

Improving Operational Efficiency of a Small Manufacturing Maintenance Organization

by

Colin Poler

S.B., Mechanical Engineering
Massachusetts Institute of Technology, 2018

Submitted to the MIT Sloan School of Management and
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Master of Business Administration

and

Master of Science in Electrical Engineering and Computer Science
in conjunction with the Leaders for Global Operations program

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

© Colin Poler, 2022. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part in any
medium now known or hereafter created.

Author
MIT Sloan School of Management and
Department of Electrical Engineering and Computer Science
May 6, 2022

Certified by
Nelson Repenning
Professor of Management Science and Organization Studies
Thesis Supervisor

Certified by
Duane Boning
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Maura Herson
Assistant Dean, MBA Program
MIT Sloan School of Management

Accepted by
Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Improving Operational Efficiency of a Small Manufacturing Maintenance Organization

by

Colin Poler

Submitted to the MIT Sloan School of Management and
Department of Electrical Engineering and Computer Science
on May 6, 2022, in partial fulfillment of the
requirements for the degrees of
Master of Business Administration
and
Master of Science in Electrical Engineering and Computer Science

Abstract

Many small manufacturers struggle with poor maintenance efficiency, resulting in high maintenance costs and/or frequent equipment breakdowns. Existing literature addresses which tasks to prioritize and how to measure results, but there is little prior work on how to accomplish more maintenance work overall with the same resources and reduce maintenance wastes. We develop a framework for conceptualizing maintenance operational efficiency as a complement to maintenance strategy, focusing on the primary maintenance process: backlog, diagnosis, planning, getting parts, executing, and observing effects. We apply this framework to a small Michigan manufacturing facility. We estimate the cost of equipment breakdowns at the facility using a novel cross-referencing between maintenance breakdowns and production bottlenecks. Finally, we propose several improvements to target wastes in each step of the primary maintenance process: shared ownership of equipment between maintenance and production, more accessible documentation, a work order system, proximal spare parts storage, and solving problems permanently.

Thesis Supervisor: Nelson Repenning

Title: Professor of Management Science and Organization Studies

Thesis Supervisor: Duane Boning

Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

This project has been an incredible educational experience for me, on an academic, a professional and a personal level. I would like to thank Heartland Steel Products for giving me the opportunity to work on a production floor, and to work directly with the maintenance team for research.

To Patrick Johnson, my supervisor at Heartland Steel Products, thank you for your guidance and support throughout this research project. You provided astute feedback for navigating the Marysville facility, and for developing a research project that became my LGO thesis. You supported my focus on the maintenance organization, and supported the pilot implementations that demonstrated the proposed improvements. Most of all, you were a friend while I worked and lived in Michigan during a pandemic, far from my support network in Boston.

To the team at Heartland Steel Products, thank you for all of your help and cooperation with this project. From the hourly workforce to administration to supervisors, I appreciate how you opened up to me and helped me to better understand the company's situation. I appreciate the expertise you shared with me on topics ranging from maintenance to welding to scheduling to safety. Thank you for working with me even when I erred. I sincerely hope that my work in this thesis makes the Marysville facility a better place for all employees to work in.

To my thesis advisors, Duane Boning and Nelson Repenning, thank you for your guidance throughout this project. You helped me to develop a research project that was feasible given the challenges in Michigan, and you helped me to make sense of difficult situations that I faced on the production floor.

To my Leaders for Global Operations (LGO) classmates, thank you for all of your support throughout. Each of you, on different internships, were an invaluable sounding board for my thesis. Most importantly, you were my support network throughout, helping me through difficult situations I faced. I want to give a special thank you to Kenny Groszman, Luke Higgins, Lauren Sakerka and Connor Stehr for joining a weekly video call to discuss building good jobs for the community.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction	13
1.1	Company Background	13
1.2	Project Goals	14
1.3	Thesis Overview	14
2	Organizational Context	17
2.1	Industry Overview	17
2.2	Corporate Context	19
2.3	Manufacturing Overview	20
2.4	Details of Manufacturing Equipment	21
2.4.1	Press	23
2.4.2	Mills	23
2.4.3	Welding	25
2.4.4	Painting	26
2.5	Existing Maintenance System	27
3	Literature Review	33
3.1	Total Productive Maintenance	34
3.2	Reliability Centered Maintenance	36
3.3	Outsourced Maintenance	37
3.4	Maintenance Performance Management	38
4	Framework for Analyzing Maintenance	41

4.1	Objectives of a Maintenance System	41
4.2	Relationship to Production	44
4.3	The Primary Maintenance Process	46
4.3.1	Incoming Tasks	47
4.3.2	Diagnose	49
4.3.3	Plan	51
4.3.4	Get Parts	53
4.3.5	Executing Physical Work on Machines	55
4.3.6	Observing Effects	56
4.4	Other Maintenance Processes	56
4.5	Wastes in Maintenance	57
4.5.1	Incorrect Diagnosis	57
4.5.2	Bad Planning	58
4.5.3	Waiting for Parts	59
4.5.4	Missing Necessary Parts During Execution	60
4.5.5	Excessive Spare Parts Inventory	60
4.5.6	Motion to Retrieve Parts	62
4.5.7	Waiting for Maintenance	63
4.5.8	Fixing Avoidable Problems	63
5	Maintenance Problems in Marysville	65
5.1	Problems with Incoming Tasks	65
5.2	Problems in Diagnosis	67
5.3	Problems in Planning	69
5.4	Problems in Getting Parts	72
5.5	Problems in Executing Work	74
5.6	Problems in Observing Effects	76
6	Estimating the Impact of Maintenance Problems	77
6.1	Available Data	78
6.2	Breakdowns	79

6.3	Bottlenecks	80
7	Targeted Improvements and Results	89
7.1	Shared Ownership of Equipment	89
7.2	More Accessible Documentation	91
7.3	Work Order System	93
7.4	Proximal Spare Parts Storage	95
7.5	Solving Problems Permanently	99
8	Conclusion	101

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

2-1	Annotated picture of pallet racking product	18
2-2	Cost structure breakdown	19
2-3	Rapidly increasing steel prices	20
2-4	Production process diagram	22
2-5	Paint system diagram	27
2-6	Illustrative map of campus	29
4-1	Agents in the maintenance system	43
4-2	Repairperson includes multiple groups of personnel	44
4-3	Repair process steps, with feedback and failures	47
4-4	Categorization of machinery documentation	50
4-5	Virtuous cycle of improving diagnosis	51
4-6	Virtuous cycle of improving planning	53
4-7	Incorrect diagnosis	58
4-8	Bad planning	59
4-9	Waiting for parts	59
4-10	Missing necessary parts during execution	60
4-11	Excessive spare parts inventory	62
4-12	Motion to retrieve parts	62
4-13	Waiting for maintenance	63
4-14	Fixing avoidable problems	64
5-1	Vicious cycle of worsening diagnosis	69
5-2	Vicious cycle of worsening planning	71

5-3	Prior state of spare parts room	73
6-1	Rows of cleaned-up MES data	79
6-2	Breakdown time grouped by machine	80
6-3	Breakdown time grouped by machine and by month	81
6-4	Breakdown time grouped by month	82
6-5	Illustration of active periods	83
6-6	Illustration of which active periods are bottlenecks	83
6-7	Example of active periods at Heartland	84
6-8	Bottleneck time grouped by machine	85
6-9	Bottleneck time grouped by machine compared to breakdown time . .	85
6-10	Bottleneck time grouped by machine and month compared to subset that was breakdown time	87
6-11	Breakdown time grouped by machine and month compared to subset that was bottleneck time	88
7-1	Pilot of spare parts cabinet	97

Chapter 1

Introduction

This thesis serves as a case study for improving maintenance efficiency, specifically improving the ability of a maintenance system to handle tasks more quickly and more efficiently. This work is complementary to the typical maintenance management problem of prioritizing which tasks to handle first. The case study occurs in the context of a small manufacturer, struggling with a mandate to increase production capacity rapidly. The unique set of constraints on this business and the vast need for maintenance work requires a highly tactical approach to accomplish maintenance tasks more efficiently.

1.1 Company Background

Heartland Steel Products is a manufacturer of interior steel structures. The firm has three manufacturing facilities: Marysville, Michigan; Harrison, Ohio; and Lodi, California. Each facility has different equipment and capabilities with a small amount of redundancy between facilities. This case study focuses on the Marysville, Michigan facility, which produces pallet racking. The primary manufacturing processes are punching steel coil, roll-forming, metal inert gas (MIG) welding and then painting.

The Marysville facility struggles with frequent equipment breakdowns, which results in poor on-time delivery and poor labor utilization among other problems. The equipment is in poor condition because of historical under-maintenance and under-

investment. The Marysville maintenance organization is comparably resourced to other similar manufacturers, so chapter 5 shows that under-maintenance is largely caused by operational inefficiencies within the maintenance system.

1.2 Project Goals

Heartland wants to reduce the incidence and unpredictability of equipment downtime by making the maintenance system more efficient. The goal of this project is to develop specific operational improvements that improve the efficiency of the Marysville maintenance system. The improvements proposed are specific enough to be immediately implemented, acceptable to the unionized workforce and have minimal impact on production efficiency. Moreover, the improvements proposed should have expected results within 1-2 years to provide value for the company owners, yet promote long-term efficiency to provide value for long-time employees. This project outlines said improvements, and implements pilot programs to demonstrate efficacy.

1.3 Thesis Overview

In chapter 1, we provide background on the company and the problems experienced. We develop goals for the on-site project, and then give an overview of this corresponding academic thesis.

We describe in chapter 2 the detailed situation of the Marysville facility. We start by providing an overview of the roll-formed pallet racking industry, then lay out the broader corporate context where relevant to Marysville's situation. We then outline the manufacturing process at Marysville, and give a detailed description of the primary manufacturing equipment. We finally describe the existing maintenance system prior to this project.

We review the literature on maintenance systems in chapter 3. We cover Total Productive Maintenance developed by Toyota, Reliability Centered Maintenance developed by United Airlines, strategies involving Outsourced Maintenance and other

strategies for Performance Management in Maintenance.

In chapter 4, we develop a framework for analyzing the primary operation of a maintenance system at a facility like Marysville. We start by outlining the objectives of a maintenance system, and consider the relationship between maintenance and production. We then break down in detail the primary maintenance process into component steps: incoming tasks, diagnosis, planning, getting parts, executing work and observing effects. Finally, we enumerate several “wastes” in the primary maintenance process.

With this generalized framework in mind, we apply said framework to the particular situation of the Marysville facility in chapter 5. We enumerate problems in each step of the primary maintenance process, and analyze the root causes of these operational problems.

In chapter 6, we quantitatively estimate the costs to Heartland of the maintenance problems at the Marysville facility to underline the importance of the problem and provide a viable range for solution costs. We start by reviewing available data, and then quantify breakdowns in the Marysville facility. We then cross references breakdowns to bottlenecks to arrive at a conservative estimate of the cost of equipment breakdowns.

We propose five specific improvements to the maintenance system in chapter 7. We propose shared ownership of equipment between maintenance and production in section 7.1. Then in section 7.2, we propose more accessible documentation to improve diagnosis. Next, in section 7.3 we propose work orders to improve planning. In section 7.4 we propose proximal spare parts storage to improve getting parts. Finally, in section 7.5 we propose solving problems permanently to address incoming tasks.

We conclude in chapter 8 by summarizing the approach, suggesting further work and underlining the central idea of this thesis: efficient maintenance operations are an important complement to maintenance strategy.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Organizational Context

In this chapter, we give detailed background on the situation of the Marysville facility. In section 2.1, we describe the nature of the pallet-racking industry, which informs the design of the manufacturing process. We describe the situation of Heartland Steel Products in section 2.2, which gives important context for aligning maintenance with corporate strategy. We provide an overview for the Marysville manufacturing process in section 2.3, and then we enumerate major pieces of equipment in detail in section 2.4. Finally, in section 2.5 we describe the maintenance system prior to our interventions.

2.1 Industry Overview

The Marysville facility participates in the market for roll-formed pallet racking.¹ Roll-formed racking consists of *frames* and *beams* (illustrated in figure 2-1), and sometimes a number of less frequent specialty parts. The beams are assembled onto the frames at the customer’s facility by professional installers.

Roll-formed pallet racking is used in warehouses (e.g., Amazon fulfillment centers), factories (e.g., parts storage at Boeing) and some retail stores (e.g., Home Depot). It is generally the most economical shelving for moderate loads (hundreds of pounds),

¹So named because the components are cold-rolled from steel coil, in contrast to structural pallet racking which is hot-rolled from steel ingot.




Figure 2-1: An annotated picture of pallet racking product, which allows customers to increase areal storage density by storing inventory above the ground.

while wire racking dominates for light loads and channel racking dominates for heavy loads. The market for all kinds of pallet racking is highly competitive, and there is an active market for reselling both new and used pallet racking.

Roll-formed pallet racking frames have a customer-specified height, depth, and sheet metal thickness, as well as customer-specified footplates or various structural reinforcements. Roll-formed pallet racking beams have customer-specified length, cross-sectional height, and sheet metal thickness. Both frames and beams have customer-specified paint colors, or in some cases are galvanized to leave a shiny finish.

The cost structure (figure 2-2) of the product is dominated (50%-70%) by the cost of the underlying steel, procured from a competitive base of steel rolling mills. Direct labor (welding and painting) is a large cost. Variable overhead includes other materials (e.g., paint) and material handling, and is a large component. The fixed overhead (management, utilities) are quite small, as is depreciation of capital assets. Because prices depend primarily upon the amount of steel, customers typically order the lightest beams that can support their expected weight loading, so almost all product is made-to-order.²

²Technically, when a large corporation orders many of the same item to be shipped to several different facilities, the company “makes to stock” and fulfills orders from that stock. However, Heartland Steel Products does not begin producing such shelving until the order is received, so for the purposes of this thesis it is also made-to-order.




Figure 2-2: The cost structure of the product at Marysville. The majority of the cost is the underlying steel coil. Direct labor is a major component and target of cost cutting. Variable overhead includes other materials and consumables (e.g., paint) and material handling. Fixed overhead includes administration and maintenance costs. Depreciation represents the accounting cost of the manufacturing equipment used.

2.2 Corporate Context

Heartland Steel Products started as Eugene Welding Company in 1954, participating primarily in the market for custom steel pipes. It entered the market for pallet racking in the 1980s with the SpaceRak brand, and later exited the pipe market. Under previous ownership, the company acquired the facility in Harrison, Ohio in 2010 and the facility in Lodi, California in 2011. Previous ownership soured labor relations by eliminating the quality department and slashing maintenance costs. Most recently, Heartland Steel Products was acquired by a US private equity firm in 2017.

The Heartland sales organization is decentralized and incentivized with a revenue-sharing commission structure. The sales price is determined from a historical price book minus a discount at the discretion of the salesperson. The nominal price is computed primarily based on the required amount of steel at historical prices plus the total amount of welding. The discounts are negotiated based on volume, relationship, and competitive pressure, and can be quite significant. After the product is shipped, the charged prices are compared to the actual costs attributed to each job by a cost accountant. In some cases, the discount is so significant that the price only just breaks even over the actual costs. It is thought that much of this problem is caused by the salespeople having incentives for revenue rather than profit, leading to excessive sales *volume* rather than *margin*.

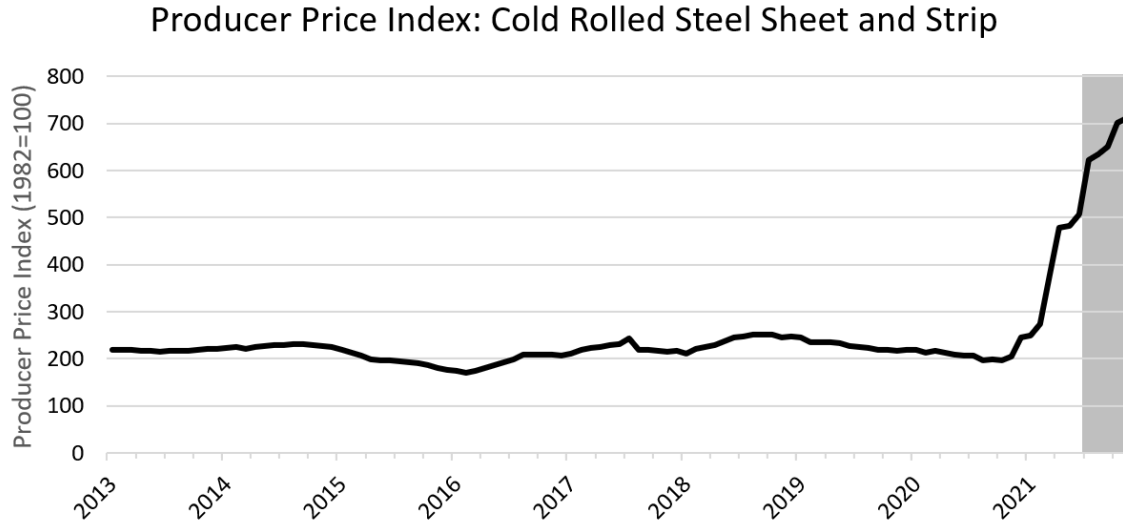


Figure 2-3: Market price index for cold-rolled steel from 2007 until 2021, showing the rapid price increase in 2021 caused by pandemic supply chain failures. The duration of the research is shown by the grey highlighted region. Data is retrieved from U.S. Bureau of Labor Statistics [10].

The business is heavily exposed to the raw material costs of steel. Until mid-2021, the company quoted the customer a fixed price, which was dominated by the cost of raw material steel. In 2021, the price of steel surged much more than ever observed in recent history (figure 2-3), forcing the company to purchase steel at a much higher price than it had quoted the customer. This significantly eroded profitability until about July 2021, at which point the CEO insisted that quotes include a clause where significant increases in the cost of steel could be passed onto the customer. This resulted in a business context where little cash is available and management is focused on aggressive cost cutting.

2.3 Manufacturing Overview

In Marysville, the production process for frames and beams depends on the specific order, but generally can be simplified to a few key processes shown in figure 2-4 that represent most of the time and cost. Coiled steel specific to the order is purchased about four weeks prior to the requested ship date. To produce *frames*, mounting holes are punched into the steel coil, the coil is rolled into a channel shape, and then braces

and footplates are welded onto the posts by manual welders. To produce *beams*, the coil is rolled and seam-welded into step-tube, and then the end-connectors are welded on by automatic beam welders. Finally, *frames* and *beams* are painted on a conveyor system, then packaged to be shipped to customers.

