
	66.46	483
	71.2	561
	75.95	683
	80.7	738
	85.44	955
	90.51	1138
	94.94	1462
	97.47	1985
	100	3287

sealer_ttr	%	Minutes
	0.93	1
	5.69	2
	14.53	3
	18.35	4
	46.64	5
	49.53	6
	53.44	7
	56.92	8
	58.03	9
	73.58	10
	76.13	12
	84.45	15
	85.39	18
	90.23	22
	95.58	40
	96.69	50
	97.79	60
	98.22	70
	98.73	80
	99.24	90
	99.75	120
	99.92	150
	100	185

	%	Minutes
unwind_ttr	1.29	1
	13.45	2
	32.28	3
	40.47	4
	75.56	5
	77.89	6
	81.75	7
	85.5	8
	86.32	9
	93.57	10
	96.73	15
	98.95	30
	99.3	45
	99.53	60
	99.65	70
	99.77	100
	99.88	120
	100	170

	%	Minutes
web_ttf	5.22	3.3
	10.09	9
	15.2	15
	20.19	23
	25.64	30
	30.05	39
	34.92	51
	39.91	63
	44.9	76
	49.88	92
	54.87	113
	59.86	138
	64.85	161
	69.84	199
	74.83	230
	79.81	284
	84.8	340
	89.91	425
	94.78	589
	99.77	1013
	100	1602

	%	Minutes
web_ttr	0.93	1
	5.68	2
	11.5	3
	14.09	4
	39.11	5
	41.62	6
	46.08	7
	49.68	8
	50.83	9
	69.3	10
	71.82	12
	81.67	15
	86.7	20
	91.95	30
	95.97	50
	98.63	100
	99.35	150
	99.57	200
	99.64	250
	99.78	300
	99.93	350
	100	390

	%	Minutes
carrier_unwind_ttf	11.59	254
	27.54	739
	33.33	812
	57.97	1870
	63.77	2271
	75.36	2859
	76.81	2861
	95.65	6041
	97.1	7415
	100	8870

	%	Minutes
anode_unwind_ttf	8.93	152.5
	17.86	387
	26.79	575
	35.71	889
	44.64	1202
	53.57	1757
	62.5	2429
	71.43	3140
	80.36	3942
	89.29	5709
	100	14891

	%	Minutes
cathode_unwind_ttf	5.82	50
	10.58	90
	15.87	112
	21.16	152
	26.46	175
	31.75	211
	37.04	260
	42.33	315
	47.62	367
	52.91	467
	58.2	597
	64.02	752
	68.78	845
	74.07	987
	79.37	1222
	84.66	1450
	89.95	1795
	95.24	2430
	100	4483

	%	Minutes
composite_unwind_ttf	10	219
	20	490
	30	771
	40	1090
	50	1402
	60	1934
	70	2442
	80	3321
	90	4531
	100	14967

	%	Minutes
mother_hopper_ttf	19.44	17
	38.89	601
	58.33	1535
	77.78	4950
	100	34519

	%	Minutes
mother_hopper_ttr	22.22	10
	52.78	20
	58.33	25
	77.78	50
	100	150

	%	Minutes
baby_hopper_ttf	7.27	5
	12.27	20
	17.5	45
	22.95	86
	28.18	133
	33.41	188
	38.64	247
	43.86	313
	49.09	445
	54.32	585
	59.55	674
	64.77	856
	70	1105
	75.23	1386
	80.45	1717
	85.68	2282
	90.91	2858
	96.14	3759
	99.77	5189
	100	11225

	%	Minutes
baby_hopper_ttr	2.2	1
	8.79	2
	18.68	3
	20.6	4
	36.54	5
	37.36	6
	39.84	7
	42.58	8
	43.41	9
	57.42	10
	68.41	15
	81.04	20
	86.26	25
	90.38	30
	92.58	35
	93.96	40
	96.98	45
	98.63	60
	99.18	75
	99.73	120
	100	164