IMPROVING RAILROAD FREIGHT CAR RELIABILITY USING A NEW OPPORTUNISTIC MAINTENANCE HEURISTIC AND OTHER INFORMATION SYSTEM IMPROVEMENTS

by

PATRICK LITTLE

B.A., St. John's University (1977)
M.S.T., Massachusetts Institute of Technology (1982)

SUBMITTED TO THE DEPARTMENT OF
CIVIL ENGINEERING
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF

DOCTOR OF SCIENCE IN CIVIL ENGINEERING

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

January, 1991

Signature of Author . . . .
Department of Civil Engineering
January, 1991

Certified by . . . .
Joseph M. Sussman
Professor, Department of Civil Engineering
Thesis Supervisor

Accepted by . . . .
Professor Ole S. Madsen
Chairman, Departmental Committee on Graduate Studies
Department of Civil Engineering

IMPROVING RAILROAD FREIGHT CAR RELIABILITY USING A NEW OPPORTUNISTIC MAINTENANCE HEURISTIC AND OTHER INFORMATION SYSTEM IMPROVEMENTS

by

PATRICK LITTLE

Submitted to the Department of Civil Engineering, January, 1991
in partial fulfillment of the requirements for the degree of
Doctor of Science in Civil Engineering

ABSTRACT

Railroads and other car owners spend billions of dollars annually to maintain the more than 1.2 million freight cars used to provide service in the United States. In spite of the considerable investment this equipment represents, maintenance policies and practices are quite simple, in part because of a complex operating environment, and in part because of lack of feasible policies which can be shown to work effectively. In this thesis, the current situation regarding freight car maintenance is critically examined, and new car maintenance policies are proposed and tested.

Reviewing the current situation leads to several important conclusions. First, the current policies being followed are suboptimal. They fail to incorporate potentially useful information regarding the reliability and costs of components and maintenance activities. They also fail to exploit potential economies of scale in maintenance. To address this, a new opportunistic maintenance heuristic is developed and tested. Second, the measures being used to monitor the maintenance function could be enhanced by use of measures which are consistent with reliability theory. Two such measures, miles per in service failure and miles per maintenance event are proposed. Finally, the information systems used to support the car maintenance function produce large amounts of data, but need considerable reorganization to be readily useful to managers. The use of a "structured car history" is suggested and demonstrated.

The opportunistic heuristic, along the currently followed policies and several other alternatives, is tested using a simulation model. The simulation showed that a number of alternatives are more attractive than the present practices. In particular, the opportunistic policy performs well and is robust over a wide range of circumstances. So called "far sighted" hard time policies, in which the car is brought into the shop at fixed intervals and then subjected to aggressive replacement of parts can also perform well, although such policies are sensitive to the failure distributions of the parts included.

Thesis Supervisor: Dr. Joseph M. Sussman

        Title: Professor of Civil Engineering

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

LIST OF TABLES

LIST OF FIGURES

IMPROVING RAILROAD FREIGHT CAR RELIABILITY USING A NEW OPPORTUNISTIC MAINTENANCE HEURISTIC AND OTHER INFORMATION SYSTEMS ENHANCEMENTS

Chapter 1
Introduction

1.1. Problem Statement

Freight cars are an integral part of the railroad system of the United States and Canada. To enable the movement of goods, the U.S. railroads and shippers own and maintain a fleet of more than 1.2 million freight cars. It is not surprising that huge sums of money are spent to keep these cars in working order. In spite of these expenditures, the maintenance policies followed by the owners of many of these cars could be improved in a way which would both reduce maintenance costs and improve car reliability. That is the subject of this thesis.

In the following chapters, it will be shown that current railroad car maintenance practice can best be described as falling under two categories, "on-condition" maintenance, and to a lesser degree, "hard-time" planned maintenance. "On- condition" maintenance calls for components to be replaced when they fail or when they reach a given condition, known as the "condemnation limit". "Hard time" policies provide for cars to be brought into repair shops at fixed intervals of time or mileage, and most of the car's systems restored to like-new condition. It will be shown that both approaches to maintenance are subject to serious problems. "On-condition" policies generally fail to utilize information about such things as modes of failure, costs of failure, and potential economies of scale in maintenance, resulting in practices which are expensive without necessarily improving the reliability of the car as an entire system. "Hard time" policies, unless very carefully developed and implemented, are inconsistent with both theoretical results and the experience of other modes, and, like "on-condition" policies, can be expensive without gaining positive results.

That railroad car owners have adopted these types of policies reflects both the

unique environment of railroad operations and the lack of attractive alternatives which can be implemented. The railroad maintenance environment is unique among modes (and production systems in general) in that the car is often in use on another railroad (interchange service), and subject to the maintenance policies mandated by semi-regulatory bodies (primarily the Association of American Railroads).

It will be shown in later chapters that a desirable alternative might be to implement what are known as opportunistic maintenance policies. Such policies treat the failure of one component or system as a potential opportunity to replace other components if there are sufficient economies of scale in performing the additional maintenance. Unfortunately, while such policies are attractive at this level of description, they are quite difficult to implement in practice. A solution to this is proposed in this thesis, and a heuristic is developed. Basically, this approach first schedules each individual component as if it were the only component in the system. When any component fails, all the other components are treated as candidates for preemptive replacement. The decision on which unfailed components to replace is based on an assessment of the costs of performing the replacement at this time versus the expected costs if the component is left in service. While not optimal, simulation results show that such an approach can outperform the current practice. This opportunistic approach should be particularly useful in the railroad environment since it does not depend on the same policy having been followed while the car is off line.

There remain, however, other barriers to implementing a more complex approach to maintenance. The most significant of these is that the volume of data and the structure of the information systems tends to discourage managers from trying new approaches to car repair. This problem not only manifests itself in the context of developing new approaches to car maintenance planning, but also in managing the current practices. A solution to the general problem is proposed in the thesis, the "structured history". The structured history has its roots in two recent developments in computer software: knowledge engineering and object-oriented programming. In both cases, the intent is to place the focus on knowledge of the problem at hand (the car, in this case), rather than on data processing issues. A structured history provides a way of organizing the data

9

available regarding a car, its use, and its mechanical reliability to correspond to the way that railroad car maintenance managers actually think about cars. This approach should prove useful not only for collecting and organizing data to support maintenance policies, but also as a basic resource for managing operations. It has already been used as part of an expert system for diagnosing cars which are experiencing excessive consumption of wheels and brake shoes. Such cars result in high repair costs and low reliability, and are found among most railroad car fleets. It was found that by using the structured histories as part of an "expert system", these cars could be detected, diagnosed, and repaired, resulting in substantial savings[1].

1.2. Research Contributions

This research makes three contributions to the state of knowledge about transportation systems and vehicle maintenance. First, it critically documents current practices and policies in the area of railroad car maintenance. It does this both in general terms by presenting an overview of the industry and in more detail in a series of three case studies of railroad car owners. These include a small regional railroad, a large (Class I) carrier, and a private company with its own fleet of specialized cars. This review of the current practice leads to three important conclusions:

- The maintenance policies currently being followed by railroad car owners should be reconsidered in light of the theory of reliability.

- In many cases, the measures of maintenance effectiveness being used are not appropriate and could be replaced with indices that better represent railroads' concerns with reliability and cost control.

- The current methods of organizing the information used to manage car maintenance activities are more reflective of historical concerns with billing and accounting than the information needed to support better car repair decisions.

The second, and potentially the most widely applicable contribution is the

---

[1] Little and Martland (1989).

development of a theory-based approach to opportunistic maintenance, a heuristic for opportunistic running repairs. This approach uses information currently available to railroad car owners and permits the car owner to make real-time decisions about maintenance actions while the car is on the repair track. Using a simulation model it is shown that use of the heuristic can reduce maintenance costs significantly and increase car reliability.

The third contribution is a set of practical tools to assist railroad car maintenance managers (and supporting information services professionals). These include the development and demonstration of the structured history, a simulation model for analyzing alternative maintenance policies, and a methodology for analyzing the tradeoffs between competing policy options.

1.3. Structure of the Thesis

The thesis is organized along the lines of the contributions and can be thought of as following the flow chart shown in Figure 1.1. After a chapter which defines basic terms, concepts, and performance measures for maintenance (Chapter 2), the balance of the research consists of two parts. The first part presents and assesses the current state of freight car maintenance; the second part develops and tests alternative methods which resolve the problems found in that assessment.

Part 1 begins with an overview of the freight car and the car maintenance process (Chapter 3). An important part of that presentation is a discussion of the problem of having cars moved and maintained under the control of other carriers (interchange). This aspect makes the environment of railroad car maintenance unique in the transportation industry, and makes the problem of maintenance planning much more complex than in the single company case. Chapter 4 examines the structure, practices and policies of the case study companies, including their information systems used to support car maintenance activities. Chapter 5 assesses car maintenance practice in the U.S. and Canada, evaluating the policies currently being followed and the information systems used to support car maintenance.

Part 2 presents opportunistic maintenance policies as an alternative to the current

```
┌──────────────────────────────────────────────────────────────────────┐
│                                                                        │
│  ┌──────────────┐          ┌──────────────────┐                        │
│  │Definitions,etc.│────────▶│   Overview of    │                        │
│  │ (Chapter 2)  │          │   Freight Car    │                        │
│  │              │          │ Operations and   │                        │
│  └──────────────┘          │  Maintenance     │                        │
│         │                  │                  │                        │
│         │                  │   (Chapter 3)    │                        │
│         │                  └──────────────────┘                        │
│         │                           │                                  │
│         │                           ▼                                  │
│         │                  ┌──────────────────┐                        │
│         │                  │  Case Studies    │                        │
│         │                  │  (Chapter 4)     │                        │
│         │                  └──────────────────┘                        │
│         │                           │                                  │
│         │                           ▼                                  │
│         │                  ┌──────────────────────────┐                │
│         │                  │      RESULT:             │                │
│         │                  │     (Chapter 5)          │                │
│         │                  ├────────────┬─────────────┤                │
│         │                  │ Current    │ Information  │                │
│         │                  │ Practice   │ Systems Used │                │
│         │                  │ is Flawed  │ by Car Owners│                │
│         │                  │     •      │ Too Complex  │                │
│         │                  │ Car Owners │ for Effective│                │
│         │                  │ Need Better│ Maintenance  │                │
│         ▼                  │ Approach   │ Management   │                │
│  ┌──────────────────┐      └────────────┴─────────────┘                │
│  │ Opportunistic    │           │              │                       │
│  │ Policies         │◀──────────┘              │                       │
│  │ Appear Well-Suited,│                         │                       │
│  │ but Few Practical│                          │                       │
│  │ Implementations  │                          │                       │
│  │ Available        │                          │                       │
│  └──────────────────┘                          │                       │
│         │                                      │                       │
│         ▼                                      ▼                       │
│  ┌──────────────────┐              ┌──────────────────┐                │
│  │   Solution:      │              │   Solution:      │                │
│  │ Opportunistic    │              │ Structured History│               │
│  │ Maintenance      │              │                  │                │
│  │ Heuristic        │              │  (Chapter 9)     │                │
│  │ (Chapters 6,7&8) │              └──────────────────┘                │
│  └──────────────────┘                                                  │
│                                                                        │
│                      Figure  1.1                                       │
│                                                                        │
│                   Basic Structure                                      │
│                   of  the  Thesis                                      │
│                                                                        │
└──────────────────────────────────────────────────────────────────────┘
```

Figure 1.1

Basic Structure
of the Thesis

approaches to car maintenance (Chapter 6). Because there are serious technical problems which restrict the use of opportunistic policies, a theory-based heuristic for opportunistic running repairs is developed and tested using a simulation model (Chapters 7 and 8). Other alternatives to the current practice are also tested. Chapter 9 addresses the problem of the information systems being difficult to use to support maintenance management, and presents an approach to unifying the data, the structured history. It also discusses potential uses of that approach both for building maintenance policies and for managing day to day operations.

Finally, some conclusions and directions for future research in the area of railroad car maintenance are presented.

Chapter 2

Basic Concepts of Maintenance and Reliability

2.1. Introduction

This thesis is concerned throughout with maintenance and reliability. While most readers already have some conception of what is meant by these and other related terms, it is useful to state some definitions both for the sake of precision and to avoid confusion in later chapters. Similarly, most readers have some sense of *why* managers are concerned with these topics, but it is worthwhile to examine these motivations in detail. A natural extension of these matters is the definition and measurement of maintenance quality. These are the topics addressed by this chapter.

The first section of the chapter presents definitions that will be used throughout the rest of the thesis. Most of these definitions are drawn from standard works in the literature of reliability and maintenance (and the appropriate standard setting bodies), but in a few cases new terms have been introduced, since the railroad environment presents some unique circumstances for the maintenance manager.

The second section presents an overview of the basic reasons why maintenance is performed in general and on railroads in particular. An understanding of the motivation for vehicle maintenance is an important prerequisite to discussion of the current practices of railroads and other car owners, since the "quality" of maintenance is best measured not so much in terms of abstract calculations as in how well the maintenance plan meets the car owner's goals. Four basic reasons why car owners undertake to maintain their fleets are presented and discussed.

Finally, appropriate performance measures for evaluating maintenance are discussed. Sound measures must address both the cost and effectiveness of maintenance policies and actions. Three measures are proposed for assessing railroad car maintenance activities.

14

2.2. Definitions

Everyone who has owned a broken down car, radio, or toaster has practical knowledge of reliability and, if the broken item must be repaired (or brought to a shop for regular servicing) of maintenance. In the following sections, these and other terms which are used throughout the thesis are defined.

First, the basic terms of maintenance and some maintenance policies are defined. Because these are implemented using tools derived from reliability engineering, some of the central concepts of that discipline are then presented. Maintenance and reliability are then united by examining replacement policies, which formalize the decision of when, if ever, to replace a component which is still serviceable.

2.2.1. Maintenance Terms

The 1974 British Standards for Terotechnology[1] [BS 3811: 1974, quoted in Corder (1978)] define a number of terms relating to maintenance and repair. **Maintenance** is defined as "a combination of any actions carried out to retain an item in, or restore it to, an acceptable condition". **Planned maintenance** is defined as "maintenance organized and carried out with forethought, control, and records to a predetermined plan". **Preventive maintenance** is "maintenance carried out at predetermined intervals, or to other prescribed criteria, and intended to reduce the likelihood of an item not meeting an acceptable condition".

Maintenance activities are generally directed in response to a **maintenance policy**. McCall (1965) offers the following definition:

> Any rule that assigns a specific action to any realization in the equipment's life is called a maintenance policy. More precisely, a maintenance policy is a function whose range is the set of possible maintenance actions, the action space, and whose domain is the set of possible realizations in the equipment's history, the information space.

In other words, a maintenance policy is an explicit plan which directs what action, if any, is to be taken when the equipment reaches any particular condition.

Carter (1986) makes the distinction among maintenance policies as being either

---

[1] Terotechnology is the formal name sometimes applied to the study of maintenance.

scheduled or unscheduled, and splits unscheduled maintenance into two further categories: repair maintenance and on condition maintenance. He defines three standard maintenance policies to go with these:

> **Scheduled maintenance** (Preventive maintenance) is carried out to keep equipment in a satisfactory operational condition by providing systematic replacement of components before they are expected to fail - it may include some inspection activity.
>
> **Repair maintenance** is carried out on a non-scheduled basis to restore an item to a satisfactory condition by providing immediate correction of a failure after it has occurred.
>
> **On condition maintenance** is carried out before an item fails, but only when its condition, established by continuous monitoring, indicates that failure is imminent.

In a later chapter it will be seen that the third of these, on condition maintenance, describes the maintenance policies of many of the U.S. freight railroads.

Several terms which are frequently used to describe maintenance activities are repair and replacement. A **repair** is the same as what the 1974 British Standards define as corrective maintenance, namely "maintenance activity carried out to restore ... an item which has ceased to meet an acceptable condition". This definition does not require that the item be restored to "good-as-new" condition. When, on the other hand, an item is removed and a new or "good-as-new" item installed, that is referred to as a **replacement**, regardless of whether the action was performed in response to planned maintenance or to the item reaching an unacceptable condition. It is important to note that a repair may be of a minimal nature (i.e., the item is restored to its condition immediately prior to failure), while a replacement resets the item's condition to new, (i.e., the item's "age", however measured, is reset to 0 in a replacement.) This and the following chapters are primarily concerned with replacement activities, although there has been considerable interest among theoreticians in the general implications of minimal repair strategies in recent years [See, for example, Nakagawa (1981, 1987) and Ohashi and Miyamoto (1987)].

2.2.2. Reliability Terms

In this section, some of the standard terms of reliability engineering are presented.

16

Many of these terms are of a more formal, i.e., mathematical, nature, which, in subsequent chapters permits the development of specific techniques and tools to assist the maintenance manager.

A distinction is made throughout between a system and the components which comprise it. A **component** is a unit of sufficient size or scale that it can be repaired or replaced integrally. Thus a railroad car wheelset can be a component, but an individual wheel cannot, since it can only be replaced together with the axle, bearings, and opposing wheel. Similarly, all four wheelsets on a car could be treated as a single component if the decision were made always to replace them at the same time. The decision regarding the appropriate level of detail to use in defining components must be considered a matter of judgement. A **system** is a set of components or systems which together perform a unified function. Notice that the definition of system is recursive; this is because a system such as a locomotive may be composed both of components (e.g., couplers) and lesser systems (e.g., the engine).

Typically, the **failure** of a component is defined as the state at which it is unable to provide an acceptable level of service (due, for example, to breakdown, wearout, or malfunction). British Standard 4778 defines a failure as "the termination of the ability of an item to perform a required function" [quoted in O'Conner (1985)]. It is useful to distinguish between various types of failures, as these bear upon the number of ways that components can cease to function and the differing consequences of component failure. The first distinction to be drawn is whether or not the system is in use when a component fails. An **in-service failure** is a failure of a component which is detected while the system of which it is a part is in use, and which requires repair or replacement upon detection. The alternative, **incidental failure**, is a component failure which is detected either while the system is not in use or when the system is already undergoing repairs.[2] An example of an in-service failure would be the finding that a coupler operating lever

---

[2] It is possible, for example, to detect that a component has reached a failure state while the system is already undergoing repairs. In general this is less costly than if the system is being used at the time that the failure occurs and is detected.

17

is broken during a routine inspection of a train prior to departure from a terminal. The component must be replaced prior to allowing the car to continue in service, even though it involves disruption to the shipper. An example of an incidental failure would be the detection of a broken center pin under a truck when the car has already been elevated off the trucks to replace the wheels. The broken center pin would not (indeed, could not) have been detected except incidentally to the other maintenance or repair action. If the repair can be made easily, there may be little or no additional service disruption.

Another useful distinction among failures is whether the component is unable to function or whether the component's condition is simply unacceptable because of a maintenance standard. In the case of industries subject to safety regulations, such as railroads, components may be required by law or industry agreement to meet certain wear limits or performance standards to remain in use. Failure to meet such externally imposed wear or condition limits will be referred to throughout as **regulatory failure**. Regulatory failures can be of either the in-service or incidental type, since the determination that a component fails to meet the standard is independent of whether the car is in use. It is generally the case that both the standards imposed by government regulators and those agreed to among the railroads themselves are preventive in nature, and so reflect conditions which precede the inability of the component to provide acceptable service, a state we will refer to as an **engineering failure**. From the perspective of a railroad car owner, a regulatory failure is the most frequently encountered "failure" mode for components which are subject to inspection, and can have many of the same consequences as an in-service failure in terms of service disruption and replacement cost. That is because a component which fails to meet a regulatory requirement must often be removed and replaced immediately, even though it is still physically serviceable. Further, such a repair is generally billed to the car owner at the same rate as if the component is unusable for engineering reasons. The costs of delays to traffic or lost customer goodwill are potentially significant, and depend on the particular circumstances of the shipment. In some cases, the various regulations provide that components which exceed mandated standards but do not present a high risk to the car or other cars in the train may be permitted to complete their loaded trip and be repaired when empty, or returned to the

18

car owner after the current loaded trip and then repaired.

Catastrophic failure is the worst sort of in-service failure; it occurs when a component ceases to provide an acceptable level of service while in operation such that the system it is part of *immediately* ceases to function. If a railroad car is in motion, the catastrophic failure of a component can result in the derailment of the car and of other cars in the train. The results can range from the stopping of the train while the car is set off or repaired to a derailment involving loss of life and property damage in the millions of dollars. It is in the hope of avoiding catastrophic failures of components that government and industry groups establish standards for removal of worn or damaged components prior to the engineering failure of the component.

The various types of failures can be represented using a set diagram as in figure 2.1. Notice that the regulatory "failures" may be in either the failed or unfailed category, that is, the component may or may not be unserviceable from an engineering standpoint. This reflects some of the ambiguity which is necessarily introduced by the mandating of replacements prior to actual failure. In general, the standards set by the various regulatory bodies are the "condition" used in "on condition" policies applied to railroad cars.

2.2.2.1. Formal Representations of Reliability

When analyzing component failures, a particularly useful tool is the **failure distribution**, F(t), which is the probability that component x fails at or before time t (Figure 2.2). More formally,

$$F(t) = P\{t \leq t\} \tag{2.1}$$

where t is the time to failure for component x. F(t) is a cumulative distribution function, so naturally, F(∞) = 1, and F(t) = 0 for t < 0. The corresponding **failure density** function is denoted f(t) (Figure 2.3), and is defined as

$$f(t) = \frac{dF(t)}{dt} \qquad \cdot \tag{2.2}$$

19

Unfailed Components

Regulatory "Failures"

Incidental

Failures

In-Service

Catastrophic Failures
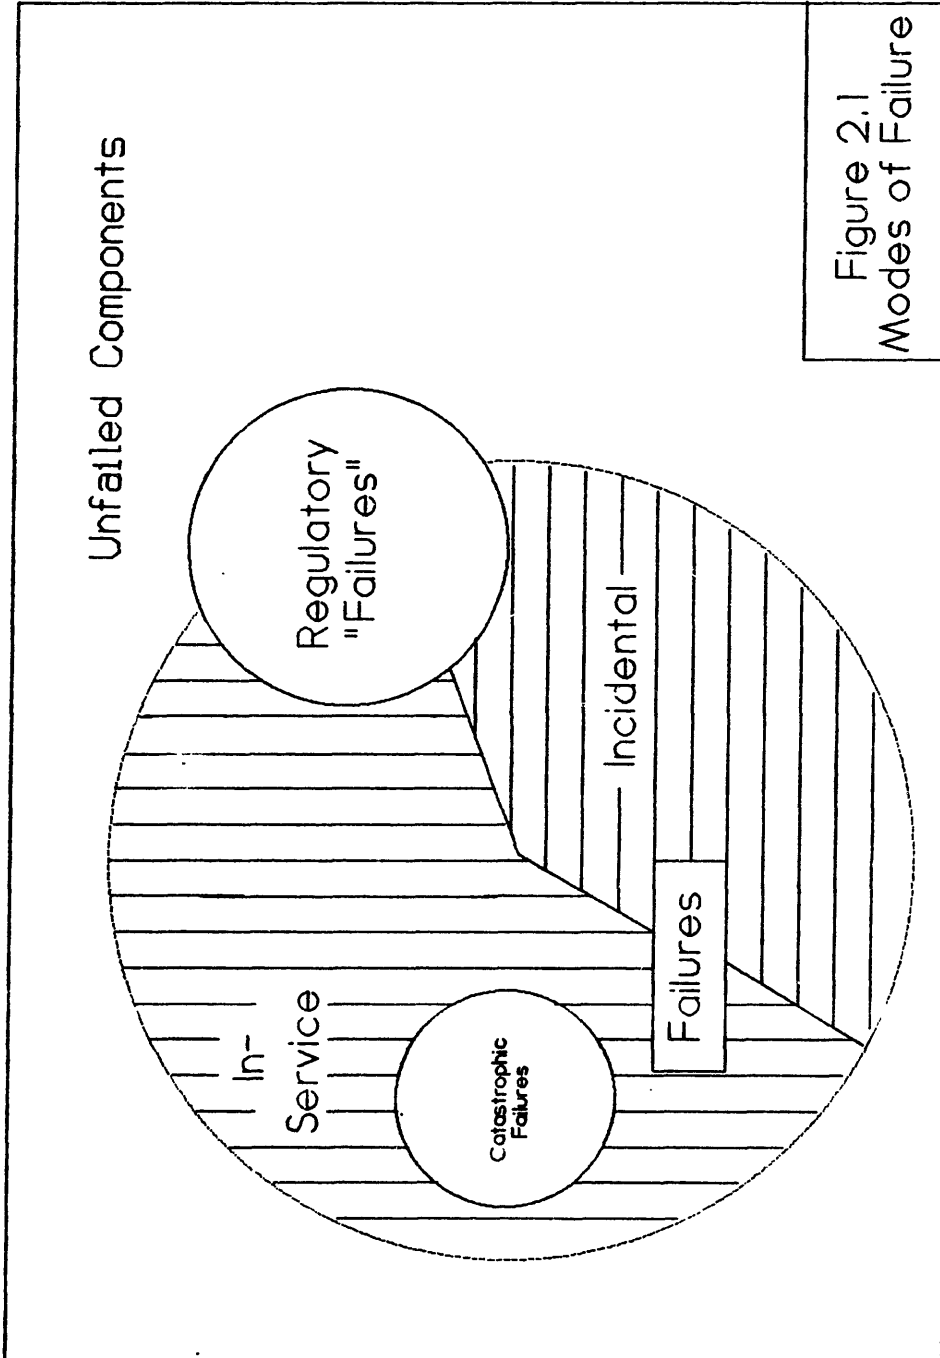
Figure 2.1
Modes of Failure
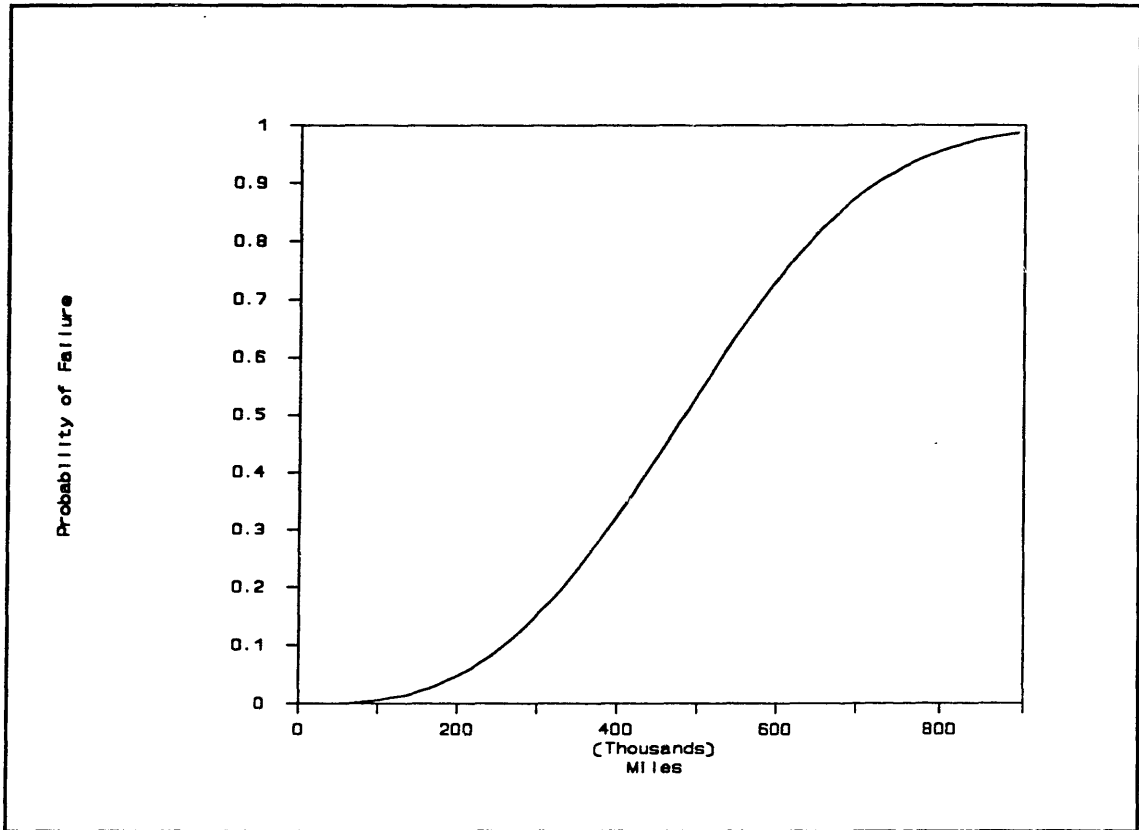
Figure 2.2
Failure Distribution
(Weibull Example)

This can be thought of as the rate of failures occuring at time t.

The **failure rate**, r(t), (Figure 2.4) is the conditional probability that x will fail, given that it has survived to time t, and is given by:

$$r(t) = \frac{f(t)}{[1 - F(t)]} \qquad (2.3)$$

for values of F(t) < 1. The failure rate is also referred to as the **hazard rate** and the **force of mortality**. Distributions for which r(t) is increasing in t are said to be **IFR (increasing failure rate)**. Similarly, where r(t) is decreasing in t, the distribution is said to be **DFR (decreasing failure rate)**. The exponential distribution exhibits a **CFR (constant failure rate)**, and has been proven to be the only distribution to do so [Barlow and Proschan (1965)].

21
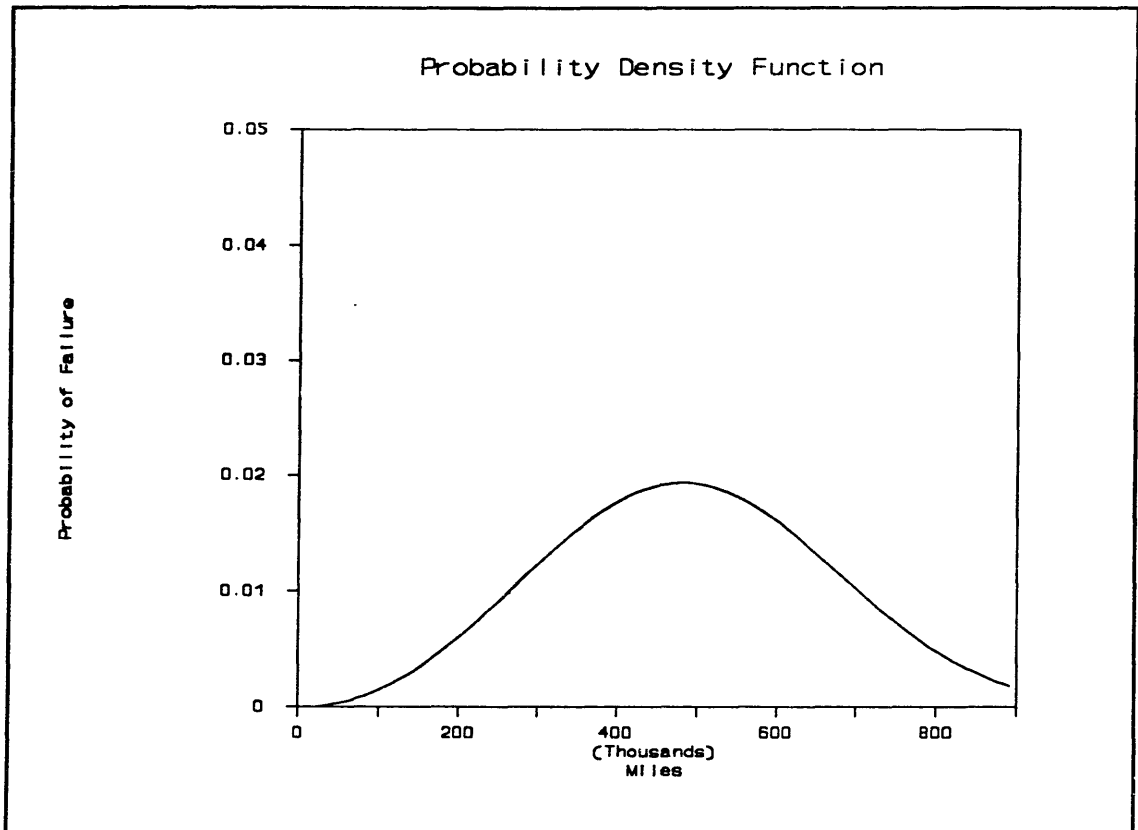
Figure 2.3
Density Function
(Weibull Example)

Many items have been studied and their failure distributions determined[3]. The failure of most components seems to be characterized by the so-called "bathtub curve", where the component exhibits a decreasing failure rate in the early life (infant mortality due to improper installation or manufacturing defects), a more or less constant failure rate for some period of time when the component may be subject to random shocks or other external influences, and finally an IFR period, when the component exhibits wearout. In such cases the distribution is often modelled by treating the three periods as separate curves. When manufacturers carefully control their processes or subject components to a period of "burn in", most mechanical components exhibit IFR distributions. Items which are not IFR are certain "wear-hardened" materials (DFR), electronic components

---

[3] Some of these studies are summarized in Jorgenson, McCall, and Radner (1967).
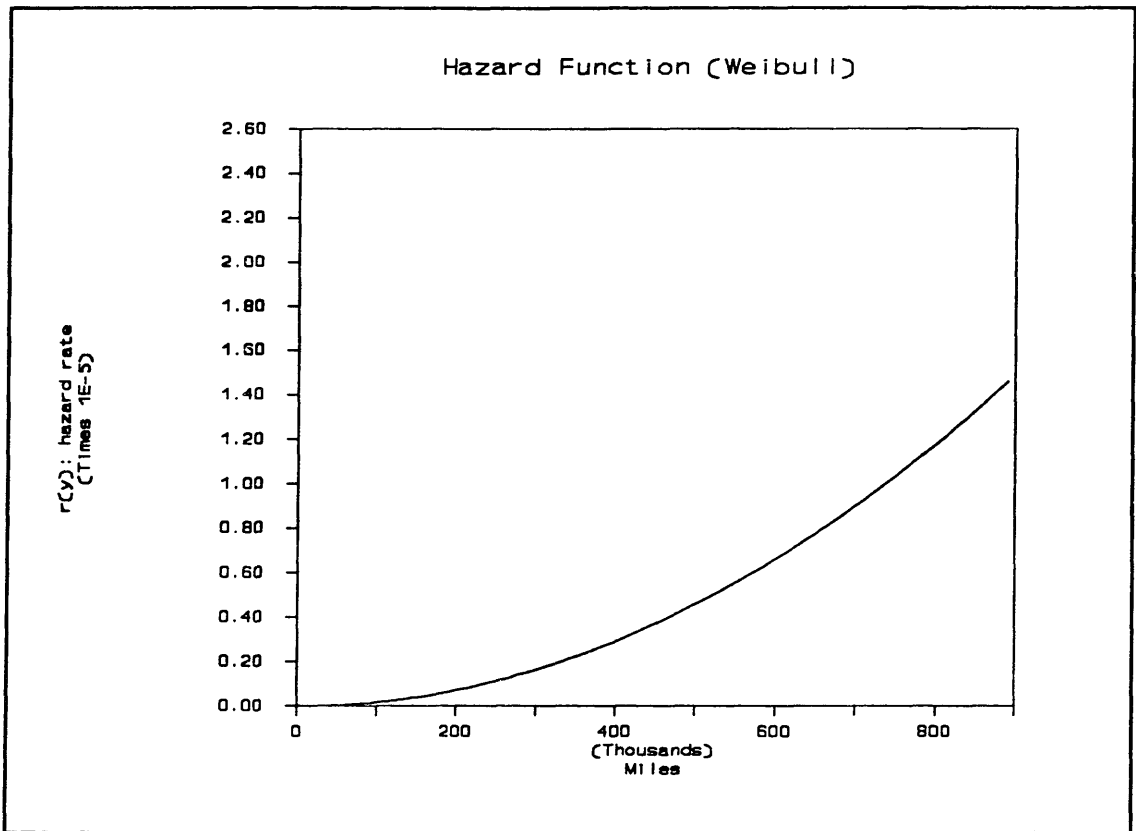
Figure 2.4
Failure Rate
(Weibull Example)

(exponential), and software (DFR, after debugging). Barlow and Proschan (1965) show that complex systems of IFR components generally tend to the exponential at the system level, an important result we will encounter again in discussions of so called "hard-time" policies used by railroads.

Reliability has been given a large number of definitions over the years. These include "the probability that a system survives for some specified period of time" [Lewis (1987)], and "the probability of a device performing its purpose adequately for the period of time intended under the operating conditions encountered"[Radio-Electronic-Television Manufacturers Association, quoted in Barlow and Proschan (1965)]. British Standard 4778 defines it as "the ability of an item to perform a required function under stated conditions for a stated period of time" [O'Conner (1985)]. Notice that while the general sense is the same in these definitions, there is a difference, namely whether or not to

define reliability as a probability. Throughout, the working definition used will be that of the U.S. Military Standards Handbook 217B (MIL-STD-217B, 1970), which defines reliability as "the probability that an item will perform its function under stated conditions of use and maintenance for a stated measure of the variate (time, distance, etc.)" [quoted in Carter (1986)]. This definition points up several matters that are relevant to railroad car maintenance and reliability:

- the reliability of an item is properly measured under the assumption that it has been under some particular set of usage and maintenance conditions;
- the appropriate measure of use may well be something other than time (e.g., miles);
- reliability can be treated as a probability, and hence, characterized by a distribution.

The **reliability distribution**, R(t), the probability of survival beyond t, can be defined in terms of its obverse, the failure distribution, and is given by

$$R(t) = 1 - F(t) \tag{2.4}$$

2.2.3. Replacement Policies

Maintenance and reliability theory come together in the area known as **replacement policies**. Replacement policies (Barlow and Proschan [1965]) are maintenance policies in which a decision is made whether or not to replace a component or system while useful life remains, and, if so, at what time to perform the replacement. After examining the conditions under which it makes sense to consider such preventive replacement of a component, we first examine the most common case of a single component, age replacement. Since most systems of interest are made up of multiple components, however, we shift our attention to multicomponent replacement problems, and discuss two well-studied approaches, block replacement and opportunistic maintenance policies.

In order for a preventive maintenance approach to be optimal, two conditions must be met. The first is that the component must be subject to some sort of worsening with time (or usage). Barlow (1963) has shown that unless a component has an increasing

failure rate (IFR), preventive periodic maintenance policies cannot be optimal and should not be considered. If the component is not IFR, one is generally replacing a component which is either "improving" or, at the least, not deteriorating, with another which is no more reliable[4]. As stated earlier, most mechanical components are IFR over some or all of their life, although some are not, usually when the smallest integrally replaceable part is itself composed of many mechanical parts, such as roller bearings [Guins (1987)].[5] The second condition which must be met for preventive maintenance to be worthwhile is that an in-service, or unscheduled, replacement must be more costly (however measured) than a scheduled one. This is because a lower cost for scheduled replacement is necessary to compensate for the foregone component life or usage until failure.