The Marysville facility is unionized with the International Brotherhood of Teamsters. The collective bargaining agreement is due to be re-negotiated in February 2022, about three months after the conclusion of the internship. Labor relations are generally strained, primarily because of decades of extreme cost-cutting measures such as layoffs of entire departments. More recently, the COVID-19 mask mandates strained relationships further, and the union feels their requests are not sufficiently prioritized. The demographics of the hourly workforce in Marysville are different from the other facilities. The entire hourly workforce is male (several women work in management), and in recent history women only stay briefly before leaving. The hourly workforce is largely but not exclusively white, with some Hispanic and some Black workers. The age distribution is bimodal, with one mode approaching retirement while another mode is in their 30s and generally started 3-5 years ago.

Inside the factory environment, most surfaces are covered in a layer of dust, which is primarily composed of an iron oxide spatter from the MIG welding operations but also partially particulate smoke from the paint oven and dirt blown in from outside. However, the floor is well swept. The factory is quite dark because most surfaces are either unpainted or painted dark colors, and the dust also absorbs a lot of light. The factory is not temperature controlled, so in the summer it can exceed 100F near the oven, while in the winter it can be well below freezing. In a few select areas, water ingress from outdoor precipitation can get inside the factory.

2.4 Details of Manufacturing Equipment

Marysville has a large array of varied equipment; in this section we cover the four most utilized, and most critical, types of equipment that production depends upon. In the descriptions here and in section 2.5, we present the prior state, before the

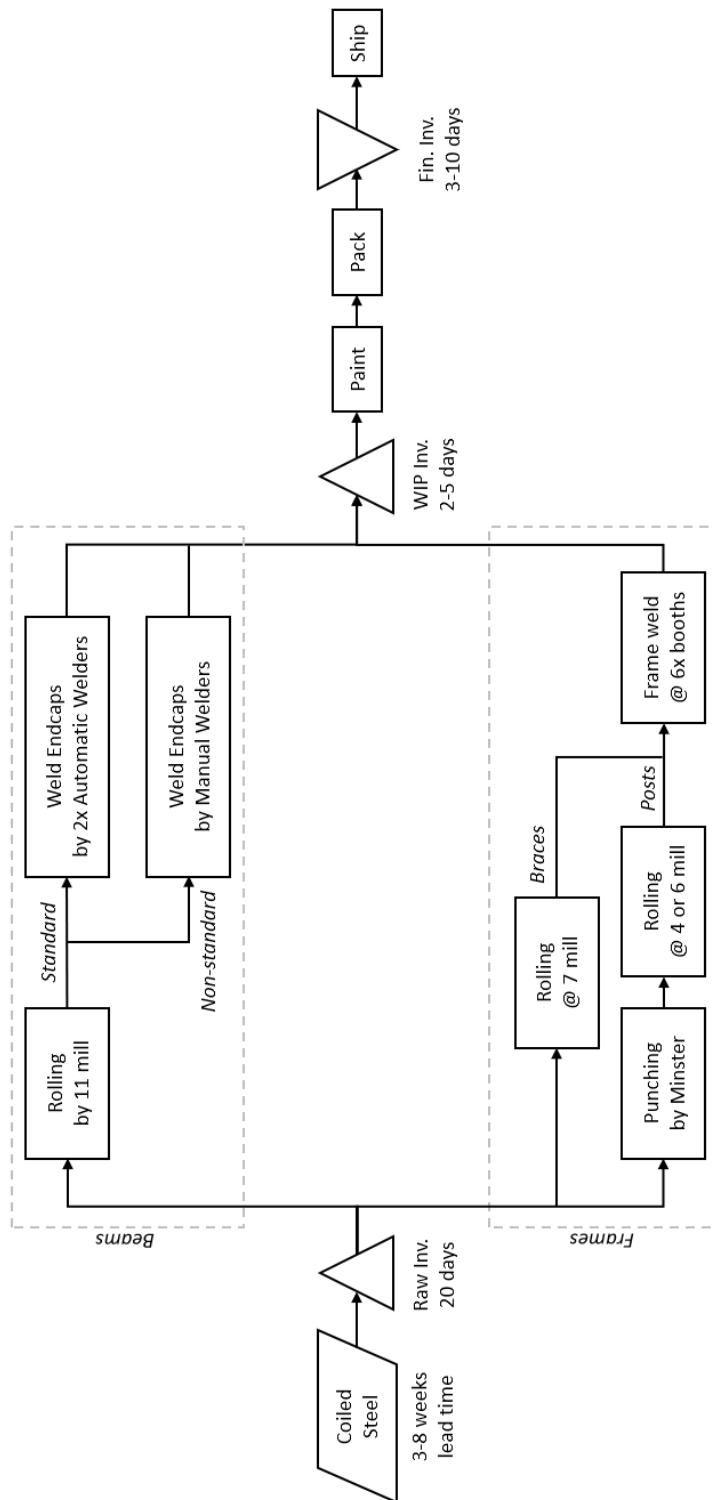


Figure 2-4: The roll-formed steel shelving process at Marysville. For simplicity, excludes a variety of equipment used to produce footplates and various structural reinforcements.

interventions proposed and summarized in chapter 7.

2.4.1 Press

Marysville has one reciprocating mechanical press, responsible for punching the mounting holes in the coils that eventually become posts. The mechanical press uses an electrical motor to drive a flywheel and crankshaft, which in turn punches through the steel coil below. Additionally, there is a servo motor that advances the steel coil the correct amount each time the punches retract. The mechanical press is from the 1970s, having recently replaced a World War II era mechanical press, but the tooling does not incorporate tooling advances made in the 1950s.

Because the press has been recently installed, the issues experienced are primarily resolving mistakes made or damage done during installation. The press has several issues with lubrication and with the reliable control of the servo.

Finally, the press has tooling that wore out rapidly, and more so because the tooling is not designed with common designs like a stripper plate.³ Therefore, the press has tooling that needs to be sharpened daily.

2.4.2 Mills

The mills are responsible for bending coiled steel into a shaped profile that is rigid enough not to buckle. The mills are a large system composed of a linear arrangement of subsystems. Each mill has a decoiler that rotates the coils of steel to release one thickness at a time, and then a station where the end of one coil is welded to the start of the next one to speed up changeovers. Then, each mill has a roll-forming mill, that uses rotating tooling to bend the coiled steel while it rolls through. One of the mills produces a closed profile and has a subsystem that would produce a seam weld where the previous loose edges met. Then each mill has one or more “turksheads,” that can be manipulated to straighten out small bends or twists in the outgoing profile. Then,

³A stripper plate serves to clamp the material and laterally support the punches to reduce lateral loads on the tooling.

each mill has a cutoff press that cut the steel at a pre-determined length, and finally, a station where the outgoing material can be stacked for transportation.

Many of the components on the mill experience significant abrasive wear. The cutoff tooling needs to be sharpened every few days. The roll-form tooling needs to be completely re-profiled once or twice a year. Other internal components wear out eventually because of continuous motion, most notably the bearings. Wear is exacerbated by the presence of “mill-sludge,” which is a thick, dark accumulation of lubricant oil, oxides from the steel, slivers of steel and other dirt; the sludge acts as an abrasive that reduces the lifespan of components.

The mills are powered by electric motors, some hydraulics, and some pneumatics. One of the mills is controlled by a PLC, while the other mills use an in-house hodge-podge of off-the-shelf controllers. The only digital sensor on each mill is the encoder that measures the output length and instructs the cutoff press to activate.

No one on site understood how these systems work; the operators are given extremely specific instructions (press button A once, button B six times, then button A again), maintenance is limited mostly to looking for loose wires, and anything more complex requires a contractor to come in. A common issue is that the operator performs the instructions wrong and does not know how to recover the machine.

The operator frequently adjusts the configuration of the tooling by tightening or loosening heavy duty screws in the arbor blocks supporting the axles. Because of the aforementioned sludge, this requires significant torque, and in some cases the torque causes solid load-bearing components to break.

Finally, the electronics requires air cooling to dissipate waste heat. In particular, the seam-welding operation has a high frequency transformer that generate substantial waste heat. The electric motors and their drivers also generate substantial waste heat. Because of environmental dust and biological fouling (anemochoric seeds from local trees or birds nests), many of these cooling systems are not getting the airflow and heat transfer that they require, leading to premature electronic failure.

2.4.3 Welding

Welding consists of nine manual booths and two automatic beam welders. Manual and automatic both use commercial MIG welding units, which consist of a power supply, a wire feeder and a gun that the welder or machine holds. The welders are fed gas from a central tank of CO₂ and of Argon, and are fed wire from large bins of coiled MIG wire. The manual booths are operated by people who lift components onto the table to be welded, then who eventually remove the finished product with a small crane. The automatic beam welders load and unload material with pneumatics and move the welding guns vertically using a leadscrew. All of the above is coordinated by a PLC, which is informed by a large number of inductive proximity sensors.

The power supplies contain transformers that generate waste heat, and because of the dust, many of these power supplies are not getting the cooling that they require. Therefore, a common problem is the solid-state relays failing due to chronic overheating. Another common failure is the wire feeders failing, either because of wear from accumulated dust or because frustrated operators press buttons too hard.

The cranes are electrically powered and operated by a battery-powered remote control. Because they are constantly in motion and flexing the power cables back and forth, the wires eventually come loose, causing sparking. The remote controls eventually run out of batteries and need replacing.

The automatic beam welders use inductive proximity sensors to determine when to actuate various functions, such as dropping in a new end connector or beginning the weld. Many of these inductive sensors are mounted on thin brackets that vibrate violently when the abrupt pneumatics accelerate or decelerate them quickly, and therefore, the inductive sensors frequently require adjustment or replacement, typically because the mounting nuts shift or because the connecting wire breaks. In other cases, unfinished parts or other parts of the machine impact the sensors and destroy them.

Also, in the automatic beam welders, the leadscrews are rotated by an electric motor at a speed controlled by the operator. However, this system was designed

before widespread use of servo motors, so the speed is governed by the electric motor resisted by a solenoid-actuated brake pad. Specifically, the motor always receives full power from the PLC, but the operator can change the amount of current that goes to the solenoid. This in turn adjusts the normal force on the brake pad, which changes the frictional torque, which creates a new equilibrium speed for the motor-brake system. However, the brake pads wear thin over time, and the friction changes with temperature, so the operators are constantly adjusting the current to the solenoid to try to maintain the same speed. The operators spend significant energy and create significant scrap trying to adjust for constant speed, and sometimes their adjustments create so much friction that the motor cannot move, which leads to the motor using too much current and blowing a fuse that needs to be replaced. Eventually, the brake pads are replaced when they wear too thin.

Finally, the end connectors are loaded in through a steel track that they ride upon while being pulled by pneumatics. The end connectors are unlubricated, so over time they wear away at both the track and the pneumatic puller, causing both components to operate less reliably.

2.4.4 Painting

Painting utilizes an overhead conveyor to carry beams and frames through the various stages of painting. As represented by figure 2-5, the parts are carried through several machines to be cleaned, painted and cured. The acid washer, rinse washer, dryer, and oven are all heated using natural gas burners, each of which is controlled by a commercial burner controller. The wash and rinse have tanks of solution from which they pump water through a series of pipes with various nozzles to spray the parts. The oscillating paint sprayer uses an electrostatic rotary sprayer that oscillate up and down through a mechanical linkage.

The most common issue is for the nozzles in the wash or the rinse to get clogged up, and then sometimes pop off entirely. This ultimately happens because the phosphoric acid forms a sludge of metallic precipitates that are pumped through the pipes and collect at the narrowest orifice. This requires turning off the washer, and having a




Figure 2-5: Parts to be painted are loaded onto an overhead conveyor by workers. Parts go through a phosphoric acid washer, which removes dirt and rust, and promotes later paint adhesion. Parts then go through a rinse washer, which removes residual detergent. Parts next proceed through a dryer, which removes excess water and warms the parts. Parts are coated by automatic paint sprayers, which use electrostatic attraction between the parts and water-based paint droplets to deposit primarily on the part. Unpainted spots are touched up by manual paint sprayers, which spray a wide fan-shaped blade with compressed air. Parts are then baked in an oven, which evaporates water from the paint and cures the polymer. Finally, parts are unloaded by workers and packed for eventual shipping. The paint will continue to cure over the next 24 hours.

person go inside to replace the nozzle. Alternatively, sometimes frames or beams in the washer swing and catch one of the PVC pipes. Because the PVC is operating at about 60C, at which temperature it is much weaker than at room temperature[4], the PVC pipes break under the load and need to be replaced.

Another issue in this system is various electronic parts failing due to chronic overheating; they are next to hot components with no effective cooling. This is particularly noticeable in the summer months.

2.5 Existing Maintenance System

To keep all of the above equipment running smoothly, the maintenance team in Marysville has personnel, facilities, inventory, and information available.

From a personnel perspective, the Marysville maintenance team has nine full-time employees:

- 6 maintenance technicians
- 2 machinists
- 1 maintenance supervisor

The maintenance technicians are union employees, typically with high-school education. Technicians have varying levels of pay and rank based on the skills that they possess. The primary responsibility of technicians is to execute the repair tasks their supervisor gave them, and indeed the technicians are the only employees formally permitted by the Collective Bargaining Agreement to execute repair tasks. The technicians are also responsible for organizing the spare parts rooms detailed below. The most senior technician work 3:00-13:00. The three next most senior technicians work 5:00-15:00 on the day shift, and the most junior technicians work 15:00-1:00 on the night shift.

The maintenance technicians are trained by shadowing another technician for about three weeks. During those three weeks, they observe and help their mentor perform any repair work that arose. In practice, the three weeks of training focus heavily on reactive maintenance work on the machinery that commonly need repairs. There is no structured training in how to diagnose issues or plan repairs, and training does not cover spare parts inventory management.

Before the interventions in chapter 7, the maintenance technicians did not all get along with each other. The two most senior maintenance technicians did not believe the junior technicians could stick to established processes, while the four more junior technicians believed the senior technicians were lazy and did poor quality work. Even within the four more junior technicians, one of the most junior technicians was very frustrated with the combative attitudes of both junior and senior technicians, and planned to leave the maintenance team as soon as he was able to.

The machinists are also union employees, with an associate's degree. Both machinists are quite new to the company. The primary job of the machinists is to sharpen the cutting tools; they spend about half their time sharpening the tooling for the mechanical press, and most of the remaining time they spend sharpening cutoff tooling or rotary tooling for the mills. In the remaining time, they also machine some custom spare parts. Sometimes machine parts are "rushed" because they are needed in the middle of a repair job.

The maintenance supervisor was recruited from an hourly, union maintenance

position but is no longer part of the union. The supervisor performs much of the support work for the technicians: retrieving parts for them, sourcing parts from external vendors, and coordinating with external contractors such as electricians. The supervisor also observes maintenance work, assigns tasks to the technicians, and reports to plant management. The supervisor aims to be liked by each maintenance technician.

From a facility perspective, the maintenance team looks after machinery in all five of the major Marysville buildings shown in figure 2-6. The maintenance team have exclusive use of the machine shop in the southern group, as well as two workshops and three spare parts rooms spread across the campus.

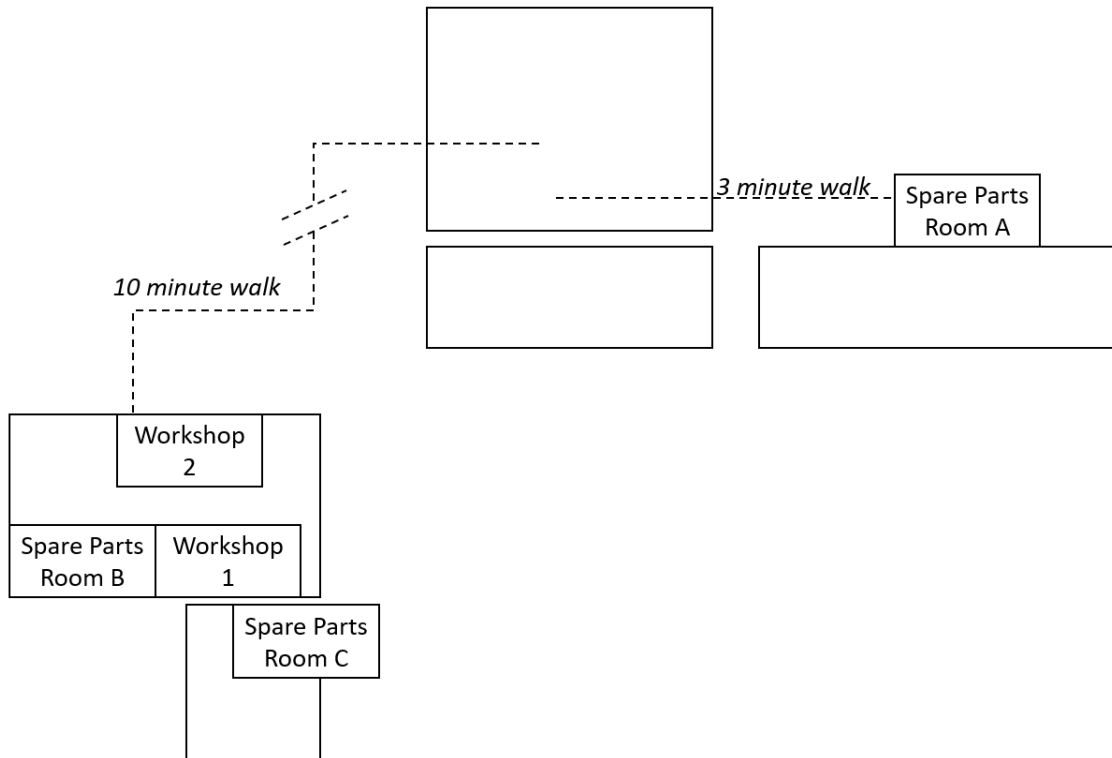


Figure 2-6: There is a northern group of three buildings, and a southern group of two buildings about a 10 minute walk away. Spare Parts Room A is in the northern group. Spare Parts Rooms B and C are in the southern group, along with both workshops.

The vast majority of employees and production equipment is located in the northern three buildings, including all the roll-form mills, all the welding, and the paint line. A few employees work in the southern buildings, primarily with the press and other

ancillary processes, as well as a large amount of raw materials inventory. Therefore, the majority of the maintenance workload is located in the northern three buildings.

The maintenance team operates three spare parts rooms across the campus. Each spare parts room is nominally common, but each has a de facto owner who made decisions about organization in that space. Spare Parts Room A is about 1500 sq. ft. and houses unlabeled shelving with parts haphazardly placed and generally dusty. Spare Parts Room B is about 1500 sq. ft. on two floors and houses labeled and numbered shelving with parts reasonably well organized, but many of the parts are obsolete or meant for machines that have been sold off. Spare Parts Room C is about 100 sq. ft. and contains unlabeled machined parts. Rooms B and C are about a 10 minute walk away from the majority of maintenance work. Prior to the improvements in chapter 7, there was no paper or digital system to track the inventory levels of any spare parts, nor what parts were supposed to be present.

The maintenance team operates two workshops where the maintenance technicians has space to perform ancillary repair processes. For some repair processes, the maintenance team will transport a small assembly to the workshop and work on the repair there. The workshops also function as storage for tools that the technicians might need for repair work.

Finally, the machine shop is the primary workplace for the two machinists and houses about 20 distinct machines for sharpening tools and producing machined parts. Tools are sharpened on one of five grinding machines, or the CNC lathe; for most tooling, the machinists sharpen one set of tooling while the other set of tooling is being used in production. Machined parts can be produced on one of the four vertical mills, the lathe, or some of the more specialty equipment.

From an informational perspective, the maintenance team keeps some paper-based information about the facility's equipment in Spare Parts Rooms A and B. About half of the machines have no information available at all. A few have comprehensive manuals available. Some of the most critical machines have binders that contain the printed datasheets of every off-the-shelf part in the machine and a printout of the entire PLC program (typically about 50-80 pages). Notably, these binders are

missing operating instructions and are missing mechanical drawings of the machined components.⁴

⁴Missing documentation is no longer available because the OEM went out of business circa 2005.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Literature Review

In this chapter, we seek to review some successful philosophies about maintenance management from the literature. These maintenance management philosophies were generally developed at large manufacturers with great success. However, some adaptation might be required for success at a smaller manufacturer.

Large businesses tend to have large maintenance organizations, in some cases hundreds or thousands of employees performing maintenance. Generally, there are several levels of hierarchy: a maintenance director, several maintenance managers, several maintenance supervisors, and many maintenance technicians. In addition to technicians, there is an additional group of support staff. Many organizations will have a dedicated planner to each support 20-30 maintenance technicians. They will also have dedicated personnel to manage spare parts inventory and will delegate much of their spare parts purchasing needs to the company's purchasing department [11].

By comparison, a small facility like Marysville has less than ten employees performing maintenance, and only one level of maintenance supervision. There are no dedicated planners nor dedicated spare parts personnel. Because of these differences, the below philosophies will require some adaptation to be applied to a facility like Marysville.

3.1 Total Productive Maintenance

Total Productive Maintenance (TPM) is a key pillar of the Toyota Production System (TPS), and it is to maintenance as Lean Manufacturing is to production [9]. The overarching philosophy is to create a partnership between maintenance and production, aimed at the ambitious goal of 100% availability, 100% performance and 100% quality.¹ It asks operators to share ownership over their equipment, performing some first-line maintenance when they are able. It creates a team including production workers from an area (workstations with related functions) and including relevant maintenance personnel, and gives the team authority to direct second-line maintenance work and improvements within limits. In advanced implementation, the team can even have input into hiring and capital expenditure [8]. However, Ichniowski et al. show that an organization must make some assurance of employment security to secure cooperation from employees [5].

Kramer [7] relates a case study about implementing Total Productive Maintenance as a maintenance turnaround. Kramer notes that urgent maintenance has become an ingrained part of the culture, for several reasons:

- No one else is trusted, which is a dignity for mechanics
- The goal is no more ambitious than “keep it running”
- Recurring failures are fixed but not root-caused nor communicated to others
- Extraordinary heroism to fix equipment is recognized, but simple planned maintenance is rarely recognized
- Operators don’t help, partially because they don’t want to be responsible for old, dirty junk equipment
- Attempts towards planned maintenance weren’t given funds, time, or slack from management

¹Availability refers to how often the machine is operational when it is requested, performance refers to how much product is manufactured compared to the standard rate, and quality refers to the fraction of output product that is acceptable.

- No training for maintenance and little cross-training
- Worry that losing the fire-fighting² will impact job security
- No engineering support to improve the design side of reliability

Kramer suggests that TPM implementations set a mission for maintenance that goes well beyond “just keeping it running,” initially boosting morale with outsourced cleaning; giving operators ownership and accountability for the performance of their equipment; providing the tools, parts, funds and support for first-line maintenance; and finally locking in the change with training to elevate the employees’ pride in their new roles.

TPM was previously understood to be difficult to implement in union environments, because since WWII [16] unions have generally resisted changes to job content such as sharing ownership of maintenance. Indeed, Kramer [7] notes that in 1990 GM wanted to empower its production operators to perform first-line maintenance, but its union (UAW) was unwilling to authorize such a provision. However, by 2007, UAW had relented and agreed that operators should perform first-line maintenance [1]. This shows that unions are reluctant to accept such a change of job responsibility, but in time they can recognize the benefits of a successful TPM program to the workforce.

Implementations of TPM typically promote a focus on preventative and aggressive maintenance. Swanson performed a large-scale survey factor analysis, and found that aggressive maintenance (focusing on improving equipment from design through purchase through operation) was correlated most with improved quality and reduced costs, and proactive maintenance (focusing on monitoring equipment, analyzing data and predictive maintenance) was correlated with improved availability, while reactive maintenance (focusing on restoring equipment) was correlated with worse performance in quality, availability and costs [14]. Feliciano found that downtime of spar machines at Boeing was correlated with overtime expense, and showed through simulation that doing so would significantly increase output of wing spars [3].

²Kramer uses this term to refer to urgent repair work, which is often where maintenance technicians receive the most praise.

3.2 Reliability Centered Maintenance

Reliability Centered Maintenance (RCM) is a prioritization technique for maintenance. The overarching philosophy is that maintenance is performed to maximize the reliability of systems. It is specifically developed to eliminate preventative maintenance that does not change, or in some cases decreases, the reliability of the system.

RCM excels in very complicated systems. RCM starts with a Failure Modes, Effects, and Consequences Analysis. FMECA involves estimating the failure rate of every asset, then what fraction of each asset's failures are due to each specific root cause, and finally the probability of business consequences arising from each root cause [2]. For example, if a machine has a 2% chance of failure, 60% of such failures are due to the widget breaking and a broken widget has a 30% chance of leading to a shutdown, then that consequence gets a score of $0.02 \cdot 0.6 \cdot 0.3 = 0.0036$ (to which one can add some measure of the importance of the consequence or utilize less-precise math to reflect uncertainty in the estimations). These scores allow the maintenance team to focus their efforts.

Tsang notes that RCM is a high-overhead analysis, particularly if there are many failure modes and they are not well understood. Tsang argues RCM will be most effective for complex and high-risk systems (e.g., where the primary risk is not loss of production but loss of life) because day-to-day personnel might struggle to weigh the relative consequence of different failure modes. Tsang also argues it will be effective where there is grossly over-maintained equipment by revealing which maintenance procedures are not justified by the consequences [15].

Kelly [6] explains why more efficient maintenance results in lower costs and better reliability. Kelly uses an illustration to show that the optimal level of maintenance is where the marginal cost of maintenance equals the marginal benefit of reliability. Kelly therefore argues that if maintenance activities are poorly chosen or resources are used inefficiently, the marginal cost of maintenance increases, so reliability decreases and costs increase.

3.3 Outsourced Maintenance

Maintenance does not necessarily need to be executed entirely in-house; manufacturers can outsource some or all of their maintenance work. At one extreme, a firm can keep all maintenance activities in-house. This carries several drawbacks including having inflexible capacity and requiring skills and capabilities scantily related to the core business. At the other extreme, a firm can essentially outsource all of its maintenance. This has several drawbacks including lack of responsiveness, misaligned incentives, and slowly losing key organizational knowledge about the machinery. Most firms choose some level in the middle.

Tsang [15] reviews several aspects of an outsourced maintenance relationship. Tsang notes that the contractual relationship can be structured in terms of a *work-package contract* (a specific set of activities), a *performance contract* (a specific level of performance from the equipment) or a *facilitator contract* (where the maintenance contractor owns the physical assets).

Ye [17] relates a case study of a unionized paper mill that used a sudden step-change in outsourcing maintenance to transform an uncooperative maintenance workforce. The paper mill had a maintenance team billing excessive overtime hours, despite doing little valuable work, and management had initially reacted to the rising maintenance costs by cutting preventative maintenance, which led to higher breakdowns and a return to higher maintenance costs. Therefore, management signed a contract for ABB to provide full-service maintenance work with a profit incentive towards higher equipment performance (this was unsuccessfully challenged in arbitration by the union). After an initial decline in productivity, which was partially masked by a labor strike, the firm recovered performance at a much lower maintenance cost. Interestingly, maintenance workers had been so focused on extracting overtime hours that they were not *happy*, and morale actually improved once the overtime was removed.

Kelly [6] recommends that maintenance workload be split up into first-line, second-line and third-line. First-line consists of minor or unplanned maintenance tasks during

the week. Second-line consists of corrective actions, inspections and scheduled services during the weekend. Third-line consists of major corrective actions and major scheduled services during an annual shutdown. Kelly recommends staffing for all of the first-line and most of the second-line maintenance, and hiring flexible contract labor to handle remaining second-line and third-line maintenance tasks.

3.4 Maintenance Performance Management

Managing the performance of a maintenance team is a difficult task for many businesses because the health of the equipment is difficult to measure; failure data is messy and sparse, particularly when equipment is operating well, and the space of outcomes is complex and hard to characterize. Unlike production where the outcome is a number of units, perhaps with some quality dimensions, the outcome space of maintenance is an almost infinite number of equipment defects.

Parida [12] categorizes all Maintenance Performance Indicators (MPI) into the following categories:

- Customer (is the end customer satisfied with the product?)
- Cost (direct costs? labor costs? costs to production?)
- Equipment (how is the equipment performing? is there an industry-standard metric for the equipment?)
- Tasks (completion of preventative maintenance tasks?)
- Learning (skills developed? knowledge documented?)
- Health, Safety and Environment (safety incidents? pollution levels?)
- Employee satisfaction (how do mechanics feel? how do production workers feel?)

A widely-used MPI is Overall Equipment Effectiveness, defined as:

$$\begin{aligned} OEE &= Availability \times Performance \times Quality \\ &= \frac{\text{Utilized Time}}{\text{Scheduled Time}} \times \frac{\text{Actual Rate}}{\text{Expected Rate}} \times \frac{\text{Good Units Produced}}{\text{Units Produced}} \end{aligned}$$

OEE is the main objective of typical TPM implementations, and teams are charged with maximizing OEE, approaching 100% in the long-term.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Framework for Analyzing Maintenance

This chapter develops a framework for analyzing a maintenance system at a small manufacturer like the Marysville facility; the framework is intended to be generalizable so it can be applied to any similar manufacturer with inefficient maintenance. This chapter starts by outlining the objectives of a maintenance system: preventing breakdowns, resolving breakdowns and aligning with production's needs. Then this chapter breaks down the primary maintenance process that largely determines the effectiveness of maintenance, and briefly reviews ancillary maintenance processes. Finally, this chapter enumerates several "wastes" in the primary maintenance process that can dramatically reduce maintenance efficiency.

4.1 Objectives of a Maintenance System

Most manufacturers have capital equipment that needs to be repaired in case of failure, therefore having an effective maintenance system is an imperative. The maintenance system is generally responsible for maintaining, recovering, or improving the productivity of capital equipment, but the specific goal of a maintenance program is contested between maintenance philosophies as discussed in the literature review.

However, we make a pair of observations that are valid within any maintenance

philosophy and may seem obvious but are key to decomposing the maintenance problem. First, all else being equal, the business always benefits from not needing to do a given maintenance task in the first place; maintenance is inherently a non-value-added activity from the business's perspective, because customers do not derive value from maintenance being performed. Second, all else being equal, the business always benefits from finishing a given maintenance task more quickly and efficiently. With this in mind, we arrive at an intuitive decomposition of the problem: (a) how to minimize the volume of maintenance tasks that need to be done, (b) how to streamline the maintenance process to finish all tasks quickly and efficiently, and (c) how to balance priorities and resources among several tasks. Any maintenance system must address these three problems.

Considering problem (a), a maintenance system is more effective if it can prevent problems from occurring in the first place, before work is needed. A business is more profitable and more rewarding for its employees if machines simply continue to work productively and safely rather than breaking down. Hence, businesses employ preventative maintenance to eliminate problems. This has the dual benefit of keeping production running smoothly and producing value-added work for customers, as well as eliminating work and cost from the maintenance system. However, businesses face the trade-off of investing initial labor or expenses to eliminate future problems, or more often, to reduce the likelihood of future problems. Even in cases where the benefits clearly outweigh the costs, businesses are often unable to make the initial investment because of constraints on available time or cash.

Regarding problem (b), a maintenance system is more effective if it can resolve problems more quickly and efficiently. For example, the business is more profitable and rewarding for its employees if machines can be restored to operating conditions more quickly, so the business can return to value-added work for customers. The business is also more profitable if machines can be restored using fewer labor hours and lower material costs. Although it is important to note that if fewer labor hours translates to smaller paychecks, this may not be more rewarding for the maintenance personnel. Another difficulty is that speed and efficiency are commonly a trade-

off. For example, increased technician capacity typically allows for faster responses to broken machinery but decreases efficiency through higher labor costs. The optimal balance between repair speed and efficiency depends upon the relative cost of downtime against the cost of maintenance resources. But most importantly, win-win solutions that improve both speed of the repair and the efficiency of resources are always good for the business.

Addressing problem (c), a maintenance system is more effective if it balances priorities and resources in alignment with the needs of production. This is ultimately an optimization problem, allocating limited resources to achieve the best outcome for production. An effective maintenance system strikes a balance between tasks that benefit production immediately and tasks that benefit production far in the future, while postponing those tasks that can be postponed without consequence. In some businesses, this is an optimization problem that can be tackled quantitatively. However, at Heartland, these decisions are generally made by individuals using unwritten experience. Due to lack of data, this thesis can only give high-level guidance about this problem.

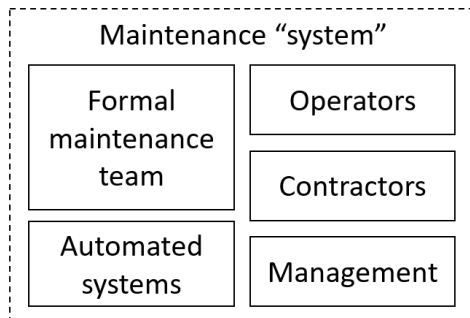


Figure 4-1: The maintenance system encompasses not only the formal maintenance team, but also operators, contractors, automated systems and management.

Finally, it is important to note that a maintenance “system” is necessarily more broad than just the employees that are formally on the maintenance “team,” as illustrated in figure 4-1. The actions of several external groups have an enormous impact on the effectiveness of the maintenance system. Perhaps most important are the actions of production: does production operate machinery as recommended, or do they

push it harder? Do all operators operate the machinery correctly? Does production prevent maintenance issues from occurring, or do they resolve some by themselves?

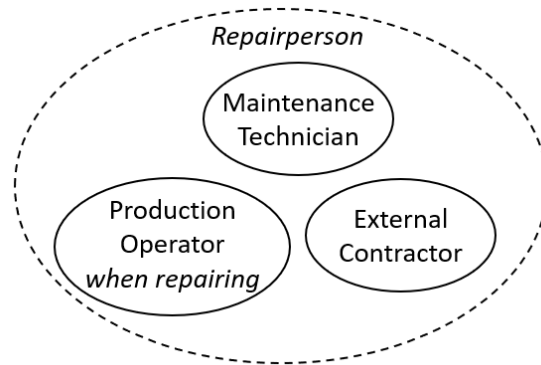


Figure 4-2: Repairperson could refer to a maintenance technician, but also to a production operator when repairing equipment or to an external contractor.

Also, within the maintenance system are several actors that complete work but are not part of the team: external contractors that are brought in temporarily and some automated systems like automatic lubricators. And finally, plant management can play a role in prioritizing maintenance tasks, as well as assigning additional work (e.g., move a machine from location A to B) or sometimes borrowing resources (e.g., maintenance technician filling in for an operator). The technicians and supervisor play a key role in the system, but their role is only a part of the overall system effectiveness. Therefore, in the remainder of this thesis, we will use repairperson as a term to refer to not only maintenance technicians formally assigned to the maintenance team, but also to anyone involved in making a repair as illustrated in figure 4-2. This can include production operators while they are performing repairs and/or external contractors.

4.2 Relationship to Production

In all factories, maintenance and production must work closely together to succeed, but it is worth noting that only production creates value for customers. Therefore, maintenance should conceive of production as its own customer: maintenance should always aim to help production succeed in the long term.

Production absolutely relies upon maintenance for day to day operations. First and foremost, the equipment that production depends upon to manufacture goods is kept operational by maintenance. If the equipment breaks down, and particularly if it is not promptly repaired, then production will find it near impossible to meet its production goals.

Second, production relies upon its workers, and those workers can easily be frustrated by difficult machinery. If equipment is difficult to use, constantly requires attention, or causes workers to be disciplined by management, workers will lose morale and either be less effective or leave the company. Moreover, poorly maintained equipment may become unsafe, which will further sap morale and could lead to injuries. Therefore, production relies in part upon maintenance to maintain morale.

Third, production relies upon the predictability of its equipment to operate efficiently. By this we mean that the long-term accumulated cost of unpredictably losing production time is much higher than simply the lost profits or the overtime costs incurred. Consider a factory that faces significant unpredictable downtime. This factory has two choices: (a) accept significant and unpredictable delays in manufacturing or (b) build up buffers to protect against unpredictable delays. The first option disappoints customers, and could erode the price premium the business might have commanded due to customer loyalty or could reduce sales dramatically. The second option requires keeping significant WIP inventory so that the business can continue producing while broken machines are repaired. This leads to the costs of excessive WIP inventory highlighted by the literature on lean manufacturing: aside from the direct costs of holding the inventory, excessive inventory also conceals inefficiencies and creates confusion. Furthermore, where a reliable business might be able to accurately optimize the schedule of labor to minimize wasted labor and overtime, an unreliable business cannot achieve those optimizations.

Therefore, production demands from maintenance a set of machinery that is operational as much as possible, easy to use, and predictable.