Given these general considerations, we now focus on the most commonly studied single component approach to maintenance, age replacement policies.

2.2.3.1. Age Replacement Policies

An age (or usage) replacement policy consists of replacing a component after a predetermined interval of time (or other measure of usage, such as miles), or upon failure, whichever occurs first.

If we let C_1 and C_2 be the costs of replacing a component after failure and prior to failure, respectively, then C(t), the maintenance costs associated with the component in the interval of length t are

$$C(t) = C_1 N_1(t) + C_2 N_2(t) \tag{2.5}$$

where $N_1(t)$ is the number of in-service failures in the interval, and $N_2(t)$ is the number

---

[4] While generally the case, this statement is not strictly true. If either replacement components or installation techniques have improved since the original component was placed in service, it may be the case that replacement of a non-IFR component will still result in a "better" component.

[5] It has also been shown that when the failure distribution is unknown and cannot be reasonably assumed to be IFR, then even if the mean failure time and the corresponding variance are known, the optimal maintenance policy is to replace only upon failure, i.e., not to maintain the component preventively [Barlow (1963)].

of premature or preventive replacements. Similarly, C(T), the expected cost per unit time of operating the component when replaced at failure or time T (whichever occurs first) up to the time of the first replacement for either reason is given by[6]

$$C(T) = \frac{C_1 F(T) + C_2 [1 - F(T)]}{\int_0^T [1 - F(t)] dt} \tag{2.6}$$

The denominator of equation 2.6 is the mean time between replacements for a component which is never allowed to remain in service beyond T. The numerator is composed of 2 parts; $C_1 F(T)$ is the cost associated with a failed component times the probability that a component fails prior to time T. The second part of the numerator, $C_2 [1 - F(T)]$, is the cost of a scheduled replacement at time T times the probability that the component survives until T.

The optimal replacement interval, T, is found by taking the derivative of the above and setting it equal to zero, which yields

$$r(t) \int_0^T [1 - F(t)] dt - F(T) = \frac{C_2}{(C_1 - C_2)} \tag{2.7}$$

The solution to this is unique if $r(T)$ is strictly increasing, which is the case if and only if log $[1-F(T)]$ is concave [Barlow (1963)] . In the same reference, Barlow has shown that when a component is IFR then the optimal replacement interval T will always satisfy the inequality

$$T \geq \frac{C_2}{C_1} \mu_1 \tag{2.8}$$

where μ_1 is the mean failure time of the density function $f(t)$. That is, we should never replace more frequently than $(C_2/C_1) \mu_1$. This result was perhaps more important as a computational tool when computing resources were scarce, but it still is useful to guide

---

[6] This result, which may not be obvious upon inspection, is given in Barlow (1963). An equivalent result, with full derivation is presented in Lewis (1987).

our intuition regarding replacement policies. In particular, we can manipulate the inequality so that

$$\frac{C_2}{T} \leq \frac{C_1}{\mu_1} \qquad (2.9)$$

which says simply that T must be selected so that the average cost per time (or usage period) associated with preventive replacement (the left hand side of equation 2.9) does not exceed the expected cost per unit time associated with a replace upon failure policy.

While methods are available to compute the optimal replacement time for a few selected distributions, numerical methods must often be used. Examples of computations for which explicit formulations are available are given for the truncated normal and gamma distributions in Barlow and Proschan (1965). Lewis (1987) gives a useful approximation for the 2-parameter Weibull distribution when $C_1 \gg C_2$. Barlow (1978) has also demonstrated a graphical method using total time on test plots, which was used by Guins, et.al. (1984) to estimate optimal wheelset replacement intervals for 28 inch wheels used on auto rack cars. The graphical method is potentially quite useful since it is nonparametric, i.e., it requires no prior information regarding the shape of the distribution. Jorgenson, McCall, and Radner (1967) present a complicated algebraic solution for Weibull distributions, which they then simplify with a graphical approximation. Niño (1974) gives a combined algebraic and graphical method for calculating the optimal interval when the failure distribution is approximately normal.

Some relevant characteristics of age replacements given by Barlow and Proschan (1965) include:

(1) Age replacement policies of IFR components increase the survival probability over replacement upon failure. The survival probability is defined as the probability that an item does not fail in service before time t.

(2) As might be expected, the more often replacement is scheduled, the longer the expected time to an in-service failure.

In some cases, the administration of these policies can be excessively burdensome, since they require keeping track of the actual age (usage) of each item. This is particularly so for inexpensive parts (bolts, etc.). One way of reducing the administrative

burden is to use block replacement policies, which is the first of the multicomponent policies presented in the following section.

2.2.3.2 Multicomponent Replacement Policies

Most complex systems are composed of many components and lesser systems. Unfortunately, finding the optimal maintenance policy for such systems is considerably more difficult than simply scheduling each of the parts separately. In the following sections, we consider two approaches to multicomponent maintenance. The first approach, block replacement, is relatively simple to implement, but inefficient if the components being replaced are expensive. The second approach, opportunistic maintenance, is based on the notion that each maintenance event, whether scheduled or not presents an opportunity to consider the condition of other components as well.

2.2.3.2.1. Block Replacement Policies

Block policies are typically applied to a component or a set of components within a larger piece of equipment. The basic strategy is to replace all components immediately upon failure and at some time T or integral multiples of T (2T, 3T,...,NT), without regard for the failure history of the component. Consider the situation in Figure 2.5. Under no preventive maintenance, failures occur at the indicated times, i.e., f_1, f_2, etc. Under an age replacement policy, items are replaced when the item fails or when it reaches some age T, whichever occurs first. Thus parts are replaced at times T_1, f_1, T_2, T_3, and f_2. Under a block replacement policy, items are replaced when they fail *and* at times T, 2T, 3T, etc., regardless of when they were replaced. Thus in Figure 2.5, items are replaced at time T, f_1, 2T, 3T, 4T, and f_2.

Block replacement policies were studied extensively by Drenick (1960) and Flehinger(1962). The selection of the optimal time to replace is not as straightforward as that of selecting the optimal time for an age based system. Indeed Drenick characterized it as "quite complicated" and suggests the use of iterative procedures to estimate it. Once found, however, the solutions themselves are quite simple to put into practice, since they consist of simply a fixed interval of time at which components are replaced without recourse to the part's history. The primary advantage of such policies are the simplified administrative procedures. No failure histories need to be recorded and

No Preventive Maintenance

Age Replacement (Intervals of Length T)

Block Replacement (Intervals of Length T)
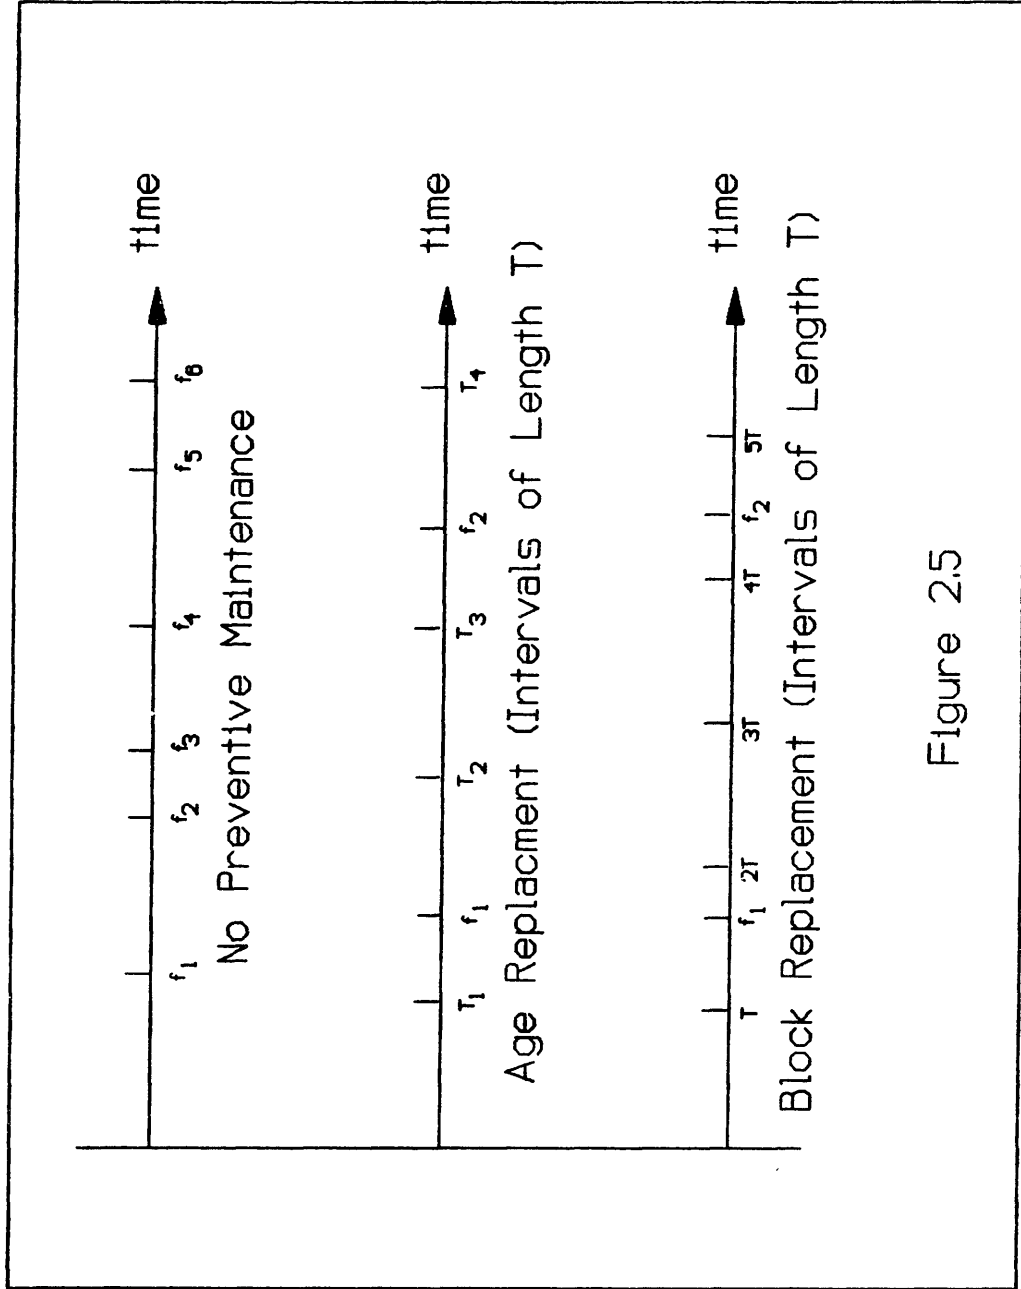
Figure 2.5

instructions to maintenance personnel can be quite simple. Some of the operating characteristics of block replacement polices relative to age replacement for the same replacement interval T are [Barlow and Proschan (1965)]:

(1) Since some parts are replaced under maintenance which were very recently replaced under failure, block maintenance policies are more wasteful (in terms of components used, and foregone wear life) than age replacement.

(2) The expected total number of removals under block replacement is higher than under age replacement.

(3) The expected number of replacements due to in-service failure is less under a block policy since the replacement interval is shorter.

Because of the computational difficulty in calculating the optimum interval and their wasteful tendencies, block replacement policies have generally been applied in cases where the value of the component to be replaced is relatively low, or the cost of monitoring and keeping usage records is high relative to the components value. Examples might include automobile spark plugs during a tune up, or the replacement of "all rubber parts" as a standard procedure of a railroad car brake system overhaul (Clean, Oil, Test and Stencil per Rule 2, AAR(c)(1987)).

In cases where the components have significantly different characteristic lives or high value, a more desirable approach is to establish policies which take into account the characteristics of the components and their relationships with other components. Most noteworthy among these relationships is the possibility of economies of scale in maintenance, in which the cost of jointly replacing two components is less than the cost of replacing the two parts separately. That is the focus of the following section.

2.2.3.2.2. Opportunistic Maintenance Policies

Policies in which the replacement decision of a component is based on the state of the rest of the system are generally called opportunistic maintenance policies. At the conceptual level opportunistic maintenance is not a new idea. As George, et.al. (1979) state,

[o]pportunistic maintenance is probably as old as man's history of the use of tools. Cave man probably replaced the thongs binding a broken axe head to the shaft at the same time he fitted a new head. Mechanics replace parts that have not failed at the time of an engine overhaul because

they are probably well used and, as long as the engine is under repair, it is cheaper to replace functioning parts rather than wait until they fail. The key element that calls for opportunistic replacement is a multicomponent system with economies of scale for replacing parts at the same time.

Sethi (1977) also summarizes the central notion of such policies:

[o]pportunistic policies pertain to situations where it costs less to replace two or more units concurrently than to replace them at different times. These cost advantages may be due to lower overhead costs and economies of scale. Thus the necessary replacement of a unit upon its failure may also justify the replacement of some other units whose failure seems imminent.

The primary goal of research into these policies has been to find a set of optimal control limits, that is, particular age or usage levels after which a component is either replaced independently or in conjunction with another failed component. The literature is summarized by Thomas (1986) and by Özekici (1988). Sasieni (1956) was among the first to examine these policies, looking at the decision whether or not to replace either of 2 identical machines based on a usage standard or when one failed and the other had reached a lesser usage level. These policies were further developed by Radner and Jorgenson (1963) and are often referred to by the name (n_i, N), reflecting the control limit notation. Because of the prominent place in the literature of their work, it is worthwhile to examine their approach more fully.

They consider a system composed of one uninspected IFR component designated part 0, and M independent inspected components subject to exponential failure rates $(1,2,...,M)$ and with economies of scale in the joint replacement of parts 0 and any of the M parts[7]. One can solve as a dynamic program to find a set of times $n_1, n_2, ..., n_m$, and N, such that:

i) If part i fails at a time when part 0 is between 0 and n_i then replace part i alone;

---

[7] More formally, let the cost of replacement of each part be designated as C_0, C_1, $C_2,...,C_M$ and C_{ij} the cost of replacing any two parts, i and j, at the same time. If $C_{0i} < C_0 + C_i$ then we say that there are economies of scale in the joint replacement of parts 0 and i.

ii) If part i fails at a time between n_i and N then replace both part i and 0 together;

iii) If part 0 reaches age N when all monitored parts are good, then replace part 0 alone.

In other words, in addition to scheduling the replacement of part 0 at some time N, one may replace part 0 at some earlier time if the failure of part i occurs "near enough" to part 0's scheduled replacement time. McCall (1963) applies the policy to a hypothetical ballistic missile system composed of one IFR component and 3 exponentially distributed components.

There are a number of problems which restrict the usefulness of Radner and Jorgenson's approach in practical situations. In most settings the system to be maintained is simply not comprised of many exponential systems and a single IFR component, nor are the components likely to be independently distributed, since they share the same operating environment and are subject to the same external shocks. In any case, Vergin (1968) showed that if failure is not costless, the control limits derived by Radner and Jorgenson are not necessarily optimal, and that the solution of the optimal problem in such cases may not have a straightforward control limit structure

L'Ecuyer and Haurie (1983) used dynamic programming in a case of 4 independent IFR components with instantaneous "good-as-new" replacement to find a strategy for deciding each time a failure occurs whether to replace any components in addition to those failed. They conclude

> An optimal strategy...has been obtained and happens to be quite complex.
> It shows, among other things, that the well known control limit rule, valid
> for a 1 or 2 component system, cannot be readily generalized to larger
> systems.

They go on to recommend that in practice one should look for appropriate suboptimal strategies, using a set of control limits found using simulation.

The difficulty in finding practical closed form solutions appears to be inherent in the problem itself. The basic approach used to find control limits is dynamic programming, in which one must specify the state of the system (and each of the components) over the entire life of the system. As the number of components increases,

the number of possible states and the outcomes associated with the states simply become too large. Liang(1985), in discussing the two component problem sums up the problem by noting, "further extensions are very involved, if not intractable."

Because of the complexity in finding optimal solutions and the difficulty of defining them in terms of a practical rule or recipe, most recent work has focussed on finding suboptimal policies with an "all-or-nothing" solution, whereby a failure leads to a decision either to replace all the other components or none of them. Liang(1985) presents an approach known as piggyback maintenance used by Xerox to maintain copy machines, in which no preventive maintenance is performed unless a part fails, in which case all parts subject to preventive maintenance are replaced. He also examines a case in which a percent of the mean life of a component is selected, and, if the component is older than that age when another part fails then the first component is replaced.

A variant on all or nothing policies, known also as "screen" policies, have been attempted for aircraft engine maintenance [Blundell and Beard (1985)] and for a proposed nuclear fusion test facility [(Day and George (1982)]. In both cases, the goal was to find a single value (the "screen") and replace components which exceed that age limit. For the aircraft engines, the basic approach was first to find the optimal replacement interval for each of the individual components. Then, using simulation modelling, a series of "screen values" were tested with the screen representing a percentage of the scheduled life. For each screen, the economic consequences of replacing all parts which have been in use for a period greater than or equal to the screen value times the scheduled life were predicted. If, for example, the screen value was 95%, then all components which have been in service for 95% of their scheduled life are replaced if for some reason, the system is in a maintenance shop. The fusion case was simpler because there were only two parts, so that the optimal value could be found by enumeration. A single optimal screen could not be found when higher numbers of parts were included.

There are a number of problems with the screen approach. If the relative costs of the various components change over time, then the screen value chosen by the simulation must be recalculated, making implementation difficult. More importantly, if the typical lifetimes of the components being studied are quite different, then the short-

lived components may never be replaced early, or the long-lived components may be replaced while considerable life remains and when more opportunities can be expected to present themselves. If the consequences of failures of the various components differ as well, the trade off that is made between the failure costs and the distributions of the components to find a single value may be far from optimal.

Ozekici[19] addressed the problem of dependence and, using an insightful formulation, shows that an optimal policy exists, but is quite counterintuitive as the system and its components age. Unfortunately, his approach also fails to yield an explicit solution, which is required in practice. He therefore suggests that one can use numerical methods to approximate the optimal control limits, provided the problem is not too large.

In sum, then, opportunistic maintenance policies have a great appeal. They are based on the recognition that actual equipment is often made up of many different components, that costs of maintaining and repairing complex equipment are subject to interactions and economies, and that maintenance policies should somehow reflect these relationships. Unfortunately, their technical complexity and dependence on numerical solutions has limited their use in practice, and when applied using the so-called screens, are subject to serious objections. In a later chapter an approach to opportunistic maintenance in the railroad environment will be presented without resorting to the use of screens.

2.3. Motivations for Maintenance

Having defined some of the key terms associated with maintenance, it is useful to turn our attention to the reasons why firms perform maintenance in general, and particularly why transportation firms maintain their vehicle fleets.

Like transportation, maintenance has the character of a "derived demand". That is, it "is not desirable in itself, but as a means [to other ends], ... and this goal is itself derived from the desire to undertake certain patterns of activities" [Manheim (1979)]. But maintenance in the transportation industry, and service industries in general, differs from the maintenance of most production (i.e., manufacturing) industries. This is because transportation company maintenance has a direct bearing on the customer's perception of

34

the quality of what he buys, since the customer is purchasing the process of production as well as a finished good. That is, in the manufacturing industries, the customer is satisfied with a high-quality good, even if the means of its production were unreliable; in the case of transportation, if the means of production are unreliable, the good purchased, the trip, is also unreliable. As a result, maintenance take on a special significance for transportation firms.

In the following sections, we look at the motivations for maintenance, focussing particularly on those aspects which follow from the special nature of transportation maintenance discussed above. It is shown that maintenance is performed in support of a broad range of goals and objectives, ranging from the straightforward support of operations to the rather subtle notion of signalling the company's actual and potential rivals as part of the strategic plan. Although the focus of this section is the activities of private sector companies, almost all of the points could be applied in other institutional environments or modes.

There are a number of concerns that a firm seeks to address in setting a maintenance plan, including ethical issues, legal concerns (liability), insurance costs, and safety. When summed up, however, it can be argued that there are 4 basic reasons why all companies undertake maintenance:

1. To support the ongoing operations of the firm;
2. As an alternative to capital investment;
3. In response to regulatory considerations; and
4. As a strategic tool.

The first three concerns are shared in common with most maintenance activities by other firms, such as manufacturers or producers of durable goods. The fourth, that of using maintenance as a strategic tool, while not necessarily unique to transportation firms, depends on the customer being interested not only in an end product or good, but in how the good is produced.

Each of the basic reasons is examined in detail below.

2.3.1. Maintenance In Support of Operations

The provision of productive capacity necessary to carry out the operating plan

must be considered the most significant reason for planning and performing maintenance. In the case of transportation firms, the productive capacity generally takes the form of vehicles, terminals and, in some cases, guideways. This aspect of maintenance is summed up well by Haven (1979), when he states

> The primary goal of any vehicle maintenance department can be stated rather simply: to supply a fleet of safe and reliable vehicles of sufficient size to meet the needs of the transportation department at the least possible cost. This goal highlights four major objectives: safety, supply, fleet reliability, and cost minimization.

Support of the operating plan can be broken down into three specific functions:

1. Insuring an adequate supply of vehicles;
2. Specifying and insuring an appropriate level of reliability for those vehicles;
3. Relating maintenance practices to the performance of the vehicle when in use (i.e., efficiently using resources to fulfill the tasks).

One of the crucial issues in meeting the supply requirements of the operating plan is determining what constitutes an "adequate" supply. Defining an "adequate" supply is essentially a fleet planning exercise. Manheim (1979) proposes that the appropriate tool for this analysis is the vehicle cycle. Under this approach, the analyst determines the entire process by which vehicles are used, beginning with a long term (e.g. annual) cycle, which can be decomposed into in-service time, maintenance time, and idle time. This long term cycle is further decomposed into service and operating cycles, including loading, travel, inspections, unloading, and other processes. In practice, the definition of an adequate supply is likely to be set jointly by the transportation, marketing, and mechanical departments.

In setting an appropriate level of reliability, many factors are important, of which safety is one of the most important[8]. Maintenance standards, for example, are generally higher where the risks are higher. If not properly maintained, equipment may endanger personnel who work around it. Further, in a transportation environment, an unsafe

---

[8] One need not be convinced of the altruism of businessmen to accept this. Blundell and Beard (1985) show the importance of well performed aircraft maintenance by citing an airline accident which resulted in a $70 million loss for the airline.

condition may also put at risk the residents who happen to live along the right of way, particularly when hazardous materials are involved.

It is useful to note the great importance attached to reliability when analyzing transportation systems compared with manufacturing firms. In the manufacturing case, the customer buys a product which, from the consumer's point of view, is independent of the process by which it is produced. Thus, if the production process is unreliable, the producer may face higher production costs and lower profits, but the consumer utility associated with the product is not directly affected so long as the product itself is reliable. In the case of transportation firms, however, the reliability of the production process is, in some sense, directly related to the consumer's utility, since it is the process itself which is purchased. While the importance of reliability to railroad operations has been studied extensively [see, for example, Martland (ed.) SROE], there has been relatively little analysis of the impact of mechanical reliability on service reliability and virtually none on railroad profitability. Dingle (1977) reviewed the experience of several railroads with respect to rejected cars and concluded that the number of cars rejected could be reduced. The Freight Car Utilization Program (1980) held that railroads could and should analyze the impact of out-of-service time on car utilization rates. Other researchers have examined the impact of vehicle reliability on train delays, and found them to be quite modest [Belovaric and Kneafsey (1972) and Sheaffer and Stern (1986)], although there are reasons to suspect that these studies understate the consequences of delays to other trains on the line. In any event, it seems apparent that maintenance plans which seek to minimize the cost of repairs without evaluating the broader implications of reliability may seriously understate the actual cost, and attempts at cost minimization based on such estimates may result in sub-optimal policies.

2.3.2. Maintenance in Lieu of Capital Investment

The second role typically performed by maintenance is to defer or otherwise avoid capital expenditures. Watson (1970) states this argument clearly when he says

> Maintenance expenditure buys production time on existing plant so maintenance should be considered as one of a number of ways of spending money to increase production capacity. By accepting this philosophy it

37

becomes possible to compare maintenance expenditures with alternative proposals designed to achieve an increase in capacity.

Clearly there are times when it is less expensive (or more timely) to incur even substantial maintenance expenses rather than purchase new equipment. The decision of whether to maintain and continue using older equipment or acquire new equipment is generally known as **replacement analysis**. Such analysis is conducted by determining the net present value of each of the options (i.e., repair and continue using, replace by lease, and replace by purchase), using the full costs and benefits over the life of the asset, and selecting the alternative with the highest net present value. (If one is dealing strictly with costs, one selects the alternative with the lowest fully discounted present costs.) An equivalent and often easier calculation is to determine the equivalent uniform annual costs, which restructures the stream of costs into a set of equal annual payments over the life of the asset. Finance and engineering economics texts demonstrate how to analyze cases to support such decisions [e.g., Riggs and West (1986)], and no further elaboration is needed here, except to note that when the discount rate is high or the economic life of the asset is uncertain, maintaining the existing asset generally becomes more attractive. In many cases (particularly in a deregulated environment) the latter is precisely the situation faced by the railroad, unless contractual arrangements can be made with the shipper.

2.3.3. Maintenance in Response to Regulation

The third rationale for maintenance is the satisfaction of requirements imposed by regulatory bodies, i.e., legal requirements. In the case of railcars, these come primarily from the Association of American Railroads (AAR), the industry's trade association, and the Federal Railroad Administration of the United States Department of Transportation (FRA). The AAR provides minimum standards for equipment which is to be used in interchange service, i.e., from one railroad to another. These rules are intended to protect the member roads from the externality[9] of accidents on the receiving road caused by

---

[9] An externality occurs when a firm or other agent produces by-products for which, if they benefit others, it cannot charge, or, if it harms others, it need not pay compensation. In the present case, the by-product is unreliability or reduced safety which

inadequate maintenance or design standards on the sending railroad. The Interchange Rules, published annually in a pocket-sized book, generally referred to as the **Field Manual**, are developed by committees composed of railroad mechanical officers, consultants to the industry, and manufacturers of equipment and components. Although the rules are not binding on railroads for the movement of their own equipment, and members are free to negotiate alternative arrangements among themselves, most railroads maintain their entire fleet to the standard of the Interchange Rules[10].

The Federal Railroad Administration mandates particular maintenance practices pursuant to the Federal Railroad Safety Act of 1970 and other legislation. These rules range from banning certain practices or equipment (some dating back to the Ash Pan Act of 1908), to mandating that components be inspected and replaced on a periodic basis (as in the case of the air brake system), to setting maximum wear levels for certain components. Unlike the Interchange Rules, the Federal standards cannot be simply ignored or negotiated away.

The two reasons generally given for the Federal standards are the safety of workers and the impact of accidents on communities along railroad rights-of-way. These concerns are not without foundation. Even a cursory examination of the accident rates of the early (unregulated) era of railroads suggests that railroads were willing to allow unsafe conditions to exist far longer than would today be considered acceptable. For our purposes, the Federal Safety Standards have the effect of reshaping any question of "optimal" maintenance levels to one of whether a car owner should engage in maintenance beyond that prescribed by law.[11]

---

a railroad can impose on another's trains by failing to perform adequate maintenance on a car prior to interchange.

[10] Prior to the adoption of the Federal standards discussed below, many railroads maintained cars which they knew would not be used in interchange service to substantially lower standards than those in the **Field Manual**.

[11] It is interesting to ask why an industry which went to considerable effort and expense to reduce the level of economic regulation during the 1970's sat passively by while the Federal government imposed a new set of potentially costly set of regulations

2.3.4. Maintenance as a Strategic Tool

To understand the use of maintenance as a strategic tool for the firm, it is necessary to review some of the basic tools of strategic management used by firms in general. Businesses naturally seek to restrict the access of other firms to their markets, since this gives them a stronger and potentially more profitable position in setting price or service levels with their customers. Among the means of controlling access to markets are the creation of high entry costs to potential competitors and the signalling that a market will not be as profitable to other firms as it is to the current participants. Other valuable tools are the ability to legally indicate to other firms already in the market an

---

over equipment and track. A staff member from the Federal Railroad Administration at the time the standards were adopted stated that the FRA not only sought comments through the standard rulemaking process, but actively requested railroads and the AAR to comment on the proposed standards. He indicated that at least some of the FRA staff was disappointed by the paucity of comments received. One possible explanation is that the industry believed that opposition to the pending safety regulations would compromise their case for economic deregulation by making railroads appear to be self-serving and even dangerous to workers and adjacent communities. Another possibility is simply that the industry was preoccupied with one regulatory battle and was unwilling to invest in a second one, particularly since the proposed standards conformed more or less with the existing interchange rules, which were believed to be widely followed. A third explanation would suggest that the railroads themselves, in some sense, wanted these regulations to better define markets and costs! It has been argued by Stigler (1971) and by Ulen (1980) that transportation companies themselves seek at least a certain level of economic regulation, primarily to police "renegade" members who might seek to develop alternative cost structures, and to insure that the cost functions of their competitors are easily estimated for pricing purposes. This argument could easily be applied by extension to the equipment and other regulation. A further argument along the same lines can be made for the notion of "raising rivals costs", based on a paper by Salop and Scheffman (1983). The basis of this argument is that a firm which has some cost advantage already in the production of maintenance can exploit that advantage by requiring less efficient companies to maintain their equipment to a higher standard, thus incurring higher costs and a worse position in contested markets. This type of behavior has been used to explain the seeming generosity of large capital-intensive firms in settling labor contracts in nationally negotiated agreements [Williamson (1968)]. Regardless of their motivation, the railroad industry permitted an extensive body of equipment maintenance rules to take effect without objection, and today these rules provide minimum standards to which components and equipment must be maintained and, as such, a minimum level of maintenance activity which a railroad car owner must sustain.

40

equitable division of customers and markets by signalling which markets are assigned to each participant.

Of particular interest, then, for maintenance planners are ways to use capacity to signal market intentions and to protect important markets or lines of business. In this section it is argued that maintenance permits the car owner to provide a continuum of capacity which can be used to restrict potential entrants. Announcements of these maintenance decisions can also be used to define particular markets which the car owner is willing to invest in and will fight to protect from potential rivals. Because detailed strategic information is not generally available, this section depends heavily on other areas of maintenance and on some degree of speculation about car owner's motives. While institutional reasons make it difficult to demonstrate that car owners engage explicitly in the practices described, there seem to be good reasons to believe they *ought* to as profit maximizers.

The use of investments in capacity as a weapon in the arsenal of the strategic planner has been studied extensively in recent years in the branch of economics known as industrial organization. In general the literature surrounding strategic use of capacity has focused on manufacturers who seek to use investment in capacity either as a barrier to entry by potential rivals, or as a signal to actual or potential rivals as to which markets or segments are unlikely to be profitable to invest in. A subtler issue is the use of capacity to "cement" relationships with customers, particularly if the costs of the capacity can be somehow borne by, or at least shared with, the customer. Use of capacity for strategic purposes has been studied in several industries, with titanium dioxide (i.e., paint) among the most widely reported [Ghemewat (1984)].

Dixit (1980) discusses the entry deterrence argument in a context of game theory, and concludes that in certain restricted markets the result is not so much to exclude competitors as to make clear the conditions they can reasonably expect to encounter if they enter. He draws the distinction between the rules of the game (i.e., how firms respond to the environment to set prices and levels of output) and the initial conditions of the game (i.e., the environment the competitors face). He concludes that "...the role of an irrevocable commitment of investment in entry deterrence is to alter the initial

conditions of the post-entry game to the advantage of the established firm, for any fixed rule under which the game is to be played."

The second strategic use of capacity is that of signalling the importance of a market to potential new entrants. This notion is built on the idea of the "credible threat" to remain in the market at the current level of production or service, even if a rival attempts to enter, which would have a potentially destructive effect on the price faced by the new entrant. As Dixit (1982), points out, for a threat to be credible "(t)here are two essential requirements: the commitment should be made (and made known prior to the entrant's decision), and it should be irreversible".

In transportation firms, capacity can be thought of as existing in 2 forms:

1. Physical capacity (track and equipment): the ability to move larger or heavier or more frequent loadings; and

2. Operational capacity: the ability to offer faster or "larger" service as a result of more effective or efficient use of the physical resources. This obviously has an upper bound imposed by the physical capacity.

Capacity in the transportation industries can take the form of additional miles of track or pieces of equipment, but it can also be increased by maintaining them at a higher level of reliability so that they can be used more intensively. Track, for example can be maintained at a level such that trains can be run at higher speed, giving more effective capacity over the same right-of-way. Similarly, equipment can be maintained with more thorough overhauls, or higher parts inventories can be held to increase the reliability and reduce the service losses associated with failures. What gives maintenance an interesting position in this regard is that it can provide some of the aspects of strategic investment without all the expenses and costs normally associated with capital investment. Recall that in order for a strategic use of capacity to be successful, several things must occur:

1. The investment must be observed by the relevant parties;
2. The investment must be large enough to give pause to other potential entrants into the market; and
3. The investment must be irreversible.

Each of these conditions can hold in the case of capacity expansion by railroads. The

information is communicated to the relevant parties via the trade press[12], often with a fanfare far in excess of that normally appropriate for maintenance activities. The investments are large enough to insure that both the customer and rivals recognize that the carrier is "serious" about the business. Finally, the investment in equipment or track is often sufficiently specialized to insure that it cannot be easily used in another market without a loss by the railroad.

A similar case ensues when the railroad and a shipper agree on shipper- owned and railroad maintained equipment. In this case, the shipper's agreement to provide the equipment and the railroad's committment to a level of maintenance has the effect of "cementing" the relationship and signalling to potential rivals that this market is effectively closed to them.

Maintenance used in these ways not only has the effect of acting like a capital expenditure, but also, by its nature allows a continuum along which to place the size and type of the investment. While one cannot build 40% of a titanium dioxide plant, one can establish and announce a new and costly maintenance plan or facility at almost any level of expenditure desired. The evidence that this occurs is anecdotal at best, but one does see some instances which appear to confirm that strategic planners use maintenance. An official of one regional railroad told the author that whenever a larger carrier negotiates divisions of revenues they always preface the meeting with announcements regarding their plans to upgrade tracks parallel to the regional's right-of-way. In another case, an official of a large railroad indicated that his railroad did not want to see a uniform standard for intermodal flatcars adopted since it would undermine their attempts to link themselves tightly with their customers.

Unfortunately, the extent to which railroads use capital investments for strategic purposes does not lend itself to formal documentation, and the use of maintenance in such

---

[12] As an example, note Conrail's announcement of a $3.5 million upgrade to its facilities between Ypsilanti and Jackson, Michigan in **Modern Railroads**, September, 1988. This is a minuscule amount compared to Conrail's overall track maintenance budget, and seems to have been timed and publicized to insure that rivals would recognize their commitment to certain auto related markets.

circumstances is understood even less. Still, it seems clear that this is one of the underlying motivations for at least some maintenance actions, particularly when either considerable fanfare is involved or when the maintenance level seems much higher than that called for to provide the current service.

Having looked at the various motivations for maintenance of freight cars, we turn our attention to the problem of measuring performance. Of particular interest is the issue of devising measures that relate maintenance activity with its intended purpose.

2.4. Measuring maintenance quality

Having presented some of the basic concepts associated with maintenance, it is now possible to focus on of the most important aspects of maintenance, defining appropriate measures of maintenance quality.

Armitage (1970) presents 7 desirable properties that a measure of maintenance effectiveness should possess.

1. The measure should be relatively easy to calculate and use.
2. It should be easy to interpret from the information provided.
3. It should be reflective of management's subjective notions of maintenance performance and organizational objectives.
4. It should indicate when something has gone wrong with past decisions.
5. Ideally, it should indicate what action to take when something has gone wrong.
6. Limits on the use of ratios should be recognized: they give relative, not absolute values, so they can be misinterpreted.
7. The measure should reflect all the relevant consequences which effect performance.

He also notes the distinction between efficiency and effectiveness. **Efficiency** measures generally relate the level of output to the levels of inputs. In the case of maintenance it can be quite difficult to define the outputs in a way that is easily measured and is in comparable units to the inputs. **Effectiveness** measures compare actual performance against planned performance. The central concept is that a 'goal' is set and, at the end of each time period, the extent to which the goal has been met is measured. Crucial to this is the setting of a goal which is measurable and is consistent with the firm's goals.

44

If we exclude the strategic aspects of the maintenance plan and take the regulatory environment as a minimum safety standard, the maintenance goals of the car owner can be generally reduced to minimizing maintenance costs while providing equipment which meets a desired level of reliability.

Turning first to costs, McCall (1965) discusses cost-based measures for evaluating maintenance policies as follows:

> A policy's performance can be measured in cost by assigning an occupancy cost to each state [of the equipment] and an intervention cost to each action. These costs are calculated to measure the money and downtime costs of each maintenance action as well as the downtime costs associated with each operational state. *The objective of the decision-maker is to choose maintenance actions so that the cost per unit time of operating the equipment is minimized* (italics added).

In the case of railroad cars, miles rather than "time" is the appropriate index for measuring the operation of the equipment, so that an appropriate cost-related measure would be cost per mile operated. This measure is similar to that proposed by Jorgenson, et. al.(1967), which was the ratio of the expected total time available to the expected total discounted cost of the equipment, although their measure includes acquisition cost[13].

The other aspect of maintenance performance measures that must be addressed is reliability or availability. Clearly a policy which reduces the maintenance cost per mile, but only leaves the equipment available for a small number of miles generally will not be acceptable for achieving the goals of providing cars to meet the operating plan.[14] To

---

[13] The reader will note that these are ratios, which Armitage gives specific warnings about (Property 7). In this case, the ratio is used simply to insure that the maintenance expenditures are "normalized" to a unit of output, miles operated. Caution should be exercised in using this measure, however, if the basis for the costs changes over time (e.g., due to inflation or other adjustments in the costs of labor or materials), or in using the measure to compare across widely differing circustances. When used in the thesis as a measure, it will be assumed that we are referring to real cost per mile.

[14] In principle, if the costs are calculated to include all those faced by the car owner, including opportunity costs from unreliable equipment and the discounted value of foregone shipments due to unreliability, then costs per mile can capture both the costs and reliability aspects. In practice, however, those costs are difficult to calculate, so a

45

measure the effect of the maintenance plan on the reliability of the equipment, one can look to some of the traditional measures of reliability. The most commonly used measure is the mean time between failures (MTBF). This is defined by British standard 4778 [quoted in O'Connor (1985)] as

> For a stated period in the life of an item, the mean value of the length of time between consecutive failures computed as the ratio of the cumulative observed time to the number of failures under stated conditions.