On the other hand, production can take actions to help maintenance succeed, which in turn helps production to do better. In particular, production can do its

best to minimize the amount of work that maintenance needs to do by operating machines correctly, safely, and cleanly. Production can also contribute directly to the maintenance effort, by providing accurate diagnoses, and sometimes by performing simple maintenance work. In some cases, production can complete entire repairs themselves, while in other cases they can help maintenance get access to the relevant subsystems or simply lend a helping hand lifting, holding, or observing.

4.3 The Primary Maintenance Process

At the company, the maintenance system is responsible for a number of interrelated activities. The system's core activity is repairing or improving the machinery that production depends on to create valuable products. However, the distinction between "repair" and "improvement" can be fuzzy. In many cases, the machinery is not operating at the manufacturer's promised productivity, yet after decades of underperforming, the expectations of the machinery have gradually relaxed to match the lower productivity state. Improving the productivity of such machines requires engineering expertise somewhere in between traditional repair and improvement. However, for the sake of simplicity, we will refer to the above work as "repair."

Although repair work is highly variable depending on the machinery to be repaired and the nature of the damage or defect, at a high level of abstraction all repairs follow roughly the same operational process shown in figure 4-3. The process starts with a stock of known tasks to be accomplished, and the team can prioritize in which order to work on these tasks and which tasks it can handle simultaneously.

The first step for the team is to diagnose the cause of the issue and formulate a plan to fix the issue. This might be trivial in the case of simple replacements, or might require significant exploration, deduction or even guesswork. The second step is to gather the materials needed to accomplish the planned repair work, perhaps from the on-site spare parts room or perhaps from an external vendor. The third step is where the technicians are doing physical work on the machine, such as replacing parts, cleaning, or modifying. In the literature, this step is affectionately known as

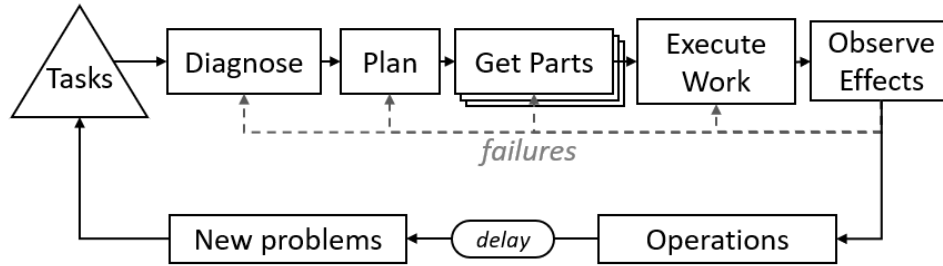


Figure 4-3: The repair process starts with a buffer of incoming tasks. The maintenance system must develop a diagnosis, then develop a plan. Then, the maintenance system must procure all necessary parts, before executing the physical work. Finally, the maintenance system must verify the solutions as successful and observe any follow-up problems.

“wrench time” and is typically the task most closely associated with the role of the maintenance system. The final step is evaluating the effect on the machinery, which could consist of a brief test of the equipment, or following up with the operators after some time, or in some cases might be skipped entirely.

With this in mind, the value created by the maintenance system depends upon (a) the value and urgency of the tasks that the team accomplishes, and (b) the speed and efficiency with which the team can move through this process.

4.3.1 Incoming Tasks

In the incoming tasks step of the process, the maintenance system collects tasks that need to be completed to support production. The desired outcome is that arriving tasks are reliably handled and that urgent tasks are handled appropriately quickly.

This part of the process is a richly studied area because in many businesses there is room to prioritize and schedule maintenance tasks more efficiently. However, in this project we consider a business struggling with effective maintenance, and one effect of their struggle is that the maintenance team was operating at maximum capacity just to handle the urgent tasks, prior to our interventions. So, we will not dive deep into maintenance task prioritization in this section.

However, there is another implicit aspect of the incoming tasks that makes a major impact, and that is the nature and rate of the incoming tasks. The design of

the maintenance system can have a significant impact on the incoming work that the maintenance team proper is asked to perform and thereby affect the performance of the maintenance team and the maintenance system as a whole.

First, some issues are simply not reported at all. They can be fixed silently by production, and this has the effect of reducing the workload on the maintenance team proper. This is a huge improvement for a maintenance team that is struggling to keep up with urgent tasks. Production operators are mechanically competent to some degree, and they are certainly capable of performing simple and surface-level tasks. In some cases, they interact with the machinery so much more frequently than maintenance that they might have a better mental model of the machinery. On the other hand, an antagonistic production operator can create a headache for a maintenance team by demanding help with extremely simple tasks.

Second, urgent and unexpected issues are inherently more difficult to deal with than non-urgent issues. Urgent issues generally require technicians to abandon work that they are already in the middle of. Such interruptions incur large frictional costs: travel from prior work to the urgent task, time getting oriented, and time gathering materials. Interruptions also affect the morale of the technicians by making them feel like they are not getting things done. Urgent issues can be minimized by preventing them in the first place and also by implementing contingency plans to forestall their urgency.

Third, the requestor of a maintenance task has significant discretion to either include or withhold valuable information about the symptoms and cause of the issue. A cooperative requestor can include a detailed list of symptoms and what the requestor thinks is the actual cause of the issue. An uncooperative requestor might instead choose to be vague about the symptoms, or even misleading. Such valuable information can be supported by educating production operators and their supervisors about the mechanical functioning of the systems that they operate and by maintaining good relations with the operators so they are willing to provide specific reports.

4.3.2 Diagnose

In the diagnose step of the process, the repairperson needs to observe the symptoms of the issue and determine the cause of the issue. The desired outcome is an actionable diagnosis for the cause of the problem, which, when remedied, will alleviate the symptoms of the problem.

The criteria for a good diagnosis are that it is correct and made quickly. The most important criterion is a correct diagnosis; in order to resolve the problem and allow production to resume using the machinery, the diagnosis needs to accurately determine the cause of the problem. The secondary criterion is that the diagnosis is made as quickly as possible. The earlier the cause is correctly diagnosed, the earlier the repair can be completed and the machine returned to production.

Cognitively, the process begins by observing the symptoms, and the repairperson builds a mental model for how the machinery is supposed to operate. The repairperson forms several hypotheses about faults that could have conceivably occurred; they typically focus on faults that are conceptually related to the symptoms that are observed. Then, the repairperson eliminates any hypotheses that are incompatible with the symptoms that they observe. If they are still left with multiple plausible hypotheses, an effective repairperson can devise measurements or tests that can narrow down the range of plausible hypotheses.

The diagnose step is highly variable, depending upon the experience of the repairperson, the information available to the repairperson, and the approach that the repairperson can take. For simple problems, diagnosis can take less than a minute. For more complex problems, but those where there is sufficient information and a methodical approach, diagnosis might take minutes or a few hours. However, when the information needed is not available or the approach taken is not methodical, diagnosis can take many hours or days. Therefore, the effectiveness of the diagnosis step depends upon the available information, the experience of the repairperson, and the reasoning the repairperson can employ.

Considering available information, a repairperson will need information to support




Figure 4-4: A repairperson starts at the root of knowledge with a mental model for how the system works. Then the repairperson might need to access details about the functioning of components, or history about the machinery.

the above cognitive process shown in figure 4-4. First, the repairperson will need to refresh themselves on an accurate mental model of the machinery. The key criteria for this information is that it is accurate for the purposes of operating and repairing the machinery, while remaining at a summary level so it is reasonably understandable to the technician. Second, the repairperson will need to eliminate hypotheses at a detailed level, and they will need access to specific details about the functioning of individual components. Third, the repairperson will often benefit from reviewing a history of the machinery to learn if similar problems have happened before or if recent changes might have caused the observed symptoms.

Elaborating upon the reasoning of the repairperson, the repairperson can be either highly methodical about making the diagnosis or highly intuitive. Highly methodical repairpersons are explicit about their hypotheses, can devise effective measurements to narrow down their hypotheses, and make sure to do so before finalizing a diagnosis. The benefits of being so methodical are that the repairperson is much less likely to make an incorrect diagnosis and they are able to avoid getting stuck on difficult problems. However, acting methodically requires personnel training and an organizational structure that supports such skills. Some repairpersons might do this naturally, but the organization cannot reliably expect every new hire to do so without training.

Contemplating the level of repairperson experience, experienced repairpersons can usually make diagnoses more quickly and more correctly than inexperienced repairpersons. They are better at generating plausible hypotheses and better at narrowing down hypotheses. Gaining experience is a learning process, and therefore can either be taught or self-learned. During training, a repairperson can be taught better men-

tal models and shown how to take a more methodical approach, but misconceptions will undoubtedly remain. Therefore, a repairperson requires trial and error with the machinery to challenge their own mental models and requires corrections to their misconceptions to become more effective.




Figure 4-5: A virtuous cycle of learning better diagnosis skills. Useful, accurate information informs a methodological approach. A methodological approach helps a repairperson to learn. After some time, a repairperson that learns can contribute to more useful and accurate documentation.

The available information, repairperson experience, and the reasoning employed form a self-reinforcing loop as shown in figure 4-5. Having useful and accurate information enables a repairperson to take a methodical approach that might not be possible with less information. Taking a methodical approach encourages a repairperson to challenge their own mental model more critically, which makes the repairperson more effective. Finally, a more effective repairperson can make improvements to the information shared with their team, and those insights can help the rest of the team in the future.

4.3.3 Plan

In the plan step of the process, the repairperson needs to determine how to address the diagnosis. First, the repairperson needs to identify what intervention is needed. Which part needs to be serviced? Does it need to be replaced, or just adjusted or cleaned? Second, once the intervention is known, the repairperson needs to conceive of an order of operations that will accomplish the intervention. Can the part be accessed safely? How will the machine be restored to working order after the part is serviced?

The criteria for a good plan are that the intervention and associated plan can be accomplished quickly and efficiently. First, the intervention must be achievable quickly to restore the machine to working order. The parts needed for the intervention must be obtainable quickly, and the process needed must also be achievable quickly. The difficulty here is that some replacement parts might require weeks to be delivered or that some procedures like letting concrete cure might take days. Second, the plan must be efficient to restore the machine quickly and minimize maintenance resources. This is complicated because certain parts might require some disassembly to access and one plan for disassembly might require hundreds of components to be removed while another plan might require only a handful of components to be removed.

The worst possible outcome is that the plan puts the machinery in a non-operational state and the machinery cannot be restored to operations. This could happen because the actions planned unintentionally damage the machine or because the plan is underspecified and the repairperson does not know how to reassemble what they have disassembled.

With this in mind, a final criterion for an effective plan is that it tells production in advance how long the process is likely to take. With a reliable estimate about the length of the procedure, production can reduce a number of the wastes that they face with repairs. First, production can reassign the labor that would be working on that machine somewhere else and bring that labor back to the repaired machine just in time. Second, production can plan contingencies to handle the non-operational machine, perhaps building up a buffer of completed work before the repair begins or perhaps offloading the demand for that machine to another machine or another facility.

Cognitively, the planning process works by proposing a particular plan, simulating the execution of the plan mentally or virtually, and then predicting which problems might occur. Therefore, as with diagnosis, good planning is enabled by effective mental models, accurate information, a methodical approach, and experienced technicians (figure 4-6). Historical records can inform new plans, making new plans more effective. A methodical approach can increase the likelihood of predicting problems before




Figure 4-6: A virtuous cycle of learning better planning skills. Historical planning records and good information inform a methodological approach to planning. A methodological approach helps a repairperson to learn. After some time, a repairperson that learns can contribute to historical planning records.

they happen and addressing them. A methodical approach helps a repairperson to learn more quickly, which leads to better records about historical plans.

4.3.4 Get Parts

In the getting parts step of the process, the repairperson needs to bring materials and tools to the site of the repair. With a good plan in mind, they know which parts they need, but they also need to find out where those parts are, retrieve them, and transport them to the site of the repair. The time spent on this step is only increased if an inefficient plan means that parts are initially forgotten, and parts need to be retrieved with several trips or several rounds of orders.

The criteria for getting parts efficiently are that it is done quickly and with minimal labor. Additionally, the inventory itself has a significant cost to the business, so the maintenance system should minimize the cost of its inventory system. Some practices can lead to getting universally all parts more quickly and more efficiently. But more often, the criteria are based upon the speed and efficiency of the average case, which affects the workload as a whole and is neatly aligned with existing inventory management practices. A good inventory system ensures that the most likely parts needed can be retrieved efficiently, while the least likely parts need not be retrieved efficiently so that storage costs can be minimized.

When a given part is needed, it is generally retrieved via one of two paths: (a) the part is stocked somewhere on-premises, or (b) the part is not stocked on-premises

and needs to be procured from an external vendor.

In (a), the part is stocked on-premises. The repairperson must first locate the part, retrieve the part, and transport the part to the site of repair. This can take anywhere from a few minutes if the part is easy to find, to hours if the part is not where it is supposed to be. Note that if multiple parts are retrieved in a single trip, the transportation costs are only incurred once.

Furthermore, repairpersons typically find that the time to locate multiple parts at once is less than the sum of the times to locate each part individually. This is because, in a typical spare parts room, a repairperson is traversing through all the labelled parts and retrieving items on the list as they move across them. Therefore, the cost of retrieving a part is roughly $O(N) + O(n)$, where n is the number of parts that the repairperson must retrieve and N is the total number of parts examined in the relevant area of the spare parts room. This provides the further insight that parts can be retrieved more efficiently if the repairperson can localize the part to a smaller set of relevant parts.

In (b), the part is not on-premises and must be procured externally. Someone must then search for a vendor that has the part available for order on the required timeframe and at acceptable cost. Commodity parts, such as standard fasteners, can usually be sourced quickly and delivered within days. On the other hand, specialized or obsolete parts might be difficult to source; they may only be available second-hand, which carries an additional risk that the purchased part is not in good condition. Getting specialized parts delivered to the factory can take weeks and/or be extremely inefficient with plant resources.

As mentioned above, for every unique type of part, the inventory system needs to balance how efficiently that part can be retrieved against the costs of storing the part. Storing spare parts incurs a number of costs, including the risk of losing the part, the costs of the storage space and the amount that it slows down retrieval of other parts. Therefore, an inventory system should choose how many of each part to stock based on the marginal expected benefit of the part being needed against the predicted marginal costs of storing it.

4.3.5 Executing Physical Work on Machines

In the execution step, the repairperson needs to actually perform the planned work with the procured parts. A good performance of the planned work executes the required tasks quickly and with minimal labor. There are several relevant variables, such as how many maintenance technicians are assigned, their skill levels, their levels of effort, and their levels of focus. Judging whether a repairperson is giving sufficient effort or quality is highly task-specific and interpersonal, and is generally a responsibility of the maintenance supervisor.

Considering whether the work is completed quickly, it is important for the supervisor to set a reasonable expectation for how long the work will take. This expectation provides feedback to the repairperson and the supervisor about whether the task was performed efficiently. If the task takes longer than expected, the repairperson and the supervisor must then diagnose whether the inefficiency was because of lack of resources, lack of skill, lack of focus, lack of effort or perhaps unreasonable expectations; often it is not the repairperson's fault that the work took longer than expected. Future maintenance efficiency can benefit from resolving these execution problems.

However, in order to set reasonable expectations for how long the work will take, the supervisor must reference previous iterations of similar work. The supervisor could reference written records of similar work, or utilize personal memories of similar work, or both, but new supervisors will need to rely more heavily upon written records.

Contemplating the amount of maintenance labor, this criterion is largely determined by the assignments of the supervisor: the supervisor can assign any number of maintenance technicians, and can encourage parallel work while execution progresses passively (e.g., a technician can do something else while concrete cures). One should note that sometimes operators can contribute some or all of the labor, particularly if the operators are unable to create value for production because their machine is down; this can allow the execution to use the same number of hands with less wasted labor.

4.3.6 Observing Effects

In the observation step, the repairperson needs to verify that the issue was actually solved. Typically, this can be done by either testing the machinery themselves, or by allowing an operator to test it and either watching or following up with them. Like diagnosis, this can be done most efficiently if the repairperson has a good mental model of the system.

Additionally, at this stage, the repairperson might start to notice recurring patterns of failure that indicate a deeper root cause. When the maintenance team identifies an actionable pattern, it is often able to propose new interventions that address the root cause and eliminate a large amount of work in the future. For example, a repairperson that observes they are refilling oil much more often than they expect might deduce that there is a hidden oil leak. By addressing the oil leak, the repairperson will dramatically reduce the amount of oil refilling in the future.

4.4 Other Maintenance Processes

A maintenance system also needs to accomplish a number of ancillary tasks:

- Managing inventory of spare parts and tools
- Sourcing and procuring replacement parts from external vendors
- Producing some subset of replacement parts, typically in a machine shop
- Training new maintenance technicians
- Sharpening production tools that have worn
- Managing relationships with contractors and vendors

Large maintenance organizations are able to assign dedicated personnel to managing inventory, sourcing replacement parts, and producing parts. However, for small maintenance departments where dedicated personnel are not staffed, it is critical to

minimize the workload that these activities create. Even more importantly, the maintenance system must minimize variability in workload. Therefore, small maintenance departments need to make it easy to check inventory of spare parts and easy to source replacement parts. The number of replacement parts produced internally should be minimized because producing such low volume parts is a highly variable workload and the machinery required is generally poorly utilized.

Training new maintenance technicians is critical for long-term success when accepting that technicians will eventually leave or retire. An excellent maintenance department should have a structured curriculum for ensuring technicians learn all the information they will need for their jobs, or at least learn where to find that information. Equally important, the curriculum should impart an attitude of methodically seeking reliability so that the above virtuous cycles are reinforced.

4.5 Wastes in Maintenance

Maintenance is by definition not a value-add activity because it does not help the customer, so, by the strictest definition in Lean Manufacturing, all maintenance is a waste to be minimized. However true that observation is, it is not a useful razor for identifying the most wasteful parts of maintenance. Therefore, in this section, we take it as axiomatic that maintenance must be done, and we identify several kinds of waste that do not add value to operations by performing the above processes.

4.5.1 Incorrect Diagnosis

When a technician makes an incorrect diagnosis, they believe they know the cause of the problem and move forward with addressing it, but in fact they have the incorrect cause of the problem. In most cases, the technician only discovers the mistake at the very end of the maintenance process when observing the effect on equipment health. Sometimes, the technician might discover the mistake themselves from new information, or might be corrected by a colleague based on experience. When an incorrect diagnosis is made, the process must return to the diagnosis step as shown

in figure 4-7.




Figure 4-7: An incorrect diagnosis is discovered during the observing effects stage. It requires the process to return to the diagnosis stage, wasting all the effort on intermediate steps.