MTBF appears to be the single most widely used reliability measure in articles and published studies. This widespread adoption appears to be more than a convention. MTBF cuts to the heart of one of the primary interests of the reliability analyst, namely, finding means to ensure failure free operation of a machine or system. An analogous definition for non-repairable items is the mean time to failure (MTTF). Maintenance actions are sometimes measured in terms of the mean time to repair (MTTR), which is the total corrective maintenance time divided by the number of maintenance actions. An extension of the above definitions is the availability of an item, which is the ratio of the MTBF to the sum of the MTBF and the MTTR. This ratio measures, in effect the percent of total time that the equipment could be used in service.

For the purposes of measuring railroad car maintenance policies and the corresponding effects on reliability, we will examine two measures that seem appropriate. The first is miles per in-service failure. This measure provides an index of the extent to which a maintenance policy is able to provide equipment to meet the needs of the operating and service plan. Miles per in-service failure is essentially a mileage-based version of MTBF. A second reliability measure we will develop and use is miles per maintenance event, where a maintenance event is any time the car 's sent to a maintenance facility or repair track, regardless of whether for planned maintenance or corrective maintenance. This measure is particularly useful to determine whether a planned maintenance policy itself is causing disruptions in the vehicle cycle even though the in-service aspect (i.e., revenue trips) of the cycle is reliable.

---

separate measure for reliability is desirable.

2.5. Conclusions

In this chapter, the basic terms, concepts, and measures which will be used throughout the remainder of the thesis have been introduced. In the following chapter, an overview of railroad car operations and maintenance is presented. That chapter also begins the process of applying the language of reliability and maintenance theory to evaluate the car repair process.

Chapter 3

Railroad Freight Car Operations and Maintenance

3.1 Introduction

In the preceding chapter some of the basic concepts of maintenance and reliability were introduced. In this chapter, we begin to apply those terms to the particular piece of equipment that this thesis is concerned with, the railroad freight car. In the first section, the basic size and composition of the fleet of cars used to serve U.S. rail markets is discussed, after which some of the standard features and attributes of the freight car are presented. This is then followed with a discussion of the typical freight car operations, inspections, and maintenance activities that railroads perform. This is done in two parts. First, we examine the circumstances of cars used solely on a single railroad; the second, more complex case, is that of cars which operate over several railroads, generally known as interchange service. The use of cars in interchange service presents some special problems of operations and maintenance, for which institutional solutions have been devised by the industry. Finally, the general organization and practice of maintenance by the freight railroads is examined.

3.2 Railroad Freight Cars

In this section, some basic facts about freight cars are presented. Attention is first focussed on the size and composition of the overall fleet of cars used on the U.S. railroads. This is followed by a discussion of some of the components and systems which are common to all freight cars. Finally, some of the design aspects that differentiate car types are discussed.

3.2.1. Basic Statistics of the Railroad Car Fleet

In 1988, there were more than 1.23 million railroad freight cars in service in the United States[1]. While this represents a considerable decline from 1980, when there were

---

[1] Unless otherwise noted, all data used in this section are from "Freight Car Statistics" (1989).

more than 1.71 million cars in use, it still represents a large capital asset. Not only has the size of the fleet been changing, but there has also been a significant shift in its ownership, composition, age, and utilization.

The most important change in the ownership of the freight car fleet has been the decrease in the share owned by Class I railroads[2]. In 1980, the 1.168 million cars owned by the Class I carriers represented 68% of the fleet. By 1988, that number had fallen to 652,123 representing 52% of the total fleet [Association of American Railroads (b, 1980,1988)]. The number of cars owned by smaller railroads, while varying by as much as 10% from year to year, has hovered about an overall level of approximately 100,000 cars over the past decade; the number of privately owned cars has remained at approximately 440,000 throughout the period [Association of American Railroads (a, 1987)]. Thus the proportional share of the non-Class I car owners has increased, but this increase simply reflects the decline in the Class I fleets.

The composition of the fleet has also changed. While there has been a decrease in virtually all types of cars used, the number of boxcars, both plain and equipped with special devices, has decreased most dramatically. In 1980, for example, there were 251,420 plain boxcars in service. By 1988, that number had decreased to 104,195 cars. The cars which exhibited the smallest decline were tank cars (down from 183,989 to 177,997), covered hoppers (299,986 to 284,556), and flatcars (152,661 to 132,365). These changes in the composition of the fleet represent the movement of the industry away from providing a broad base of customers with general transportation service to offering the movement of specialized trains, especially bulk goods movement, to a smaller number of larger shippers.

Along with the decline in the size of the fleet, it is not surprising that there has been an overall aging of the cars. In 1980, the average age of a freight car was 14.9 years. By 1988, that number had climbed to 17.7 years. Once again, that increase in age

---

[2] The U.S. railroads are organized into "classes", as defined by the Interstate Commerce Commission, based on their annual operating revenues. Class I railroads are the largest railroads. In 1986, for example, the basis for the Class I railroads was $88.6 million in annual operating revenues [Association of American Railroads (a, 1987)].

came primarily in the fleets of railroad owned equipment. The average age of the private car fleet remained at approximately 14 years over the entire period.

Despite the increased age and decreased size of the fleet of cars providing service in the U.S., the utilization of those cars improved throughout the 1980s. The number of annual carloadings remained steady at approximately 22 million carloadings between 1980 and 1988, even though the number of cars available for loading decreased by about 28%. Thus the average loadings per car increased from 13.1 to 17.9. The average length of haul of these loadings also increased, from 626 miles in 1980 to 697 miles in 1988 [Association of American Railroads (b, 1980,1988)].

In other words, a smaller, older fleet of cars is being used more intensively by railroad shippers. This puts a great burden on maintenance managers to make the fleets more reliable, and to perform maintenance quickly and efficiently. Indeed, railroad maintenance has grown considerably in importance relative to other railroad function in recent years as a direct response to the need to use fewer resources in better ways. It is in this context that we now shift from the general numbers surrounding freight car movements to the structure and use of the cars themselves, ultimately to the policies and practices for maintaining them.

3.2.2. The Freight Car Itself

The modern freight car is composed of many different components and systems. Some of these are common to all freight cars, while others are unique to particular designs created for specific uses. Consider the car shown in Figure 3.1, a typical boxcar. Virtually all railroad freight cars use steel wheels joined by an axle into a *wheelset*, which is held in position by a *truck*. The wheelset is able to turn freely by virtue of devices known as *bearings*, (typically fully enclosed roller bearings), on which the truck rests. The truck supports the rest of the car, which is, in effect, the actual container for shipments moved. Since the container part of the car (the box, in the figure) must rest on the two trucks, an important part of the car is the *center sill*, which is actually a large beam. Freight cars are equipped with what are usually called *safety appliances*. These include various ladders and handholds; one of the first Federal laws applied to the

50
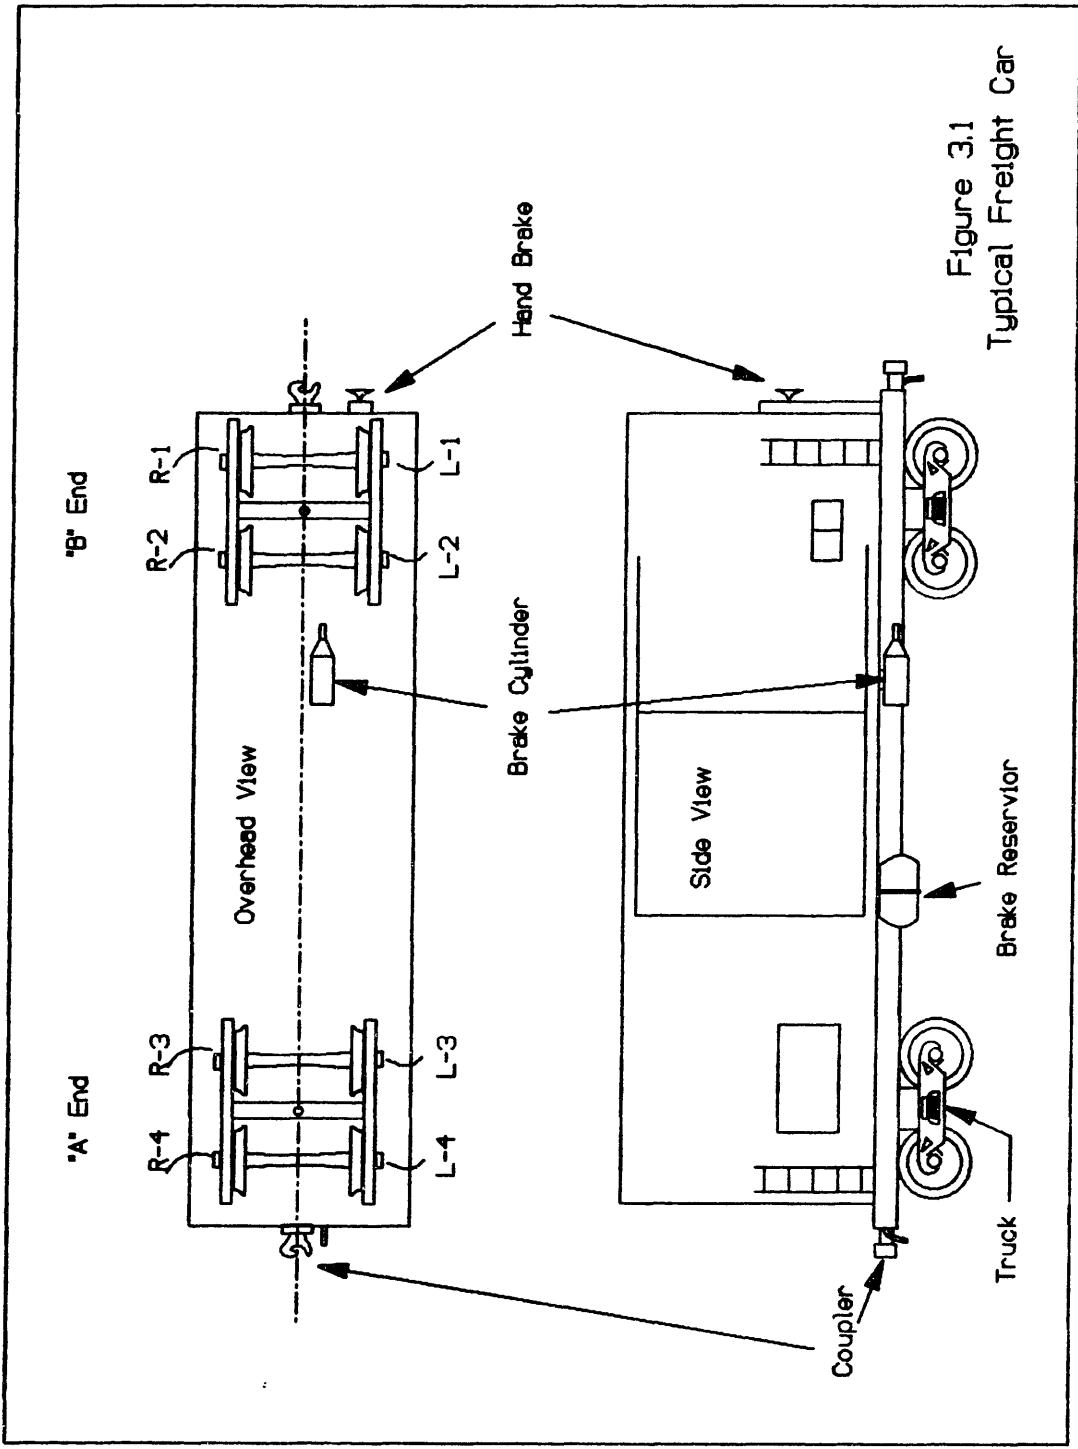
Figure 3.1
Typical Freight Car

railroad industry, the Safety Appliances Act of 1893 specified that the placement of safety appliances be uniform on all cars to insure worker safety. The brake system used on all freight cars in the U.S. and throughout most of the world relies on a system known as the *automatic air brake*. This system uses a continuous flow of compressed air to keep the brakes released. In the event of any disruption in the flow, the tension on a spring in the *brake cylinder* is released, and the *brake shoes* press against the wheels and cause the car to stop. Because of the importance of the brake system, much of the attention of regulatory and industry bodies has focussed on when and how the brakes should be inspected and maintained. Finally, to join the car together to other cars to form a train, the car is equipped with coupling devices, including the *coupler*, which actually joins the cars together, and other devices for reducing the impact of train forces on the lading in the car, including the *draft gear* and *end of car cushioning devices*.

As indicated above, there are also many unique components in railroad cars. Among these, the most obvious is the design of the containers for carrying shipments. These can range from general purpose designs such as the boxcar, which can carry a wide variety of commodities to tank cars, which are used to carry liquids, to gondolas and hoppers which carry dry bulk commodities such as coal, scrap metal, and when covered, grain. Other designs are specialized to gain operational efficiencies, such as "double stack" intermodal cars which can carry one container on top of another, or auto part cars, which are very long boxcars equipped with bins for carrying a number of different components to auto assembly plants. In general, the maintenance of specialized cars or components present special challenges to railroad managers, since they require unique inventories of parts and special skills by maintenance workers. Because they are generally assigned to particular customers and routes, however, these inventories and skills can often be assigned to particular repair facilities.

Clearly, some components are of critical importance in terms of safe operations. Certain components, such as wheels or brake systems, are such that a complete failure can result in a disaster (i.e., catastrophic failure), while the failure of others such as roofs or doors may subject the commodity being carried to risks (e.g., leakage). As a result of their higher potential for risk, any failure of some safety related components to meet basic

wear standards mandates an immediate removal of the car from service until the component is repaired or replaced.

In terms of maintenance expenditures, these critical components are responsible for the largest share. Guins and Hargrove (1980), for example, found that brakes and wheel systems (not including bearings) accounted for 62.7% of all repairs dollars billed under the car repair billing system in 1977. More recently, data from one of the case study companies reported in the next chapter shows that over a three year period, wheel, brake shoe and periodic brake system repairs were approximately 55% of the annual car maintenance expenditures. In both cases this is, however, a reflection not only of these components' importance in terms of safety, but is also due to the fact that these components are common to all cars.

Having looked briefly at the cars themselves, we now examine the ways they are used in operations, and the implications of the uses of the cars on maintenance practices.

3.3. Freight Car Operations and Maintenance

In this section, we look at the way in which freight operations are performed, with an eye toward how these operations impact upon the maintenance of freight cars. Naturally, one of the most important issues in this is how and when cars are inspected for defects. Many of these inspections are mandated by Federal regulations, and are primarily concerned with insuring that the braking system and various safety appliances are in working order. We will look first at the case of railroad operations on a single railroad, and then at the more complicated case of cars moved on two or more companies.

3.3.1. Freight Car Operations: The Single Line Case

Cars can be considered as moving according to a *vehicle cycle*[3], which includes placement at the customer siding for loading, return to the railroad yard, classification, line haul movement to the destination yard (including intermediate classification), delivery to the consignee, and unloading, followed by return to the railroad yard for movement to

---

[3] These processes are described in general terms in Manheim (1979), and more specifically for railroads in Messner (1980).
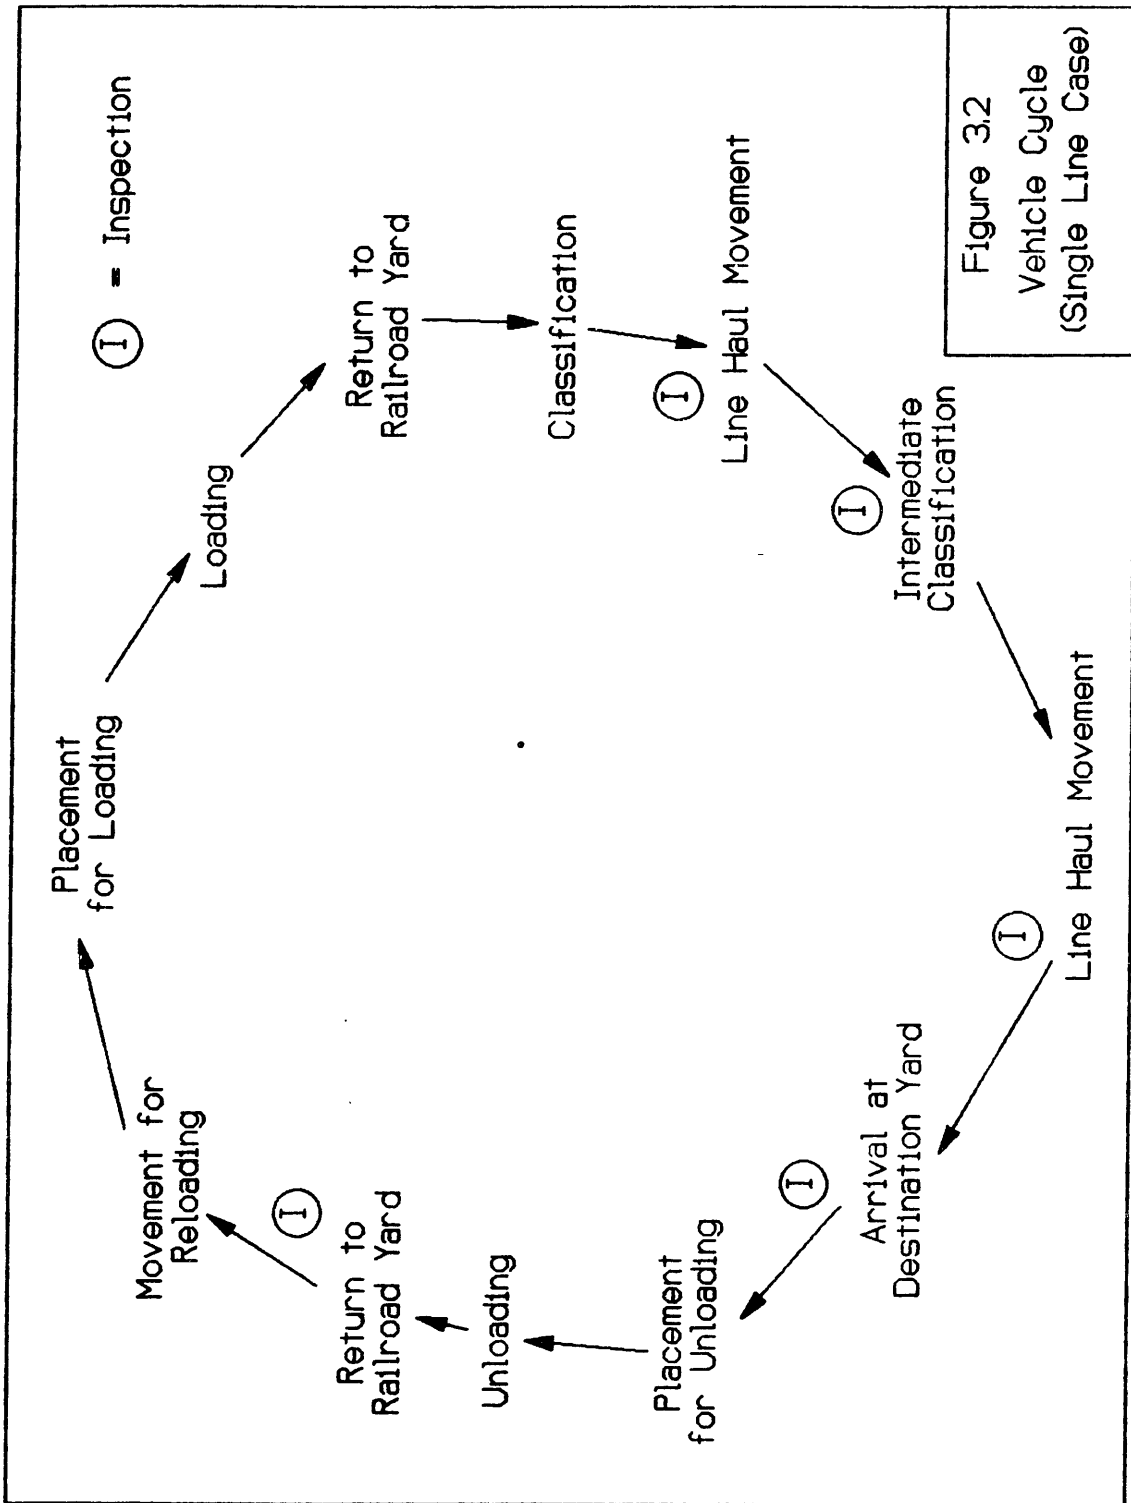
Figure 3.2
Vehicle Cycle
(Single Line Case)

$\textcircled{I}$ = Inspection

Placement for Loading

Loading

Return to Railroad Yard

Classification

Line Haul Movement

Intermediate Classification

Line Haul Movement

Arrival at Destination Yard

Placement for Unloading

Unloading

Return to Railroad Yard

Movement for Reloading

the next customer. These operations are shown in Figure 3.2.

The car is inspected at various point in the vehicle cycle, and, if defects are found they must be addressed. Some of these inspections are marked in Figure 3.2, and include:

- Inbound inspection upon arrival at yards: each car in a train is inspected whenever a train arrives at a yard for classification or delivery to customers. If the car is empty, part of this inspection typically includes *grading*, in which the suitability of the car for hauling various commodities is rated.

- Inspection prior to movement in train service (brake test): Each car must be tested to insure that the brakes and other safety systems are operable prior to departure from a yard in a train. Most railroads also require that the personnel performing inspections look for defects which might cause the car's contents to be damaged (e.g., open or defective doors), and for other defects.

- Intermediate brake tests: If a train is moved more than a fixed number of miles, the brakes must be tested as if the train were departing a yard. Formerly, this was required every 500 miles, but has now been extended to 1000 miles, and so is less significant in current operations.

To control these inspections, most railroads provide their car inspectors with written guidelines for what to look for during inspections. The actual conditions used to define unacceptable components are typically those given in **Field Manual** [Association of American Railroads (c)], which is discussed below.

In the course of these inspections, defects are naturally found. Depending on the nature and severity of the defects, there are three possible outcomes.

1. Some defects which are easily accessible, do not require specialized equipment, and do not take long to fix are repaired on the spot. For example, brake shoes or air hoses can be fixed *in situ* by the carman who discovers the defect. This type of repair is the least disruptive of the overall operation of the train, and some railroads attempt to perform as many repairs as possible in the train, especially when the train is a more or less integral unit such as a specialized piggyback train.

2. Other defects can safely be allowed to remain unrepaired until the car is unloaded, either for the convenience of the shipper, or to facilitate the repair, or both. An example of this would be a defective door mechanism, when detected at the final destination yard. Clearly the customer benefits from

access to the material in the car, and the repair is generally easier if the lading is removed.

3. Some defects must be repaired before the car can be allowed to continue in service, and can only be repaired at a fully equipped repair facility (in railroad parlance, a *rip track*). An example (at most yards) would be a car with defective air brakes, or with a defective wheel. This is the most disruptive case, since all the other cars in the train are delayed while the defective car is removed from the train, and the shipment is delayed until the car can be repaired.

When the third case occurs, the process followed is that shown in Figure 3.3. The car first must be removed from the train or block of cars attached to it. The car is then sent to a holding track, awaiting scheduling into the repair track. This scheduling process reflects the priorities of the railroad performing the repairs. Normally, railroads give a higher priority to loaded cars than empty ones, and to cars belonging to others than to their own (reflecting the fact that they must pay for the time when other railroads' cars are on their line). Other factors influencing which cars are scheduled into the repair shops include the demand for a car type, the availability of parts and labor, and the suitability of a shop to perform a particular type of repair. Once ordered into the shop for repairs, the car must be moved in by crews from the Transportation Department (i.e., the car must be switched), and the repairs are made. After the car is repaired, it must be "switched" out of the shop and returned to the yard, where it is grouped with other cars by destination and moved on a subsequent train.

Notice the many opportunities for a car to be delayed in this process. If, for example, the shop has cars with higher priorities, the car is delayed. If the parts needed are unavailable, the car is delayed. If the Transportation Department makes an error and brings in the wrong cars (or spots the car at the wrong place in the rip track) the car may be delayed. If the car is not finished until after the switch is made, the car must wait until the track is switched again (daily on most facilities, several times each day at the largest shops). After being repaired and returned to service, the car still must wait for the next available connecting train to its destination and must make that connection. If there are other cars found to be defective in that train, the car is delayed while those cars are

56
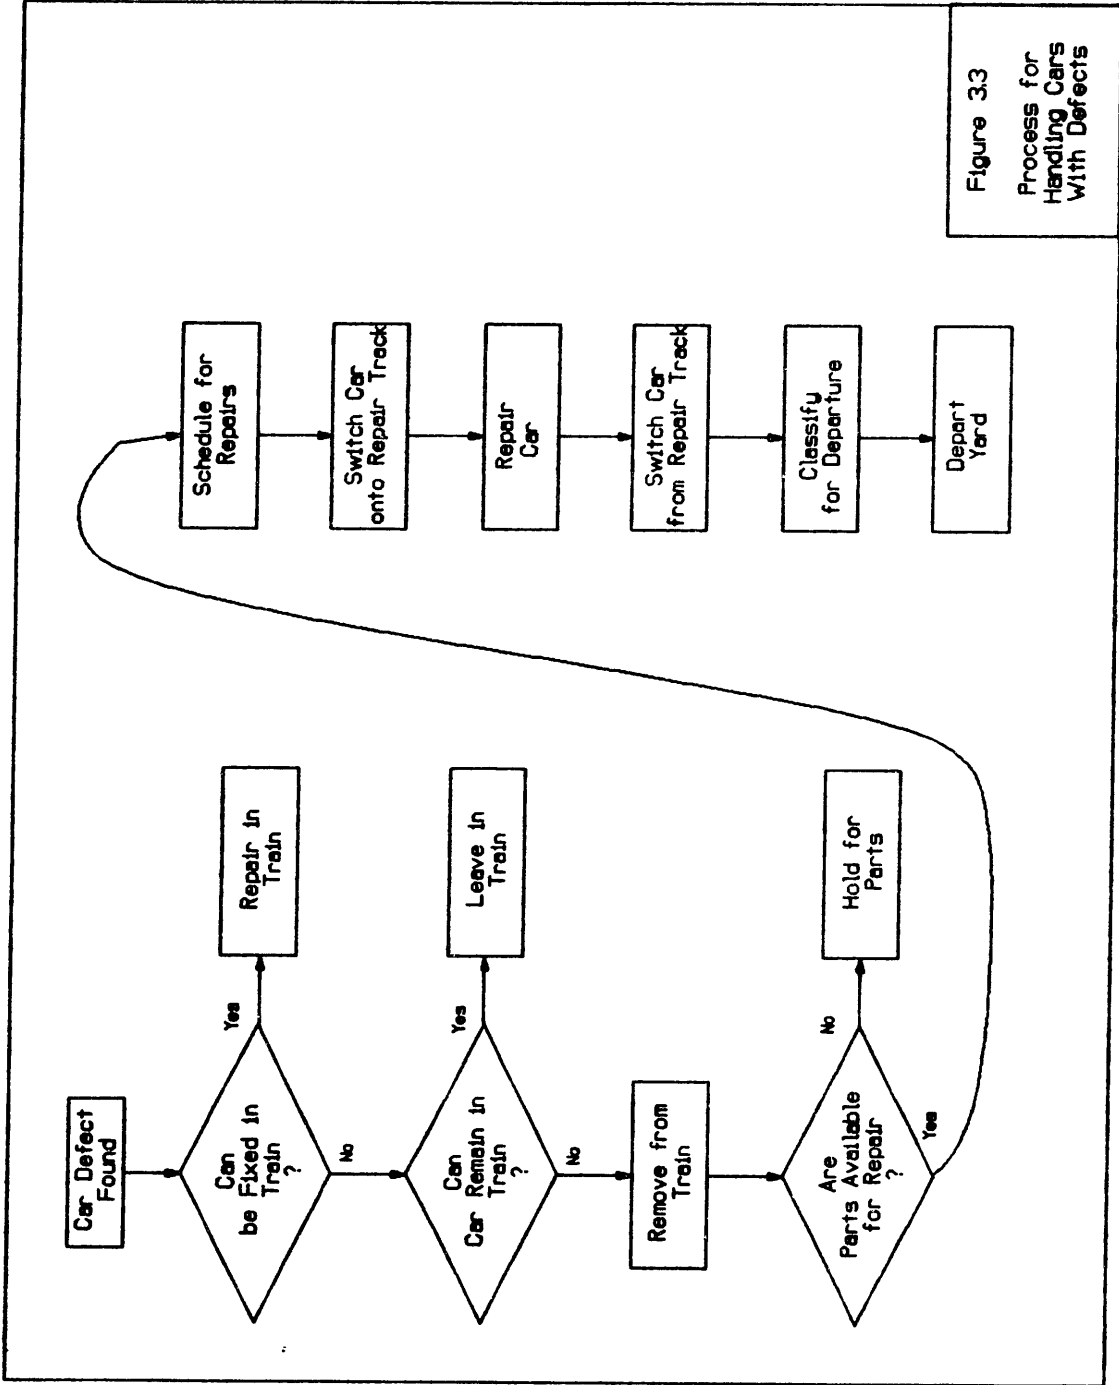
Figure 3.3

Process for
Handling Cars
With Defects

switched out of the train. In short, once a car is found to be defective and require repairs, there are a great many things which can lead to that car and its contents being delayed.

If the car is moving on more than one railroad, the number of ways the car can be delayed only increases. That is the subject of the following section.

3.3.2. Freight Car Operations: The Interline Case

Shippers desire to move shipments as part of their production and delivery processes, without regard for the particular networks of individual railroads. But railroads are limited to the particular regions or parts of the country where they have trackage. The result is *interchange service*, whereby cars and shipments are allowed to move freely from one carrier to another at specified junction points, in accordance with rules agreed upon by the various roads. Interchange service has implications for the design, operation, and maintenance of freight cars. In the absence of agreed upon standards and procedures for interchange service, railroads with inadequately maintained equipment could cause derailments on other railroads with no means to compensate them for the consequences. Similarly, if cars broke down on another railroad, each railroad would need to negotiate the cost and quality of the repairs. Some equipment might not even be able to operate over other railroads because of clearance restrictions.

To resolve these problems and to establish basic standards for the design, operation, repair, and billing of interchange service, the railroad industry's trade association, the Association of American Railroads, has set various voluntary standards and established common formats for the exchange of information.

Consider the case of a car moved according to the cycle in Figure 3.4. Note that part of the line haul portion of the car cycle, and the customer delivery occur on a different railroad than the loading of the car. The car now receives an additional set of inspections, at the junctions where the car is interchanged. These inspections insure that the car is in accordance the standards set forth in the **Field Manual of the AAR Interchange Rules**, which state exactly what condition the components and systems of a car must meet to be accepted by another carrier. Like the single line case, there are several possible outcomes when a defect is found, including a spot repair, allowing the car to complete the trip (then returned to the owner for repairs), and sending the car to
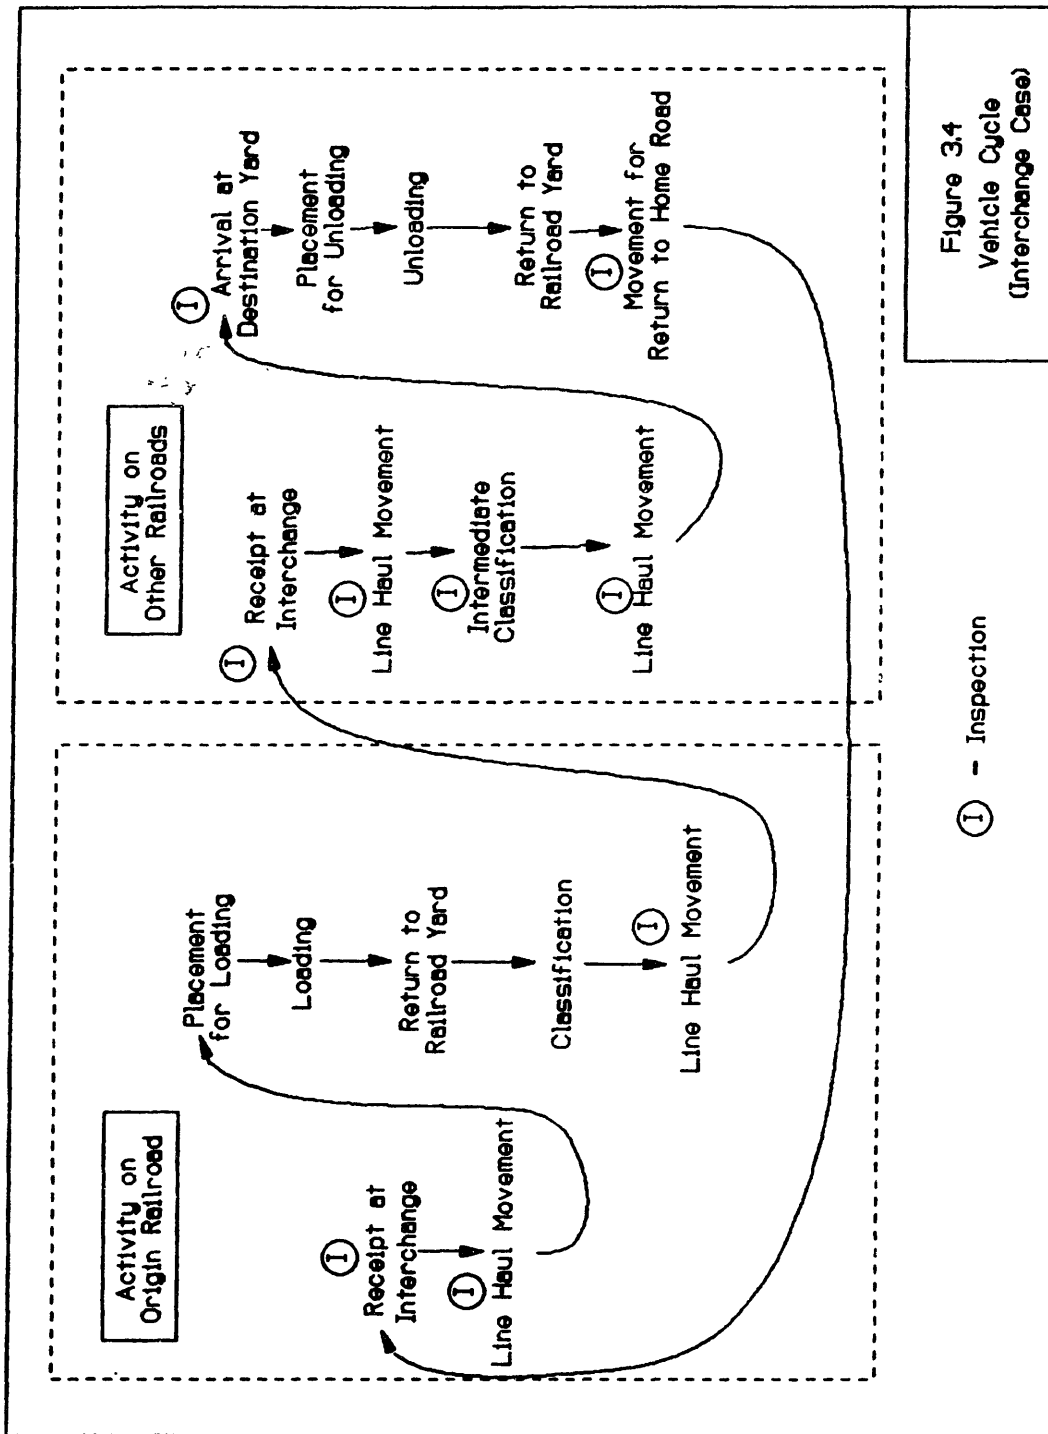
Figure 3.4
Vehicle Cycle
(Interchange Case)

Activity on Other Railroads

Arrival at Destination Yard → Placement for Unloading → Unloading → Return to Railroad Yard → Movement for Return to Home Road

Receipt at Interchange → Line Haul Movement → Intermediate Classification → Line Haul Movement

Activity on Origin Railroad

Placement for Loading → Loading → Return to Railroad Yard → Classification → Line Haul Movement

Receipt at Interchange → Line Haul Movement

(I) – Inspection

a repair track. To these is added the additional possibility of rejection of the car. The receiving railroad is not required to accept a car at interchange which fails to meet the standards, although most railroads have agreements that provide for the car to be repaired at the nearest rip track of the receiving road. Two key issues follow from the above discussion; the source of the standards and compensation for repairs to cars which fail to meet the interchange standards.

Turning first to the source of the standards, the standards are set forth by the Mechanical Division of the AAR. Most divisions of the AAR are organized into sections which are supervised by committees, and the Mechanical Division is no exception. The most powerful committee in this area is the General Committee, which is composed of the Chief Mechanical Officers of the AAR's member roads. This committee sets policy for all the other committees concerned with car design, maintenance, repair, and billing issues, and must approve all significant changes in the interchange rules and associated practices. Other committees oversee particular AAR activities, including certification of new car or component designs, approved maintenance practices, and auditing the quality of repairs of interchanged cars.

The second issue, the billing of repairs, is also addressed by the AAR's Mechanical Division through the Car Repair Billing (CRB) system. This system provides for a standard set of data formats, billing procedures, and uniform pricing for repairs performed on cars in interchange service. The data formats insure that sufficient machine readable information is provided to permit car repair bills to be properly audited. The information includes what was repaired, why it was repaired, who performed the repair, and when the work was done. One result of the CRB system's standard formats and universal application is that virtually all the railroads in the U.S. and Canada follow the CRB formats, or very close variants, in their internal car maintenance management information systems.

The CRB system also provides standards for determining the responsibility for paying for repairs. Generally, if a component fails to meet the standards due to normal "wear and tear", the owner of the car is responsible for the repair. If, on the other hand, the defect is the result of damage or rough handling, the handling railroad (i.e., the one

performing the repair) is responsible.

The pricing of repairs is based upon ongoing studies of maintenance practices by the railroads. The prices of components are updated every 3 to 6 months, based on averages of prices paid by member roads. The pricing of labor is based on time and motion studies performed by AAR staff at field sites, and the average wages and other compensation being paid to car repair personnel. An important component of the pricing in the CRB system is the assignment of overheads and fixed costs to repair activities.

The Car Repair Billing system is subject to a problem which affects the management of car maintenance. This is the permitted lag between the performing and the billing (i.e., reporting) of repairs. The CRB system requires only that a repair be billed within 12 months of the date it is performed. The result is that some railroads fail to bill repairs in a timely manner, and on many railroads the clerks who process the car repair bills are vulnerable to layoffs during cash shortages. The effect of this is that the data regarding car repairs and maintenance activities is often out of date, creating uncertainty about the age and quality of the components on railroad cars. A car returning from interchange service may be in a very different condition than what the car owner believes. This uncertainty about the true condition of a car returning from interchange service encourages the use of on condition maintenance[4], since that depends only on detecting the actual state of a part during an inspection[5].

There are also a number of consequences of the widespread application of the interchange rules. The first is that most railroads tend to adopt the standards in the **Field Manual** as their own internal maintenance standard. This has several practical advantages:

- Use of the AAR standard simplifies the supervision of car inspectors and repair personnel by limiting their activities to one standard for all cars.

- Liability from the consequences of using wrong or inappropriate standards is

---

[4] See p.11.

[5] It will be shown in Chapter 6 that it is possible to confront and resolve this matter directly as part of an opportunistic heuristic.

spread among all the member roads of the AAR.

- It sets a uniform standard for cars that cars are maintained offline, so that when they return the general scope of maintenance activities applied to them are known, although the specifics are unknown for some time.

Use of the **Field Manual** as a standard also has its disadvantages, however:

- The AAR standards may not be appropriate for particular groups of cars. Consider, for example, cars which are used only on a slow speed line in very short trains. It may be needlessly expensive to maintain these cars as if they were going to be used in high speed trains in interchange service.

- Use of the AAR standards creates another incentive for managers to follow the "on condition" maintenance policies described in the previous chapter, even if there are factors such as economies of scale in maintenance which could lead to a better approach.

Interchange and the supporting institutional structures present the railroad car maintenance manager with a situation unique among transportation modes. Not only must the car be maintained to Federal safety standards, but in order to be fully used in the railroad environment it must meet a second set of standards which reflect the economic concerns of the industry as a whole. The use of the car in interchange service also effectively insures that the car owner faces uncertainty about what recent maintenance actions have been performed on the car. In a subsequent chapter, an alternative approach will be presented which recognizes and addresses the unique circumstances which the interchange of cars presents to the maintenance manager. In the next section, however, we look at how railroads have organized their maintenance functions to meet this unique environment.

3.4. Organizational Issues and Maintenance

Having shown some of the unusual circumstances that railroad operations create for the maintenance manager, we now turn our attention to how railroad car owners, and particularly the Class I railroads, typically organize the car repair and maintenance function. In this section we look at the general approach, drawing upon articles from the trade press and conversations with various railroad mechanical officers. In the next chapter, case studies of three companies are presented.

Given that the motivations for maintenance are so varied, it is clear that without direction from senior managers the activities undertaken by one part of an organization may well be in conflict with the overall objectives of the company. It is therefore worthwhile to focus briefly on the various levels of the typical railroad company and how the maintenance program fits with each level's perspective.

The organizational model used is that presented by C.K Mao (1988), i.e., the dual-system control paradigm[6]. Mao argued that transportation systems can best be understood as 2 systems: **the controlling system**, which is essentially organizational in nature, and **the system being controlled**, which is primarily technological. This approach argues that the physical and technological units are at the lowest levels of the organization, and the highest levels seek to develop general policies for harnessing technology in support of broad goals and objectives. The functions of the intermediate structures, such as departments or regions is to provide **linkages** between the two, by developing implementation plans and creating feedback mechanisms. Most railroads' maintenance activities are arranged along functional lines into a hierarchy of three levels (Figure 3.5):

- **Executive level**, the highest officers of the organization, with responsibility for strategic plans, corporate policies, and market goals.
- **Departmental level**, in this case the level including and immediately below the Chief Mechanical Officer, with responsibility for defining, implementing, managing and monitoring programs which translate the policies developed by the Executive level into actions followed at the shop level.
- **Shop level**, with responsibility for actually carrying out the maintenance program developed at the department level.

This description is something of a simplification of what is in practice a complex structure, but is illustrative of the general manner in which most railroad companies organize their maintenance functions. In the following sections, each level of the organization is discussed in terms of its areas of interest and responsibilities, and the

---

[6] A number of alternative approaches are available to analyze the activities of large organizations. Mao's approach has been adopted primarily because he specifically addresses the case of railroad vehicle maintenance, albeit for locomotives rather than freight cars. Also, the author finds his approach convincing.
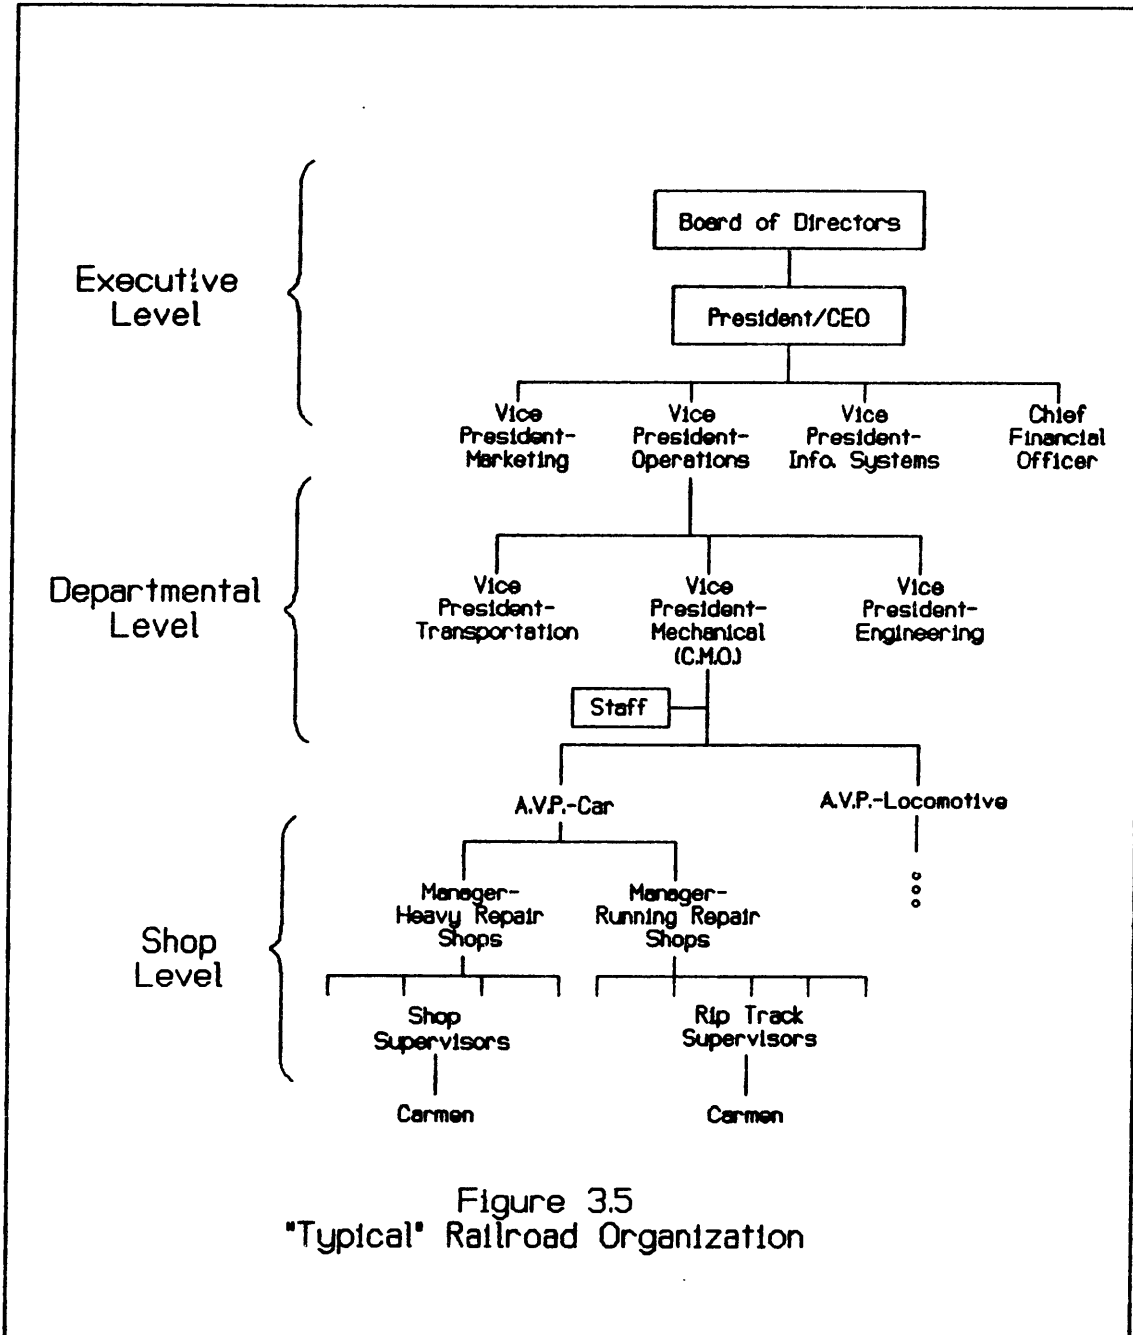
Figure 3.5
"Typical" Railroad Organization

supporting information that can be expected to be supplied to them.

3.4.1. The Executive Level

At the executive level, maintenance can be properly viewed as something to be integrated into other organizational plans. The company will have a group of overall plans for defining and achieving its objectives. These include

- *a strategic plan*, which defines markets the company intends to serve and a competitor analysis;
- *a capital plan* which outlines resources available and how these will be used in support of the objectives;
- *a marketing or service plan* which specifies the markets to be served and performance standards required to serve them; and
- *an operating plan*, which clearly states the activities to be performed by the various operating departments (transportation, engineering, equipment maintenance) to meet the requirements of the service plan.

In each case, equipment maintenance should be an integral part of the larger plan. In the strategic case, for example, maintenance is a tool which can appropriately be used to signal commitment to some markets. In the capital plan, maintenance is an integral part of fleet and facility planning, serving as a form of capital investment and as a tool for addressing uncertainty in some markets where investment may or may not be desirable (with deferred maintenance as a further option to generate short term cash). Finally, in the case of the marketing and operating plans, maintenance performs its traditional role of insuring an adequate supply of vehicles at minimum cost (and, in companies with well-formed "feedback loops", of indicating markets where the reliability and cost requirements may exceed the organization's capability).

Because maintenance is such an important component of railroad operations[7], both in terms of resources consumed and its potential impact on the provision of services, the information passed up to the executive level is much more extensive than in most corporate environments. Still, it is likely that the information is mostly of a financial or budgetary nature in the largest railroads, with a limited amount of exception reporting

---

[7] In Appendix A, for example, it is shown that freight car repairs constitute more than 10% of all railway operating expenses in the U.S.

when either service or cost targets are not being met.

3.4.2. The Departmental Level

At the departmental level, the maintenance plan must actually implement the "big picture" given above. The issues include how to deliver equipment according to standard at minimum cost, including fleet or series level planning, resource planning, and performance monitoring.

Departmental planning is generally performed at the level of fleets or series of cars, called *fleet management*. Fleets or series are groups of cars with common characteristics, both as matters of design and utilization. Each railroad generally seeks to define series in a way which is both natural (grouping boxcars with other boxcars, for example), and in a way which contributes to managing the maintenance function. A key issue in fleet management includes how best to assign cars to particular users so as to provide an adequate supply, while not giving excessively high cost cars to low revenue shipments, or incurring unnecessary wear on cars which are difficult to maintain. Another issue in this area is the decision of when to maintain, rebuild, and replace cars in a series, and what to replace them with.

Resource planning is the process of deciding what level of manpower, facilities, and money to apply to the maintenance of equipment. Effective resource planning involves both deciding on the level of resources that will be made available and how best to organize those resources to deliver the desired level of maintenance at minimum cost. The first matter, the overall level of resources, will necessarily involve some form of negotiation with other departments and upper management as part of a budgeting process. The second, the application of resources to specific tasks, can range from the large scale organization of the maintenance function down to the assignment of particular workloads and personnel to facilities.

The third aspect of the departmental level maintenance planning process is *performance monitoring*. Ideally, the maintenance function is organized in such a way that it is possible to track the performance of both the equipment being maintained and of the resources being used to carry out the maintenance. These will include

- shops and facilities;

- fleets and series of cars;
- individual cars;
- components.

In the next section we review some of the measures of performance currently being used, as reported in the trade press and in interviews.

3.4.2.1. Performance Measures Being Used By Railroads

In the previous chapter, some of the basic concepts of performance measures for maintenance were discussed. In particular, a good measure of maintenance effectiveness should exhibit several key properties [Armitage (1970)], including:

- ease of calculation, use, and interpretation;
- consistency with organizational objectives;
- an indication when something has gone wrong with decisions;
- the use of ratios should be avoided;
- the measure should reflect all the relevant consequences which effect performance.

The distinction between efficiency (resource control) and effectiveness (goal achievement) was also noted.

Given the importance of measuring the performance of maintenance activities, the number and quality of measures which railroads indicate they use is surprisingly low. Indeed, after reviewing the trade literature and interviewing a number of Chief Mechanical Officers (CMO), only the following performance measures were uncovered:

1) **The percent of the fleet bad ordered.** This number simply takes the number of cars listed as being out of service because of repairs divided by the total size of the fleet. A ratio, this measure is problematic for a number of reasons. If the size of the fleet is being reduced by retiring old or unreliable cars this is, at best, a measure of the effectiveness of the capital policy, not the maintenance policy.

2) **The number of cars set out on trains for bad orders.** This measure is also somewhat problematic, since it presumes that the quality of car inspections is constant over time. (If the number or quality of car inspectors declines, the apparent quality of the fleet goes up.) The index also fails to reflect changes in the number of inspections a car undergoes as the average length of haul changes. Perhaps most disturbingly, it punishes the Mechanical Department which finds and fixes defects, while assigning no penalty to the car with a severe defect which causes a catastrophe.

67

3) **The cost per loaded car-mile** (or per day, in some cases). While superior to the above measures, this measure must be adjusted for changes in the AAR billing rates for offline repairs; it also makes no distinction between the car with inexpensive repairs which "nickel and dime" the owner throughout a trip versus the car with a single expensive defect.

In the past, the cost and difficulty of developing appropriate information systems to carry out the performance monitoring function has been viewed as a critical restraint on implementing sophisticated maintenance systems. Because of the widespread availability of computerized information systems, and the need for systems to bill other railroads for repairs of foreign cars, most railroads now have the data required for monitoring performance available.

3.4.3. The Shop Level

At the shop level, the most important issue is the effective utilization of human resources and materials in order to carry out an assigned maintenance plan or approach. Management attention is likely to be focused on the selection of an appropriately sized workforce, determination of levels of parts inventories, and assignment of work in a way which realizes the company's business plan.

A crucial issue which is derived from upper levels of management but implemented at the lowest levels is the relationship between the maintenance plan and the conduct of labor-management relations. If there is a high level of tension between management and labor it will manifest itself most profoundly at the shop level in such forms as strict adherence to craft and seniority constraints, insistence on written directives for all tasks, poor safety records, and generally low productivity.

Labor-management relations have a long and sometimes bitter history in the railroad industry. Much of the history has been characterized by management seeking to change work practices in order to reduce the number of employees, and labor unions seeking to protect the jobs of their members. One of the ways that labor unions have been able to protect their members' jobs is through the creation of separate "crafts", or skills, and "scope" agreements, which provide that only members of certain crafts may perform certain tasks. This has had several effects in the area of car maintenance. First, it has raised the cost of work performed at union shops relative to otherwise identical

non-union facilities by introducing an artificial restriction on use of human resources. A second effect has been the need to define explicitly all tasks and operations associated with car repair, so that an agreement on which craft performs which task can be drawn up. This has the effect of making railroad operated shops much more resistant to changing procedures and practices. A third effect of scope agreements has been in the area of contracting out work. Most scope agreements have been written so that most, if not all, work on the equipment belonging to railroads must be performed by railroad employed workers. This has limited the ability of railroad maintenance managers to direct work to contract shops, which are often staffed by lower cost non-union workers. In the short run the scope provisions have tended to save jobs, but a review of railroad employment statistics suggests that this may be an unsuccessful long-term strategy for labor[8].

A further issue at the shop level is the relationship between the level of skill and experience of the workforce and the level of complexity of maintenance plans. If the level of skill is low, then no amount of sophistication in the maintenance plan can lead to a well-maintained fleet, and if the skill levels are high, then failure to use the skills and experience of the workforce essentially wastes the firm's human capital.

The information made available to the shop level will depend to a great extent on the company's particular information systems and on the maintenance plan. If maintenance is restricted simply to carrying out a single centrally defined set of maintenance actions (e.g., following the **Field Manual** for all cars), then the information provided to shop managers is likely to be quite limited, and almost exclusively related to parts inventories and workload reports. If, on the other hand, the railroad has either a sophisticated plan for deciding what repairs to perform or allows a certain level of local discretion, information regarding a car's history and future activities may also be provided.

---

[8] Consider, for example, that from 1980 to 1988 the number of workers on Class I railroads employed in the I.C.C.'s maintenance of equipment and stores category (Form A, line 904) has declined from 99,487 to 45,209 [AAR(b,1980,1988)].

3.5. Conclusion

In this chapter, we have begun the movement away from the general terms of maintenance to the maintenance of the railroad freight car. In order to make this transition, first some of the basic facts about the fleet of freight cars in use on U.S. railroads were presented. It was found that the number of cars is decreasing, and the average age of the cars is increasing. This older, smaller fleet was shown to be used more intensively, however, making the potential costs of breakdowns more significant. In a subsequent section, the freight car itself was introduced, and some of the ways that cars are used in railroad operations were presented. The use of cars on many railroads (interchange service) has had several effects. These include the development of uniform maintenance, repair and billing standards, and indirectly, the adoption of the AAR standards by railroads for their own fleets.

In the final section, we looked in general terms at the organization of railroad car maintenance activities by the large railroads. One of the findings was that the performance measures appeared to be weak. In the next chapter, we look in detail at the approaches taken by three companies.

Chapter 4

Railroad Car Maintenance: Case Studies

4.1. Introduction

In this chapter, the general environment described in the Chapter 3 is made concrete by examining the maintenance policies and practices of three companies who own and manage freight car fleets. The companies were selected not only because they represent current practices (as described in a series of interviews with various industry officials), but also because they have a genuine and demonstrated interest in improving their maintenance programs. Although quite diverse in terms of size, fleet composition, transportation activities, and maintenance policies, the three companies have important similarities that are central to this thesis. Each of them uses information systems based around the Car Repair Billing System described in the previous chapter. The measures of maintenance effectiveness used by each company could be improved to facilitate management of car repair activities. As stated above, the maintenance managers of each company seem to be genuinely interested in improving the reliability of their fleet. Most importantly, the maintenance policies of each company are relatively simple, often have only a limited basis in reliability studies, and fail to exploit all the information the companies have regarding costs, failure modes, or potential economies of scale in maintenance. Their policies accurately reflect the best of current practice; like the rest of the industry, the policies fall into one of two categories, either "on condition" policies, based more or less on the AAR standards, or "hard time" maintenance, with maintenance activities scheduled at fixed intervals. In a later chapter, the problems with each approach are discussed in detail, and an alternative policy is proposed.

In examining each company, the basic approach will be to focus upon several key matters:

- the company itself, including the size, structure, markets served, and the relationship of freight car maintenance to overall activities;
- the freight car fleet owned and operated by the company;
- the organization and staffing of the car maintenance function, including the resources applied to car repair and maintenance activities;

- the information systems used to support the maintenance function, including the types of data available to maintenance managers;
- the maintenance policy (or policies) followed, and the company's assessment of the policy.
- the measures used to evaluate maintenance effectiveness within the company.

Because some of the materials provided by the companies are proprietary, they are referred to throughout as companies A, B, and C. Company A is a regional railroad with a relatively small fleet, which is used extensively in interline movements. Because its cars are offline so much of the time and it believes it can perform repairs at lower costs than the standard rates, Company A's current maintenance focus is on ways to reduce car repair billing payments to other carriers. Company B is a very large railroad, with a large traffic base, including substantial international movements. Company B is presently in the process of attempting to institute a large scale program of planned maintenance, primarily to increase service reliability, and as part of an overall drive to let marketing considerations determine operations activities. Company C is a chemical manufacturer, with its own fleet of specialized equipment used to support manufacturing and sales objectives. Transportation expenses represent only a small proportion of overall sales revenues, and in light of the hazardous nature of many of their products, Company C's railcar maintenance strategy is driven almost exclusively by mechanical reliability and safety concerns.

In the following sections, each company is examined individually. After presenting a brief overview of the company, the maintenance policies are described in detail. Particular attention is then paid to the information systems used to support the car maintenance function. Finally, we consider the direction each company's managers would like to see their maintenance policies and programs take. Following the presentation of the three cases, some observations about car maintenance practices are made and some conclusions drawn.[1]

---

[1]In order to honor the companies' requests for anonymity, some of the numbers presented are given in general ranges rather than actual values. A few other numbers were simply not made available by the companies. Only Company A made full copies of all the reports used by their managers available to the author. In no case has data been

4.2. Company A - A Regional Railroad

Company A is a comparatively small carrier primarily engaged in providing service to the paper and paper products industry in a northern state. Company A maintains an active membership in the Association of American Railroads (AAR), the Regional Railroads of America and the American Short Line Association, the three trade associations which represent the interests of railroads.

Company A operates approximately 450 miles of track, with a fleet of more than 40 locomotives and 3000 freight cars in active service. Some 250 work equipment cars are operated by Company A. Approximately 400 people are employed by the railroad. The company provides no passenger service.

Annual freight revenues are in excess of $30 million. Car loadings in 1988 were approximately 55,000, up from 51,000 in 1987. The average net load in 1988 weighed 57.6 tons. Although no data was available regarding on-line car miles, off-line mileage grew over the past few years from just under 32 million miles in 1986 to more than 36 million miles in 1988. 1989 mileage extrapolated from the first 4 months of the year was expected to exceed 38 million miles.

About 55% of Company A's loadings are local moves, that is, traffic which is not handled by any other railroad. The local loadings are primarily pulpwood, logs, and wood chips used as raw materials by the paper mills, and energy products moved from an ocean port to the paper mills. Forwarded traffic (primarily paper and paper products) comprises 33% of the railroad's loadings, and the balance (12%) is traffic received from other carriers (mostly chemicals and raw materials, and seasonal agricultural products such as fertilizer). Cycle characteristics for cars forwarded from Company A were computed from data that the railroad made available. A cycle is defined as the entire period from placement of a car at a customer for loading until the next placement for another loading by one of the railroad's customers. The approximate distance travelled per cycle was found to be just less than 2000 off-line miles. This is somewhat

---

explicitly falsified; nor does the author believe he has been denied access to information which might have substantially changed the chapter's conclusions.

misleading, however, in that many cars are believed to be reloaded and used by the receiving carriers for other shipments. The mileage thus reflects not only the average length of haul on behalf of Company A's customer and return, but also any additional mileage incurred by intermediate loadings. While Company A is compensated for the use of its car, the car is completely beyond the control of the railroad at such times. The average cycle time for all forwarded loadings is 36 days, 26 days off-line and 10 days on-line.

The management structure of Company A as shown in Figure 4.1 can reasonably be characterized as lean. The company is organized along functional lines, with 4 basic groupings: human resources, financial, operations, and marketing. The Vice-President of Operations (VPO) oversees the engineering, mechanical and transportation functions, with the aim of insuring cooperation and coordination among the three operating groups. In spite of this broad grouping, the railroad's size and relatively small number of managers insures that the separation between line supervisors and the President is never more than 3 levels of the organization. This means, among other things, that a meeting to address problems in some part of the operation can easily include both policy makers and line personnel.

Senior managers indicated that one of the railroad's goals was to improve coordination between departments in both policy and operation. In the past, managers from the various departments had been extremely conscious and protective of what they viewed as their own "turf", or areas of responsibility and authority. Interviews with various personnel suggested that the desired coordination is becoming a reality. This change is relatively recent, with the example given of cooperative work between Mechanical and Marketing officials to provide input to a major customer in the design of some specialized equipment to haul tree length logs.

4.2.1. Company A's Car Fleet

Company A's car fleet is primarily composed of three groups of cars, rack cars (used in the transport of pulpwood and logs), gondolas (for hauling woodchips), and boxcars ("the paper fleet"). The paper cars are considered to be the pride of the fleet, since they are used in interline service and are what most employees and customers think

Figure 4.1
Company A

President

Vice President-Marketing
A.V.P.-Sales
A.V.P.-Marketing

Executive Department
Data Processing Department

Chief Engineer
Roadmasters

Vice President-Operations

Superintendent-Transportation
Trainmasters

Chief Mechanical Officer
Ass't. CMO
Foreman-Back Shop
General Car Foreman
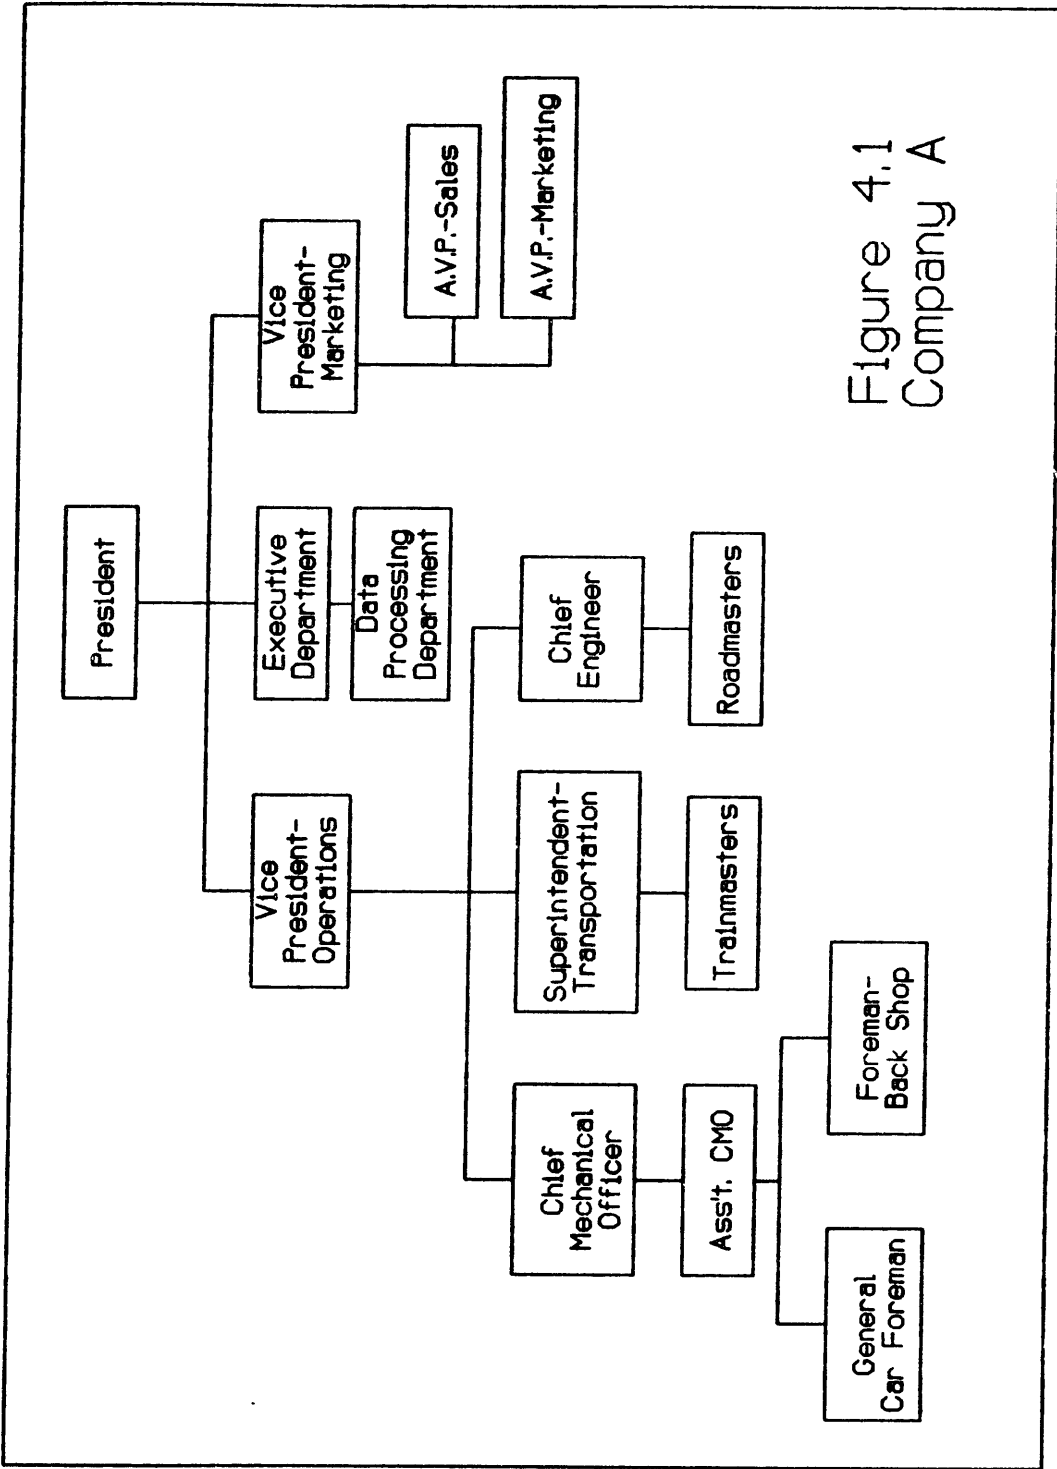
of when they think of the railroad's cars. The entire fleet is quite old; the typical boxcar is 15 years old, the rack cars more than 30. The age and condition of the gondolas is a particular source of concern to the Chief Mechanical Officer (CMO). One of the important decisions that he feels must be faced in the near future is whether to overhaul, replace, or scrap these cars. That decision must be made jointly with the Marketing department, and he expressed considerable relief that the relationship with them has improved in recent years.

Financial and other conditions have led to increased use of leased cars. About half of the paper cars are under some form of lease, with Company A responsible for maintenance costs of about half of those. The remainder can be rebilled to the owner at the AAR rates.

4.2.2. Car Maintenance on Company A

The railroad operates 4 primary maintenance facilities and has roving carmen assigned to others areas the railroad serves. One of the facilities is almost exclusively a "back shop" dedicated to heavy repair and major overhauls. It also operates the wheel shop for the railroad. (A wheel shop "turns" or renews certain types of wheels for re-use.) Another of the maintenance facilities is actually part of the mill facilities of a major customer. The work performed consists mostly of replacing brake shoes and air hoses and other minor repairs which do not involve extensive equipment. The largest repair facility on the railroad is dedicated mostly to performing "running repairs", that is, unscheduled defects which have been noticed on cars in active service or under load. Much of the work there consists of repairing safety defects (bent grab irons) and air brake work, although they are capable of extensive repairs when called upon. Because much of the workload involves cars under load, a night shift is also employed at this location. The fourth facility, located at a major interchange point, performs both running repairs and "program work", such as replacement of end-of-car cushioning devices. This is the only facility other than the back shop with an indoor work area; it also includes an air brake shop where brake components are rebuilt following the periodic servicing of brake systems known as a Clean, Oil, Test and Stencil (COT&S). None of Company A's facilities is switched more than once per day, and while each is relatively near to

76

computer terminals and facilities, none relies on computerized information in selecting cars to be repaired or to assist in diagnostics. This is probably reflective of the lack of such data in a form usable by the shop foremen.

The car maintenance activity is supervised under the structure shown in Figure 4.1. The CMO reports to the Vice-President of Operations. (The CMO is also responsible for locomotive maintenance.) He has a very small staff, consisting of an assistant CMO and a person responsible for the AAR Car Repair Billing (CRB). The AAR/CRB person also acts as a de-facto administrative assistant to the CMO, providing him with special data when requested and pointing out any anomalies in the car repair process he observes. This person also shares responsibility for entering repair data into the computer system. (None of this data is currently input at field sites where the work is done, although there are plans to do so at one location in the near future.) The chief field supervisor is the General Car Foreman-System. In addition to providing day to day supervision of the facility near the interchange, he also supervises the General Foremen at the other two running repair locations.

Annual maintenance activity measured in terms of costs and patterns is shown in Table 4.1. The online expenses are somewhat misleading because they reflect Company A's billing of itself at AAR rates, even when the actual out-of-pocket cost may be different. There are several reasons to believe that this substantially overstates online repair costs. The AAR labor rates have increased in the past several years (Table 4.2), while Company A's hourly wage rates have remained unchanged for more than 5 years due to the failure to reach a labor agreement during that period. The railroad also believes that it is also more efficient than most other carriers, and the AAR rate is an average of all carriers which is used in computing the AAR rates. The general car foreman cited the case of end of car cushioning device replacement as one example. The AAR estimates these require 7 man hours to complete, but Company A typically require only 2 man hours to perform such a repair.[2] A third reason why Company A believes

---

[2] The author observed one being completed in that time while touring one of the company's shops.

		TOTAL	REPAIRS	WHEELS	B&SHOES	COT&S	EOC
1986	TOTAL	$1,683,628	906	$649,349	$250,040	$121,130	$51,773
				(38.6%)	(14.9%)	(7.2%)	(3.1%)
	FOREIGN	$663,049	420	$297,308	$136,999	$24,520	$9,708
				(44.8%)	(20.7%)	(3.7%)	(1.5%)
	ONLINE	$1,020,579	486	$352,041	$113,041	$96,610	$42,065
				(34.5%)	(11.1%)	(9.5%)	(4.0%)
1987	TOTAL	$2,431,768	1118	$843,415	$242,954	$199,379	$143,590
				(34.7%)	(10.0%)	(8.2%)	(5.9%)
	FOREIGN	$675,635	418	$300,136	$132,803	$22,838	$12,993
				(44.4%)	(19.7%)	(3.4%)	(1.9%)
	ONLINE	$1,756,151	700	$543,279	$110,151	$176,540	$130,597
				(36.4%)	(7.4%)	(11.8%)	(8.8%)
1988	TOTAL	$2,811,328	1164	$900,576	$257,025	$313,335	$144,112
				(32.6%)	(9.3%)	(11.3%)	(5.2%)
	FOREIGN	$814,485	443	$332,073	$151,142	$60,901	$26,469
				(40.8%)	(18.6%)	(7.5%)	(3.2%)
	ONLINE	$1,996,843	721	$566,503	$105,883	$252,434	$117,643
				(28.5%)	(5.3%)	(12.6%)	(5.9%)
1989 Jan-June	TOTAL	$1,385,266	668	$537,767	$151,460	$40,319	$114,668
				(39.7%)	(11.2%)	(3.0%)	(8.5%)
	FOREIGN	$473,282	258	$206,059	$88,266	$7,228	$43,050
				(43.5%)	(18.7%)	(1.5%)	(9.1%)
	ONLINE	$911,984	410	$331,708	$63,192	$33,116	$71,618
				(36.4%)	(6.9%)	(3.6%)	(7.9%)

Table 4.1
Company A
Repair Costs, Online and Foreign, 1986-1989

January, 1986	$ 57.99
January, 1987	$ 58.81
July, 1987	$ 62.41
January, 1988	$ 63.44
July, 1988	$ 67.48
January, 1989	$ 64.36

Table 4.2
Hourly A.A.R. Labor Rates
(including overhead and indirect labor)

that it performs repairs at less than the AAR rate is the railroad's relatively low overhead. The railroad has seen substantial staff reductions in the past decades (although this is true throughout the industry), and has salaries which reflect a more modest cost of living than in much of the U.S.

In recent years, Company A has experienced increasing expenses both on and offline. The past several years have seen an increase in the ratio of offline to online expenses, and the decline in the relative cost of online labor suggests that offline expenses are increasing even faster. The total costs in 1988 ($ 2,811,328) represent a 67% increase over total costs in 1986 ($ 1,683,628) while offline mileage was up only 14% over the same period. In 1988, the total cost per offline mile was $.078. If we adjust for the estimated online mileage (10%), and the overstatement in using AAR billing for work, the cost per mile is probably closer to $0.06[3]. The cost per car in 1988 was approximately $875, which compares favorably with the U.S. average of $1,131.

When repair activity is examined by key component groups, we find that wheels and brakeshoes constitute approximately 44% of total repair costs in 1988, and almost 60% of offline repair costs. The CMO furnished the author with a report issued to CRB

[3] The overall cost per mile for all cars in the U.S. is approximately $0.07. See Appendix A.

system participants by the AAR which showed that typical AAR Car Repair Billing (i.e., average off-line) figures for the same two components for 1988 were less than 50%[4]. The higher than average offline costs are believed to reflect that Company A's offline mileage as a percentage of total mileage is substantially above the national average. Company A's high COT&S expenses probably reflect several series coming due for that maintenance in this period.

Several long term trends are noteworthy. While the percentage of wheel and brakeshoe expenses are relatively stable, there has been a marked increase in the number of end-of-car (EOC) cushioning device repairs in recent years. Such repairs have more than tripled as a percent of total costs, and the number of replacements appears to be growing at an even higher rate. Some of this increase is simply that units are wearing out at the end of an approximately 10 year wear cycle. Company A officials believe, however, that the dramatic increase in foreign repairs of EOC units reflects overpricing by the AAR and a new awareness on the part of other roads that this is a profitable repair.

4.2.3. Information Systems Supporting Car Maintenance

To support the car maintenance function, Company A has a computerized information system. This system was developed in the early 1980's by an individual who had worked in the Mechanical Department prior to his transfer to the Data Processing (DP) Department. The system has been changed over the years when time was available, but most of the system remains as originally designed. The programmer indicated in an interview that there are a number of changes that could be made to improve the system, most notably allowing on-line access to the data by Mechanical Department officials, but the current data processing workload simply does not permit such an effort. Although

---

[4] These figures are consistent with earlier studies. In 1977, for example, wheels and brake repairs were found to constitute 43.5% of all repair costs for the railroads participating in the A.A.R.'s Car Maintenance Cost Data Base, and offline repairs (i.e., CRB system billed repairs) of these components were 62.5% of total costs [Guins and Hargrove (1980)]. In that study of more than 200,000 freight cars owned by four railroads, the discrepancy was assumed to represent the inclusion of heavy repairs in the online expenditures.

developed primarily to assist in producing the CRB system bills, the system is also used to generate a number of daily, monthly, and ad hoc reports.

Reports used by the Mechanical Department include:

(1) Daily Statistics Report: This report lists all the bad ordered cars by location, giving such information as the car type, repair status (light, heavy, or derailed), loaded or empty status, date when bad ordered, hourly car hire rate, and total amount of car hire lost while car has been on bad order status. It is used by the CMO and his staff to monitor the status and performance of the various facilities and fleet, and by Marketing officials to track loaded cars for customers.