An incorrect diagnosis wastes resources in several ways. First, the incorrect diagnosis wastes all the maintenance labor that was spent planning, getting the parts, performing the work, and observing the effects. Second, the incorrect diagnosis wastes the material costs of any parts that were consumed, which can be significant if they were expensive. Third, and usually most importantly, the downtime when the machine is not producing between incorrect diagnosis and eventual correction is not creating any value, and so wastes production capacity that the business depends upon.

4.5.2 Bad Planning

When a technician makes a bad plan, they choose an inefficient or infeasible method to accomplish a given intervention. In many cases, the technician remains unaware that there was a more efficient way to accomplish the intervention, so they struggle through and the waste is often never identified. Otherwise, the technician generally finds out during the performing-work stage that the plan they have selected is either infeasible or much less efficient than the alternative that is then apparent to them. At this point, they can choose to continue on with the previous plan or backtrack and try a new plan as shown in figure 4-8.

A bad plan wastes resources in several ways. First, a bad plan wastes all the additional maintenance labor used in executing the bad plan relative to the more efficient plan. Second, a bad plan might waste some material costs of spare parts if they cannot be recovered. Third, and most importantly, a bad plan wastes the




Figure 4-8: A poor plan is discovered during the execution stage. The repair can proceed with an inefficient execution, or return to the planning stage.

production capacity corresponding to the additional time required for the bad plan relative to the efficient plan.

4.5.3 Waiting for Parts

When a technician is waiting for parts, they know what part they need but cannot find it themselves, so they are waiting for an external source or sometimes a colleague to deliver the part (figure 4-9). In some cases, the delivery can take place in minutes or hours, while in other cases, the delivery can take days or weeks. The technician may choose to perform other work while waiting, but in many important cases they will not be able to work efficiently on other tasks because pressure from management to fix the original problem is so great.




Figure 4-9: Waiting for parts introduces a delay in the getting parts step, from hours to weeks. As a result, the repair takes much longer than expected.

Waiting for parts wastes resources in several ways. First, waiting wastes maintenance labor, specifically the difference between what the labor could have achieved with the part versus what was actually achieved during that time. This can be a large difference if the technician has spent most of that time deflecting management pressure. Second, waiting wastes the productive capacity that could have been used if the machine was repaired sooner.

4.5.4 Missing Necessary Parts During Execution

One particular variant combines a bad plan with waiting for parts, where the technician begins physical work on the machinery when in fact they do not have the parts that they need as shown in figure 4-10. This waste is particularly noticeable during planned maintenance when the machinery could have been operational if it were not taken down intentionally. In particular, the machinery is taken down with the intention of finishing the work quickly, but because some part is not present, the procedure takes much longer than scheduled.




Figure 4-10: Missing parts during execution forces a return to getting parts. Work that was planned to take minutes might need to wait days for a part to be delivered.

Not having necessary parts wastes resources primarily by wasting the productive capacity that could have been used if the repair had been completed on schedule. It is usually better to postpone the entire procedure until parts are available and the procedure can be completed quickly rather than start the procedure and have to stop in the middle. Furthermore, not having necessary parts wastes labor, and usually wastes money, trying to get the required parts urgently in the middle of the repair.

4.5.5 Excessive Spare Parts Inventory

When a maintenance team has excessive spare parts inventory, it is holding spare parts that are not necessary and incurs several forms of holding costs. There are a few reasons that a given spare part can have unnecessary stocking:

- Too many of the part are stocked, to the extent that the expected benefit of having the marginal unit on hand is less than the holding costs. For example, if the need for a part is ten units per year, with an uncertainty of about five units, having 100 units on hand means that the business is vanishingly unlikely to utilize the marginal 100th unit.

- The part is very unlikely to need replacing, to the extent that the expected benefit of having even one on hand is less than the holding costs. For example, a single part that is unlikely to fail and services a machine that is low criticality need not be stocked, particularly if the part can be quickly and easily sourced externally. For these parts, if any units are kept on hand, there is unnecessary stocking.
- The need for the part can be adequately foreseen to find the part externally before the part is actually needed. For example, a part that is replaced at regular intervals or for which the operator can recognize the need several weeks out can be procured externally and used on arrival instead of stocking. For these parts, if the part is kept in stock when there is no immediate reason to keep it in stock, there is unnecessary stocking.

Having unnecessary inventory is not an uncommon problem for maintenance organizations, but having excessive inventory incurs several forms of holding costs:

- More inventory in the same amount of space generally leads to disorganization, which means it takes longer to find any given part in the spare parts room.
- Purchasing the inventory requires upfront expenses, which incurs a cost of capital to fund the purchase. For businesses that do not capitalize spare parts in balance sheet inventory but instead expense it, this can be particularly costly for management by depressing net income.
- More inventory requires more storage space, and space is costly. The space must either be purchased with a cost of capital or rented. Furthermore, the space is generally furnished with shelving that must be funded. Finally, the space needs lighting, heating, and its own building maintenance.
- Most businesses will need to check spare parts inventory periodically to confirm that the right parts are stocked. Counting inventory requires labor, and so the cost of checking inventory rises linearly with the amount of inventory.

Fortunately these holding costs are small if the degree of unnecessary stocking is small. However, if a large quantity of unnecessary parts are stocked, then the costs mount and become significant. Costs accumulate to delay in getting parts and higher costs of getting parts, as shown in figure 4-11.




Figure 4-11: Missing parts during execution forces a return to getting parts. Work that was planned to take minutes might need to wait days for a part to be delivered.

4.5.6 Motion to Retrieve Parts

The time that a maintenance team spends retrieving parts and bringing them to the machine do not contribute to the machine being fixed, and therefore constitute a waste to be minimized (figure 4-12). As above, the amount of inventory contributes to the time it takes to search for the correct part. Moreover, typically it is more efficient to search for list of parts all at once, rather than one part at a time.




Figure 4-12: Motion when retrieving parts requires time and energy during the getting parts and execution steps. As a result, it takes longer to resolve an equipment problem.

Additionally, travel to and from the spare parts room can be a significant waste. If done multiple times, even a short travel time can add up to a significant time sink. Moreover, leaving the site of the repair means the technician must reorient and refocus once they arrive. In fact, the technician might also stop during travel either to chat with colleagues, or even to warm up if the weather outside is poor. A good maintenance team therefore minimizes the number of times that parts must be retrieved but also minimizes the amount of time it takes to retrieve parts.

4.5.7 Waiting for Maintenance

Waiting for maintenance refers to when the machine in question is no longer operable and needs repair, but the operators, supervisors, or technicians linger nearby the machine without performing productive work to fix it (figure 4-13). Often, the operators do not relocate to another workstation, either because they are instructed not to or because they are reluctant to do so. They might wait nearby the broken machinery until maintenance has fixed it. Therefore, the labor of the operator waiting for the repair to be completed is wasted. If the supervisor is also waiting for it to be completed, the supervisor's labor is also wasted.

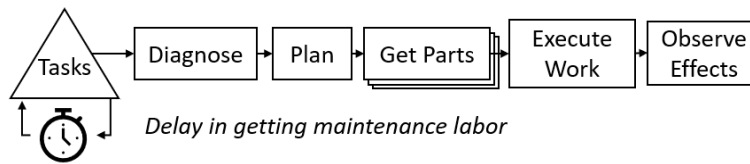


Figure 4-13: Waiting for maintenance is wasting resources waiting for maintenance to begin on a problematic resource. As a result, production resources are wasted not performing any value-added work.

Similarly, maintenance technicians can also waste time waiting for labor. For example, consider that some jobs require two or three people to be completed or might require specialized electrical skills. When a technician is simply waiting for peers or an electrician to show up, their maintenance labor is wasted.

Another variant of this waste is when production labor is assigned to operate the machinery after the repair procedure is expected to be completed but the repair procedure is delayed. This is typically during planned maintenance, where the maintenance team expects to work for a limited time and production assigns labor after that time. Therefore, if the maintenance team is not ready by the expected time, the production labor has nothing to do. This is therefore a waste of production labor.

4.5.8 Fixing Avoidable Problems

Finally, remember that all maintenance is inherently a non-value-add activity. Therefore, any problem that was fixed but was avoidable is a waste of resource. Typically,

many problems can be prevented at a cost much less than the cost of fixing them, particularly when considering the lost productive capacity and the increased efficiency of performing well planned maintenance.




Figure 4-14: If root causes are not addressed, unnecessary tasks are added to the incoming tasks. As a result, maintenance resources are wasted fixing problems that were avoidable.

As above, fixing avoidable problems has wastes that show up in several areas. First, the productive capacity that is lost because of the problem is wasted, compared to a typically small loss of capacity to prevent the problem. Second, the maintenance labor and material costs of executing the repair are wasted. Third, the least obvious cost is the impact the unpredictability has in preventing management from optimizing production. For example, in a reliable factory production can proceed just-in-time and minimize production costs, but an unreliable factory necessitates high levels of buffer and suffers from confusion.

Chapter 5

Maintenance Problems in Marysville

In this section, we apply the above framework to Heartland's Marysville facility and highlight specific problems, prior to the application of specific interventions proposed in chapter 7. This section enumerates problems following the steps outlined in figure 4-3: problems with incoming tasks, problems with diagnosis, problems with planning, problems with getting parts, problems executing work and problems observing effects.

5.1 Problems with Incoming Tasks

An effective maintenance system is structured to create a desirable rate and nature of incoming tasks, but the maintenance system at Heartland was overwhelming the maintenance team with menial tasks, urgent tasks, and vague descriptions. This led to a maintenance system that was falling behind on incoming tasks, which led to decreasing reliability of the equipment for production.

The most common problem with incoming tasks was that they are menial tasks but were assigned to valuable maintenance technicians. This wasted labor because the production operator was not creating any value for the business while their machine was broken, and the maintenance technician was not performing other repair

work. The most obvious cause for this misallocation was that the Collective Bargaining Agreement specified that only maintenance technicians can perform repair work, therefore supervisors could not ask operators to repair equipment, and many operators wanted to follow the rules in the CBA.

However, the underlying cause of this misallocation was an established antipathy between maintenance and production. Maintenance technicians, particularly the more senior ones, believed that the production operators do not know how to operate their machines, and that the production operators were lying about many problems that they report. Correspondingly, these technicians often dismissed concerns from the production operators and did not trust the operators to perform repairs. In response, the production operators saw that their concerns were not being addressed, so they believed that the maintenance team was being lazy and not doing its job. This was compounded by the fact that maintenance was simply too busy with more urgent work, so often they did not get around to less critical repairs. Therefore, the production operators often refused to do any repair work, believing that they could punish the maintenance team and instill some work ethic by making a maintenance technician perform the work.

For example, when a production operator discovered oil dripping from a machine and causing an oil spill, the operator reported it to maintenance. The maintenance technician refused to fix the issue, claiming that he could not see anything leaking. Weeks later, the production operator needed an oil filter replaced. The operator acknowledges that he was capable of replacing the filter, but he refused to do so; he insisted that maintenance was not doing its job, and he wanted to punish the maintenance team with a dirty job.

Other menial jobs were not being handled efficiently. Maintenance technicians were spending time applying regular lubrication to continuously moving parts, which could be accomplished by the production operator or could be done even better by an automated drip lubricator. Maintenance technicians were spending a lot of time maintaining overhead cranes, which does not require any Heartland-specific knowledge and can be accomplished more efficiently by external contractors.

Finally, many of the incoming tasks for the maintenance team were vague. Vague issues took the maintenance team longer to diagnose, did not permit spare parts to be organized ahead of time, and were difficult to verify as fixed. Many of the vague issues reported were quite visibly antagonistic, falling along the lines of punishing maintenance for perceived slights.

For example, a production operator reported “fix the fan ITS BROKEN do your job.” The reporter entered the name “concerned worker.” When a maintenance technician visited the fan, the fan was spinning as it was meant to. The maintenance technician spent an hour investigating the fan to discover a problem but never found one. The maintenance technician was unable to follow up with the worker to ask for clarification.

5.2 Problems in Diagnosis

An effective maintenance system needs to make correct and quick diagnoses, but the maintenance system at Heartland was frequently making diagnoses that were slow or incorrect. This led to excessive loss of production on the machinery, unpredictable repair times, and wasted maintenance resources.

The most common problem was that a diagnosis was not made quickly, resulting in the completion time of the repair being unnecessarily late and unpredictable. This was typically caused by a combination of an inexperienced repairperson and a lack of useful information available to the repairperson. Heartland has a number of more experienced maintenance technicians, but they only work on day shift, so the only available maintenance technicians on night shift are much less experienced. Compounding this problem, there is little system summary information available about the machinery. There are no written system summaries to communicate a mental model of the systems, and the technician training includes little in the way of communicating mental models of the systems.

For example, one evening there was only one maintenance technician on night shift. This technician had only been on the job for about five weeks, three of which

were shadowing a senior technician to perform repairs. On this night, one of the mills broke down suddenly and was no longer cutting consistent lengths of steel. The maintenance technician had only done minimal work with this machine and did not have a good mental model for how the length-control subsystem worked. The technician opened the electronics box and looked for loose connections, but did not see anything obviously wrong. The technician was overwhelmed by the number of possible things that could have been broken and thought it would get worse if he started replacing components at random. The technician wrote a note for the more senior maintenance technicians on day shift. The machine was not able to function for the rest of night shift, resulting in one workstation being starved for materials the next day.

Another common problem was that the repairperson would make an incorrect diagnosis, resulting in the subsequent steps of planning, getting parts, and performing work not actually fixing the problem. This wasted maintenance resources and delayed the resumption of operations. This was typically caused by not following a methodical approach to making the diagnosis: the repairperson would either not consider some plausible hypotheses, or would neglect to carry out an appropriate test to narrow down the hypotheses. Senior technicians would make this error about as often as junior technicians.

For example, one day the oven broke down, and the temperature was exceeding the setpoint. There were three maintenance technicians and the supervisor on duty, including the most senior technician. The technicians came up with a hypothesis that an actuator that manipulates a gas valve had failed and gotten stuck open. The supervisor and the plant manager purchased a new actuator; they spent a premium to get a replacement quickly, and spent time picking up the actuator personally to deliver that night. The night shift technician on duty at that time was not given instructions about how to implement the fix, so he left it for day shift to complete. When the day shift returned, they replaced the actuator, but the problem was not fixed. The technicians realized there was an alternative hypothesis that the electronic controller had failed, and they had not ruled that out. So, the supervisor purchased

a new controller to be delivered the next day and hired an electrical contractor to install the new controller. The new controller fixed the oven.




Figure 5-1: A vicious cycle of worsening diagnosis skills. Lack of information prevents a methodological approach, so an intuitive approach is used. An intuitive approach slows down a repairperson’s learning. A repairperson that does not learn about the systems cannot contribute to useful and accurate information.

As in our framework, there should be a self-reinforcing loop of useful information enabling a methodical approach, which enables a repairperson to learn and subsequently contribute to the information. The opposite loop is also possible: a lack of information encourages a repairperson to use a more intuitive approach, which leads to the repairperson not learning as fast and not contributing to the information available. It is this opposite loop that had taken hold at Heartland, as shown in figure 5-1. Most of the systems had no system summaries available to the repairperson to convey a mental model, and significant gaps existed in the component details that enable eliminating hypotheses. Therefore, each repairperson had grown to rely on intuitive diagnoses, which were error-prone and limited to the most senior technicians. Finally, by not challenging their own mental models, technicians were not learning quickly and were not contributing to the organizational knowledge store.

5.3 Problems in Planning

An effective maintenance system needs to make efficient plans for accomplishing repairs, but the maintenance system at Heartland was frequently making plans that were inefficient or was not able to generate a plan at all. This led to delays in returning the machines to operations and production working around unpredictable repair procedures.

The most common problem was that a repairperson did not know how to plan a repair at all, and therefore could not proceed. A junior repairperson with a correct diagnosis would still not understand enough about the system to feel comfortable starting a repair and did not have the experience to remember doing the repair before. This is a symptom of lack of information: in addition to the missing mental models of the whole system, the technicians lacked access to records of how the procedure had been performed in the past. Such records are invaluable, because they enable any repairperson to replicate a similar procedure and, in fact, vastly speed up the planning process.

For example, a bushing on one of the automatic beam welders needed to be replaced because it was wobbly. Only the most senior maintenance technician had ever performed this repair before, and this senior technician was absent due to COVID-19. The procedure had not been written down anywhere, so the more junior technicians were uncomfortable starting the repair. The machine was operating at a high level of defects, but it was operating, so production did not want to shut it down. Several days later, when the senior technician returned, he was able to perform the repair quickly with help from another technician, who thereby learned how to perform it, but other technicians remained unaware of the procedure.

Less commonly, but more seriously, the repairperson proceeds with a repair with an insufficient plan that leads to the machine being non-operational for far longer than anticipated. The plan would have serious but unrecognized flaws, including necessary parts that were not available or a lack of recorded detail about how to restore the machine's operational condition. This could lead to the machine being non-operational for days or weeks, while the maintenance team wasted days or weeks of maintenance labor that could have been used elsewhere. These problems could have been avoided by making a clear plan and making sure all prerequisites were in place before proceeding.

For example, one week the maintenance team planned to perform some preventative maintenance on one of the automatic beam welders. They scheduled two days to replace a couple of key components. After they had disassembled a large section

of the machine, they went to go look for the replacement parts and discovered that there were none left. The parts would need to be machined from stock, which would take days or weeks. The maintenance team initially intended to wait for the parts, but when production heard about the delay, they insisted on abandoning the procedure so they could use the machine again. However, the maintenance team had not recorded the procedure for disassembly or reassembly, so the technicians were unsure about how to reassemble the machine correctly. After seven days of work, the technicians were finally able to reassemble the machine so that it worked again. No valuable work had been done to improve the equipment, so this work wasted about a week of maintenance labor and five days of productive capacity that the business had planned to utilize.




Figure 5-2: A vicious cycle of worsening planning skills. A lack of planning records or information result in a haphazard approach to planning. A haphazard approach leads to slower learning, particularly without reflection. A repairperson that is not learning will likely not keep records, or keep very sparse records.

As before, a vicious cycle had taken hold at Heartland as shown in figure 5-2. The repairperson was using a haphazard approach to planning, usually to catastrophic effect. This meant that the repairperson was not challenging their mental models and, furthermore, not applying what they did learn in previous experiences. Their experiences were not recorded well, which contributed to the lack of attention paid to planning repair processes.

5.4 Problems in Getting Parts

An effective maintenance system needs to efficiently bring parts to the site of a repair, but the maintenance system at Heartland was frequently spending a lot of time retrieving parts or unexpectedly scrambling to order a part externally.

The most common problem was that a repairperson would waste time and resources retrieving parts. As shown in figure 2-6, the spare parts rooms are located relatively far away from the machines that they service. Furthermore, as illustrated in figure 5-3, the spare parts rooms were cluttered and difficult to navigate. For some of the most common repair processes, more than half of the time and labor required to fix the issue was spent retrieving parts. This also frustrated the maintenance technicians who were regularly spending a significant amount of their work day travelling back and forth to the spare parts rooms to search through dusty parts.

For example, the automatic beam welders frequently needed to have a “cam follower” replaced, which is a small bearing and is stocked in Spare Parts Room A. We observed the whole repair process taking 18-25 minutes: walking 3-4 minutes to the spare parts room, searching 2-5 minutes for the necessary part, and walking back for 3-4 minutes. As another example, the automatic beam welders periodically needed a fastener to be replaced. The fasteners are stocked in Spare Parts Room B. We observed the whole repair process taking 45-55 minutes: walking 10-12 minutes to the south building, searching 10-15 minutes for the correct fastener because there are hundreds of available fasteners in that room, and returning for 10-12 minutes. In one case the technician had to make the walk twice because the first fastener they brought was not the correct one.

The less common, but much more severe problem, was that a repairperson discovered that the necessary part was out of stock. In most of these cases, no one was previously aware that the part had been out of stock. This is because no one was regularly counting how many of the parts were present and discovering that the number was too low. Indeed, the inventory system had no definition for what was too low. It should be noted that there was a nominal process to reorder more of the



Figure 5-3: Shelves in the spare parts room, labeled A in figure 2-6. Cardboard boxes containing spare parts are stacked on pallet racking. Some boxes are labeled with permanent markets, but there are no part numbers. Many parts are covered in a layer of dust. Parts are not counted regularly.

part whenever it was used (by submitting the empty packaging to the supervisor to reference), but this process was error-prone because a repairperson might forget to submit the empty packaging, or because a repairperson might submit only one empty package, which the supervisor interpreted as only to reorder one unit, or because the packaging got lost before the supervisor found it. The effect of a part being unexpectedly out of stock is that the machine cannot be operated by production for days or weeks until the part is delivered.

For example, the automatic beam welders often wore out their cables, which underwent violent accelerations. In one instance of the cable breaking, the technician went to look where the cables are usually stacked on the shelves. He discovered that there were no more of the cables remaining: there was no label indicating the specific part number nor the quantity that should be present. Because the specific part number was not recorded there, the supervisor searched through a large binder to discover the specific part number and eventually ordered more from an online marketplace to be delivered as fast as possible. The cables arrived two days later and were quickly installed to return the machine to working order.

5.5 Problems in Executing Work

An effective maintenance system sets clear expectations for how long maintenance tasks should take and assigns maintenance resources efficiently. However, the maintenance system at Heartland had unclear or no expectations for how long maintenance tasks should take, and did not make efficient use of maintenance resources.

Considering expectations for task completion times, the maintenance supervisor at Heartland has practically no records of previous maintenance work and how long each should take. When technicians are asked how long it would take them, they generally inflate their estimates of how long tasks would take; inflation is an understandable hedging behavior, because a previous maintenance supervisor tended to yell at technicians if they did not meet expectations. Inflated estimates give the technicians time to deal with unexpected delays due to incorrect diagnoses or missing

parts as mentioned above. However, because expectations are relaxed, the causes of these delays are consistently not addressed. Certainly, there are instances where a repairperson does not sufficiently focus on the task, but more often the technicians are dealing with many setbacks that had never been addressed after previous experiences.

Again, the root cause of unclear expectations is a lack of accessible records about previous maintenance tasks. The maintenance supervisor is new to the position, and there are few written records of maintenance performed. The records that do exist are handwritten on loose leaf paper without organization, so it is difficult to locate relevant records.

For example, management asked a technician how long it would take to fix a problem with the automated beam welder. The technician first got angry about being asked to make such an estimate. Then the technician said about six hours, but not to hold them to that number. The procedure in question should take about fifteen minutes if all parts were available and no complications were encountered. Eventually, the technician was assigned to do the work, and finished in about two hours. The technician took a few minutes break to smoke, and then returned to the workshop since there was no more work assigned that night.

Moving on to efficient assignments of maintenance resources, the maintenance supervisor at Heartland tends to assign two or three maintenance technicians to almost all tasks. The rationale is that multiple technicians can help each other out, keep each other safe, and hopefully share experience to solve problems. However, this policy reduces the number of maintenance tasks that can be simultaneously addressed by a factor of two or three.

Consider here that a production operator can provide extra hands, keep a lookout for hazards, and contribute their knowledge of the machine (operators are often more experienced than the maintenance technicians). Furthermore, assigning maintenance technicians to work with operators forces socializing between the two groups, which maintenance technicians try to avoid to some extent by travelling in pairs or trios.

For example, two technicians were assigned to replace the die on one of the mills. The production operator stood and watched from afar (which is wasted labor); the

production operator has worked with this machine for more than twenty years. The maintenance technicians (both with less than three years' experience) had a couple of brusque conversations with the operator, but all the jovial workplace conversation was between the maintenance technicians. Consider that the production operator had the most experience and could have created a social bond with the maintenance technicians if he were working closely with them.

5.6 Problems in Observing Effects

An effective maintenance system identifies actionable patterns and addresses root causes, but the maintenance system at Heartland was making few improvements to reduce their future workload. This led to an ever-increasing amount of maintenance work without long-term improvements to reliability.

For example, a recurring problem was that the nozzles would get clogged in the washer and blow off. This required a brief but dirty repair process. Sometimes the entire pipe would break off and need to be replaced. These issues were caused by the dirtiness of the wash water, and the inappropriate choice of PVC for a high-temperature structural application.

The primary reason for not tackling the patterns was that the team was not having discussions at all. For much of the time, the maintenance team existed in two factions that disliked each other and refused to talk to each other in a civil manner. Moreover, even within the factions, some technicians felt that their peers were being unnecessarily toxic, and they planned to leave the team as soon as they were allowed.

Another reason was that when such ideas were proposed, they were usually shot down promptly by a peer. The senior technicians had gotten so used to having budgets cut that they instinctively thought solutions were too expensive. Other technicians believed that their technician colleagues did not know anything about maintenance, and therefore any idea that their peers had was necessarily a bad idea and poorly thought through.

Chapter 6

Estimating the Impact of Maintenance Problems

In this section, we estimate the impact of the maintenance problems at Heartland. Understanding the magnitude of the impact is important for two reasons. First, knowing the scale informs the range of appropriate responses to the problem: large problems can justify large investments to improve. Second, knowing the scale helps to impress the need for change upon all employees from whom support is needed: a large impact is more likely to align diverse groups such as maintenance, the union and management.

To estimate the impact of equipment breakdowns, we review the available data from the Marysville facility, and identify all of the equipment breakdowns within the data. Then we cross-reference those breakdowns to production records of which equipment was a bottleneck at the time, and therefore has a direct impact on the profitability of the company.

Notably, this estimate does not include the cost of wasted maintenance labor, wasted maintenance parts, wasted external contractors, nor wasted management labor due to confusion. Additionally, we leave the above estimates in units of hours rather than dollars so as to protect proprietary competitive information about Heartland.

6.1 Available Data

The data for this estimate originates from Heartland’s Manufacturing Execution System (MES). The MES is available to production workers via several computer terminals spread across the various workstations. The MES prints out inventory tags that the workers use to label their output parts for later retrieval. The MES also collects information from the production workers. For each job, in which a single workstation makes some quantity of a single part number, the MES records the start time, end time, how many parts were produced, and which production workers were working during that time. In many cases, the MES also records time when the workers were not productively working on a job under a special job identifier “INDIRECT.” In such cases, the workers can report a specific reason such as waiting for materials, waiting for orders, being pulled away for training or a meeting, or the machine needing repair. However, not all repairs or other interruptions are logged in this system. In particular, nothing is logged on the MES if no employees are at the workstation, even if the machine is broken.

As part of this project, we manually cleaned up this data to the best of our knowledge. To a local copy of the MES data, we added appropriate records to detail repairs when we knew positively that the machine was broken. To detail repairs that took place prior to the thesis project start, we manually catalogued breakdowns referenced by a 12 month historical archive of emails. Some repairs are certainly missing from these records, but manually cleaning up the data greatly improves the quality of our estimates. The cleaned data covers the time period between June 1, 2020 and September 30, 2021, a period of 16 months.

Heartland did not maintain any significant digital or paper records of maintenance work performed, at least not until just before the completion of this project at the site. Therefore, we cannot reference data about exactly where maintenance personnel spent their labor hours, which parts were used, nor exactly what problems were being fixed.

For the sake of simplicity, we limit this analysis to only the most critical processes

and exclude a litany of minor machines. In fact, the data for the minor machines is lower quality because there is less management pressure to report breakdowns, and indeed it has little impact because they are working well below capacity.

6.2 Breakdowns

From the cleaned up MES data, it is relatively simple to identify the time intervals labeled as breakdowns: they are reported to “INDIRECT,” have zero parts produced, and give a reason of “MAINT / REPAIR.”

	Start Time	End Time	working	total_time	Machine
32	2020-06-01 05:03:00	2020-06-08 07:39:00	False	170.60	Automatic Beam Welders
1789	2020-06-01 06:58:48	2020-06-01 07:31:12	False	0.54	Mill 4
2169	2020-06-01 16:34:48	2020-06-01 16:36:00	False	0.02	Paint Line
1817	2020-06-02 06:57:00	2020-06-02 07:28:12	False	0.52	Mill 4
1830	2020-06-03 06:58:12	2020-06-03 07:28:48	False	0.51	Mill 4
...
2086	2021-09-28 05:16:48	2021-09-28 05:25:12	False	0.14	Minster
298	2021-09-28 13:03:00	2021-09-30 06:15:00	False	41.20	Automatic Beam Welders
1606	2021-09-29 07:16:12	2021-09-29 08:10:48	False	0.91	Minster
890	2021-09-30 06:15:00	2021-09-30 15:28:12	False	9.22	Automatic Beam Welders
1389	2021-09-30 13:21:00	2021-09-30 15:10:48	False	1.83	Minster

Figure 6-1: Rows of cleaned-up MES data. Contains the start time, end time, number of hours, machine identity and whether the machine was working (i.e., not broken). This table highlights only the rows where the machines were broken and not working.

From the data in figure 6-1, we can quickly visualize the magnitude of the downtime between the various machines in the process. In figure 6-2, we immediately see that the most downtime was contributed by the Minster press (punches hole in post coil), which is well known to management. The beam welders, two of the mills, and the paint line all have comparable downtime, while Mill 7 (braces) has much less downtime.

By adding a time dimension, we can get a sense for the historical issues with each machine (figure 6-3). The Minster press had significant issues in the winter of 2020-2021 and needed to be entirely replaced in February 2021. The mills and paint line had a fairly steady stream of problems over the 16 month period. The beam welders

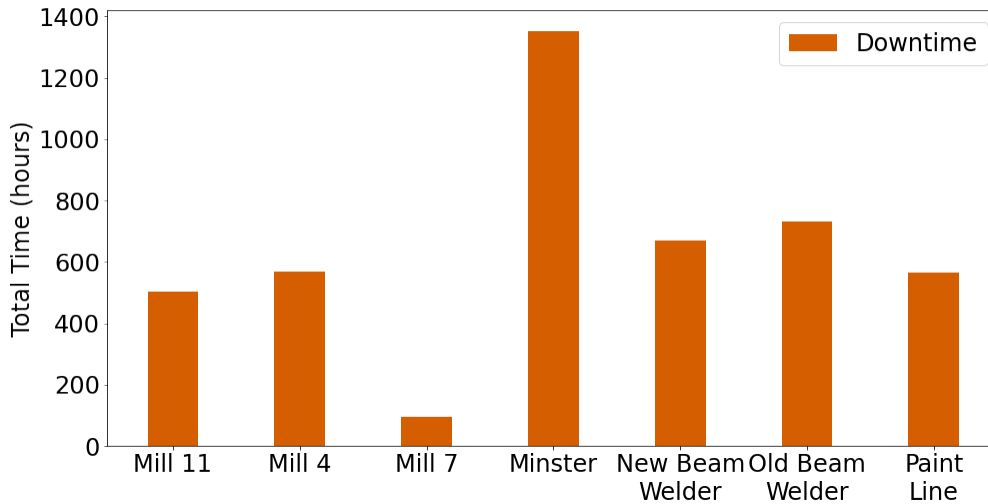


Figure 6-2: Breakdown time in hours, grouped by machine.

are temperamental, in that they can work well some months and have significant problems in other months.

Finally, it is useful to look at the monthly breakdown but stacked vertically (figure 6-4). This makes it harder to gauge the individual problems with each machine, but it makes intuitive the overall workload and stress levels of the maintenance team. Clearly, the maintenance team was very busy with the Minster being replaced in February 2021 through April 2021. The peak in October 2020 is not as high, but is more impactful than it looks because no individual machine was a clear focus.

Overall, figure 6-4 represents about 4500 hours of downtime. Making a conservative estimate that an average of two maintenance technicians were assigned to each broken machine and that the cost of employment for the technician is about \$30/hr (Marysville average), this represents a direct maintenance labor cost of \$270,000.

6.3 Bottlenecks

In order to estimate the lost opportunity to produce saleable product, we first must develop which breakdowns represent lost opportunity and which breakdowns do not hold back the total production. In particular, if a given machine is not one of the




Figure 6-3: Breakdown time in hours, grouped by machine and by month, showing the patterns of downtime for each machine.

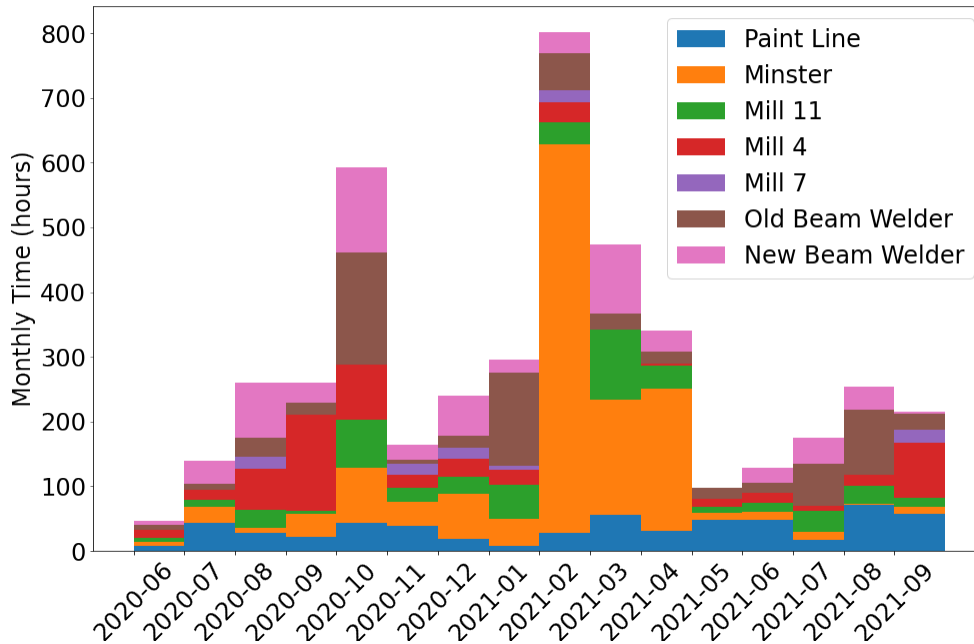


Figure 6-4: Breakdown time in hours, grouped by machine and by month, stacked vertically, showing workload of the maintenance team.

limiting factors to production, then having a breakdown will not reduce the net saleable product produced. In fact, this is the same concept as a machine being the production bottleneck.

Identifying the bottleneck at a facility like Heartland in Marysville is a non-trivial task. As detailed in the introduction, customers order different relative quantities of and dimensions of beams and frames. Therefore, one week the factory might be producing a vast number of short beams and short frames, which management knows will be a tough week for the beam welders. Another week, the factory might be producing tall frames and fewer long beams, which management knows will be a tough week for the mills and manual welding. Therefore, for our purposes of measuring lost opportunity, we must develop an empirical measure for the instantaneous bottleneck.

In this section, we utilize a methodology developed by Roser et al. [13] to determine the momentary bottleneck at all moments. Roser presents a methodology that is both theoretically a good fit to determine whether a machine is instantaneously losing opportunity to produce, and is also convenient given the nature of Heartland’s

available data.




Figure 6-5: A diagram showing which activities are considered active periods. Diagram follows a similar diagram from Roser et al. [13].

First, Roser defines whether a workstation is utilized at any given point in time by an “active period.” Specifically, a workstation is active if it is producing parts and also if there is any interruption preventing the workstation from producing parts. A workstation is only inactive if it is either starved for material or blocked from producing more output. These interruptions include repairs, tool changes, trainings, or quality issues if those interruptions happen when the workstation otherwise could be making parts. Intuitively, this makes sense: for example, even if a machine would not otherwise be the bottleneck, if it is prevented from producing parts long enough it will eventually become the bottleneck.




Figure 6-6: A diagram illustrating which active periods are considered bottlenecks. Diagram follows a similar diagram from Roser et al. [13].

Second, Roser defines the bottleneck(s) at any given moment. Loosely, the bottleneck is defined as the machine which has the longest uninterrupted active periods at any given point. However, Roser notes that because the bottleneck shifts from one machine to another and the active periods of those machines might overlap, there exist periods where it cannot be determined which of the two is really the bottleneck.

Therefore, more strictly the bottleneck is a one-to-many mapping from time t to one or more machines m_i , which has values of all machines m_i which are at time t in active periods $p_{i,j}$ such that $p_{i,j}$ is the longest active period for at least one value of time t' .




Figure 6-7: Active periods of various machines during a two week period.

We implement this test for Heartland's data. Figure 6-7 is a zoomed in view of about two weeks at Heartland, where each rectangle represents an active period on the corresponding machine. The fully saturated bars are the bottleneck active periods, and the less saturated bars are not bottlenecks. In this example, the paint line is the bottleneck for most of the time.

With all active periods classified as bottleneck or not, it is fairly simple to compute the total time when each major machine was the bottleneck. In figure 6-8, we can see that the paint line and the automatic beam welders were the only sole bottlenecks, and the vast majority of the time the bottleneck is shared. This is indicative of the fact that the Marysville facility serves an extremely high mix of products, and the bottleneck is almost constantly shifting.