(2) Daily Car Hire Report: This report, which is generated primarily for the Transportation officers, shows the car hire associated with cars currently in a bad-order status, and gives an indication of the opportunity cost associated with shop backlogs.

(3) Monthly Car Repair Statistical Report: This report gives summary and detailed information on car repairs reported in the previous month. It is divided into 3 parts.

> Part A: lists by car series the number of cars billed, the labor, material and total costs, and the cost per repaired car.
> Part B: lists on and offline costs by general mechanical component systems (brakes, wheels, trucks, etc.), in terms of Company A repairs of foreign cars, foreign repairs of Company A cars, and repairs by Company A of its own cars, all costed at AAR repair rates.
> Part C: lists disbursements per Rule 23, i.e., who owes Company A money for repairs performed on their cars.

The report is used by the CMO and VPO to keep track of car repair costs (vs. budget) and to monitor changes in the number and cost of particular component repairs.

(4) Audit Summary of Exchange Tape: Reports the number and cost of repairs (labor, material, and total) billed by or to Company A in the previous month (including internal repairs.) This is, in effect, the summary of the CRB

system bill to be paid by Company A. The railroad is a "car repair debtor". That is, it owes more money than it receives each month because of the nature of its traffic. This report, like the Statistical Report, is used mostly to monitor unusual expenses and to determine if a pattern of increased expenses owed a particular railroad is developing. The report is also used by the Financial group to organize payments to and from other roads through the AAR.

(5) Monthly Car Movement Report: This is one of a series of reports developed to assist the person responsible for the AAR/CRB system to audit the bills received from other railroads. It generates all movements of cars belonging to Company A that are reported through the AAR. These include bad ordering, interchanges, arrival at final destination, placement at customer siding, etc. The report is used to insure that a car which is billed for a repair by a railroad was actually on that railroad and at the repair location at the time the repair was said to have occurred.

(6) Monthly Exception Report of Foreign Car Repair Billings: This is another report designed to assist in auditing the CRB data. Any car which is billed is checked against a file for irregularities, including:

- low usage job codes
- unusual why-made codes
- inappropriate or non-standard job codes relative to the design of the car
- non-existent or restricted car number
- unusually expensive repairs
- duplicate bills
- identical repair in the past year (certain codes only).

This report depends on knowledge built into the system at its design, when the programmer responsible was a mechanical engineer with considerable experience in the Mechanical Department. The knowledge base has not been maintained over the past several years, so some of the information is not entirely reliable, and particular design features of newer cars have not been

added to the system. In addition to its use in auditing CRB data, the report is also used to bring expensive repairs to the attention of the CMO and Assistant CMO.

There are also a number of ad hoc reports which are generated when needed:

(7) <u>Selected Car Repair Billing Items From Car History</u>: This report can be generated upon request of the CMO or staff. It lists all repairs made to a car or group of cars for a range of repair codes for the previous 5-6 years. It includes who performed the repair, the accounting and repair dates, the job code, the location on the car (if applicable), the responsibility code (i.e., who paid for the repair), and a description of the defect. The report is used to examine individual cars which may be experiencing undue repair expenses or to look at groups of cars which have become a source of concern. The report is not generated interactively by the user, and is only available in printed form from the DP Department, which sometimes results in time lags between the time the report is requested and when the data is actually received. The report is not currently made available to repair track personnel.

(8) <u>Derailed Cars</u>: This is simply a list of all cars which are known to have experienced a derailment in the past 5-6 years (based on certain why made codes in CRB data), the date the flagged repair was made, and the railroad which billed the flagged repair. It is used as part of the auditing process to determine repairs which may be the responsibility of another railroad.

In addition to these reports, the AAR/CRB monitor also manually maintains logs of various car repair and usage data, including reports of offline mileage and repairs by month, total online and offline repairs to Company A's cars, number of carmen by month, the number of cars receiving repairs, and number and cost of repairs by major components. These data are entered into a personal computer used by the CMO and used to generate graphics and charts to assist in tracking and understanding unusual or disturbing trends. These charts are the only reports that have been designed specifically to assist the CMO in monitoring performance.

To sum up, Company A maintains an extensive data base of the repairs performed

on their cars, the activities of their shop facilities, and the costs of various maintenance actions. From this data base, a large number of reports are generated, presenting the data in ways that help monitor the car repair expenses and the level of activities at their shops. Because substantial cash payments are involved, foreign car repair billings are carefully checked for possible exceptions. Trends in overall performance of cars or components are examined only with direct intervention and action by the CMO.

Not surprisingly, managers of Company A see the need for changes in the information they receive. One of the themes voiced most often by the various individuals interviewed was the belief that the appropriate information is not getting to the right people. This concern almost surely translates into reduced efficiency in managing car maintenance. The General Car Foreman-System, for example, indicated a number of uses that he could make of the a car's repair history. He stated he would use the histories to look for excessive wear or parts consumption, and to monitor repairs billed but not satisfactorily performed. He indicated that a particularly useful fact would be a car's derailment history, since a derailed and repaired car may be subject to a number of potentially expensive defects.

At the other end of the chain of command, the CMO acknowledged that he could easily "spend a week just looking over the reports and data to get a better handle on what [they] ought to be doing, if I had a week". The real premise here seems to be that the volume of data is too great for an official at the level in the decision hierarchy of the CMO. The lack of on-line computer access to the data by the field personnel and the auditor seems to restrict spontaneous inquiries and cause a disruption in the process of auditing improper repairs. One of the more interesting items learned in the interviews was that DP maintains a car master file which includes various standard information about the car, the date of the most recent repair of the car for several components, and recent car movement data. While this data is not complete, it could serve as a starting point for developing a more sophisticated car history such as that discussed in Chapter 9.

4.2.4. Car Maintenance Policies

The current policy being followed by Company A can best be described as "on condition" maintenance, i.e., replace components upon failure or when they reach the

84

condemnation standards set in the AAR **Field Manual**, except in a few noteworthy cases. The two primary exceptions to this are early performance of COT&S (scheduled brake system overhaul) and limited early replacement of end-of-car cushioning devices.

In the case of early COT&S, the railroad is considering instituting a program based on the impending due dates for about 200 boxcars 2 years from now. Because the cars are leased with a provision for billing the lessor for repairs, such a program would have to be approved by the car's owner. The General Car Foreman believes that the COT&S can be performed for less cost by his workers than offline, and that they can therefore afford to offer a limited discount to the car owner, if necessary. Of more concern to him is that if the program is not instituted soon, the brake shop will not be able to keep up with the number of cars which have exceeded the due date and it will be necessary to buy rather than rebuild parts (at higher costs) to complete the repairs. While no definite plans have been made, the General Foreman indicated that they would initially limit early COT&S to those cars which were shopped for other defects; as the mandated date for the overhaul approached, cars would be scheduled into the shops specifically for the COT&S. This can be thought of as a form of opportunistic maintenance[5], where the failure of another component creates an opportunity to perform an earlier COT&S than scheduled or required.

The end of car cushioning device (EOC) problem is relatively new, with an increasing number of failures seeming to occur on cars 10 to 12 years old. The expectation is that some sort of preemptive maintenance program will be undertaken in the next few months to exploit the lower costs of online replacements compared to offline failures. The company recently hired an outside consultant to study the problem and develop a planned maintenance program. Again, one of the options the company is considering is to perform these replacements based on an opportunistic approach, i.e., replace EOC's earlier than the scheduled date (as developed by the consultant) if another

---

[5] See pp. 29-33.

part fails.[6]

In both cases, the managers of Company A are willing to schedule components for replacement prior to reaching the AAR condition and to preempt that schedule opportunistically. They expressed concern, however, about the need for sound economic criteria upon which to base decisions regarding when it is appropriate to perform opportunistic maintenance actions. In particular, they indicated they would like to have a formal basis for making the decision. Such a basis is presented in Chapter 6.

The CMO indicated that there are no formal guidelines or explicit rate of return targets currently used for approving planned maintenance programs, although he would have to justify the economic advantages to the President. Similarly, if his department was over the budget levels he was expected to explain what had happened and why. He indicated that the senior management was generally receptive to changes in his budget or operations when the change was a good investment. He stated that he believed that more planned maintenance instead of emergency repairs could significantly reduce the long term maintenance costs. The VPO, in a separate interview, said that he believed they could reduce the offline costs by as much as 50% by appropriate planned programs.

One of the potential areas for planned maintenance discussed in depth was wheels. Because these are such a significant part of the total budget, it would be desirable to reduce costs in this area. There are several problems with a planned approach in this area at this time. The General Foreman indicated that the wheel facility at the back shop is a bottleneck at the present time, and that he must sometimes allow wheels to remain in service which are very near the condemnation limit for want of a set on hand to replace them with. The other problem is that because of recent rises in the price of wheels and lags in the AAR pricing system, it has been as cheap to allow other railroads to assume the costs of wheels at AAR rates as to repair them online. This is a transient phenomena

---

[6] Already the railroad is performing some of these opportunistically. In particular, if the EOC is "about 10 years old" and the EOC on the other end of the car fails, and the crew has time, the General Foreman indicated he will sometimes direct them to replace both. No formal rule is in place to determine what the minimum age might be for replacing EOC's.

by all accounts, and, in any case, the other roads are probably doing the same thing which simply resulted in a lot of deferred maintenance until the billing system caught up with actual prices. There was some interest in a preventive or opportunistic program in this area as well if appropriate condemnation standards and decision rules could be defined.

4.2.5. Company A: Conclusions

Company A has organized itself in a way which takes advantage of its small size and well defined business plans. The car maintenance function seems well run by people interested in improving their department's performance over time. The most significant shortfall in the information system is the need to provide appropriate access to people at various functional levels in the organization. The shop level supervisors are willing and able to make use of detailed information regarding car repair and usage histories. The CMO is looking for reports which assist him in defining and measuring performance of the facilities and people under his direction. The entire organization is willing to undertake planned maintenance (either preventive or opportunistic) to reduce offline costs and improve fleet reliability. Although some of the necessary expertise to develop such reports and programs appears to be available in-house, the staff appears to lack the necessary time to develop such programs entirely on their own. (Certain analytic skills necessary to properly develop programs may be missing at this time as well.) The railroad may thus be forced to go outside for support in addressing their concerns. The company has been willing to use outside consultants and university researchers in the past to meet technical shortfalls. Of more concern is the need to develop a workable set of decision rules for opportunistic maintenance. With the current staff and information systems, the company has the resources to implement a program of opportunistic maintenance such as that presented later in the thesis.

4.3. Company B: A Large Railroad

Company B is one of the largest railroads in North America, operating over more than 20,000 miles of track, including that of subsidiary companies . The Mechanical Department is responsible for a fleet of more than 1000 locomotives and 50,000 freight cars. The fleet is composed of a wide variety of car types, including general purpose

boxcars, covered hoppers used in grain and other bulk commodity movements and an extensive fleet of gondolas used in coal train service. Company B is involved in significant intermodal movements and owns and operates a dedicated fleet of flat cars for this service. Recent purchases and activities have included upgrading cars used to transport automobiles from points of import or manufacture to distribution centers around North America. In all, Company B's car fleet is divided into 52 different groupings for maintenance purposes. In 1989, the fleet was operated for over 1.125 billion car miles.

4.3.1. Car Maintenance

In 1989, Company B spent approximately $90 million on the maintenance and repair of freight cars, with approximately 20% spent in heavy repairs and programs, and 80% at running and intermediate repair facilities. The maintenance cost per car was $1820, and the maintenance cost per mile was approximately $0.08. These costs are slightly higher than those for all railroads[7]. While some of the higher costs may reflect unique environmental or operating conditions, these higher costs also reflect several factors which characterize the company's maintenance policies discussed below:

(1) The company has made a conscious decision to replace components while useful life remains to insure that components can be remanufactured;

(2) The company has undertaken a program of planned maintenance for some of its fleet, which may be incurring high startup costs, and leading to further premature replacement of components;

4.3.2. Organization of Company B's Car Maintenance Activities

Company B has undergone a substantial reorganization in recent years, with the establishment of two regional management units which are responsible for day to day operations under policy direction and oversight from the company's headquarters, generally referred to as the "system" level. Each regional office makes decisions regarding budgeting, staffing, and actual activities, but must conform to the policies and standards set by the headquarters. The major impetus behind this has been the desire to

---

[7] See Appendix 1.

shift the company to being "market-driven", a term used regularly by employees at virtually every level of management. To further encourage this market orientation, Company B has recently spun off part of one of its regional units which currently hauls boxcar traffic into a separate business unit with the express mandate to initiate profitable competition with the trucking industry.

A natural consequence of the geography faced by Company B is that one of the regions is dominated by customers producing bulk commodities for unit train movements, while the other is largely oriented toward more specialized movements such as auto traffic, intermodal, chemical, and merchandise traffic. One of the primary roles of the headquarters offices is to insure coordination across the regions in terms of operations and service, since the equipment and trains are allowed to move freely across regional lines without interference or obstruction.

Unlike most railroads, a majority of the traffic hauled on any of Company B's trains is moved on equipment owned or controlled by the company. Part of this reflects the company's emphasis on developing unit train movements, but it is primarily an accident of geography. The railroad is large enough to complete many of its full shipments without interline service, and it runs parallel to competitors so that shippers generally choose either a full routing over Company B or over another carrier. As will be noted below, this has the effect of making Company B much more interested in developing comprehensive maintenance plans. They are seeking, as the Chief Mechanical Officer indicated in an interview with the author, to "capture all the benefits of reliability and to reduce the risk of someone else's bad car undermining their efforts to establish a reliable fleet".

Car maintenance on Company B is organized along the same lines as the overall company structure (Figure 4.2). There is a system level staff which reports to the Chief Mechanical Officer (CMO), who is responsible for all policies, standards, regulatory matters, and labor relations relating to equipment maintenance for the entire company. The CMO reports to the Vice President - System Operations, who reports to the President and Chief Executive Officer (CEO). Each of the regional groups has its own staff responsible for budgeting, staffing, and hands-on management of the shops. The regional
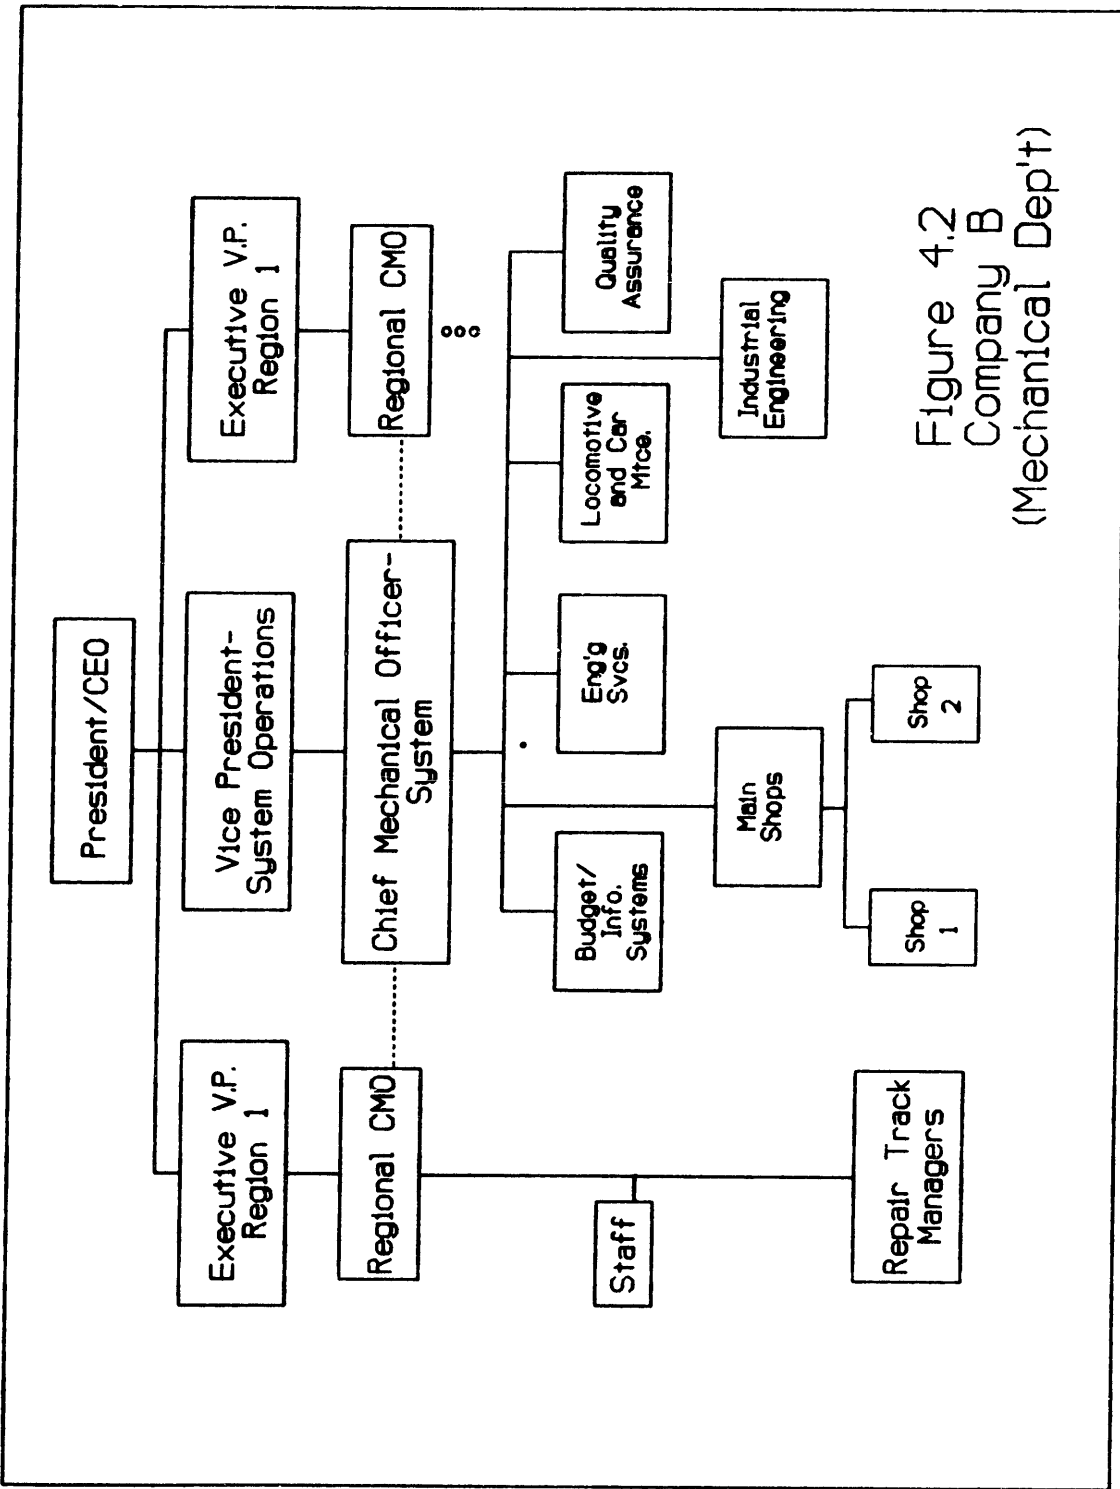
Figure 4.2
Company B
(Mechanical Dep't)

President/CEO

Vice President-
System Operations

Executive V.P.
Region 1

Executive V.P.
Region 1

Regional CMO

Chief Mechanical Officer-
System

Regional CMO

Budget/
Info.
Systems

Eng'g
Svcs.

Locomotive
and Car
Mtce.

Quality
Assurance

Industrial
Engineering

Main
Shops

Shop
1

Shop
2

Staff

Repair Track
Managers

mechanical function is supervised by the regional-level CMO, who reports to the Vice President for that business unit. The Vice President for the region reports to the President and CEO (Figure 4.2).

The headquarters staff is divided along functional lines into six groups who report to the CMO:

(1) Mechanical-Electrical Engineering Services is responsible for design specifications for cars and locomotives, data gathering and analysis with respect to the functioning of engineered systems, and monitoring equipment performance.

(2) Locomotive and Car Maintenance reviews and issues maintenance instructions issued to the regional units, and acts as a "firefighter" for any concerns or issues raised by the business units.

(3) Main Shops staff manages the company's heavy repair facilities.

(4) Industrial Engineering performs traditional industrial engineering studies, facilities planning, and productivity analyses. In recent years this group has taken the lead role in developing the company's planned maintenance program.

(5) Budget Planning and Information Services provides budget planning and oversight, and acts as liaison with the company's Computer and Communications Department (C&C), in all stages of information system design. This group has been responsible for the car and locomotive maintenance information systems since the systems' earliest design stages.

(6) Quality Assurance performs internal audits over the regional and main shops, remanufacturing centers, and over vendors. An important function has been to insure both that standards issued at the system level are being met at the regional level, and that instructions from the headquarters are clear and useful to shop-level personnel.

An especially important actor in the determination of car maintenance policy is the Freight Car Repair Committee. This is an interdepartmental group which meets every six weeks to review the car needs of the company and to assign priorities for car types which are then passed along to the various maintenance facilities. This committee includes representatives of the Transportation and Marketing departments, and both the system and regional level Mechanical departments. This group attempts to match market forecasts for equipment demand with shop capacities to plan workloads for the next month. The Car Repair Committee also monitors any complaints or concerns raised by

customers. This group performs a particularly important role in encouraging the company's "market-driven" approach to maintenance.

One of the effects of this interdepartmental approach has been an increased concern by the Mechanical Department for those aspects of the car of most concern to the railroad's customers. Several persons indicated that they had recently become more aware of and sensitive to the fact that the customers' requirements go well beyond the simple mechanical reliability of the car (i.e., the running gear), and extend to things such as the doors, linings, and floors. This has caused the railroad to pay more attention to standards in these areas than in the past.

4.3.3. Car Maintenance Facilities

The company operates three types of maintenance facilities: back shops, intermediate shops, and running shops. The back shops are responsible for so-called "heavy repairs", including rebuilding programs, repairing cars extensively damaged in wrecks, and in major capital-intensive programs such as modifications and painting. The back shops are under the direction of the system level management. There is an increasing movement by the company toward using the back shops more for specialized repairs (painting cars or relining tank cars) or for reconstruction of component assemblies (such as truck remanufacturing or turning wheels).

The intermediate and running shops are supervised at the regional level; the intermediate shops are capable of extensive repairs, and are dispersed geographically to allow various car fleets and locations access to extensive facilities which do not require special equipment. The intermediate shops can undertake, for example, complete overhauls of the running gear of the car (e.g., trucks, wheels, brakes, and draft and coupler systems).

The running shops generally perform spot repairs of cars which experience component-level failures (including failure to meet limits for allowing a component to remain in service although still operative). Work performed at the running shops usually reflects failures of components or cars while in revenue service, which is most undesirable from the customer's perspective. One of the goals of the company over time is to shift the work load from the back shops and the running shops to the intermediate shops.

4.3.4. Information Systems Used in Support of Car Maintenance

As is the case with most large companies, Company B has an extensive computer-based information system. Financial and operating data are maintained by the company's Computer and Communications (C & C) Department. Large information systems, such as those supporting car maintenance are designed and built by C & C in response to (and in coordination with) the user department. Once built, maintenance of the system and supporting data bases is also the responsibility of C & C, which has the effect of making changes to the data bases or the systems difficult and expensive.

In the past 2 years, the company has developed the Car Maintenance Information System (CMIS), which is used both for car accounting and maintenance management purposes. This system was modelled after the company's highly successful Locomotive Maintenance Information System (LMIS), which company managers indicate has significantly changed the type and productivity of locomotive maintenance. LMIS allowed the railroad to monitor the usage, failure patterns, and costs of various locomotive components. Prior to the development of LMIS, locomotive maintenance was performed based on preventive maintenance schedules recommended by component manufacturers. Using LMIS data, it was learned that some components which were being replaced preventively were, in fact, subject to random failures rather than wearing out, and so were not good candidates for preventive maintenance. By adding some components to the preventive maintenance plan, and removing others, it is generally agreed that the locomotive fleet has become more reliable at the same or lower cost. This success inspired the development of such a system for freight cars (CMIS).

Like virtually all railroads in North America, Company B participates actively in the AAR's Car Repair Billing System. This system provided a basic framework of data structures, component names, repair codes, and cost structures to use in developing CMIS. CMIS currently includes general information about a car, such as the car type, built date, and specialized components, and the complete repair record of the car. Like most railroads' car maintenance systems, revenue and usage information is recorded and stored in other files as part of other information systems. While there is considerable repair data stored in CMIS, the lack of readily available car usage information and the system's

93

adherence to CRB formats tend to make inquiries regarding car or maintenance performance difficult.

CMIS and other information systems are currently used by Company B to generate a number of reports on a daily, weekly, and monthly basis. These include:

(1) A report comparing the budget and actual expenditures for car maintenance.

(2) Equipment availability, measured in terms of the number of cars in each of the 52 car groups available and the number in bad order status.

(3) Manpower levels, for each region and for each facility, the number of person employed by craft, matched against the actual production of the facility, measured in cars repaired (weekly and monthly).

(4) Component usage/consumption - A monthly report showing the number of cars in each group, the number of miles the cars ran, the labor hours spent in maintenance on the cars and the number of each of the major components (wheels, draft gear, brakes, etc.) consumed by the car group. Component usage is compared with the recent month and the same month in the preceding year. In addition to a summary report by groups, the detailed information regarding each group by car numbers (i.e., particular series' of cars) is also provided.

(5) Performance by facility reports, which show the number of cars assigned to be rebuilt or repaired at each of the main shops, and the number of cars actually completed. This report is used by the Freight Car Committee to monitor fleet availability and to assign priorities to the main shops.

(6) Cars set out from trains for mechanical problems: This daily report is used as a proxy for the fleet's in service reliability by the CMO and others. This is considered to indicate the degree to which the Mechanical Department is inhibiting the marketability of the company's service.

All the reports are available using computer terminals on the various managers' desks, and are only printed when specifically requested by the user. Each report begins with a page of summary information, an information policy imposed by the CMO to encourage managers to focus their attention on problem areas rather than dwelling on details which

94

may not be very relevant to the management of the car fleet.

Repair data is input by clerks using terminals at field locations (i.e., shops), but the shop personnel themselves do not generally have access to a car's repair history.

4.3.5. Maintenance Policies and Standards

Company B is currently in the process of changing its car maintenance policy from what it calls "bad order" to "planned maintenance". The current system, based upon an "on condition" policy, provides for components to be replaced on a car when a component either fails or exceeds the company's condemnation standards. Condemnation standards come from both the general sources discussed in Chapter 3 (i.e., government and industry standards), and from the Engineering Services group.

Company B believes that the AAR's interchange standards are not rigorous enough in many cases, so their wear limits are generally more restrictive than those found in the **Field Manual.** This is not so much out of a desire to avoid in-service failure as a conscious concern with rebuilding or remanufacturing parts and components. The company's Quality Assurance group has undertaken to measure components received from vendors, and has found that as many as 80% of certain components delivered are not in compliance with the specifications ordered. After attempting unsuccessfully to improve vendor compliance, Company B established remanufacturing facilities for certain high value components when they are worn out, particularly trucks and truck components. In order to be remanufactured, however, these components must be removed from service before they reach the AAR condemnation limits. The difference in cost between a remanufactured truck and a new one is more than $1000 according to the CMO, so there is a clear incentive to maintain them to higher than interchange standards.

The condemnation standards currently used by the company are the product of "informed estimation" by experienced engineers and maintenance supervisors. A program to measure the actual wear of components under field conditions has been underway for several years, and is expected to result in some modification of the standards for the removal of components, although there is as yet no uniform agreement as to how to proceed with the analysis of the component wear measurements to develop new standards.

The current inspection practice is to examine cars as part of the train brake tests

and upon arrival and departure from the yards with running shops. If a component is found to be failed or beyond the condemnation limit, the car is removed from service and the component replaced. If the component cannot be replaced at the running shop, the car is sent to an intermediate shop (when empty, if the defective part is not safety-related). If the repairs are estimated to require more than 150 man hours to complete, the car is sent to the nearest back shop. This approach has resulted in some cars remaining in continuous service until components which are not easily inspected are worn beyond the limits for remanufacture. More significantly, one of the Mechanical Department managers stated that the Marketing department has indicated through the company's Freight Car Repair Committee that there may be substantial losses in customer goodwill in using cars with a high degree of unreliability. One measure of the overall reliability used by Company B is the percent of the fleet "bad ordered" at any time. Currently, the number of cars subject to the "on condition" policy in a bad order status is reported in the range of 4-5% of the fleet.

4.3.5.1. Planned Maintenance on Company B

It is the consensus of the company's managers, and the decision by the CMO, that the losses in revenue and long term business due to equipment breakdowns while in revenue service are such that the company will benefit from a program of planned maintenance, even if the costs associated with such a program are as high or higher than the out-of-pocket costs currently incurred. This conclusion came on the heels of the company's recent experiment in planned maintenance undertaken at a new maintenance facility built to service unit coal trains.

4.3.5.1.a. The Coal Car Experiment

In 1986, Company B built a new maintenance facility for coal cars used in unit trains in its Western region. Some of these cars are owned by private companies and others by Company B. All the cars were in captive service over two fixed routes from mines to export ports or electric utilities. The railroad was operating 20 train sets of 114 cars each, and had to maintain enough reserve equipment on hand to fill out the trains to replace bad ordered equipment. The cars typically ran 1400 miles per cycle, and were operated 100,000 miles per year. The cars were experiencing a bad order rate of 5

cars/114 car train/ 1400 mile trip, for a wide variety of causes. The new facility was located at a point through which virtually all of the cars pass on every cycle, both loaded and unloaded.

When the new coal car maintenance facility was put on line, the Industrial Engineering Group was asked to develop a new approach to planned maintenance in order to reduce the number of bad-ordered cars in the coal fleet. The manager of the Industrial Engineering group indicated that he decided to model the system on that being followed by a contract shop handling maintenance for a captive fleet of coal cars owned by an electric utility company. A series of 4 types of inspections were devised, labeled A, B, C, and D. An A inspection is basically an examination of the car's running gear after stopping the train during an empty return trip from the port or generating plant. A inspections are to be completed in less than 2 hours, including setting off any cars which are found to have failed components. Type B and C inspections involve increasing degrees of dismantling and inspection of components, including trucks, springs, side frames, and bolsters. During Type B and C inspections, various components were to be measured and records kept of the wear rate. Type D inspection involves a complete disassembly of the car's undercarriage, and replacement of any parts which exhibit a high degree of wear.

Originally, cars received the following series of inspections[8]:

Miles	Inspection Type
41,500	B
83,000	C
124,500	B
166,000	C
207,500	B
250,000	D

The planned maintenance program was implemented in 1987. It became apparent over the next several years, however, that the type B inspections were not finding many defects, so they were discontinued. Recently the owners of the private cars in the service

---

[8] The cars also receive type A inspections at the loading and unloading points.

have concluded that even the policy of inspecting the cars every 83,000 miles is not economically efficient, and so are now performing a type C inspection every 125,000 miles and a type D inspection and overhaul every 375,000 miles. The regional managers wanted to follow suit, but it was decided by headquarters personnel that there was not enough history to support such a large increase. The railroad's current policy is to examine each train with a type A inspection when it passes through the facility empty, and to perform a type C inspections at 100,000 miles and 200,000 miles, with a type D overhaul at 300,000 miles. If a car needs to be shopped and is "near" to the scheduled mileage for a PM inspection the car is sent in for planned maintenance rather than to the fast track. The criteria currently used for "near" is up to the supervisors at the repair facility, who attempt to balance the current work load of the shop and how long the interval is until a car is due for planned maintenance.

During the planned maintenance inspection, the basic rule used is that if a component can be expected to last until the next scheduled inspection time without exceeding either the failure limits or the remanufacturing standards, it is left in service. If the component is judged likely to fail or wear beyond the company standard, it is removed. Currently a set of "hard standards" are used (i.e., components worn, for example, to within 3/16" of the condemnation standard are removed.) None of the managers interviewed claimed responsibility for developing the particular standards in use, but that may reflect a significant realignment of personnel recently rather than a lack of support for the standards in use. An ongoing program involves measuring the condition of components to try and develop better wear standards.

The planned maintenance program, then,is essentially a "hard time" policy, in which cars are brought in at fixed intervals and components repaired or replaced as warranted. Noting that the times are not based on any particular analysis of the car as a system or its components' individual reliability characteristics, it is not surprising that the times chosen are subject to constant re-evaluation regarding whether they are appropriate or effective.

Although the effectiveness of the planned maintenance experiment is not entirely agreed upon, there are a number of reasons for judging it a success. The bad order rate

found is currently 2 cars/114 car train/1400 mile trip, which is a significant improvement over the previous rate of 5 cars/train/trip overall, and the more than 14 cars/train/trip reported for some of the coal equipment. The increased reliability has also allowed the railroad to reduce the number of spare cars held out as reserves for bad ordered equipment. According to the Industrial Engineering group, this improved reliability, coupled with operating improvements has allowed the railroad to operate 2 additional train sets without incurring the capital costs associated with actually acquiring them. The equivalent annual capital cost of train set is estimated by Company B to be approximately $800,000. The additional contribution associated with the train cycles such a set can perform is estimated to be in excess of $2,000,000 annually.[9] This increased reliability and associated savings in capital costs have not resulted in significant increases in labor costs, which have remained essentially constant over the life of the experiment, while the number of cars assigned to the PM program has steadily increased. Material costs have increased, however, as components are removed earlier with more residual life remaining. For many of the components, however, the earlier removals have meant that more are reclaimable at the remanufacturing facility, which offsets at least some of the costs of early replacement.

The area of costs is one which seems to provoke heated responses among the company's managers. It is generally agreed that the planned maintenance program has increased the overall out-of-pocket costs associated with maintaining the cars in the experiment, but that alone is not an appropriate measure of maintenance effectiveness. At that point the agreement ends. Some officials maintain that cost would have risen as fast for this fleet with or without the planned maintenance program, and that the increased costs reflec: the aging of the cars and the expansion of the facility. Other managers claim that costs have gone out of control, and that shop level personnel are "gold-plating" the equipment. Proponents of the planned maintenance approach point out that there has been

---

[9] Supporters of the company's planned maintenance program are quick to point out that not all these savings are the result of the car maintenance activity, but reflect a systematic approach to managing the overall transportation service, including better train operations and handling, better track maintenance, and improved train scheduling.

a decline in the number of main shop personnel, and claim that at least some of those reductions follow from the increased quality of field level maintenance. Critics suggest that the reductions at the main shops would have occurred in any event and reflect the tightening economic climate of recent years. In any event, cost savings are not the primary justification of the planned maintenance program.

Objections to the program, then, can generally be summarized as taking 2 forms:

(1) The basis for deciding on the time frames and the components to replace are essentially of an *ad hoc* nature, so that while reliability has gone up, it has done so by overmaintaining the fleet. The desire by the private owners and the regional managers to extend the time frames for inspections is taken by some persons as an indication that the cars are being "overmaintained". There is clearly some merit to this concern, particularly regarding the lack of an engineering or statistical basis for the current standards, and this is admitted by the program's defender's, who assert that when the wear measures are developed, the time frames can be adjusted to optimize the level of maintenance.

(2) Most of the real improvements in reliability may be the result of the quality of the Type A, i.e., run-through inspections, rather than the planned programs. Critics note that cars which are repaired during the type A inspections are not included in the bad order count, since the car was not removed from service; failure to count these repairs gives undue credit to the planned maintenance program. Similarly, when a car is bad ordered, if the local manager also performs a planned maintenance inspection or overhaul on the car, it is not counted as bad ordered. It seems undeniable that this explains at least part of the improvement in reliability, but no data was available to measure the degree to which the improved performance is a reporting artifact.

The experiment's critics also question the degree to which a highly structured program such as this can be extended to fleets which are more free-running, that is, which do not follow repetitive service patterns. They point out that the coal cars follow a well defined loop, and can be easily monitored. For cars which are used in interline service

100

or over more complex service patterns than the coal cars, it is questionable whether there is a single inspection and overhaul interval which will achieve the improvements in mechanical reliability claimed for the coal fleet. In the absence of the regular type A inspections, it seems likely that a free running fleet may be much more prone to components failing at seemingly random times than is currently found in the experimental fleet.

4.3.5.1.b. Extending Planned Maintenance

Notwithstanding the concerns raised in the preceding section, the general attitude of Company B's managers, and particularly the CMO, is that the planned maintenance approach developed in the coal car experiment has been a success and should be extended to other cars in the company's fleet. It is planned to implement the system wide program in stages, concentrating first on groups of cars which are operationally similar to the coal cars, that is, high mileage cars in more or less captive service. These cars will also have regular measurements of components, and it is hoped to develop wear curves and standards to derive appropriate maintenance intervals. A point of concern to some of the managers involved is that the current measurements are all logged on paper, so analysis will either be quite difficult or will require that the data be converted to machine readable format. This is actually a natural outcome from the learning process the company had to go through in developing the wear measures. Initially more than 30 measurements were taken on a coal car. It has since been established that only 7 points are needed to chart the wear pattern for the components Company B is interested in.

An integral part of the strategy to implement planned maintenance throughout the fleet is the development of a supporting information system, the Planned Car Maintenance System (PCMS). PCMS will track each car in the fleet in terms of when it is due for a planned maintenance inspection or overhaul and will order it sent to the nearest shop when due and empty. Since each car type and series can have its own schedule and preferred shops for various types of maintenance, the system may be quite complex. It will also be linked to the company's car movement system, which monitors train movements, so that when a train contains cars which should be routed to a planned maintenance facility the blocking assignments can be updated without disrupting

101

transportation operations. (Actual routing of cars to planned maintenance facilities will be performed by the Transportation Department.) PCMS is still in the design stages at this time, and there seems to be a friendly but intense debate over what should and should not be included in the system.

To date the company is performing some version of planned maintenance on approximately 10% of the car fleet. It hopes to include all cars in the fleet which travel more than 20,000 miles in the next 2 years, although some of the managers say that may be too ambitious.

4.3.6. Company B: Conclusions

The most important thing to note about Company B's car maintenance policies and practices has to be the high level of enthusiasm exhibited for developing a planned approach to maintenance. This positive attitude is clearly fueled by the attitude of the CMO, who has advocated a more systematic approach to maintenance for more than 10 years. Indeed, it is noteworthy that even those people who expressed reservations about the benefits of the coal car experiment objected on technical rather than philosophical grounds. One of the directions in which the company clearly wants to move in the next several years is to integrate the various information systems used to support car maintenance with other relevant data bases and systems in the company. Interest was expressed by a number of the company's managers in developing a more uniform and easily accessible mechanism for user queries concerning a car's usage, engineering characteristics, and maintenance history.

4.4. Company C: A Chemical Company With A Private Car Fleet

Company C is one of the world's largest industrial manufacturers and suppliers of chemicals. petroleum, and related products, including a number of what are usually characterized as hazardous materials. Annual revenues from the company's various product lines are more than $10 billion. To support the movement of both raw materials and finished goods, Company C has a large staff which oversees the transportation and logistical functions. Because of the specialized nature of the commodities shipped by Company C, a fleet of more than 5000 railroad cars is owned or leased on a continuing

basis. It is not surprising that given the large volume of sales and the hazardous nature of the materials shipped, maintenance cost control is not the primary focus of the transportation managers.

Company C is organized into approximately 10 Operating Departments, along the lines of the general products produced (petroleum products, polymers, etc.) These departments are supported by a headquarters staff, which is responsible for what could be called traditional staff functions, including accounting and auditing functions, logistics and transportation, and health and safety policies. The Logistics Department is divided into several sections along modal and functional lines, one of which is the Rail Transportation Section (Figure 4.3). The Rail Transportation Section is further divided into four groups,

- Fleet Management, which oversees the assignment of cars to particular Operating Departments, and monitors the day-to-day utilization of the fleet;
- Modifications, an engineering group responsible for changes to cars to meet specialized needs of customers or the Operating Departments;
- Acquisitions, which has responsibility for deciding on the appropriate size of the fleet, and for arranging the purchase. lease, or sale of equipment to meet the needs of the Operating Departments;
- Maintenance, with responsibility for both day-to-day repairs to and maintenance of cars and for mandated inspections and maintenance activities (such as relining tanks and painting).

4.4.1. Company C: Fleet and Shipment Characteristics

As previously stated, Company C has a fleet of over 5000 cars. In the past, maintenance of leased cars was the responsibility of the car owners. but some recent leases have provided that Company C assumes this obligation. This has been done at the company's insistence, to insure that all the cars used are maintained to Company C's high standards. The fleet is composed almost exclusively of tank cars and covered hoppers, with tanks cars comprising some two-thirds of the cars. The fleet is divided into approximately 150 pools or groups and each pool is assigned to an Operating Department for shipment of particular commodities. The average age of the fleet is 18 years, which is approximately the same as that in the overall U.S. fleet.

Figure 4.3
Company C

Chairman

HQ

Operating Department

Operating Department

Operating Department

Operating Department

Operating Department

Operating Department

Legal

Logistics

Acct'g.

Other Staff

Ocean Transportation

Rail Transportation

Other Logistics

Fleet Management

Modifications

Maintenance

Acquisitions

Quality Assurance

104
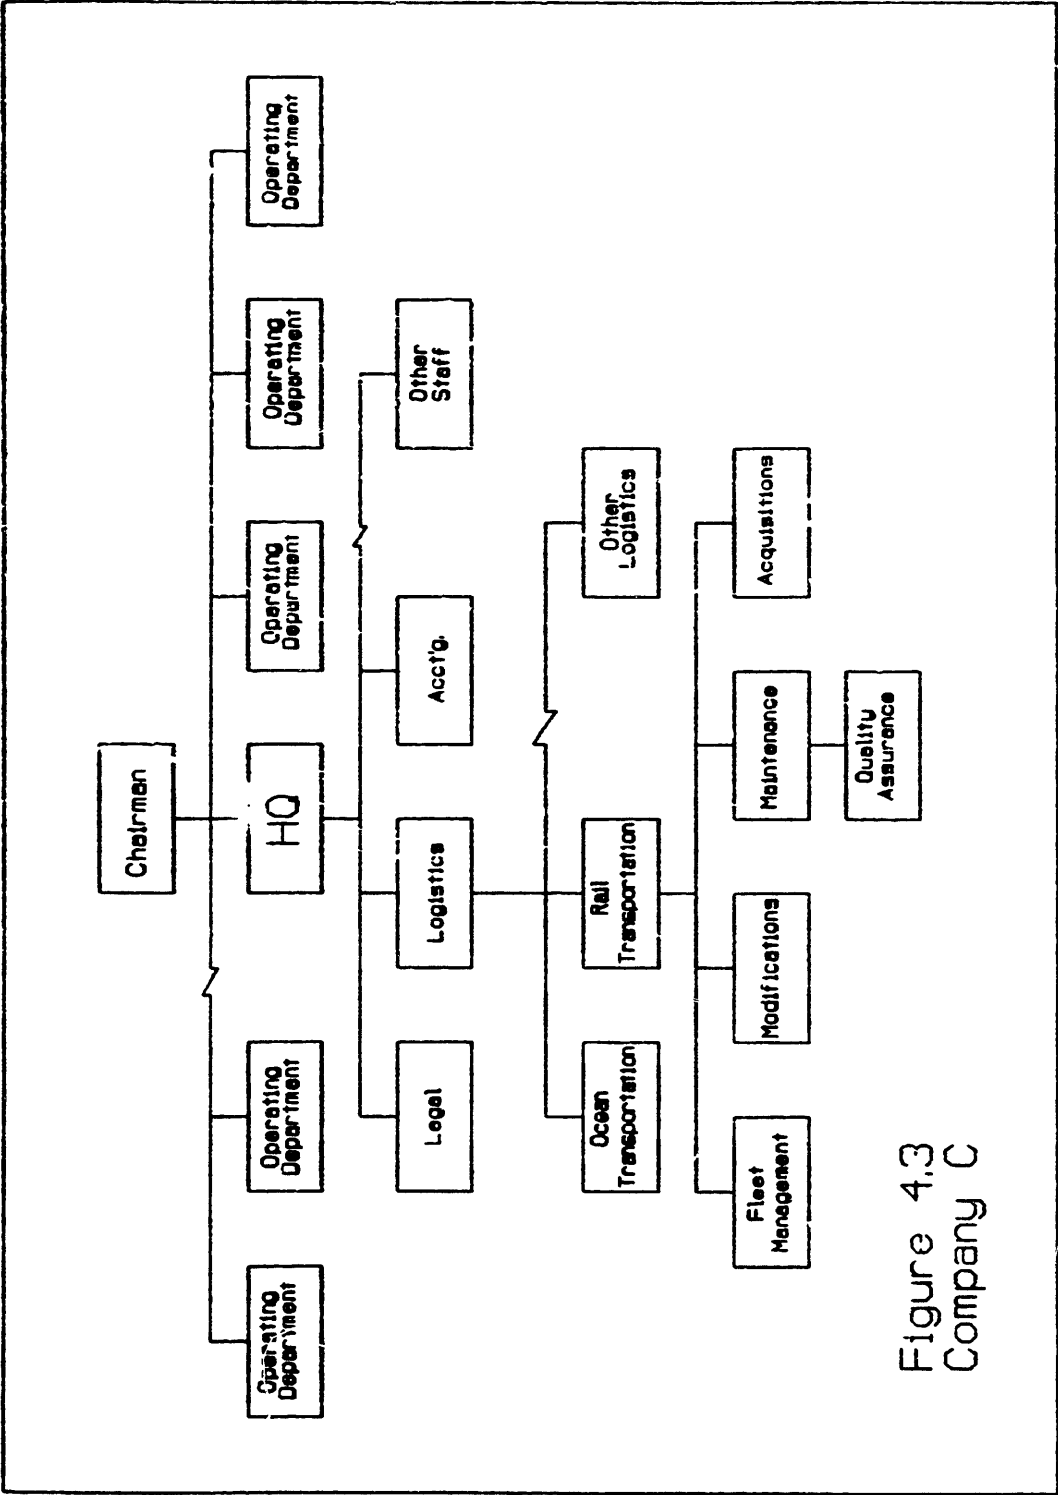
The typical length of haul is 600 to 700 miles loaded, and virtually all backhauls are empty. A typical shipment is carried on 2 or 3 railroads, with trip cycle averaging 45-60 days. This rather long cycle reflects the company's use of the equipment as holding tanks, and as marketing tools for end customers. One of the company's managers indicated that in many cases, the value of the commodity is so high that the opportunity cost of allowing the customer to hold the railroad car is insignificant in comparison. The typical car generates 10-12,000 car miles per year, although the company operates several higher mileage groups which are used as much as 30,000 miles per year. About 60% of the shipments made are considered "hazardous materials" for transportation purposes, i.e., cars with Standard Transportation Commodity Codes beginning with "49".

4.4.2. Car Maintenance by Company C

Company C spends about $15 million annually on the maintenance, repair, and modification of its fleet. Of this, $1-1.5 million is paid to the railroads for repairs to cars while in service, $10 million is spent on preventive maintenance, scheduled tests, and unscheduled maintenance at contract shops, and $2-3 million is spent on modifications to equipment to improve the safety or reliability of the fleet. This works out to some $1700 per car per year. which is higher than the overall industry amount of approximately $1155 per car cited in the previous chapter. On a per mile basis, however, it becomes clear just how high the expenditures really are. The overall figure for the industry is on the order of $.07 per mile, while for Company C, the amount is $.14 per mile. This higher expense reflects both the specialized nature of the equipment, which requires more inspections and maintenance and the particular policies followed by the company, which are discussed below. Managers indicated that AAR repair costs per mile (i.e., costs for repairs for in-service failures) have been decreasing in the past few years, and are 25% lower than five years ago, in spite of the increases in AAR rates for labor and materials in that period.

A concern which was expressed by several of the company's managers was that the cars, which are maintained to a higher than mandated standard still must travel in trains alongside cars which may be operating exactly at (or even slightly beyond) the condemnation limits. This concern suggests that the company would welcome a

105

tightening of the AAR condemnation standards, although no one expressly stated as much.

4.4.3. Maintenance Facilities

Company C owns no car repair shops. To perform maintenance, it contracts with several private car repair shops, which perform 80-90% of the company's discretionary maintenance work. Other repairs are performed at "mini-shops" run by contractors at some of the busiest plants, and by the railroads when cars are found to be defective enroute. One of the goals of the company's maintenance plan is to reduce the number of repairs performed by railroads on the fleet. This is not so much a reflection on the quality or cost of railroad repair as on the effect of removing a car from service on the customer. Company C believes that shopping a car incurs a substantial cost in lost goodwill and foregone business revenue in addition to the actual cost of the repairs. Concentration on preventive maintenance has had the result that less than 1% of all Company C cars sent to repair shops are loaded.[10] (Cars are virtually never sent to contract shops while loaded, i.e., for discretionary maintenance.)

The small number of shops currently used for planned maintenance is a sharp decrease from just six years ago when Company C used more than 30 facilities. At that time it was decided that the company could improve the quality of maintenance and better supervise the activities of the shops by reducing the number. A study was performed, which analyzed the number of cars used at each of the companies industrial locations and defined a set of key transportation corridors over which most of the rail movements occur. Shops located along these corridors were then examined in terms of their capabilities and qualifications, and they were invited to bid on various types of repairs. Shops were selected on the basis of quality of repairs and ability to perform a broad range of activities (e.g., x-rays of welds, tank cleaning, etc.). Over the past few years this number has been reduced further, as the company has moved work to the best of these shops.

---

[10] While no statistics on this are available, this is a remarkable figure in the author's experience. A more typical figure would be 20-30% of all the cars on a repair track at a given time would be loaded, excluding very heavy repairs, which are essentially in storage. This is about what was found at the various industry facilities the author has visited.

The manager of one of the contract shops was interviewed as part of the case study, and he indicated that it was made clear to him from the outset that price alone would not be the determining factor until quality control over repair work had been clearly established. He also indicated that he knows from contacts in the industry that they are not the "low bidders" among shops of comparable size.

In addition to the contract shops, mini-shops are operated under contract at several of the company's largest plants. Company C selected 6 plants which were making a large number of rail shipments (more than 400 per year), and negotiated contracts with a national car maintenance company for mini-shops. These shops perform more extensive inbound and outbound car inspections, and based on what they find, can perform a wide array of maintenance actions, including schedule preventive maintenance programs, valve tests, and standard running repairs (safety appliances, running gear, etc.) The mini-shops do not perform major work such as painting cars, lining tanks, or replacing all of a tank's seals. The mini-shops currently perform work on 40% of the cars in the company's fleet.

The performance of the various shops and mini-shops are monitored by an aggressive quality assurance program. The company has a group of full-time inspectors who monitor the repairs performed by the contract shops and mini-shops. The quality assurance program includes examining cars after the inbound inspection to the shop or plant for undetected problems or unneeded repairs, auditing the actual repairs while the shops work on cars to insure compliance with both industry standards and the company's own guidelines, and inspecting outbound cars to ascertain if all needed repairs were correctly made.

Prior to renewal of a shop's contract, Company C's Railroad Maintenance Group meets with the quality assurance inspectors responsible for that shop and determines whether the contractor's work is of high enough quality to merit renewal. The manager of one of the contract shops indicated during an interview that the quality assurance inspectors are seen so frequently that after a period of time they come to be viewed by the workers and staff as "regular fixtures rather than visitors from headquarters".

4.4.4. Information Systems to Support Car Maintenance

Company C has been in the process of upgrading its information systems over the

107

past several years, spending more than $2.25 million on systems to support the Rail Transportation Section. The current system is comprised of 3 parts:

- Car tracing, used to track shipments across the various railroads;

- Equipment Management, used to support the car maintenance function; and,

- Car Accounting, used to keep track of mileage, demurrage, and other car related expenses incurred.

The three parts were developed by an outside contractor in response to specifications developed by the users and the company's Information Services Department.

The equipment subsystem consists of a number of separate files which are unified to present the user with a group of standard screens to use when managing the various car repair processes. The system also supports inquiries to the data base using conventional computer languages and data access routines. The equipment systems includes:

(1) The physical description and characteristics of the car, including UMLER records, the car type;
(2) The fleet or group to which a car is assigned;
(3) Any modifications which have been performed (or are pending);
(4) The due dates for any tests or mandated inspections or scheduled preventive maintenance;
(5) Summary data regarding the costs of repairs in the preceding year and thus far in the current year.

The equipment system is an online system which allows the managers to make limited queries against the system and generate reports for their own use. The system is used to assign cars to particular repair shops when an inspection comes due or the car needs maintenance which can be deferred until the car is empty. The assignment of cars to particular shops is currently done by a manager who inputs the routing data into the computer as he makes the decision. It is hoped to automate this process in the future, at least for the simplest cases, such as where the nearest shop can perform the maintenance action at reasonable cost, or where there is only one possible shop which can perform the work. (Such cases are believed to constitute more than 75% of the decisions). The online system also includes access to a local area network, so that managers can download the

results of searches of mainframe data bases and then work on that data using personal computer programs such as LOTUS 1-2-3.

In addition to the online system, the car maintenance group has a contract with an outside vendor to manage the repair records received from the railroads and the shops. This data, which is input in a manner more or less compliant with the AAR's Car Repair Billing System, is processed by the vendor, who performs limited audit checks, generates reports, and sends summary information to the company's mainframe computer for use in the equipment subsystem. The decision to use an outside contractor for this was based on the desire to avoid building and maintaining a set of standard computer billing programs when there were good ones available, and the desire to free up computer resources for the equipment system.

In retrospect, the managers seem reasonably happy with the contractor's performance, but expressed some misgivings about the timeliness and accuracy of information. Car Repair Billing data sent to the company seems to lag about four to six months behind the actual repair date. Some of this is delay incurred waiting for the repair records to be keypunched by the vendor. The company is hoping to reduce the delays and improve the accuracy by arranging for repair bills at the shops to be transmitted electronically to the company's headquarters using standard EDI protocols. The bills will be audited and the equipment subsystems updated, and then will be transferred to the vendor for storage and analysis.

The two information systems are used to prepare a wide variety of reports used by managers at the Maintenance Group. Reports generated include:

(1) Maintenance Schedules: showing which cars are due for either preventive maintenance, or for mandated tests and inspections. by location, which is a proxy for which Operating Department the car is assigned to. This report is generated monthly and is passed to the Operating Departments, who are responsible for insuring that the cars are released for scheduling to a particular shop.

(2) Shop Workload Projections: each month the shops are given an estimate of the number of cars they will be receiving in the next few months so they can

plan their workloads and inventories.

(3) Maintenance Costs at Shops by General Car Categories: All the maintenance activities are assigned to general groups, such as painting, cleaning, lining, repairs, wrecked cars, and miscellaneous. The expenses incurred at each shop by category are used to monitor costs, workloads, and focus at the shops. If the miscellaneous category is too large, it usually indicates that there is a problem in either the definition or the application of the job codes.

(4) Costs Incurred by Budget Category: This monthly report includes costs for car leases, facility rents, demurrage, etc.

(5) Fleet or Car Group Summaries: these reports summarize the status of each of the 150 pools in terms of the availability, bad order rates, the costs by categories, and a summary of work done thus far and still to be done to complete applicable modifications. This report is used by the manager responsible for each fleet.

(6) Summary of Repairs: This report is generated by the vendor who manages the repair records and shows the historical repair costs by railroad, car, and car group. Because of the time lags, it is used mostly for trend analysis to see if either a particular railroad seems to have higher than normal repair activity or if avoidable repair expenses are being incurred, which would suggest that the preventive maintenance program is overlooking something.

In addition to these reports, the manager of the Maintenance Group generates a set of reports using programs he developed himself either on the mainframe or using downloaded data to the local area network. These reports include:

(7) Labor and Materials by shop, which are given to the quality assurance inspectors to use in deciding where to focus their attention;

(8) Cars Due for Maintenance and Cars Released Year to Date: this report is used to encourage the various Operating Departments to release the cars in a more timely way. It also has the effect of showing where responsibility lies for cars held in service beyond their date due for inspection.

It was indicated by several of the managers that in the future it is hoped to automate the

110

routing of cars to the shops when the car is empty and approaching a due date for inspection.

4.4.5. Maintenance Policies

As has been stated, the primary focus of car maintenance at this company is on safety and reliability rather than on controlling maintenance expenditures. This is not surprising when one considers that the costs of a single catastrophic incident involving hazardous materials could easily involve higher costs than the entire annual maintenance budget for the cars. Further, the use of the cars to enhance the marketing of certain high value commodities makes reliability an integral part of the marketing process, rather than simply the obverse of failure. To achieve the desired high level of reliability, Company C uses three approaches:

(1) Company C has an extensive program of scheduled preventive maintenance (PM), including extensive inspections (disassembly of trucks, for example).

(2) The company undertakes periodic modifications to replace components or systems which are believed to undermine the reliability of cars. An example would be the program to replace all friction bearings several years earlier than the date mandated by the AAR.

(3) Company C publishes its own manual for repairs which is considerably more stringent than the AAR's **Field Manual**. Part of the motivation for this is the desire to reduce the number of condemnations of parts while cars are under load, which can disrupt customer operations.

Each of these is discussed below.

4.4.5.1. Preventive Maintenance at Company C

Company C's Preventive Maintenance (PM) program is based on both time and miles. Depending on the commodity carried, cars are brought into a shop at fixed intervals of time or mileage (whichever occurs first), and the car is extensively reworked. Any welded parts associated with carrying hazardous material are inspected and, if deemed appropriate, X-rayed. Quality assurance and fleet managers have the company's authority to request X-rays at any time.

The PM standards provide that every car will be brought into the shop at least

every 6 years or 100,000 miles, whichever comes first. Cars used for carrying hazardous materials are brought in more frequently, either every 2 years/40,000 miles or 3 years/60,000 miles, depending on the nature of the commodity carried. When brought into the shop, the intent is to insure that all the components leave "as good as new". As an example, all cars brought in for a PM have the trucks built up to conform to the full specification for new trucks while in the shop. This is quite remarkable, since the typical life of a truck is on the order of 500,000 miles, which means that the trucks are only using about 20% of their wear life before being renewed.

This has the effect of increasing the cost per mile, but reducing the likelihood of a catastrophic failure. Setting time and mileage standards is made more difficult for Company C by concerns with liability. The time and mileage limits used are the result of a staff study of the time to failure of various components, which found that failures were quite rare up to the 40,000 mile limit. It was thus decided to put all cars on that basis initially, and then begin to push the limits out by using cars assigned to carry non-hazardous materials as test cases.

Although the PM program is a corporate initiative, and is not mandated by any regulatory body, cars are almost never allowed to exceed the limits. The reverse, however, is not true. If a car is in one of the contract shops for an unscheduled maintenance event or for a mandated inspection (e.g., a tank or valve test), and the car is due for PM within 1 year, the PM date is moved ahead and the PM performed. The choice of 1 year is acknowledged to be arbitrary. The same limit, one year, is used for deciding whether or not to perform a mandated inspection early. In other words, if a car came into one of the contract shops for an unscheduled repair, and the valve test was due in 11 months, and PM in 18 months, the valve test would be performed now, but the PM would not. These guidelines are not extremely rigid, however, and if a car came in and it was due for PM in 1 year and 1 day, the car would almost certainly be given the PM at this time, unless the shop was extremely backlogged.

4.4.5.2. Car Modification Programs

These are programs in which either the industry or the company's own engineering staff has determined that a component should be removed and replaced with a presumably

better design. In such cases, cars which are subject to the modification are "tagged" in the computer, and when the car is sent into a shop, the modification is performed. Because this work can often only be performed at a few shops, the effect is often to make the decision of where to send a car which is due for either maintenance either quite easy ("Send it to the place that can do both") or quite difficult ("Should it be sent to the place that can do the modification or the place that can do the mandated inspection?"). In any event, the company uses the system both to monitor and insure compliance with changed regulations and to improve the reliability of the fleet.

4.4.5.3. Company C's Field Manual

Because such a high percentage of Company C's fleet is tank equipment engaged in the transport of hazardous materials, the company has undertaken a number of engineering studies of its own to determine the adequacy of the interchange standards for its cars. One of the results of those studies has been the development of a special set of guidelines for use by car repair personnel in the contract shops. In some cases the guidelines used reflect proprietary car designs which would simply not be covered in the AAR Field Manual, or mandate the use of paints or liners appropriate to the particular commodities to be transported in the car. In other cases, however, the standards reflect the company's willingness to replace components earlier than the conditions mandated in the AAR's Field Manual. This is generally more of an indication of the company's concern with the costs of removing a car from service than a belief that the AAR's standards are unsafe. The net effect, however, is that components are often removed from Company C's cars while considerable service life remains, with the company's knowledge and consent. The basic rule seems to be that if the maintenance shop foreman believes that a part will not last until the next preventive maintenance cycle then the part is replaced.

4.4.6. Performance Measures

The measures used to monitor performance by Company C are quite simple and straightforward. Two measures are used to reflect control over reliability and control over costs. The first is *railroad repairs (in dollars) per loaded car mile*. This is used as a reliability measure, since it is the basic position of the company that at best a car will

never fail while on a railroad, but will be preventively maintained at the shops and mini-shops. The second measure is the *overall cost per loaded mile*. This is a straightforward cost control measure.

Overall, both measures have decreased in recent years, in spite of increases in the AAR rates. Railroad cost per loaded mile has declined from $.06 to $.035 over the past 6 years. Overall, the cost per loaded mile has fallen from $.25 to $.20. The first reduction is attributed by the company's managers to the effectiveness of the preventive maintenance program. The savings in the second measure is believed to be a result of economies generated by rationalizing the number of shops, savings at the mini-shops, increased time between preventive maintenance and, of course, the savings at the railroad shops.

The performance measures are generally applied only at the level of the overall fleet although there are analyses by product and shipping locations. There are no explicit performance measures applied to particular pools of cars, although it is likely that fleet managers are performing these calculations manually when they receive their cost reports. Performance is measured at particular locations, and when cars assigned to a shop or plant are found to be experiencing excessive railroad repair costs per loaded mile, a detailed breakout of the costs is generated. It is usually found that the higher railroad costs per mile are the result of "excessive" railroad repairs of a particular type, such as COT&S, or wheels. The typical management response to this is to notify the shop to be aware of the need to monitor the component or system, and to notify the quality assurance inspector responsible for the shop to perform outbound inspections at that facility.

It was indicated that there is interest in developing alternative reliability measures based on engineering or statistical criteria in the future, but that at this time there is no particular effort in this direction.

4.4.7. Company C: Conclusions

Like Company B, Company C views the maintenance of rail freight cars as an extension of the activities of other corporate functions more than as an end in itself. In the case of the chemical producer, freight cars have a natural logistics function within the company, and a marketing function directed at the company's consumers. A third

114

function of the maintenance activity is to protect the company from litigation in the event of an accident involving the company's cars.

The result of this view has been that Company C is more willing to spend money on car maintenance than most railroads, and this willingness to invest in reliability has been marked by a period of declining costs per loaded mile. This suggests that the long term costs of car repair and maintenance can, at least in some cases, be appropriately reduced by increasing short term expenditures, although expenses are still far higher than the industry average. To achieve reductions in costs, the money was spent at a carefully selected group of shops which were willing to implement aggressive quality control programs and on information systems to help track the condition of the cars and the activities of the shops.

4.5. Observations and Conclusions

In this section, the three cases are examined together to see the areas which the three companies have in common and where they diverge. To facilitate the discussion, some of the characteristics of the three companies are summarized in Table 4.3.

4.5.1. Common Characteristics

Perhaps the most striking common theme found among all three cases is the desire to avoid in-service failures. Company A is motivated in this concern by the desire to avoid the higher costs associated with AAR billing rates for offline repairs; Company B's concern with in-service failures is primarily based on its desire to be a more market driven firm; Company C is concerned with both marketing concerns and with safety, because of its high volume of hazardous material shipments. Notwithstanding their differing motives, each is looking for ways to increase the level of planned or preventive maintenance to avoid repairs while cars are under load. An extension of this is that each of the companies is willing to invest at least a limited amount of money at the outset to develop a planned maintenance program. Company B has undertaken a large scale experiment with its coal car fleet and, even in the face of mixed economic results, has committed itself to extending this to more of its fleet. Company C already has an extensive program of preventive maintenance, and undertook it without any particular

115

	Company A	Company B	Company C
Fleet Size	3000 Cars	50,000 Cars	5000+ Cars
Annual Revenues	$30 Million	$1 Billion	$10+ Billion
Car Maintenance Expenses	$2.8 Million	$90 Million	$15 Million
Cost per Car Mile	$0.06	$0.08	$0.14
% of Miles Of. Line	90%	20%	100%
Repair Policy	On Condition	On Condition Hard Time (for some cars)	Hard Time
Primary Reliability Measure	none	Cars Removed/ Train % Bad Ordered	Railroad Cost per Loaded Mile
Primary Cost Measure	Total Cost	Cost v. Budget	Cost per Loaded Mile

Table 4.3
Comparison of Case Study Companies

expectation of reduced out-of-pocket costs (although this appears to have occurred). Company A has hired outside consultants to help develop preventive maintenance plans, and has indicated that it is willing to share any economic benefits it obtains from leased cars with the car owners. The obvious conclusions here are that those responsible for railroad car maintenance perceive in-service breakdowns as expensive and want to undertake planned or preventive maintenance as a means of avoiding them.

A second characteristic that the companies all seem to share is that while they have large amounts of computerized data, their reporting of the maintenance function remains quite limited. In no case is the data organized or stored in such a way that the managers who are most familiar with the problems of car repair and maintenance can easily or even readily access all the information to build computer programs or reports that truly support the way they think about car repair. Company A's data is not available as an on-line resource. Company B keeps revenue and mileage information in a separate file and the manager responsible for their car information system acknowledged that only programmers and users with extensive training can easily access the data. Company C not only lacks timely car repair information, it stores the repair record in a completely different computer environment from the management and usage information. One result is that in each case. the most expert managers of the car maintenance function were unable to use the computerized databases except by requesting data from another department that controls the databases.

Another theme common to two of the three companies seems to be the difficulty in finding the "right time" to perform preventive maintenance. Companies B and C have both implemented planned maintenance programs based on bringing cars into the shop at fixed intervals, but there seems to be no underlying engineering or statistical basis for the times or mileage used, as is evident from the desire of both companies to experiment with finding a "better" time. In both cases this is not for want of studying the problem. Company B has an ongoing quite expensive program to measure and monitor wear in various components of cars subject to preventive maintenance, and Company C undertook a study of their fleet prior to extending times for cars in non-hazardous material service. In the next chapter, we will see that it is unlikely that a single number which is "best" for

117

a car exists, particularly if the car is free-running in the general car fleet. What appears to be "best" changes frequently. In the case of Company C this is not much of a problem, since the use of very frequent intervals has the effect of making the cars highly reliable, which is the primary objective. But in the case of Company B, the large railroad, there may be negative economic consequences if the attempt is made to apply this approach to cars such as boxcars which are free running and often carrying relatively low value shipments.

In all three cases, the performance measures used to evaluate the quality of the maintenance program are not complete. None of the companies uses any of the standard reliability measures such as mean time between failures or time to first failure after maintenance events. The cost measures used by Company A are dependent on the AAR rates, and give no indication of the overall reliability of the fleet or of the various facilities. The current measure for Company B, cars removed per train is a function of the size of trains and quality of inspections, and is probably subject to substantial random variation. The measure used by Company C, AAR costs per loaded mile, is strictly dependent on the AAR billing rates, and as such cannot be compared across time periods unless adjusted. (A study has been done which reviewed repair costs over several years at 1982 rates.) Further, if a car or series of cars is subject to inexpensive repairs which require removal from service, the problem will likely go undetected for a very long time.

As indicated in Chapter 2, a better set of measures for cars might be miles per in- service failure and miles per maintenance event. Miles per in-service failure is independent of both the nature of the event which caused a car to be removed from service to a shop and the cost of the repairs. The second measure indicates the overall reliability of the car, and recognizes that a car which is in the shop for planned maintenance too frequently is just as unproductive as one in the shop for unscheduled maintenance. Along with cost per car mile, these measures will be used when comparing various alternative maintenance policies in Chapter 6.

4.5.2. Differences Among the Companies

A number of the differences among these three companies stem simply from the differences in size, economic circumstances, and operating patterns. Managing freight

cars is a tiny part of Company C's overall business activity, no matter what sort of measure one wishes to use. Out-of-pocket costs associated with freight car maintenance represents, on the other hand, 10% of the total corporate revenues of Company A. It is not surprising that there is a difference in focus with respect to cost control.

One of the most notable differences among the companies studied is a natural consequence of their shipment and cycle characteristics, namely, the perceived and realized benefits of planned maintenance. Company A's cars spend 72% of their time and accrue 90% of their mileage outside the control of Company A. Company B's cars accrue most of their miles (and time) on line. Company A rightly focuses on the out-of-pocket cost consequences of maintenance programs, since the reliability of its cars may well be overwhelmed by the reliability of other railroads' cars. Consider that Company A's cars are likely to be a small part of a train consist on another carrier who may follow any number of maintenance policies. In other words, many of the non-cost related benefits of a highly reliable car will be experienced (or undermined) by another company. In the case of Company B, on the other hand, many of its cars are used in unit trains over routes fully controlled by the carrier. Thus all the benefits of operating reliable equipment accrue to the company.

4.5.3. Conclusions

In this chapter, the car maintenance practices and policies of three companies have been studied. The companies range from a small regional railroad whose cars are frequently under the control of other companies to one of the largest railroads in North America. The third case, a huge industrial company, is representative of the private car companies who depend on the railroads for transportation, but seek to be self-sufficient in the area of equipment maintenance. While each of the companies is unique, they share certain common interests and needs. One of those needs is a theoretical framework to develop, implement, and monitor maintenance plans and programs.

In the next chapter, we will critically evaluate the current maintenance practices and policies of freight car owners, and begin to develop the theoretical framework which will address some of the shortcomings we have already found.

Chapter 5

A Critique of Freight Car Maintenance in the U.S. and Canada

5.1. Introduction

In the preceding chapters, we have looked at the car maintenance practices and policies of various railroads and private car owners, focussing in detail on three representative companies. In this chapter, those maintenance activities are critiqued. Because the thrust of this thesis is on ways to improve freight car maintenance, the chapter emphasizes the negative aspects of the current policies and practices. This is not intended to convey the impression that freight car owners and maintainers do not have good reasons to be pleased with their actions in recent years. The positive aspects include willingness to implement new maintenance programs to serve their customers and meet organizational goals, and invest resources to try to improve fleet reliability and the cost effectiveness of their maintenance programs. Unfortunately, the maintenance policies, monitoring systems, and the information systems to support them suffer from serious shortcomings. The problems are discussed in terms of their nature and sources, and the specific solutions which are provided in later chapters are introduced.

The chapter first briefly highlights the encouraging facets of the current situation. The focus is then shifted to the deficiencies found. That the negatives are examined in so much greater detail than the positives might seem to denigrate the legitimate efforts of maintenance managers to improve their actions in recent years. That is not the intention of this chapter. Rather, the focus on weaknesses is because managers have demonstrated a willingness to improve the fleet and the maintenance systems when legitimate alternatives are presented. The detailed analysis of the problems with the current practices leads directly to the constructive alternatives proposed in the chapters which follow.

5.2. Current Maintenance Practices: Positive Aspects

In recent years, freight car owners have demonstrated a remarkable willingness to

120

invest in facilities, information systems, and organizational approaches to improve freight car maintenance. This willingness has included the hiring of technical staff and consultants to provide analysis for maintenance planning and the consolidation and reorganization of shop facilities to meet the equipment needs of users of particular car types. To validate this claim of generally positive attitudes toward investments in improved car maintenance, we look first at some examples from the case study railroads presented in the previous chapter, and then other examples drawn from the trade press. In each case, the motivation for the improvements seems to be a recognition that the determination of the costs of car unreliability is more than simply the costs of shopping and repairing the car. In particular, the relationship between customer demand and car reliability seems to be at the heart of many of the improvements cited, i.e., maintenance is becoming "market-driven".

Recall that Company A is a small railroad, with limited repair facilities and support staff. The typical trip is spent mostly off-line, both in terms of miles and time. Because the company has reasons to believe on-line repairs were significantly less expensive than those billed at A.A.R. rates, they desire to move as many of the maintenance actions performed on the cars as possible into their own shops. A crucial aspect of this decision is also the concern that off-line failures of components come at the expense of the service level provided to customers. To reduce costs and improve car reliability, Company A hired outside consultants to perform engineering analysis of certain components and to help plan a preemptive replacement program for end of car cushioning devices. What is noteworthy in this is that the company recognized the value of reliability analyses enough to pay to have one performed, and had sufficient confidence in the results of such analysis to preempt the A.A.R. condemnation standards. This case appears to be one of the first where a railroad has, based on its own studies, tried to move beyond the "on condition" policies to maintain equipment based on a set of economic decision rules. The case serves to demonstrate that the size of a company or its fleet is not a determinant of the willingness and ability to invest in better maintenance policies and practices.

Company B also serves to demonstrate an important point about the current

attitude toward freight car maintenance. Like Company A, they also use preemptive standards for replacing certain components, most notably trucks. In the case of Company B, however, the primary concern is with removing trucks while they can be remanufactured at the railroad's plant. The investment in that facility was motivated by a concern with insuring that parts received were of sufficient quality.

Even more noteworthy is that company's foray into planned maintenance for its coal car fleet. Notwithstanding the objections to some of the technical aspects of that program detailed later in this chapter, it is striking that this program was undertaken specifically at the direction of the Chief Mechanical Officer. He indicated that he was aware that the program might not result in reduced costs for his budget, i.e., out-of-pocket maintenance and repair costs, but believed that there would be an overall improvement in both car reliability and customer satisfaction. What makes this significant is the recognition by the most senior maintenance manager of a major railroad that car maintenance is ultimately an interdepartmental concern, and that investments of time and resources to improve the quality of the fleet are appropriate even at the expense of near-term budget issues.

Company C also serves to demonstrate the willingness of contemporary car owners to allocate financial resources to improve the reliability of the fleet. Company C has spent more than $2 million in recent years for information systems to provide computerized data concerning car maintenance and repairs, particularly through use of another company's car maintenance information systems. More significantly, the company has carefully adopted a system of continuous oversight of car repair shops, with a high emphasis on quality control over the repair process. Indeed, one of the current concerns is that some of the components are currently being maintained too frequently, resulting in excess infant mortality failures. Like Company A, they have also demonstrated a readiness to invest in technical services from outside the company to insure that their maintenance programs meet their needs.

These positive attitudes towards improving freight car maintenance are by no means unique to the case study companies. A review of the trade press reveals numerous examples of the same perspective by other railroads and private car owners.

Railroads have adopted preventive maintenance as a standard part of their vocabulary, representing a major change in the past few years. In 1985, the Chief Mechanical Officer (CMO) of one of the largest Class I railroads was quoted as saying

> The concept of preventive maintenance for freight cars is almost unheard of. It seems that most of us operate our locomotive fleet under a program of preventive maintenance, but we are unable to convince management to do this for freight cars [Brownlee (1985)].

Since that time, almost every month a CMO is quoted in one of the trade journals crediting his railroad's preventive maintenance program with the quality of their fleet[1]. Unfortunately, a careful inspection of their programs often reveals that their "preventive maintenance program" is often little more than a long term rebuilding program under a different name. Consider, for example, one of the Class I railroads whose CMO praised his preventive program for processing several thousand cars per year. Based on the size of the railroad's fleet, each car will only return for preventive maintenance every 12 years! That is not much shorter than the typical overhaul cycle. Nonetheless, the fact that CMOs view preventive maintenance programs as signs of a well maintained fleet is a significant and positive change in the industry.

Even where the details of the program is subject to skepticism, however, there are positive signs. The Burlington Northern is now performing limited opportunistic maintenance of wheels during maintenance at Havelock shops, their largest facility, so that if wheels are "close" to condemnation limits they are replaced to prevent subsequent disruptions of service.

Perhaps even more noteworthy is the shift away from the geographically based shop (a major shop every 500-1000 miles, for example) to facilities organized to meet local customer needs. Both Burlington Northern and Conrail have realigned their shops to build up expertise on special car types used by major customers in specific areas, such as auto racks, coal or steel cars ["Smaller Fleet..." (1989), Bauer (1989)].

Not only railroads have been refocussing their maintenance programs to meet changing markets. TrailerTrain, Inc. (TT) is the largest private car owner in the United

---

[1] See, for example, ["Smaller Fleet... (1989)] or Bauer (1989)

States, and has been in the forefront of supporting better approaches to car maintenance[2]. One of the largest markets they serve is the provision of railcars for hauling automobiles from assembly plants to distribution points, representing about 65% of the motor vehicles produced in the U.S. From 1979 to 1988, loss and damage to automobiles in such moves declined from $92 million annually to $28 million, which TrailerTrain attributes in part to improved maintenance of end of car cushioning devices. They have spent large sums of money on car maintenance, and car maintenance information systems (which they now offer commercially to other car owners); much of this investment has been driven by an awareness of the shipper's concerns. A recent article in the trade press stated:

> As the major supplier of cars that move the cars and trucks, TT knows what can be done [to retain motor vehicle busine.s].
> - Both cars and racks have to be better maintained.
> - Cushioning devices have to be checked more frequently, and maintained to higher standards.
> - Truck maintenance has to be improved, so that trucks and suspension systems are in tune with loads.