The true value of this methodology is that we know at every instant in time whether a machine is a bottleneck or not. Therefore, for each machine, we can now find the set of instants where the machine is both the bottleneck and is also

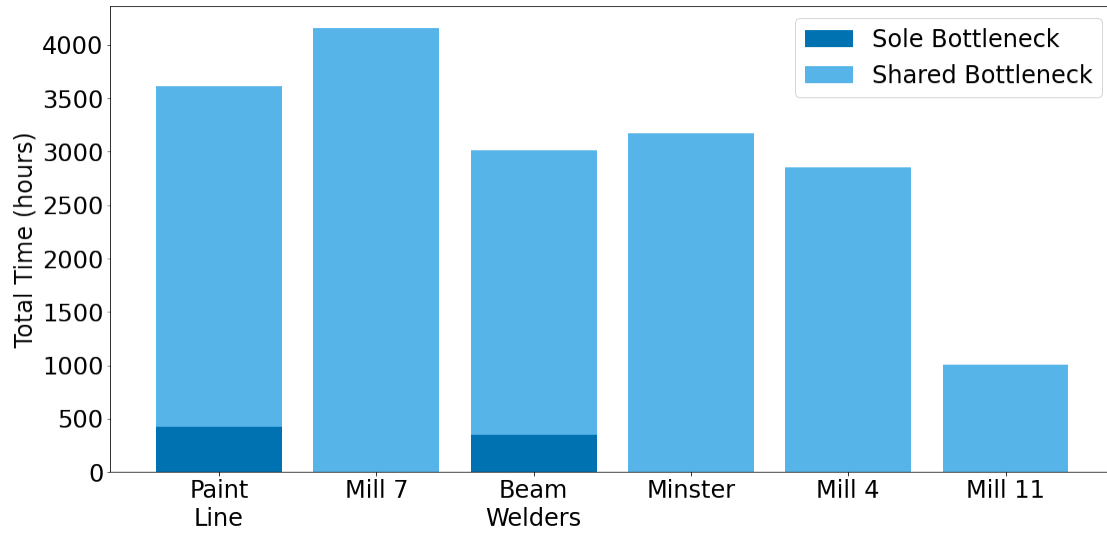


Figure 6-8: Amount of time in hours where each machine was a sole or shared bottleneck.

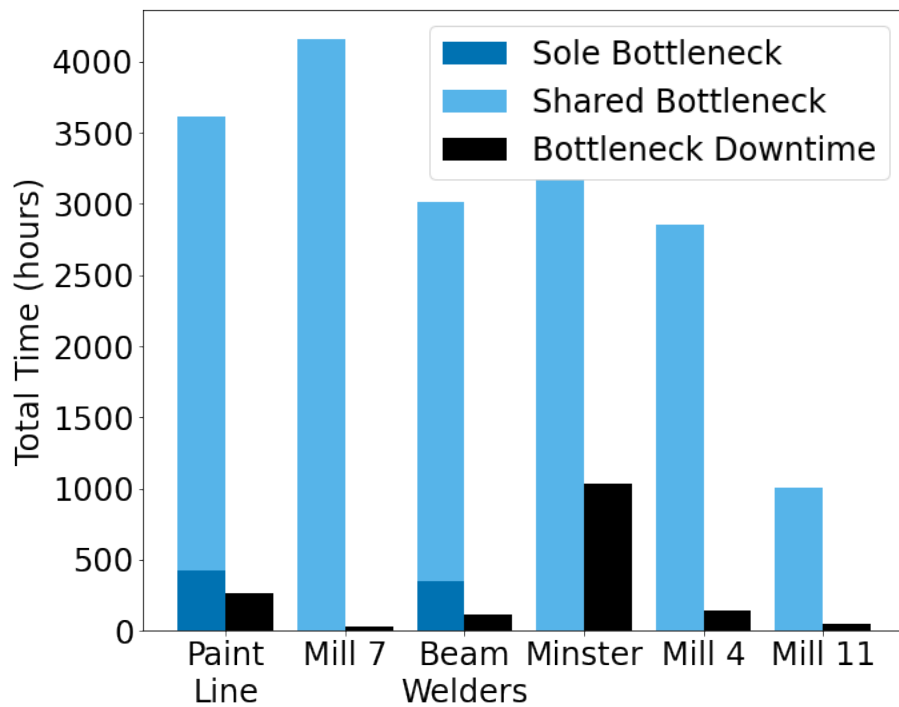


Figure 6-9: Amount of time in hours where each machine was a sole or shared bottleneck, compared with the amount of downtime while each was the bottleneck.

experiencing a breakdown. This leads to figure 6-9, where the black bars represent only the amount of time where the machine was the bottleneck and also broken. The black bars therefore represent the hours of downtime that were truly lost opportunities to produce saleable product.

We can visualize over time which machine was the bottleneck (figure 6-10 in blue) and, of that bottleneck time, how much was lost opportunity due to breakdowns (figure 6-10 in black). Clearly, a vast amount of opportunity was lost when the Minster needed to be replaced, and there was a smaller spike for Mill 4 in September 2020. Additionally, there is a steady trickle of saleable product being lost to breakdowns on the paint line, and automatic beam welders.

Additionally, we can visualize how much of the downtime ultimately resulted in lost product. In figure 6-11, the overall height of the bars represents the downtime, while the black bars represent the fraction of the downtime that ultimately resulted in less saleable product.

All told, the total downtime that occurred on the instantaneous bottlenecks sums to about 1500 hours during a 16 month period. This dominates the total cost of maintenance, costing an estimated \$500k per year in lost profits.

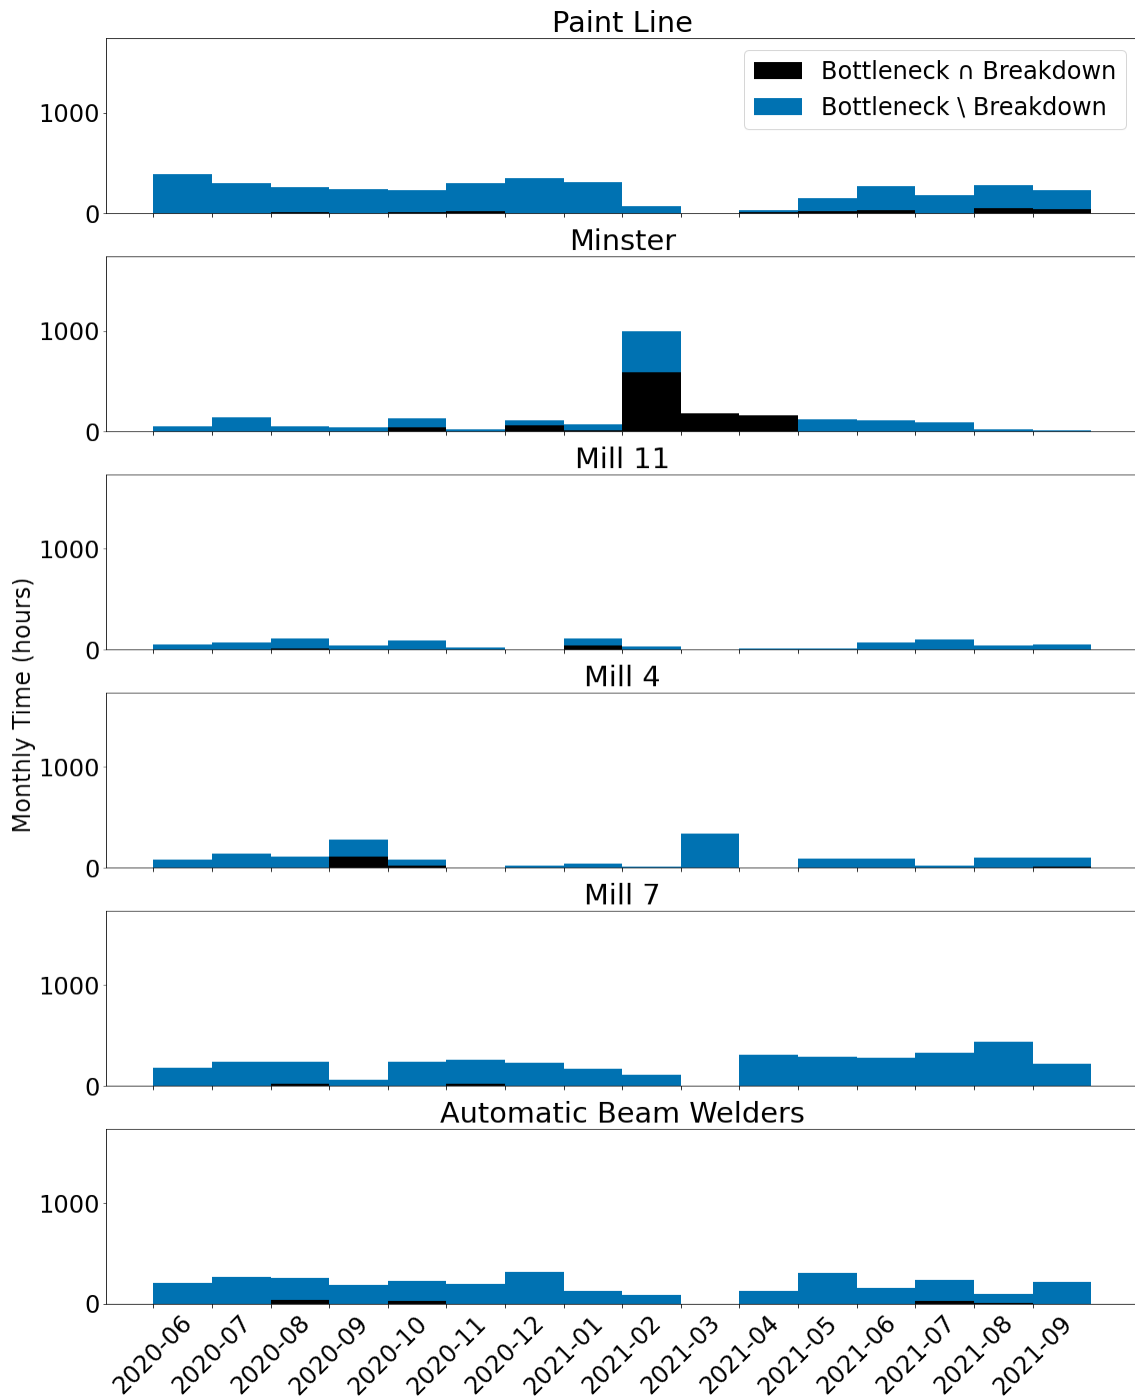


Figure 6-10: Total height of each bar represents the amount of bottleneck time in hours for that machine in that month. The portion that is black represents the amount of that where the machine was broken. In the legend, \cap indicates set union and \setminus indicates set difference.

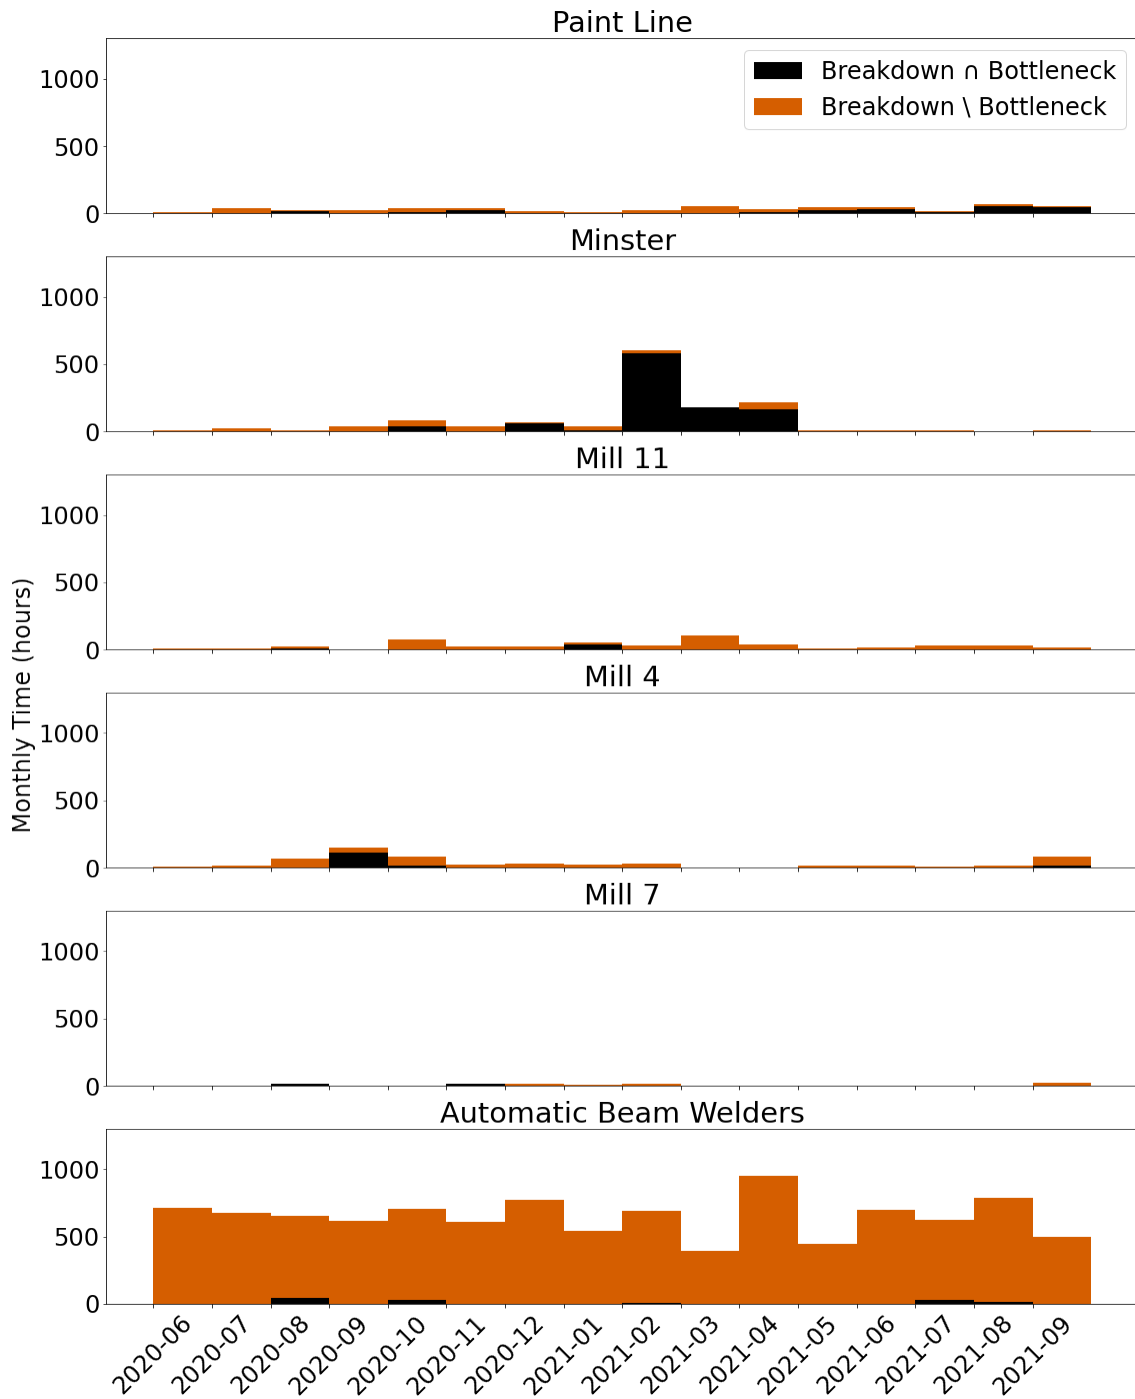


Figure 6-11: Total height of each bar represents the amount of breakdown time in hours for that machine in that month. The portion that is black represents the amount of that where the machine was also the bottleneck and resulted in lost opportunity. In the legend, \cap indicates set union and \setminus indicates set difference.

Chapter 7

Targeted Improvements and Results

In this chapter, we detail improvements to make Marysville's maintenance system more effective. Some improvements are approachable in the short-term, and we were able to pilot the improvements in small settings on the factory floor, usually with the automatic beam welders. Other improvements are not possible in the short term, and they are delivered as recommendations for the future of maintenance at Marysville.

7.1 Shared Ownership of Equipment

Recall from section 5.1 that the maintenance team was overwhelmed by menial tasks, urgent tasks and vague descriptions. Further recall problems with incoming tasks were caused by antipathy between maintenance and production, as well as production operators refusing to perform repair work.

The desired state is that no labor is wasted waiting for maintenance and that maintenance gets to focus on the tasks that most require their particular skill set. We planned an initiative to break the gridlock of antipathy and ask production operators to perform basic repairs. This would ensure that common issues are addressed more quickly, because the necessary labor is already present, and would allow the maintenance technicians to not be frequently interrupted by these issues.

We anticipated that this impasse would be difficult to address, but we recognized that the production operators for the automated beam welders were emotionally most ready to accept responsibility and in some cases were bypassing the formal policy of the union. These employees were extremely frustrated with how frequently their machines were breaking down. The automated beam welder operators started hiding a collection of electrical fuses at their workstation. When a fuse would blow, they would make sure their union stewards were not watching, then they would replace the fuse and return to work.

We embraced this group's willingness, and proposed a pilot study to the union where we would empower the operators to perform safe, minor, and recurring repairs. We defined in detail for the union which repairs were considered safe, minor, and recurring. We argued to the union that with negotiations coming up, Marysville would be able to pay operators more if they could take care of some repair work. The union agreed to allow the pilot to move forward, and the production operators were willing and, in some cases, excited to get more autonomy to repair their equipment. Finally, as detailed later in section 7.4, we made common spare parts available nearby the machines.

We administered the program with a "repair training matrix," which tracks which kinds of repairs each operator was trained to perform. If an operator was trained to perform the repair, they could go ahead and perform it. Furthermore, the supervisor knew which repair tasks they could ask each operator to perform. If the operator was not trained, then they could call over maintenance who would guide the operator through the task. The maintenance technicians were given the ability to add tasks to the list that they thought were safe, minor, and recurring, or they could subtract tasks that they noticed were not being completed correctly.

For example, with cam followers stocked nearby, the production operators were soon able to replace them without involving the maintenance technicians. The production operators were shown how to use the reorder cards. The maintenance technicians went from spending 18-25 minutes every few days on this task to performing the task very rarely. This allowed maintenance to spend more time elsewhere.