In other words, they are trying to structure their maintenance programs to meet customer needs rather than simply following the A.A.R. standards at minimum cost.

These improvements in attitude and willingness to invest financial and human resources highlight that the barriers to improved car maintenance and repair are not primarily the result of institutional resistance. Yet, in spite of the many positive signs among the various car owners, there remain serious problems, the correction of which would serve to further improve car reliability. In the following section, we will examine the three most significant, and will offer solutions to them.

5.3. Current Maintenance Policies and Practices: Problems

The most important problems in the current policies and practices can be summarized as falling into three areas:
- the policies themselves can be improved,

---

[2] The information in this section is drawn primarily from "TT Beefs Up Rack-Car Fleet" (1989).

- the performance measures used to assess maintenance effectiveness are sometimes inappropriate, and
- the information systems are often structured in a way that obstructs rather than enhances good maintenance practices.

In the following sections, we look at the nature of each problem, the sources or roots of the problem, and introduce methods of resolving the problem.

5.3.1. Policy-Related Problems

It was shown in the case studies that the maintenance policies followed by freight car owners can generally be characterized as either "on condition" maintenance or "hard time" planned maintenance. Each of these policies has, of course, some aspects which commend it to the owners, usually the ability to be implemented in the face of the rather unique environment of railroad car operations. In both cases, however, there are serious drawbacks, which can result in wasted resources and unreliable cars or even fleets of unreliable cars.

"On condition" policies are those in which a component or part is repaired or replaced whenever it fails or it reaches a certain condition, or level of wear. In general, the condition levels used in the railroad industry are those prescribed in the A.A.R.'s Field Manual, although we have seen examples where car owners follow a more demanding standard. These policies have several virtues, including their preventive nature, which presumably reduces the number of catastrophic failures the car experiences, and their uniform application across many repair shops on a railroad and across other railroads. For the car owner this means that he can be reasonably confident that his car is being maintained to the same preventive standard throughout its trips. For the railroad, it means that car shop laborers have clearly defined standards for component replacements, and that any liability for improper standards is shared with other railroads. Since no additional preventive maintenance is undertaken, recordkeeping and informational requirements can be reduced to the set of accounting transactions to compensate other car owners for maintenance actions and, if desired, a similar set of on-line repair records.

The following of "on condition" policies also creates a number of problems.

Because the components of a car can reach the condemnation limit at any time, failures and resultant visits to repair tracks occur both during loaded trips and empty ones, resulting in poor service to shippers. They are as likely to occur at expensive shops (i.e., off line or inefficient facilities) as at more efficient ones, depending on where the car travels and the inspections the car receives. In other words, the failures cannot be "concentrated" to occur at desirable points in the car cycle. Further, because each component is inspected and evaluated independently of other components, cars may be subject to repeated removals from service when individual parts reach the condemnation limit. For example, a car might be sent to the repair track for wheels on axle 1, returned to service, then, after a few miles, sent in for those on axle 2 (which is on the same end of the car), returned to service, and crippled a few miles later for the A end truck (which is supported by axles 1 and 2). This type of individualized failure led one industry manager to remark about their "grey fleet", which seems to travel from repair track to repair track, never completing entire trips without needing repairs[3]. The CMO of Company B also expressed concern that they have seen many cars which have no condemnable parts, but have many parts very near the condemnation limits, resulting in poor performance and a high degree of unreliability over time.

This sort of repeated removal from service for repairs has direct economic consequences for the car owner, the shipper, and the railroad handling the car. Even if the removal of a car from a train takes 15 minutes, which in the author's experience is quite fast, this means that a 100 car train experiences almost 25 hours of car delay because of the defective car. A.A.R. studies[4] have suggested that the typical total out-of-service time while undergoing repairs is approximately 4 days, although only 3 days for loaded cars. This time reflects the fact that most repair tracks are switched only once per

---

[3] These sorts of "problem cars" were studied by Little and Martland (1989), and are used to demonstrate benefits of information systems improvements in Chapter 8.

[4] See, for example Dingle (1977), and Guins, et.al (1984). Guins, in a telephone interview suggested that this has probably been reduced to 2-3 days in recent years. On Company A, where the author had occasion to examine some repair records, it would appear that the delay to loaded cars is about 2-3 days, and to empty cars 4-5 days.

day, so that a car will be placed on the repair track for repairs the following day (day 1), and, if parts and labor are available, and the car is of a high enough priority, it will be returned that night. It must then make connections with a train to its destination. For the shipper, this unreliability can result in high economic cost in terms of either raw materials unavailability or high inventory costs. If the variance in travel times is such that the consignee must keep half a weeks inventory of a high value product, these costs can be quite high. Consider the data in Table 5.1.

Table 5.1 compares the opportunity costs associated with being required to maintain an emergency supply of one half week's product for three commodities. The prices of the first two products are taken from the financial pages of the **Boston Globe**, April 19, 1990. The price of the specialized chemical is an estimate based on a conversation with an official from Company C. In the example, the user of the commodity receives one carload per month of the commodity in question. The point, which can be readily seen from the bottom line, is that unreliability due to repeat failures of cars has an increasing impact as the value of the commodity increases. This not only effects the costs which the shipper faces, but also his mode choice.

For the car owner, the repeated removals from service that result from an "on condition" maintenance policy mean that potential savings from economies of scale in maintenance are lost. The existence of such economies is a virtual certainty in freight car maintenance, at least in terms of the costs of removing the car from the train and switching it to a repair track. Many of the operations involved in replacing components are also the same for other nearby components. Replacing a wheelset, for example, involves jacking the car and truck, and removing the wheelset. If the wheelset is an interior one (i.e., axles 2 or 3), then the exterior ones must also be removed to gain access. Thus the cost of replacing the other wheelset at the same time is almost entirely the cost of the materials (and, of course, the foregone wear on the still serviceable component). Guins and Kyparesis (undated) estimated that in the case of wheelsets, replacing all four at the same time at an efficient facility would result in a reduction of 50 percent of the cost of replacing them separately at A.A.R. rates.

The overall effect of "on condition" maintenance policies is that while easy to

	Scrap Steel	Lead	Spec'lized Chemicals
Tons/Carload	85	85	40
Cars/Month	1	1	1
Tons/Year	1020	1020	480
Price/lb.	$0.05	$0.50	$10.00
Dollars/Ton	$100	$1000	$20,000
Dollars/Year	$102,000	$1,020,000	$9,600,000
Tons/Week	19.6	19.6	9.2
Dollars/Week	$1,962	$19,615	$184,615
Stockout Inventory (Weeks)	0.5	0.5	0.5
Discount Rate	10%	10%	10%
Annual Opportunity Cost of Excess Inventory	$ 98	$ 981	$ 9,231

Table 5.1
Inventory Costs Due to Unreliable Service

administer, they result in cars which experience frequent removals from service, and are wasteful of resources.

The other general policy followed by freight car owners is what is known as "hard time" maintenance. This policy is characterized by bringing the car into a maintenance facility at fixed (hard) intervals for inspection, repair, and replacement of components. Like "on condition" maintenance, "hard time" policies have a number of potential benefits to the car owner. The car owner can concentrate efforts to improve the efficiency of the maintenance facility and reduce maintenance expenditures. By judicious selection of which parts to repair and replace, economies of scale may be capture. If the parts which are replaced at the maintenance facility are IFR, the likelihood of in-

service or regulatory failures is reduced, making the car more reliable. Further, since the facility performing the work on the car is known, personnel there can be given different rules for assessing the condition of components without incurring labor relations problems. The key in a "hard time" policy is the appropriate selection of the interval at which to order the car into the shop. It is also here that problems arise.

Each company that adopts a "hard time" policy seems to frequently shift the time in the hope of finding a better time. Company B, for example, in their coal car experiment, originally looked to bring cars in every 41,500 miles. Within 2 years the interval had been raised, first to 83,000 miles and later to 100,000 miles. The managers of the coal car fleet wanted it even higher, but were overruled by the headquarters staff. Those cars in the experiment which belong to a private car company are now being brought in for inspection every 125,000 miles and overhauled every 375,000. Similarly, Company C has steadily increased the interval for bringing in cars for overhauls, except in those cases where the potential for a catastrophe is highest.

This extending of the maintenance interval is not unique to the freight car fleet. In the 1960s and 1970s, the airline industry experimented with "hard time" policies for aircraft engines and airframe components. Gregory (1973) reports the experience of British European Airways. They experimented with various "hard time" policies, and ultimately opted for an "on condition" policy because of findings "that the time between overhauls does not seem to be related well to reliability". He went on to state that for complex equipment

> [w]e know from mathematical analysis, and from millions of hours of
> operating experience, that the removal of components for test or overhaul
> at a fixed time limit is generally a futile method of improving reliability.

This is consistent with the finding of Nowlan (1964), who reported the experience of United Airlines that while individual components which wear out could be scheduled for replacement, complex components, including the aircraft as a whole, were not adversely affected in terms of reliability as the period of scheduled maintenance is lengthened.

This result is really not that surprising upon reflection. Recall that in Chapter 2, it was pointed out that systems of IFR components tend to the exponential [Barlow and

129

Proschan (1965)]. (Exponentially distributed systems are characterized by a constant failure rate (CFR), and are not good candidates for preventive maintenance at the system level.) This tending to the exponential is borne out in the case of freight cars. The author and T.S. Guins used the A.A.R.'s Equipment Reliability Analysis System (ERAS) to estimate the overall failure rate of three groups of freight cars. ERAS is a set of computer programs developed by Guins and Kyparesis to estimate the failure rates of components or systems which are subject to multiple censoring, i.e., some of the components are removed while still serviceable (such as an axle removed with a wheelset), and others are in life throughout the entire test[5]. ERAS fits the repair records of the components or systems in question to a two parameter Weibull distribution[6]. The groups of cars tested were hoppers used in coal or grain unit train service, built between 1977 and 1980. Each group was composed of approximately 500 cars and was homogenous in construction and usage. The shape parameters of the three groups were .941, .959, and 1.117. A shape parameter of 1 is an exponential distribution. That is, these freight cars (as entire systems) seem to be subject to more or less CFR distributions, which suggests at the level of the entire car there is no single optimal time to bring the car in for repairs and replacement.

That systems of IFR components would exhibit this tendency can be seen by considering the example in Figure 5.1, which shows what happens when several groups of components which are each IFR are superimposed onto the system as a whole. As can be seen, the individual components are each reasonably regular in their pattern of failure, but for the system as a whole, the result appears to be a more or less random sequence.

---

[5] For a full discussion of multiple censoring and its effect on estimating failure rates, see Nelson (1982) or Mann, et.al. (1974).

[6] A two parameter Weibull distribution, widely used in life studies, is described by a scale parameter and shape parameter. The scale parameter is the characteristic life of the distribution (i.e., when approximately 63.2% of the items have failed). The shape parameter is dimensionless and, as the name indicates, describes the shape of the distribution. A shape parameter of 1 describes an exponential distribution. A shape parameter of 2 is a Rayleigh distribution, and 3 is approximately a normal distribution.

130

Figure 5.1

Picking out the "best" time for the overall system is the task that the "hard time" maintenance planner faces.

An important caveat is in order here. While it is true that the system as a whole tends toward a CFR distribution, it is not the case that the components which make up the system cannot be preventively replaced to good advantage. The difficulty in picking the single best time to "do everything" should simply lead the maintenance planner to consider planned maintenance at the component level, or look for strategies which consolidate components with similar distributions, or to consolidate maintenance activities opportunistically to extend the time between failures of other components as well.

Each of the two strategies currently being followed has been shown to have serious problems. The "on condition" approach, because it is based on the condition of individual components and disregards the costs of failure or potential savings of joint maintenance activities, results in cars which are often unreliable and expensive to maintain. The "hard time" approach results in ongoing attempts to "tinker" with the maintenance interval in the vain attempt to find a single "magic number" which may not exist. The scheduled time continues to be increased until it approximates the infinite interval appropriate for CFR distributions. What is needed is a multicomponent replacement method which considers the joint costs of maintenance activities, the costs and likelihood of various component failures in the future, and is implementable under the present environment.

The clear candidate for the ideal maintenance policy would be some form of opportunistic maintenance. Recall from Chapter 2 that opportunistic maintenance policies attempt to use the failure of a component as a time to evaluate other components in the system and replace them when economically appropriate. What is required for an opportunistic maintenance policy to be appropriate is simply that there be some economic justification (usually economies of scale) for considering the joint replacement of two or more components. As indicated in Chapter 3, it seems clear that this is the case for freight cars, due to the costs of switching cars to the repair track alone. What is required to *implement* an opportunistic policy, however, is solution to the technical difficulties outlined in Chapter 2. In the next chapter, that is accomplished by the use of a heuristic

which, while not optimal, leads to significant improvements in both cost and reliability over the current maintenance policies.

5.3.2. Performance Measurement Problems

The second general area that presents problems in improved freight car maintenance is in the measurement of maintenance effectiveness. In earlier chapters, some of the characteristics of good measures of maintenance effectiveness were discussed; when the particular measures used in the rail industry and by the case study companies were compared to these, the measures were often found wanting. The lone exception was the set of cost per mile measures used by company C, and this alone is not sufficient to measure reliability or cost control over long periods of time. In this section some of the practical consequences of using inappropriate measures are discussed, and solutions presented.

Before describing the problems associated with the currently used set of measures, it is worthwhile to note the advantages afforded by these measures. Recall that the reliability measures used by various companies included the cost per loaded mile, the bad order ratio and the number of cars removed per train. The first and most obvious benefit of these measures are that they are easily calculated with existing data sets. That is, whatever efforts are required to gather the information have been made. A second, and more important benefit of these measures is that they do focus on particular problems of interest to the maintenance manager. Cost per loaded mile is important both in terms of cost control and for dire-:ting marketing managers away from markets which are expensive to serve. The bad order ratio gives a snapshot of the extent to which fleet utilization is impaired by mechanical failures or maintenance activities. The number of cars removed per train gives an insight into the extent to which mechanical problems are impinging on transportation operations. In the following discussion, it will be suggested that these measures be augmented by others, not that maintenance managers should disregard tools which have, in general, served them well.

In addition to their benefits, each of the currently used measures is prone to operational impacts which are undesirable. The most serious of these follows from the use of cars removed per train, since the obvious way to decrease this number is to reduce

133

the number or quality of inspections performed on cars in trains. This can result in components reaching conditions which prevent them from being reconditioned, or worse, may result in catastrophic failures. It is not uncommon to hear operating officials state that the best way to move trains is to "keep car knockers away from them". The use of this measure tends to encourage this approach to car inspections.

The use of the bad order ratio, which compares the number of cars in a shopped condition to the total number of cars in the fleet, can also lead to bad practices. Because the bad order ratio is a fleet level measure, the maintenance manager is really given no information on which cars are experiencing maintenance needs, or why. Certain series of cars, because of design complexities, usage characteristics, or component aging may be receiving extensive maintenance attention, which is not indicative of problems in the fleet or even the series. The bad order ratio tends to encourage managers to focus on getting these cars back into service, when more maintenance actions may be warranted. The use of the bad order ratio can also create a misleading impression that all is well simply because the car fleet is being downsized. This is because car owners will tend to retire cars which are older, or subject to higher maintenance requirements, which reduces the bad order ratio without changing the maintenance policies or practices. The manager then believes that maintenance has become more effective (or efficient), when, in fact, the fleet being maintained is all that has changed. Under this scenario, as the fleet ages, the maintenance manager is doomed to look increasingly ineffective.

As was suggested in the preceding chapter, the solution is to add measures which are more directly related to maintenance activities and their desired effect. Such measures should be easily calculated and understood, should be directed toward the organization's goals for maintenance, and should give some insight into problems as they develop. It was suggested that three measures be used together to monitor the effectiveness and efficiency of freight car maintenance. These are:

- Cost per mile,
- Miles per in-service failure,
- Miles per maintenance event.

The first of these, cost per mile, is useful as an efficiency measure, since cost can be

treated as an input, and miles can be thought of as the output that a freight car produces. The second measure is more or less a standard reliability measure, corresponding to mean time between failures, except that miles are the usage measure for freight cars. The final measure, miles per maintenance event, recognizes that each time the car is removed from service there are costs incurred for the car owner, the shipper, and the railroad handling the car. This measure tends to encourage managers to "bundle" maintenance activities to reduce the number of separate maintenance events (i.e, removals from service for maintenance).

Application of these measures at the level of the individual car allows managers to detect both problems with cars and with the maintenance policies or practices used on the car. These problems can have serious impacts on both the costs incurred by the company (detected by cost per mile), and the service offered to customers (indicated by miles per in-service failure). Consider, for example a car (or series of cars) which exhibit decreased miles per in-service failure. These cars can be readily detected by a monitoring program which is triggered to detect changes in car performance. If the car series is subject to some form of planned maintenance, the increase immediately suggests that there is a problem either in the design or the application of the maintenance policy. In the meantime, marketing managers can begin to deal with the customers who depend on that car series, either offering different cars, rates, or at least information that the problem is under review. Similarly, if the miles per maintenance event measure increases while miles per in-service failure remains the same, this may indicate that certain systems of components are requiring more attention, and the maintenance manager can begin to determine which, if any, of the increased maintenance activities should be consolidated or handled differently. The important point is that these measures correspond to the company's objectives of providing reliable vehicles at reasonable cost, and using capital efficiently. They also serve to highlight changes in maintenance performance, and, properly used, can begin to guide managers to find the sources of problems or inefficiencies.

An important aspect of these measures is that the data needed to calculate them are already available, in some form, to freight car owners. Both railroads and private car

135

owners maintain some form of mileage information (or waybill data which can be translated into accurate approximations). Indeed, private car owners are compensated, in part, for the miles their cars are moved on railroads. The costs of maintenance on cars is virtually always kept as part of the car repair billing system. The number of maintenance events (i.e., trips to repair tracks) can be easily determined from a car's repair records. The number of in-service failures, is simply all offline repairs plus all online repairs, net of scheduled maintenance events. This too can generally be determined from the car repair records. In Chapter 7, when alternative maintenance policies are presented and compared with the present practices, these measures are used as the basis for evaluation.

While the "raw materials" are available for building and using proper indices for measuring maintenance performance, direct access to that information in formats which are usable to maintenance managers is often not available. Problems with the organization, structure and access to the data needed by maintenance managers is the subject of the next section.

5.3.3. Information System Problems

A third serious problem in freight car maintenance is in the area of information systems. Quite simply, the currently used information systems are not constituted in a manner which leads to better maintenance management. Indeed, the typical structure of the data and limitations on access actually serve to restrict car maintenance to ineffective maintenance policies. Car information systems, because of their roots i 1 car accounting systems, tend to concentrate on the billing of components rather than the reliability or even the use of the car itself. Because of this component-by-component approach, car repair data sets also tend to be quite large and cumbersome. Finally, because the data has important accounting functions, and because of the large size of the data sets, access is often restricted, and requires computer skills beyond those often found in maintenance managers. The result is that the information systems are not an integral part of the car repair and maintenance process, and are viewed as administrative requirements rather than managerial assets. The practical effect is that the information systems are not used to detect cars with problems, inefficient facilities, or poor maintenance policies.

The driving force behind virtually all the car maintenance information systems in use today is the AAR's Car Repair Billing (CRB) system. The CRB system specifies a set of formats for the pricing, billing and payment of repairs. As a result, most car owners use either the CRB formats (or a superset of them) as the basis of their own data bases. The problem with this is that the CRB data is really part of an accounting system, which is organized around the particular components and billable items associated with a car repair rather than with the car itself. As a result, the replacement of a wheelset causes the creation of records for each wheel, the axle, the bearings, the associated labor, and other minor parts. Needless to say, the data sets involved can quickly grow quite large. This is an appropriate way of structuring the data as long as the maintenance policy is simply an "on condition" policy, since the concern is when and where a component reached the condemnation limit. But if another policy is to be followed, the tracking of individual components may well be inappropriate. Notice that when Company B wanted to begin a program of "hard time" maintenance they found it necessary to undertake a series of detailed component measurements in order to try to estimate wear rates. The point is that what is good for the accountants of any organization is not necessarily good for all the other managers.

This same problem is found in the other relevant data bases which freight car owners, especially railroads, typically maintain. (On one railroad, the waybill data, which tracks the routing and loading information of a car trip is stored in a file keyed by the waybill number.) More importantly, because these data sets are important to many parts of the organization, changing them can be expensive and difficult, since the programs of a broad user community must be changed.

The result is that information which should be reaching maintenance mangers frequently does not. One example of this is the existence of cars with recurrent problems. In section 5.3.1. the problem of cars which are repeatedly removed from service for distinct defects was discussed in the context of "on condition" maintenance policies. A more pathological case, however, is that of cars which are removed from service repeatedly *for the same defect*. These are cars in which some sort of underlying problem or defect (e.g., improper brake rigging) causes other components (e.g., wheels) to wear

137

out or reach condemnable conditions very rapidly. Because the repairs occur at many different points, and the data sets are too large and inaccessible to monitor them, managers have, in many cases, simply allowed these cars to continue in service (at considerable expense). These cars, known as "problem cars", have been estimated to represent an avoidable economic loss to car owners of more than $60 million annually [Little and Martland (1989)]. In Chapter 9, this problem is discussed further and a solution is presented.

Given that the present information systems are needed for various accounting and auditing functions, what is called for is not so much a redesign of the overall data sets, structures and programs currently in use, as a single, direct intelligible (and intelligent) pathway into the data sets which can be used by maintenance managers. In Chapter 9, the concept of the "structured history" is developed and demonstrated. The structured history is what is known as a knowledge base, which is organized around the car in the way that car maintenance managers think about cars, rather than the way that data processors think about data. Use of the structured history permits the maintenance managers and planners to build information systems which reflect their expertise in maintenance management, and to apply that expertise to very large fleets.

5.4. Conclusions

In this chapter, we have looked at the policies and practices followed by the case study companies and freight car owners in general and found some specific problems. While car owners have demonstrated a willingness to invest financial and human resources to improve freight car reliability and overall maintenance performance, the actual policies followed are subject to a number of drawbacks. The measures used to evaluate maintenance effectiveness have also been shown to lead to operational problems. Finally, the information systems used to support car maintenance and repair are unwieldy and difficult to use, resulting in problems even under the current policies. In each instance, specific remedies have been introduced. The balance of this thesis demonstrates that these treatments are workable and beneficial.

In the next chapter, we take up the problem of maintenance policies. It has been

138

shown that railroads are good candidates for opportunistic maintenance policies. We now focus on showing that the technical difficulties in finding a workable approach can be overcome.

Chapter 6

A Heuristic For Opportunistic Railcar Maintenance

6.1. Introduction

In the previous chapters it was established that while opportunistic maintenance may be a natural policy for the railroad car environment, the standard models from the literature cannot be directly or easily applied to the case of railroad car maintenance. In Chapter 5, it was shown that the use of opportunistic maintenance has an intuitive appeal because of economies in switching costs and in common maintenance actions such as jacking the car. The review of opportunistic maintenance policies presented in Chapter 2 showed that the use of such policies in cases where there are many components presents an unacceptable computational burden, even to the point of being intractable. The problem is further complicated by the complex cost structure faced by railroads, with various costs of failure depending on where the car is located when a part fails. Freight railroads face yet another complicating factor in that the car is subject to an on condition policy whenever it is off line, which can be a considerable fraction of the car's life. All of this suggests the use of heuristics, which, while suboptimal, may still outperform the current practice.

Only a few such heuristics are presented in the literature, and these do not appear to be readily applicable to freight car maintenance. Recall that the two approaches usually suggested are "screens", in which all components are replaced at a fixed percentage of their expected life, and "all or nothing" policies, in which the failure of any component results in the replacement of all the other components that are subject to preventive maintenance. In Chapter 2 it was pointed out that if the typical lifetimes of the various components in the maintenance program cover a wide range of values, then both approaches are likely to result in considerable waste. A further concern is that the particular values used in the "screen" approach are found as a result of simulations, and if the costs or failure distribution of any of the components are not constant over time, then the screen value selected may become inappropriate.

In this chapter, a heuristic is proposed which is tailored to the circumstances faced

by railroads, and which makes the replacement decision in a dynamic way that allows the costs and quality of parts to change while remaining robust. This method seeks to achieve three related goals:

(1) increasing the survival probability typically associated with the replacement policies discussed earlier;

(2) achieving the cost savings typically associated with scheduled rather than in-service failures;

(3) capturing some of the economies of scale associated with opportunistic maintenance policies.

The basic tradeoff that the heuristic seeks to balance is between the foregone component life that could be realized under an on condition maintenance policy and the increased reliability and economies of scale under an opportunistic policy.

An additional concern is to make certain that the approach depends only on information that is readily available to the railroad car owner. In principle, one could require the collection and processing of additional data and information, but a number of the railroad car managers interviewed in the course of this research indicated that they would consider any such information requirements excessive because of constraints they face in this area. Accepting this as a constraint, the approach presented in this chapter depends more on rep. cessing and use of existing information sources than on new data.

The key notion of the heuristics is that when the car is on the repair track for any reason, it may be desirable to perform maintenance on any of the car's components. Such maintenance of unfailed components is considered in order to increase the reliability of the car while in service, reduce the likelihood of failure at a more expensive facility in the future, or exploit economies in maintenance.

The structure of this chapter is as follows: first a conceptual description of the heuristic is given. The details of the heuristic are then presented using a simple two component model. In presenting the two component version a simple, "greedy" form of the heuristic is first described which is then "extended" to include more of the available information. This model is then expanded to include a larger number of components. Finally, the problem of imperfect imformation is examined.

In presenting the two component model a numerical example is given using

wheelsets. Wheelsets are particularly good for illustration purposes since they constitute some 30 percent of repairs reported in the Car Repair Billing System [Guins and Hargrove (1980)] and they wear out with usage. Wheelsets are also expensive as individual items, with typical replacement costs over $1000. In Chapter 7, a simulation model is used to compare both versions of the heuristic with current practice and a simple component-by-component scheduling approach. In that chapter, eight components are modeled, clearly demonstrating the multicomponent nature of the proposed methods.

Terms which were defined in Chapter 2 relating to general concepts of failure, distributions, and reliability theory are used throughout without repeating the definitions.

6.2 Conceptual Overview of the Heuristic

Before beginning a detailed exposition of the heuristic, it may be useful to discuss the underlying philosophy behind it. The basic notion of the heuristic is that one can schedule each of the components for replacement individually, and then preempt that schedule in order to collect maintenance activities together. As a starting point, one can ask the question, "what happens if each part is treated as a completely separate and independent item and maintained accordingly?" This results in a policy that might be called "naive scheduling", in which the "optimal" replacement interval for that component is calculated using the single-part age replacement policies presented in Chapter 2. No one could seriously expect that maintenance managers would follow such an approach for very many components, however, because the car would be going to the shop very often. If, for example, one component's "optimal" replacement time was after 180,000 miles, it is likely that the car's owner would also want to replace another component which is due after, say, 181,000 miles. The issue then becomes one of finding ways to combine such events together when economically justified, and allowing components to remain in service when replacement is not warranted.

It was decided that a reasonable approximation of that decision process could be developed using the notion of expected values. Expected values weigh the probability of various outcomes and the values (or costs) of the possible outcomes to create an "overall effect" of uncertain events. As the number of trials or actual outcomes is increased, the

expected value and the mean of the actual values generally approach one another. The process adopted for the heuristic is based on calculations of the expected costs of allowing the part to remain in service until the single component "optimal" interval, and comparing that cost with the cost of replacing the part at the present time. This has the effect of causing maintenance managers to combine "nearby" events (such as the 180,000 and 181,000 mile events), while leaving in service those components which still have considerable potential life remaining. Because the number of components (and the number of cars in most fleets) is high, the use of expected value calculations seems to be an appropriate method of dealing with the uncertainties of future events.

In deciding how to combine maintenance actions, one would like to include as much information as possible about the costs and consequences of each potential outcome, and select accordingly. In practice, however, the decision must be reduced to some set of reasonable alternatives for which costs and consequences can be evaluated. The heuristic presented is based on a simplification regarding future opportunities. In its simplest version, the implicit assumption is that the present opportunity created by the car being on the repair track is a one-time phenomenon. That is, the car owner decides between replacing a part now or leaving it in service until it either fails or is replaced on schedule. The possibility of another opportunity is not considered. In the "extended" version, this possibility is considered, albeit in a limited way. In both cases, the calculation of future costs assumes that a replaced component will remain in service as if a policy of "naive scheduling" were being followed. Put simply, the calculation of future costs after replacement are based on the assumption of no opportunities for replaced parts. This assumption has the effect of "isolating" the decision from future events, and simplifies the calculations considerably. The consequence of this assumption is that expected future costs following the replacement of a part at the present time are probably overestimated (since we are assuming no future inexpensive opportunities to arise and be taken), so that some opportunistic replacements are deferred to a later time.

The basic decision regarding which maintenance actions to combine at a given time are based on an analysis of the costs of replacing the part at this time versus leaving it in service until some future time. To formulate this decision, we can use the "time
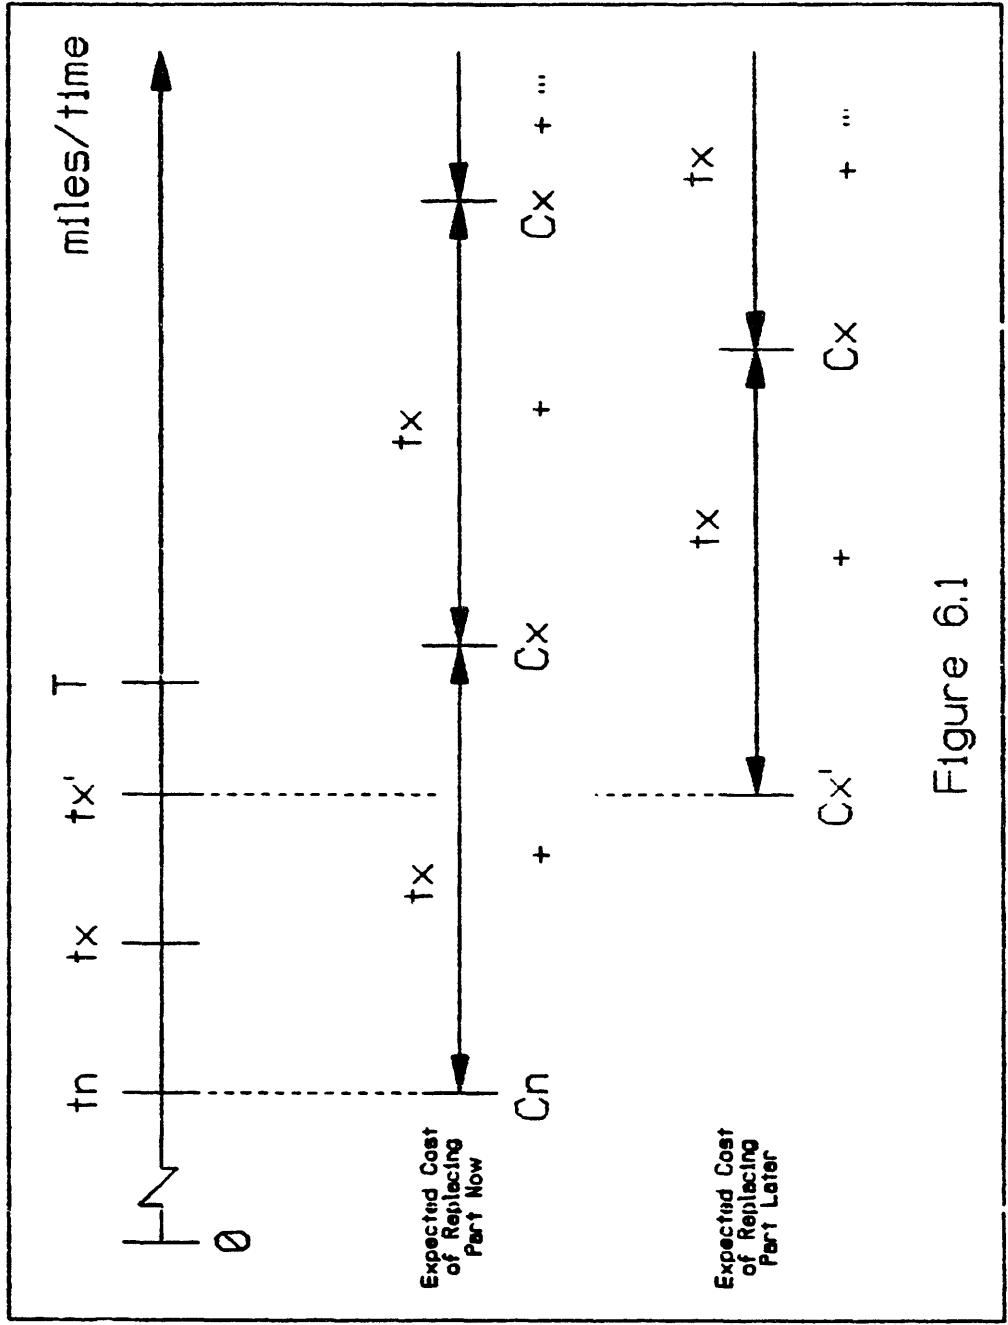
Figure 6.1

line" in Figure 6.1. (We will be using time and miles interchangably, since we are concerned with the unit of measure by which things "age". In the case of railroad freight cars, this is generally mileage.) If we designate the time at which a part is installed as time 0, then we can indicate the scheduled replacement time for that part as time T. This time can represent a mandated replacement time imposed by regulatory authority, an estimate of the expected life of the component, or, the "optimal" replacement time for the component as a stand alone system (a discussed in Chapter 2)[1]. In the case of the heuristic, we will use the latter, although the structure of the heuristic does not depend on how time T is derived. Given a failure distribution, we can also find a time t_x, which is the expected time of replacement of the part given that the part will be replaced at time T if it survives that long. This time is given by

$$t_x = \int_0^T [1 - F(t)]\, dt \qquad (6.1)$$

where F(t) is the failure distribution of the part. The expected cost of replacement of the part over the interval (0,T) is C_x, where C_x represents the costs of a failure of the part and the cost of a scheduled replacement, each weighted by the probability of that outcome. Consider the case where the part has survived until some time t_n, (representing time now); the expected time and cost of replacement must be recalculated in light of the information that the part has survived to the present time. These new values, conditioned on the survival to t_n, are designated t_x' and C_x'. The value of t_x' is given by

---

[1] The methods given in the literature for finding the "optimal" time for replacing a single component (including that used in the following sections) generally do not incorporate a discount rate. If there is a positive discount rate, such methods will tend to replace parts sooner than the true optimum, since the discounted cost of a failure at a later time will be less than the cost used, and any preventive replacement will occur earlier than a failure would be expected to. The effect of using an undiscounted T in the heuristic is probably to cause components to be replaced earlier than they would be if the true optimum were available. It is beyond the scope of this thesis to develop a new and better method for estimating the single component replacement time. If such a method becomes available, it should be used in the heuristics that follow.

$$t_x' = t_n + \frac{\int_{t_n}^{T} [1 - F(t)] \, dt}{1 - F(t_n)} \quad (6.2)$$

and the value of C_x' is again the costs of failure and scheduled survival appropriately weighted by the probabilities of each outcome.

The decision of whether or not to replace the part at the present time turns on whether the present value of the costs of replacement are less than the present value of the expected costs of leaving the part in service. We can designate the costs of replacing the part at the present time by C_P, and the costs of leaving the part in service as C_L. C_P is given by

$$C_P = C_n + \sum_{j=1}^{\infty} C_x \frac{1}{(1+i)^{jt_x}} \quad (6.3)$$

where C_n is the cost of the replacement at this time (i.e., the actual cost of installing a replacement part, net of any salvage value), and i is the discount rate[2]. The summation term on the right hand side of equation 6.3 is the expected cost of repeatedly replacing the new part every t_x miles, at an expected cost of C_x. The expected cost of leaving the part in service at the present time is given by

$$C_L = \left(\frac{1}{(1+i)^{t_x'-t_n}} \right) \left[C_x' + \sum_{j=1}^{\infty} C_x \frac{1}{(1+i)^{jt_x}} \right] \quad (6.4)$$

The term in brackets on the right hand side of equation 6.4 can be understood as the expected cost of replacing the part at time t_x', given as C_x', and the cost of replacing that part every t_x miles at a cost of C_x, discounted to time t_x'. All that is discounted to the current time (t_n) by the leftmost term on the right hand side.

---

[2] Later in the chapter it is shown that this equation can be simplified considerably.

The decision of whether or not to replace the part at the present time is thus based on whether the present cost, C_P, is lower than the expected cost at a later time, C_L.

The heuristic described in this chapter shares some features with greedy algorithms used in optimization studies. In particular, whenever an opportunity for preemptive maintenance presents itself, it is taken if the cost is lower than the expected cost of waiting until the scheduled time for that component. In the simplest version, the expected cost of waiting until the scheduled time only incorporates information about the failure distribution of the component itself; in the "extended" version, the expected costs of waiting for the scheduled replacement includes the overall failure distribution of the car as a whole.

We turn now to the implementation of the heuristic, first by examining a simple two component model, and then looking at multicomponent implementation.

6.3. The Two Component Model

The central notion behind the heuristic is that one seeks to make a "good" decision about maintenance actions by balancing information about maintenance costs (present and future), failure rates, and expected usage of the car. In particular, the present value of two streams of costs are calculated: the cost of replacing the component at this time (and the expected costs of replacing that component in the future), and the expected costs of leaving the component in service until it either fails or is scheduled for replacement under a single component replacement policy (again including the expected costs of future replacements).

The basic approach to planning opportunistic maintenance consists of three steps:

1. Given that the car is in a repair facility under the control of the car owner, the present value of the costs associated with replacing the part at this time is calculated.

2. The present value of the expected cost stream associated with allowing the part to remain in service is calculated (including the costs associated with a possible in-service failure).

3. If the present value of the costs associated with replacing the part are lower than the present value of the expected costs associated with leaving the part in service the part is replaced. If not, the part is left in service.

147

In order to demonstrate the decision process, a numerical example is presented. The data for the IFR part is for wheelsets, and is taken from Guins and Kyparesis [undated]. The overall data for the car was estimated by T.S. Guins and the author using data reported in the A.A.R.'s Car Maintenance Cost data base for a group of cars in unit train service averaging 66,000 miles per year. The cost data used approximates actual figures to the extent possible, with the data coming from the case study participants, A.A.R. studies, A.A.R. rates reflected in the Office Manual, and, to a limited extent, the author's conjecture.