This program gave maintenance technicians a small relief from the frequent repairs they were handling on the automated beam welders. We recommend that Marysville continue to empower production operators to complete safe, minor, and recurring repairs. In aggregate, we expect this will significantly reduce the workload on maintenance technicians.

This program also reversed a trend of punishing maintenance with trivial tasks, and gave the production operators a stake in improving the state of their equipment. The production operators requested tools to keep their machines cleaner and began suggesting improvements that would make their equipment more reliable.

For example, one production operator requested a pick to clean steel shavings out of a slot that regularly got clogged. The operator also became frustrated with frequently replacing cables and suggested that the machine could employ guards to protect the cables from impacts.

7.2 More Accessible Documentation

As we noted in section 5.2, maintenance diagnoses were frequently slow or incorrect. Problems with diagnosis were caused by a vicious cycle of a lack of information leading to intuitive diagnoses, which leads to slow repairperson learning.

The desired state is a virtuous cycle of useful and accurate information enabling a methodical approach to diagnosis, which enables a repairperson to learn quickly. This is the virtuous cycle shown in figure 4-5. However, the attitude and the experience of the repairpersons is not directly addressable by management, so the leverage point is to improve the availability of useful, accurate information.

We designed a program to keep focused, useful information nearby the machines that they detail. We created a pilot of this program, keeping useful information about the automatic beam welders in a binder about 20 ft. away from the machine. We seeded this binder with the written information that management had available:

- Explanatory diagrams, with diagrams communicating a mental model of the system

- Operating manual, with instructions about how to use the machine correctly
- Maintenance manual, with instructions about how to perform regular maintenance
- Reference of part numbers used in the machine, and where they were purchased

We encouraged the maintenance team to make additions and corrections to the above information, and we also made blank sections for the maintenance team to contribute the following information:

- Troubleshooting guide, including common problems and their solutions
- Records of maintenance performed and changes made

We then highlighted the usefulness of the information to the technicians, particularly the most junior technicians. They particularly appreciated having a summary of the system as a whole, which they had never seen before and allowed them to better conceptualize the role that each subsystem and component plays. They also particularly appreciated having a reference of the part numbers used in the machines, because they were able to locate replacement parts much more easily with the part number and vendor in hand. We expect that having such documentation will significantly improve the training experience of new technicians that need to learn about this system.

Finally, the junior technicians encouraged the senior technicians to contribute some of their unwritten knowledge to the binders, particularly in the troubleshooting guide. For example, the senior technicians knew how to fix issues with the wire feeding subsystem quickly, but junior technicians did not know about this particular solution. We expect that if this encouragement continues, the senior technicians will begin to record more useful information in the binders, and potentially adopt a more methodical approach themselves.

For example, we provided some details about the proper configuration of an electronic controller. This was an issue that no one was able to fix because the person

who knew had previously retired. With this information in hand, a junior maintenance technician on night shift was able to fix the electronic controller on their own. They developed a hypothesis that the controller had become misconfigured, and confirmed that the configuration no longer matched the specification. This allowed the machine to be fixed the same night, rather than waiting several days for an electrical contractor to arrive and potentially replace the whole controller.

We recommend that Marysville start assembling and providing critical information about more machines on the shop floor, starting with the machines that the maintenance team finds most confusing and most frustrating. We expect that a further roll out will substantially improve the job quality of the maintenance technicians and become a maintenance resource that the technicians value and want to improve.

7.3 Work Order System

Recall from section 5.3 that maintenance plans were often inefficient or entirely absent. Further recall that poor planning was caused by the repairperson not having access to the information to create a procedure or not sufficiently specifying and verifying the plan.

Also recall from section 5.5 that expectations for task completion were unclear or missing, and maintenance technicians were not assigned efficiently. Further recall that poor execution was caused by the supervisor lacking accessible documentation about previous maintenance work.

The desired state is that the repairperson uses historical planning records to inform a methodical approach to planning their current repair. Then the repairperson critically learns from their approach and applies the learnings across their work wherever applicable. Finally, this leads to better historical planning records. This is the virtuous cycle shown in figure 4-6. However, with no historical records available to management, the only addressable area is encouraging the methodical approach to planning.

We designed a work-order system to encourage and induce more thorough planning

for upcoming maintenance tasks. We piloted this program with a paper ticket system, primarily because the senior leadership was not able to initially spend on digital devices and software. We asked the maintenance team to document all problems that they worked on with the following information:

- Symptoms of the problem(s)
- Cause(s) of the problem(s), with encouragement to identify the root cause
- Plan to address the problem(s)
- Materials required to execute the plan
- Expected time to complete the procedure(s)
- Actual time used, which was recorded when the procedure was complete

We foresaw that the maintenance might receive this system as being micromanaged, even if it was necessary to ensure success. Therefore, we framed the need for this against a recurring concern that we were hearing from maintenance technicians: that they were never allocated enough downtime to complete the preventative maintenance that they wanted to accomplish. We helped them to understand that they were not allocated time because the production supervisors faced a lot of uncertainty about how long repairs might take, and that with more information about the planned length of repairs, supervisors would allocate appropriate time. With this in mind, the maintenance technicians were more willing to plan out the procedure that they wanted to perform and estimate how long it would take so that production supervisors would leave them alone for that long. The production supervisors also appreciated this framing, because they were able to plan around the length of the repair process.

However, the biggest benefit was that the repairperson was able to create and execute a well-planned procedure efficiently. We encouraged technicians to plan how they would accomplish the tasks listed and what materials were needed. Furthermore, we encouraged them to bring a kit of materials before they started the work, which

meant that the machinery was often operating while the repairperson collected parts, and the repairperson never proceeded to disassemble before the parts were available.

For example, we made a list of all outstanding problems known on one automatic beam welder. We estimated the time to fix each problem, and we prioritized four important tasks we could complete within eight hours. We made a detailed plan for how to accomplish each of the four tasks, using the available maintenance labor efficiently. After getting approval from production, we brought the kitted supplies to the machine and only then ceased production. The technicians completed the work efficiently, removing one part that needed to go to the machine shop, then completing three separate tasks while the machining proceeded, and finally, reinstalling the machined part. All told, the work was completed in six hours, and production was able to resume on the machinery as scheduled.

We recommend that Marysville assign maintenance technicians to work closely with production operators, and fully utilize the help and expertise that operators can offer. The maintenance supervisor should promote among technicians that operators are knowledgeable, trustworthy and friendly colleagues. We expect this will have social benefits for the company, as well as increasing the number of maintenance tasks that can be handled concurrently.

We also recommend that Marysville continue to emphasize the importance of making detailed plans for repair work and further that Marysville eventually transition to a mobile-phone based system. The mobile-phone system has the benefit that records are less likely to be lost over time and that the maintenance supervisor can easily review past and upcoming plans.

7.4 Proximal Spare Parts Storage

As discussed in section 5.4, the maintenance system spent excessive time retrieving parts, both locally and externally. Recall this was because the parts were stored far away, disorganized, and not routinely counted.

The desired state is that the most common and most relevant parts can be re-

trieved efficiently, while less common parts are primarily focused on minimizing the number of parts kept in storage. We recommended the simple ABC system that many organizations follow. The most commonly used parts are A class parts and should be retrievable within one minute. Somewhat less commonly used parts are B class parts and should be retrievable within 10 minutes. Finally, the least commonly used parts are C class parts, and the priority is really to minimize inventory of these parts, because they also tend to be the most expensive spare parts (e.g., electronics or complex subsystems).

To handle the A class parts, we developed a program to stock these parts in highly relevant spare parts cabinets shown in figure 7-1. We piloted this program with the automated beam welders, installing spare parts cabinets seconds away (20 ft.). The spare parts cabinet only included a small number of relevant parts. For example, the cabinet contained only four distinct fasteners that were likely to break on the machine, compared to hundreds of distinct fasteners in the previous spare parts room. Altogether, travel time to retrieve common parts was greatly reduced, and the difficulty of locating common parts was also greatly reduced.

For example, with the proximal spare parts cabinets, the time to replace a cam follower fell from 18-25 minutes to 8-10 minutes because the walking time was reduced from 3-4 minutes to under a minute, and searching time was reduced from 2-5 minutes to under a minute. Similarly, the time to replace a fastener fell from 45-55 minutes to just 12-15 minutes because the walking time was reduced from 10-12 minutes to under a minute, and searching time was reduced from 10-15 minutes to 1-3 minutes.

An additional benefit to moving the A class parts to a nearby cabinet was that low inventory levels were more reliably reported. Consider a particular spare part that should be reordered when levels fall below N parts so that replacements arrive before they are needed. First, we tagged the N th-to-last part with a reorder tag, which contains all the information needed to reorder the part. The repairperson simply needs to drop the reorder tag in a dedicated drop box to ensure it is reordered. Second, in case the reorder tags are lost or incorrectly used, we also posted a printed list of the different part numbers and quantities that are supposed to be in the cabinet.



Figure 7-1: A picture of the spare parts cabinets. Visible are the bins containing each unique spare part, with labels applied to the cabinet and reorder cards inside the bins. At the top left is the documentation binder, and at top right are specific tools for the machine. At far left is an listing of the parts and quantities that should be present when inventory is counted.

Thereby, anyone can check that each part is sufficiently stocked, and we recommend that management have someone perform such a check monthly.

When setting inventory levels, we did not have maintenance data available to choose optimal stocking levels with any kind of Newsvendor model. Therefore, we asked the maintenance technicians for their opinions about how many to keep in stock at all times in order to avoid running out. The maintenance technicians provided reasonable guesses and appreciated being consulted. We acknowledged that the stocking levels would be adjusted in the future by simply increasing stocking levels if the maintenance team ever runs out or by reducing the stocking levels if the parts never get close to being exhausted.

For example, the cam followers were promptly reordered when the cabinet ran low because the reorder card was submitted to the supervisor when there were only five units remaining. Because of the advance warning, the new cam followers were ordered and did not require rush shipping. They arrived a few days later when there were still three units remaining in the cabinet. The maintenance supervisor found it much faster to reorder parts with the reorder cards because the information was already there, and was able to spend more time doing supervisory tasks.

In another example, the reorder card for the second to last cable connector had fallen off, and so the second to last unit was used without being reordered as planned. To our surprise, the production operator spent one of their break times checking the contents of the cabinet against the list of what was meant to be inside, and discovered that there was only one cable where there was meant to be two. The production operator let maintenance know and maintenance quickly reordered new cables based on the list and made a new reorder card. This avoided a potential situation where the machine was inoperable because of a missing cable. The production operator was pleased to avoid a shutdown.

We recommend that the maintenance team continue to install spare parts cabinets nearby machines with A class parts. Doing so is expected to reduce travel time and searching time, and to do a better job of managing inventory levels effectively. Additionally, when widely deployed, it is expected to significantly reduce the number

of parts that are remaining in the centralized spare parts room.

To handle the B class parts, we expect that after the A class parts are eventually moved out into specialized cabinets on the shop floor, and after the obsolete or unusable parts are discarded, that the remaining parts could be stored in a single centralized spare parts room. We recommend that the remaining B class parts be clearly labeled and stored on shelves. We also recommend that B class parts be organized by part class (as opposed to organized by equipment served), primarily because many B class parts are applicable to multiple machines on the shop floor, and so the individual likelihoods of needing parts can be pooled to reduce overall stocking requirements. Moreover, organizing by parts class allows special storage needs to be addressed more efficiently. For example, all flammable materials can be stored in the same flammables locker while all steel components that are liable to rust can be stored in a dehumidified locker.

Finally, to handle the C class parts, we recommend that each part be re-evaluated to determine if it truly needs to be stocked or if it is better to just source it from external vendors when needed. For those that truly need to be stocked, perhaps because external vendors cannot deliver for many weeks, we recommend that Heartland's three facilities pool their stocking levels such that spare parts inventory can be reduced. For example, Heartland as a corporation only needs to stock one spare high-frequency transformer for seam welding (which is a large and expensive spare part with a long lead time): the spare can be shipped to any of the three facilities in the rare event that it is needed, which allows the facilities to pool their risk.

7.5 Solving Problems Permanently

As detailed in section 5.6, Marysville was making few improvements to reduce future maintenance workload. We noted that the lack of improvements was caused by the team not having discussions about recurring patterns, rejecting ideas prematurely.

We first planned to resolve the inter-faction conflicts and enable discussions. We created parallel meetings for each faction that had the same structure and the same

content, covering past and upcoming work. For each faction, we highlighted the work that the other faction had done and might not have been visible. This helped both groups to recognize that the other faction was busy doing valuable work and provided some commonality. We were later able to merge the meetings and finally observed the maintenance technicians having discussions about recurring patterns.

We also needed to change the attitude about prematurely dismissing ideas. We had several conversations with Human Resources about a couple of technicians' behaviors, and we made explicit the change in management priorities to eliminating issues even at some cost. Finally, we found that the maintenance team was more willing to pursue permanent fixes.

For example, the maintenance team identified that the issue with the clogged nozzles and breaking pipes could be addressed forever. They verified that an automated system suggested by the plant manager could be tasked with continuously cleaning the water, removing the particulate matter that caused the nozzles to clog and virtually eliminating that problem. Furthermore, they identified that the pipes could be made out of stainless steel and protected by chicken wire, which would eliminate the recurring cause of pipes breaking. This was not implemented until after the internship ended but is expected to vastly decrease the maintenance required on the washer.

Chapter 8

Conclusion

In this project, we characterize the efficiency problems of maintenance at Heartland, and develop targeted improvements to improve maintenance efficiency. The effects of improved maintenance efficiency will be critical in enabling further improvements to the business, but will likely take months or years to become evident in day-to-day operations.

Based on review of the relevant literature on maintenance management, we develop a framework with which to analyze Heartland's maintenance efficiency. This framework focuses on the primary repair process, and specifically on how that process can be efficient or wasteful. We identify several key potential sources of waste in maintenance execution.

We apply that framework to the specifics at Heartland. Doing so reveals a number of clear problems that were causing poor performance. Heartland's maintenance team had problems with the nature of the incoming tasks, making efficient and correct diagnoses, making well thought-through plans, getting parts efficiently and with avoiding issues altogether.

We then estimate the impact of these maintenance problems from empirical data collected by Heartland's MES. With this data, we estimate the hours of potential production that were lost to breakdowns.

Finally, we piloted targeted improvements to address the specific maintenance execution problems at Heartland. We proposed a new program of shared equipment

ownership between operators and maintenance to improve the nature of the incoming tasks. We wrote more accessible and useful documentation to improve the efficiency and quality of diagnoses. We initiated a more explicit process of work orders to promote better planning of repairs. We demonstrated a spare parts cabinet nearby a critical machine to make getting parts more efficient. We supported modifications that would significantly increase reliability and reduce maintenance workload.

Further work could use maintenance work order data and performance data to quantitatively estimate the impact of each individual problem and intervention. Additional work could also use optimization to allocate responsibility for various repair types or locate spare parts optimally on the shop floor to minimize both inventory and retrieval costs. Further work could also apply the maintenance framework in this thesis to other small manufacturers, and further develop the framework based on additional experience.

An efficient maintenance execution program is an important complement to a maintenance strategy. More efficient maintenance improves the amount of maintenance that can be done or reduces maintenance costs, regardless of how maintenance tasks are prioritized. Simple improvements can improve the day-to-day productivity of maintenance personnel, which is profitable for the business and creates more fulfilling manufacturing jobs for the local community.

Bibliography

- [1] Catherine Clegg. Letter of Understanding: Production Maintenance Partnership (GM), October 2015.
- [2] J. L. Coetzee and S. J. Claasen. Reliability Centered Maintenance for Industrial Use: Significant Advances for the New Millennium. *The South African Journal of Industrial Engineering*, 13(2), January 2012.
- [3] David Michael Feliciano. *Exploring Barriers, Enablers, Justification and Planning Methods for Total Productive Maintenance Implementation in Automated Production of Commercial Airplanes*. PhD thesis, Massachusetts Institute of Technology, June 2015.
- [4] D. J Hitt and M. Gilbert. Tensile properties of PVC at elevated temperatures. *Materials Science and Technology*, 8(8):739–746, August 1992.
- [5] Casey Ichniowski, Kathryn Shaw, and Giovanna Prennushi. The Effects of Human Resource Management Practices on Productivity. *National Bureau of Economic Research*, (5333), November 1995.
- [6] Anthony Kelly. *Strategic Maintenance Planning: An Essential Guide for Managers and Professionals in Engineering and Related Fields*. Plant Maintenance Management. Elsevier Butterworth-Heinemann, Amsterdam Heidelberg, 1. ed., reprinted edition, 2007.
- [7] Anthony J. Kramer. *An Integrated Approach to Maintenance in a Three-Crew, Two-Shift Environment*. PhD thesis, Massachusetts Institute of Technology, June 1997.
- [8] Bill Maggard and David Rhyne. Total Productive Maintenance: a Timely Integration of Production and Maintenance. *Production and Inventory Management Journal*, 33(4):6, 1992.
- [9] Seiichi Nakajima. *Introduction to TPM (Total Productive Maintenance)*. Productivity Press, 1988.
- [10] U.S. Bureau of Labor Statistics. Producer Price Index by Commodity: Metals and Metal Products: Cold Rolled Steel Sheet and Strip [WPU101707], March 2022.

- [11] Doc Palmer. *Maintenance Planning and Scheduling Handbook*. McGraw-Hill, New York, 2nd edition, 2006. OCLC: ocm62282525.
- [12] Aditya Parida and Uday Kumar. Maintenance Performance Measurement (MPM): Issues and Challenges. *Journal of Quality in Maintenance Engineering*, 12(3):239–251, July 2006.
- [13] C. Roser, M. Nakano, and M. Tanaka. Shifting Bottleneck Detection. In *Proceedings of the Winter Simulation Conference*, volume 2, pages 1079–1086, San Diego, CA, USA, 2002. IEEE.
- [14] Laura Swanson. Linking Maintenance Strategies to Performance. *International Journal of Production Economics*, 70(3):237–244, April 2001.
- [15] Albert H.C. Tsang. Strategic Dimensions of Maintenance Management. *Journal of Quality in Maintenance Engineering*, 8(1):7–39, March 2002.
- [16] Gene E Wilkins. The Role of Job Classification in Collective Bargaining. *Indiana Law Journal*, 32(2):173–192, 1957.
- [17] Jacqueline Ming-Shih Ye. *Improving Maintenance Operation through Transformational Outsourcing*. PhD thesis, Massachusetts Institute of Technology, June 2007.