6.3.1. The Two-Component "Greedy" Model

In this section, a model is considered in which the car is treated as composed of 2 parts. One of the components, which we will designate as part i has an IFR failure distribution and the other component, denoted part r, is exponentially distributed. Part r can be thought of as representing all the rest of the components in the car. The basic components modelled in this case are very similar to those in the standard Radner and Jorgenson opportunistic problem. Unlike the Radner and Jorgenson problem, however, the cost structure is dependent on when and where failures occur, with costly in-service failures, and with some of the maintenance decisions made outside the control of the car owner. This complicated cost structure with limited control is part of what characterizes the railroad freight car maintenance environment and makes the problem unusual. The other important aspect of the 2 component "greedy" model presented here is that, unlike the Radner and Jorgenson model, it can be extended to the multicomponent case, to which their model does not apply.

To simplify the initial analysis, several assumptions are made. We assume the owner of the car has complete and immediate information as to what repairs are performed on the car, regardless of who performs the repair. (Relaxation of this assumption is discussed later in the chapter.) We assume further that the owner of the car may not ask other railroads to replace components which are not condemnable under the A.A.R. interchange standards, i.e., the car is subject to on condition maintenance when off line or, in the case of private cars, when not in a shop controlled by the car owner. It is assumed that the owner has only one repair facility, or has the same cost structure

148

at all facilities, and that capacity is not a constraint at the on-line repair facility. (This assumption is made to restrict the costs considered to those of off and on-line repairs as opposed to the further consideration of which on-line repair point is best.) Finally, we assume that all repairs made are, in fact, replacements, which restore the component to "good-as-new" condition, and that the repairs are made perfectly.

We can formulate the decision process regarding the replacement of part i (the IFR part) as following the pattern in Figure 6.2. If the IFR part has failed, it is replaced. If the IFR part has been in service for a period of time greater than time T (as a result of the car being used off-line, for example), then it is replaced. If neither of these conditions hold, then we calculate the expected costs associated with replacing the part at the present time, C_P, and the expected costs of leaving the part in service, C_L. If the expected costs of replacing the part at the present time are less than the expected costs of leaving the part in service it is replaced; if not the part is left in service.

We can calculate C_P and C_L by applying equations 6.3 and 6.4. Prior to this, however, we can reduce the terms involving infinite series. To do this, recall that for an infinite series[3], if $|r| < 1$, then

$$\sum_{k=1}^{\infty} ar^{k-1} = \frac{a}{1-r} \qquad (6.5)$$

We can rewrite equation 6.3 as

$$C_P = C_n + C_x \sum_{j=1}^{\infty} \left(\frac{1}{(1+i)}\right)^{t_x} \left[\left(\frac{1}{(1+i)}\right)^{t_x}\right]^{j-1} \qquad (6.6)$$

which can then be reduced using equation 6.5, and, after some manipulation of terms, becomes

---

[3] See, for example, Sobel and Lerner (1983), Chapter 10.
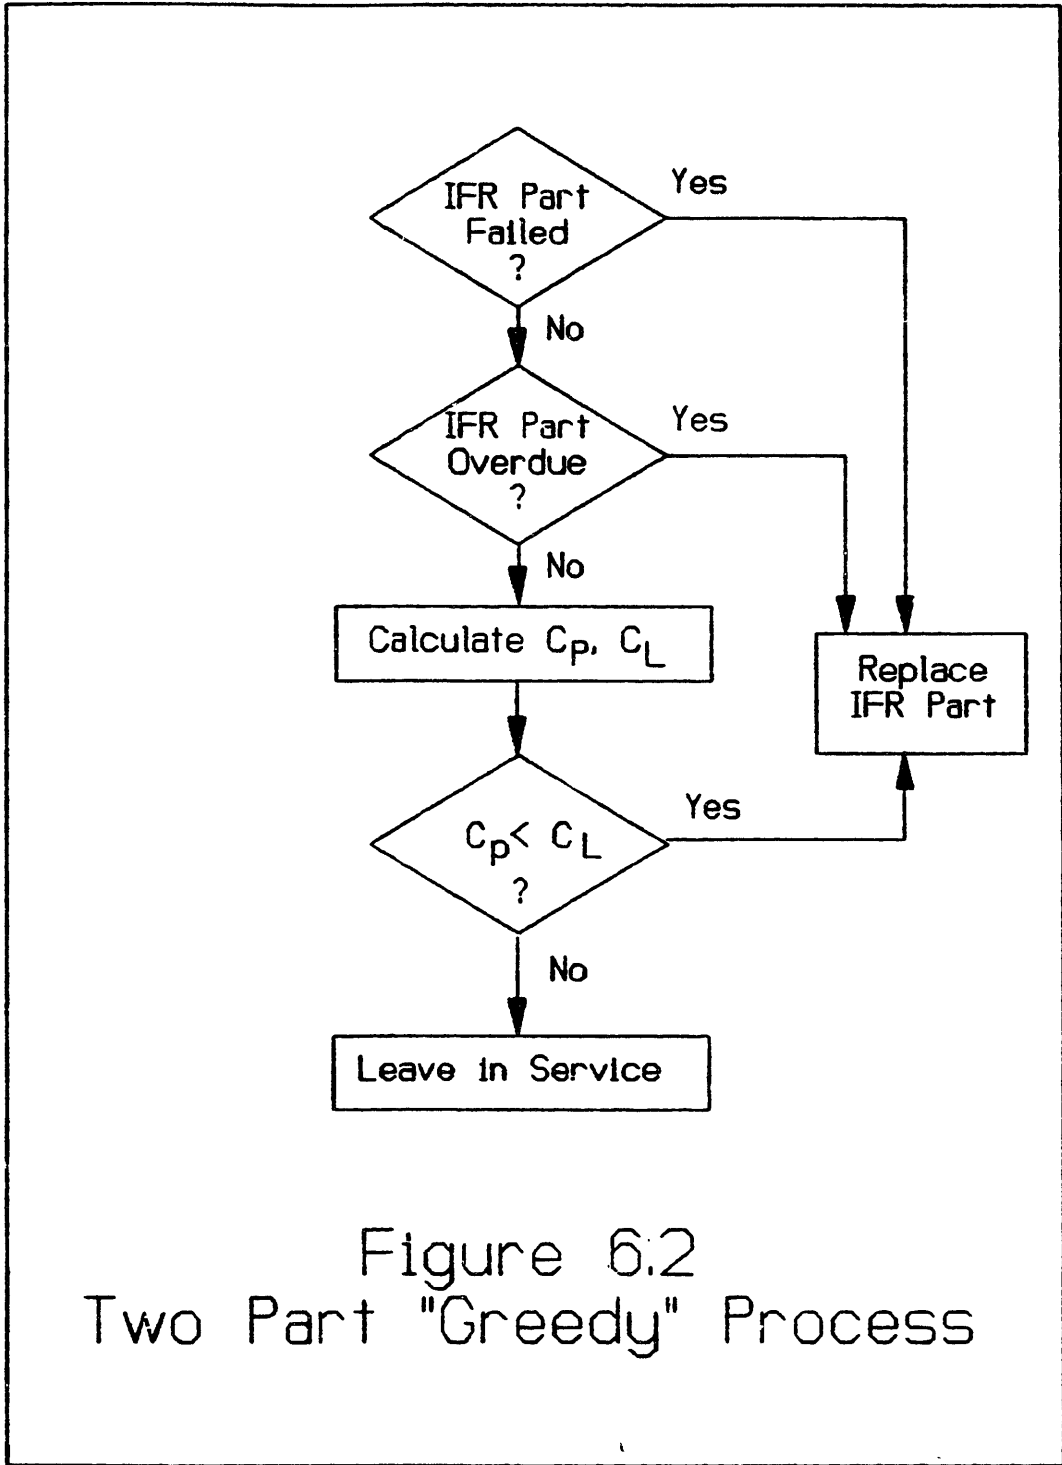
Figure 6.2
Two Part "Greedy" Process

$$C_P = C_n + \frac{C_x}{(1+i)^{t_x} - 1} \qquad (6.7)$$

Similarly, equation 6.4 becomes

$$C_{I_c} = \left(\frac{1}{(1+i)^{t_c'-t_a}}\right)\left[C_x' + \frac{C_x}{(1+i)^{t_x} - 1}\right] \qquad (6.8)$$

In order to evaluate these, it is necessary to find the terms t_x, C_x, t_x', and C_x'. As has been indicated earlier, t_x and t_x' are found by applying equations 6.1 and 6.2. (If the integrals of the failure distribution do not have closed forms, as in the case of the Weibull distribution, numerical methods must be used). C_x, the expected cost of a replacement evaluated when the component is newly installed, subject to a failure distribution F(t), and to be replaced at time T if it survives that long, is given by

$$C_x = \left[F(T) C_F \left(\frac{1}{(1+i)^{t_F-t_x}}\right)\right] + \left[(1-F(T)) C_S \left(\frac{1}{(1+i)^{T-t_x}}\right)\right] \qquad (6.9)$$

where C_F and t_F are the expected cost and time of an in-service failure, and C_S is the cost of a scheduled replacement of the part. In the case of a railroad freight car, the expected cost of an in-service failure is given by the cost of an off-line in-service failure and the cost of an on-line in-service failure, weighted by the probability of the failure occurring off- or on-line respectively, i.e.,

$$C_F = C_{off}(P(Off)) + C_{on}(1 - P(Off)) \qquad (6.10)$$

where, C_{off} and C_{on} are the costs of off-line and on-line failures, respectively, and P(Off) is the probability that the failure occurs off-line. In practice, C_{off} can be estimated by the A.A.R. rate for the repair, plus all the costs associated with unreliable service, C_{on} by the costs of repairing the car on-line plus the associated unreliable service, and the costs of

151

delay imposed on other cars in the consist. P(off) can be approximated by the share of miles that the car is operated off-line (assuming that inspections on and off-line are of comparable quality).

We can calculate t_F in several ways. Perhaps the simplest, and most intuitively appealing is to note that t_x represents a "balance point" between T and t_F and their respective probability masses, so that

$$F(T)\,(t_x - t_F) = (1 - F(T))\,(T - t_x) \tag{6.11}$$

hence

$$t_F = t_x - \left[\frac{(1 - F(T))\,(T - t_x)}{F(T)} \right] \tag{6.12}$$

C_S is the cost of a scheduled replacement of the part. This will include the cost of the repair, any switching costs associated with sending the car to a repair track, and the opportunity cost of having the car unavailable during the replacement. In solving for C_F, the C_S must be discounted back to time t_x, since a scheduled replacement will take place only after the part has been in service for T miles, and all the costs are evaluated at the expected time of replacement, t_x.

C_x', the expected cost of replacing the part later, given that the part has survived to the present time, is given by

$$C_x' = \left[\frac{F(T) - F(t_n)}{1 - F(t_n)}\, C_F \left(\frac{1}{(1+i)^{t_F' - t_x'}} \right) \right] + \left[\left[1 - \left(\frac{F(T) - F(t_n)}{1 - F(t_n)} \right) \right] C_S \left(\frac{1}{(1+i)^{T - t_x'}} \right) \right] \tag{6.13}$$

The essential change in this from equation 6.9 is in the calculation of the probabilities; in this case the probability of failure (or survival) is conditioned on the fact that the component has survived to the present time, t_n. The other changes are that the cost of a scheduled replacement is now discounted to time t_x' rather than t_x, and t_F' is evaluated as above using the conditional failure distributions. The expected time of failure over the

152

interval is given by

$$t'_F = t'_z - \left[\frac{\left(1 - \left(\dfrac{F(T) - F(t_a)}{1 - F(t_a)}\right)\right)(T - t'_2)}{\left(\dfrac{F(T) - F(t_a)}{(1 - F(t_a))}\right)} \right]$$

(6.14)

Once one has numerical values for each of the terms, one can proceed in either of two directions. One can calculate C_P and C_L, and if $C_P < C_L$ then one should replace the part because the expected costs associated with replacing are less than the expected costs associated with leaving the part in service. If $C_P > C_L$ then the part is left in service since that results in lower expected costs. If the two are equal, then we are indifferent regarding replacing the part or leaving it in service.

An alternative approach is to set C_P equal to C_L, and find the corresponding t_n. This is done by setting equations 6.7 and 6.8 and solving for t_n (recalling that C_x' in that equation is also a function of t_n). Once this value is found, which we designate t_n^*, we can simply apply the rule that if $t_n > t_n^*$, then the part should be replaced. The problem with this approach is primarily administrative. If the cost of parts used is subject to changes, then it becomes necessary to update the critical value of t_n^* whenever the inventory is changed, even though no decision is faced at the present time. In any event, the two methods are equivalent in their results.

6.3.1.a. A Numerical Example

To show the decision process, we consider an example drawn from actual railroad data. Let the IFR part, part i, be a wheelset, with a failure distribution which can be represented by a Weibull distribution with shape parameter of 3.5, and a scale parameter, or characteristic life, of 274,000 miles (Figure 6.3). Let P(Off), the probability of being off-line at any time, be 0.5; that is, the car is off-line for half of the miles operated. The cost of an off-line in-service failure, C_{off}, is $1400, and C_{on}, the cost of an on-line in-service failure is $1250 ($1100 for the failure and associated costs plus $150 in switching

153

Wheels and Rest of the Car

Probability Density Function
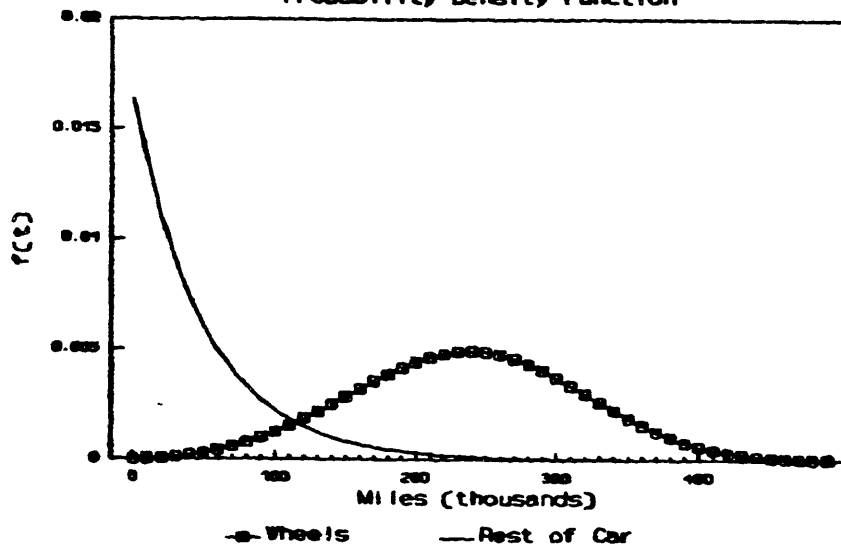
(a)

Wheels and Rest of the Car

Cumulative Distribution Function
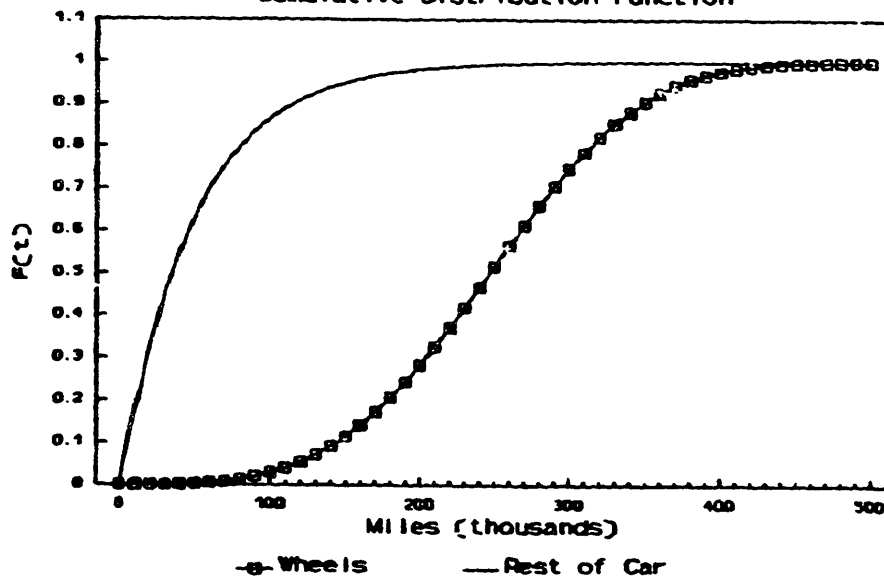
(b)

Figure 6.3

costs). From these, we can calculate C_F, the expected cost of an in-service failure, using equation 6.10 to be $1325 (= $1400 * (.5) + $1250 * (.5)). Let the cost of a scheduled replacement, C_S be $850 ($700 for the replacement plus $150 in switching costs). Let the cost of an opportunistic replacement be the same $700 as in a scheduled replacement. (There are, of course, no switching costs for an opportunistic replacement, since the car is already on the repair track.) Let the discount rate, i, be 10%[4]. All the values used in the example are summarized in Table 6.1.

Given these costs and this failure distribution, one can apply any of the methods described in Chapter 2 to determine T, the single component "optimal" replacement interval. Using the graphical method in Jorgenson, McCall, and Radner (1967), T can be estimated at 257,000 miles.

Let the present time be 180,000 miles since the wheelset was installed. If the car is on the repair track at the present time, a decision can be made whether or not to replace the wheelset.

Because the wheels are modelled with a Weibull distribution, equations 6.1 and 6.2 do not have closed form solutions, and numerical methods must be used. Using the trapezoidal method given in Press, et.al., (1989), t_x evaluates to 220,000 miles. That is, the expected replacement time at time 0 with a scheduled replacement time of 257,000 miles and the indicated failure distribution is 220,000 miles. Given that the part has survived to 180,000 miles, t_x' evaluates to

$$t_x' = 180,000 + 49,000/.79 = 242,000 \text{ miles.}$$

The expected time of failure if a failure occurs in the interval (0,T), t_F, is given by equation 6.12,

$$t_F = 220,000 - (((.45)(257,000 - 220,000)) / (.55))$$
$$= 190,000 \text{ miles.}$$

C_x, the expected cost of replacing a wheelset assessed at the time of installation

---

[4] Since discount rates are generally applied over periods of time, and we are concerned with periods of usage, i.e., mileage, the discount rate is divided by the annual mileage of the car. Thus a .1 annual discount rate for a car with an annual mileage of 66,000 miles becomes .1/66,000 = .001515 per thousand miles.

Variable	Meaning	Value
P(Off)	Probability that an in-service failure occurs off-line	.5
t_a	The present time (measures in miles)	180,000 miles
i	The discount rate	.1
Variables Related to the IFR Part		
β	Shape parameter	3.5
α	Characteristic life	274,000 miles
C_{off}	Cost of an in-service failure that occurs off-line	$1400
C_{on}	Cost of an in-service failure that occurs on-line (includes switching costs)	$1250
C_s	Cost of a scheduled replacement (includes switching costs)	$850
---	Cost of replacing when an opportunity arises	$700
T	Optimal replacement interval as a stand alone system	257,000 miles
t_I	Expected time to first failure estimated at time of installation	220,000 miles
Rest of System (used in later examples)		
---	Mean time between random failures	50,000 miles

Table 6.1
Values Used in Numerical Examples

is given by equation 6.9,

$$C_x = (.55) (1325) (1/(1.0015)^{220-180}) + (.45) (850) (1/(1.0015)^{257-220})$$

$$= \$1125.$$

$$t_F' = 242,000 - (((.57)(257,000 - 242,000)) / (.43))$$

$$\approx 222,000 \text{ miles.}$$

C_x', the expected cost of replacing the wheelset currently in use at time t_o, is

$$C_x' = (.43) (1325) (1/(1.0015)^{-20}) + (.57) (850) (1/(1.0015)^{15})$$

$$= \$1060.$$

With these values, we can now calculate C_P and C_L and make the replacement decision. C_P, the present value of the stream of expected future costs associated with replacing the wheelset at the present time, is given by equation 6.7,

$$C_P = 700 + 1125/[(1.0015)^{220} -i] = 700 + 2846 = \$3546.$$

C_L, the present value of the expected future costs of replacing the wheelset at a later time, is given by equation 6.8,

$$C_L = (1/(1.0015)^{(242-180)}) (1060 + 1125/[(1.0015)^{220} -1]) = \$3556.$$

Since $C_P < C_L$, we replace the wheelset at this time.

We now turn to the matter of including information regarding the overall system reliability in the decision process.

6.3.2. The Two Component "Extended" Model

In this section the model is extended by considering the possibility of another, better opportunity to replace the part in the interval between the present time and the scheduled time. The notion here is that if the car as a whole is subject to frequent failures, there may be an advantage in waiting until a later time to make an otherwise acceptable component replacement, since another opportunity is "likely", and will result in lower overall costs. The tradeoff comes in the possibility that the part will fail while awaiting the next opportunity, or another opportunity may simply not arise before the scheduled replacement time. What we are interested in, then, is the possibility of a "random" failure of another component in the interval (t_o, T), as measured from the time of last replacement.

To frame the decision process, we can again construct a flow chart like that in

Figure 6.2. In this case, we follow the process in Figure 6.4. If the car is on the repair track for a failure of part i, the IFR part, then we replace the part at this time. Similarly, if the part is "overdue", in the sense that it has been in service for more than T miles, it must be replaced. The part is then tested according to the "greedy" approach described above. If the part would not be replaced under that scheme, then it will surely not be replaced under this, since the extended version is only concerned with leaving parts in service for a longer time when the possibility of better opportunities warrants it. Finally, if the part would be replaced under the "greedy" version, we calculate C_L', the expected cost of replacing the part at a later time given that it has survived to this time and that the car as a whole is subject to a failure distribution.

To understand and formulate a decision rule, consider first the costs if another opportunity presents itself. If we assume that the salvage value associated with a replaced part is the same regardless of when the part is replaced[5], then the cost of an opportunistic replacement, C_{OR} will be the same as the cost of replacing the part now, except that it will be discounted back to the present,

$$C_{OR} = \left(\frac{1}{(1+i)^{t_{oc}-t_n}} \right) C_P \qquad (6.15)$$

where t_{oc} is the time at which the new opportunity arises. For all $t_{oc} > t_n$, $C_{OR} < C_P$, since we defer an identical payment to a future time.

As before, we can solve for the expected costs of leaving the part in service, which we will designate as C_L'. These are given by

---

[5] This is not an unreasonable assumption given that there is no market for used parts other than for scrap or remanufacture at the present time. If companies were to begin to remove unfailed parts before reaching the condemnation limits set by the A.A.R. or the F.R.A., we might expect that a market for such parts would come into being, and railroads would consider "cascading" used parts in a manner similar to that for used rail (see, for example, Martland, et. al. (1990)).
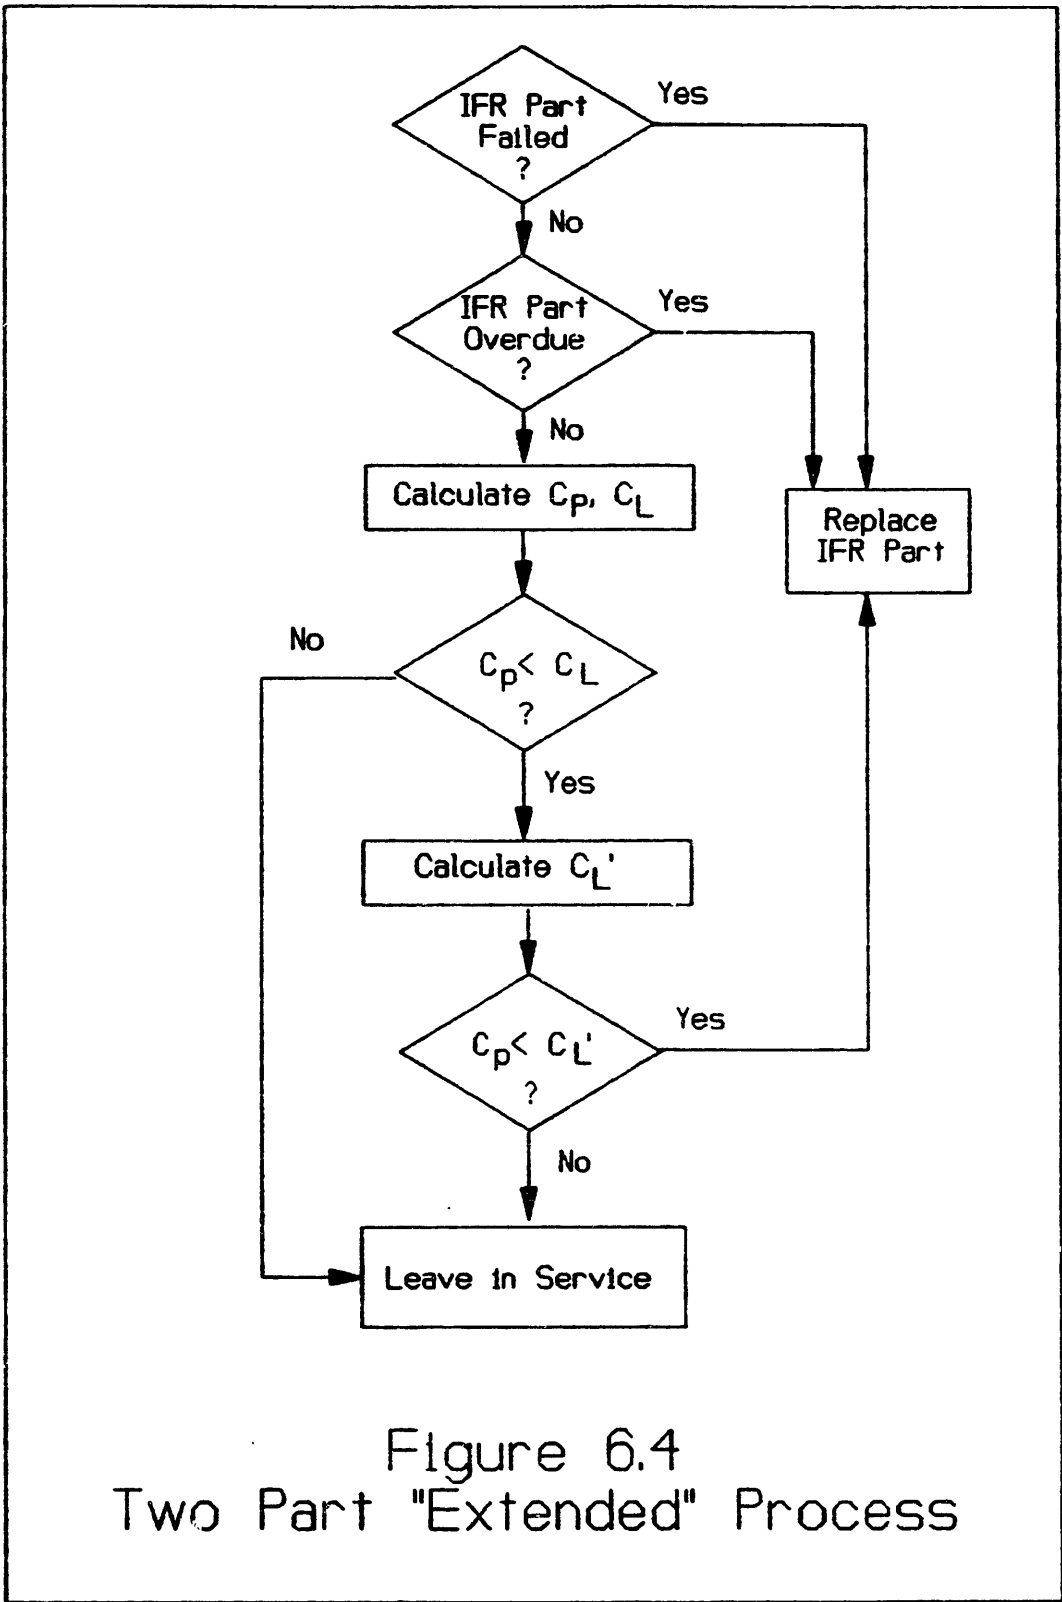
Figure 6.4
Two Part "Extended" Process

$$C_L' = \left(\frac{1}{(1+i)^{t_x''-t_n}}\right)\left[C_x'' + \frac{C_x}{(1+i)^{t_x}-1}\right] \qquad (6.16)$$

where t_x'' is the expected time of replacing the current part given that it has survived to this time and the possibility of another opportunity, and C_x'' is the expected cost of that replacement.

To find t_x'', we consider that the probability that we will replace a part prior to time T corresponds to the probability that the part fails plus the probability that some other part of the system fails (creating an opportunity). If we designate the event of failure of the IFR part as event I, and the event of some other part failing as event R, then we have

$$P(I \cup R) = P(I) + P(R) - P(I \cap R) \qquad (6.17)$$

The probability that we will replace the IFR part prior to some time τ we will designate by

$$F_R(\tau) = P\{t \le \tau\} \qquad (6.18)$$

which corresponds to

$$F_R(\tau) = F_I(\tau) + F_O(\tau) - F_I(\tau) F_O(\tau) \qquad (6.19)$$

In the case of the IFR part, we are interested in the conditional probability given survival to time t_n, while for the "other" part (or the system as a whole), the unconditional probability must be used. This is a result of two things; the fact that we are on a repair track assessing the decision of replacing the IFR part indicates that some other part has failed (or been scheduled in, which constitutes the same effect). Also, we will assume the system as a whole is exponentially distributed, which is consistent with both theory and the empirical work cited in Chapter 2. This means that the system is "memoryless", and makes conditioning inappropriate. Finally, the argument of the failure distribution

of the system as a whole must be adjusted by the probability that the car is on-line at the time a failure occurs, since only on-line failures create opportunities. These various assumptions mean that

$$F_R(\tau \mid t_n) = F_I(\tau \mid t_n) + F_O(\tau) - F_I(\tau \mid t_n) F_O(\tau) \qquad (6.20)$$

The expected time of replacement, t_x'', is given by

$$t_x'' = t_n + \int_{t_n}^{T} [1 - F_R(t \mid t_n)] \, dt \qquad (6.21)$$

$$= t_n + \int_{t_n}^{T} \left[1 - \frac{F_I(t)}{1 - F_I(t_n)} - F_O(t - t_n) + \frac{F_I(t) F_O(t - t_n)}{1 - F_I(t_n)} \right] dt \qquad (6.22)$$

$$= t_n + \frac{\int_{t_n}^{T} [1 - F_I(t)] \, dt}{1 - F_I(t_n)} - \int_{0}^{T - t_n} F_O(t) \, dt + \frac{\int_{t_n}^{T} F_I(t) F_O(t - t_n) \, dt}{1 - F_I(t_n)} \qquad (6.23)$$

The sum of first integral on the right hand side of equation 6.23 and t_n is t_x', so that

$$t_x'' = t_x' - \int_{0}^{T - t_n} F_O(t) \, dt + \frac{\int_{t_n}^{T} F_I(t) F_O(t - t_n) \, dt}{1 - F_I(t_n)} \qquad (6.24)$$

There is an important assumption implicit in this approach; if we are given any future opportunities for replacement in the interval (t_n, T), then we will take the first of them. More precisely, the cost of a present replacement is compared with the expected cost associated with taking the next opportunity for a replacement. A consequence of this assumption is that in some circumstances where there are very many future opportunities, we will select to replace at the present time even though one of the later ones may be

better. Such cases are not believed to be likely, since, as will be shown, we are concerned with both the timing of the next event and its probability. In cases where there are many future random events expected (which might lead us to consider the next rather than the best event), the probability of that next random event is likely to be high. As will be seen in the following section, this high probability of a future lower cost opportunity should lead the car owner to defer in the hope of capturing the savings.

To estimate C_x'', the expected cost of replacing the part currently in use at some later time (including the cost of a later opportunistic replacement if it arises), we again sum the costs of the possible outcomes weighted by the probability of that outcome. There are three possible outcomes:

1. Failure of the IFR part, with cost C_F, which we evaluated earlier;

2. Failure of some other part and subsequent replacement of the IFR part at that time, at cost C_o;

3. No failures in the interval (t_n, T), with the IFR part replaced at the scheduled cost, C_S.

Associated with these, we need to estimate three probabilities:

1. the probability that the first failure is of the IFR part;

2. the probability that th ~ - failure is of some other part (and occurs on-line);

3. the probability that no relevant failure occurs in the interval.

Beginning with the last of these, the probability that there is no failure in the interval is given by

$$P(n) = [1 - P(i)] \cdot [1 - P(r)] \tag{6.25}$$

$$= \left[1 - \left(\frac{(F_I(T) - F_I(t_n))}{1 - F_I(t_n)}\right)\right] \cdot \left[1 - (F_o(T - t_n))\right] \tag{6.26}$$

where $P(n)$ is the probability of no failures in the interval, $P(i)$ is the probability that the IFR part fails, and $P(r)$ is the probability that some other part fails. Notice that no conditioning is required for the rest of the system because of the assumption of an

exponential distribution.

The remainding of the probability must be allocated over the two types of failures. This could be done by integrating over the density function of each type of failure multiplied by one minus the probability of the other type of failure. A reasonable approximation, however, is to allocate the probability of failure according to the ratio of the distribution functions evaluated over the interval so that

$$P(i) = \left[\frac{\left(\frac{F_I(T) - F_I(t_a)}{(1 - F_I(t_a))}\right)}{\left(\frac{(F_I(T) - F_I(t_a))}{(1 - F_I(t_a))}\right) + (F_O(T-t_a))} \right] [1 - P(n)] \qquad (6.27)$$

and

$$P(r) = 1 - (P(n) + P(i)) \qquad (6.28)$$

Thus we now have the cost and the probability of each of the possible outcomes. These can be brought together to find C_x'', the expected cost of replacing the part in use at time t_x''. The result is

$$C_x'' = C_R P(r) \left(\frac{1}{(1+i)^{t_r - t_x''}}\right) + C_F P(i) \left(\frac{1}{(1+i)^{t_F'' - t_x''}}\right) + C_S i \ (n) \left(\frac{1}{(1+i)^{T - t_x''}}\right) \qquad (6.29)$$

A comment on the point of discounting for the various failures is in order. In principle, we should discount each of the costs from the time it is expected to be incurred, over to time t_x'', from which we discount back to the present time, t_n. As a practical matter, however, calculating t_F'' and t_r, the expected time of the first IFR and "random" failure is more complicated than is probably appropriate for a heuristic. Instead, we can calculate the "joint balance point", which we will designate t_j, using the same method as for calculating t_F, equation 6.12. This is the point which combines the centers of gravity

163

of the two probability masses associated with the failure densities for the IFR part and the rest of the system, and balances that associated with scheduled replacements. This is given by

$$t_j = t_x'' - \left[\frac{(1 - F_R(T \mid t_a))(T - t_x'')}{F_R(T \mid t_a)} \right] \tag{6.30}$$

Using t_j for both types of failures, we can reduce equation 6.29 to

$$C_x'' = [C_a P(r) + C_F P(i)]\left(\frac{1}{(1+i)^{t_j - t_x''}}\right) + C_s P(n)\left(\frac{1}{(1+i)^{T - t_x''}}\right) \tag{6.31}$$

These terms can be input into equation 6.16 to find C_L' the present value of the expected stream of costs associated with leaving the part in service. Once again, if the cost of replacing the part now, C_P is less than C_L', the part is replaced at this time. If it is greater, then the part is left in service. If the two costs are equal, then we are indifferent between removing the part or allowing it to remain in service.

6.3.2.a. A Numerical Example

We can continue the numerical example from before, with the additional assumption that system failures are exponentially distributed with mean frequency of 50,000 miles. Again consider the case where t_a is 180,000 miles. Recall that under the "greedy" version we replaced the part since the expected cost associated with removing the part now, C_P is less than C_L, the expected cost of replacing the part later.

The first thing that must be done is to solve for t_x'', the expected time of the next replacement when system reliability is included. Using equation 6.24, and numerical integration where needed (the exponential term has a closed form), we find (using units of thousands of miles)

$$t_x'' = 242 - 23.3 + (9.8/.79) = 230,000 \text{ miles.}$$

To find C_x'', we use equation 6.31; first, however, we need to find the probabilities of

the various possible outcomes and t_j, the joint balance point for the failure outcomes. $P(n)$, the probability of no failure is found by using equation 6.26 to be

$$P(n) = (1 - .43) (1 - .537) = .264$$

The remaining probability, 1 - $P(n)$, is allocated over the two possible failure types using equation 6.27,

$$P(i) = [.43/((.43)+(.537))] (1 - .264) = .327$$

$$P(r) = .409.$$

The point from which to discount failures, t_j, is found using equation 6.30 as

$$t_j = 230 - [(.264)(257-230)/.736] = 220,000 \text{ miles.}$$

Substituting into equation 6.31,

$$C_x'' = [(700)(.409)+(1325)(.327)](1/(1.0015)^{-10}) + (850)(.264)(1/(1.0015)^{27})$$

$$= \$945.$$

Using equation 6.16, C_L' is found to be

$$C_L' = [1/(1.0015)^{50}] [945 + 2846] = \$3514.$$

Recall that $C_p = \$3546$. Since the cost of replacing at the present time is greater than the expected cost of leaving the part in service, we do not replace the part at this time. In other words, the information regarding the overall reliability for the car has led us to a different decision than when that information is excluded. In particular, we follow a lower cost alternative.

6.4 Multicomponent Models

In this section, the heuristic is augmented to include multiple components, with potential economies in repairs among the various components. The basic approach is first to determine which components can be replaced profitably at the present time based on general economies of scale such as switching or disruption costs (which we will designate as first-order effects). Those components which are not found to be economical to replace based on first order effects are then examined to see if there are additional joint economies which accrue as a result of components already being repaired or replaced (second order effects). An example of a first order effect might be the decision to remove a wheelset on axle 1 because it is very near the condemnation limit. A second order

165

effect would be to replace a less worn wheelset on axle 2 based on the knowledge that the car must be jacked to remove wheelset 1.

The implementation of both the "greedy" and "extended" versions of the heuristic are essentially the same. In each case, all the parts are first examined individually; higher order effects are then examined for each part which is slated for leaving in service. The primary difference between the "greedy" and "extended" versions are in the degree to which the underlying assumptions are challenged by this piece-by-piece approach. In the greedy case, the notion of treating each part more or less separately is consistent whether or not the system has many components. In the "extended" version, however, the failure to consider the effects of future expected failures among the various components highlights that at least some future "better" decisions are foregone. This is a result of the compromise between tractability and efficiency in the use of a heuristic.

As regards terminology, most of the definitions of the preceding sections are left intact, with the addition of a subscripted letter "m", to denote a particular component among the M components which make up the system. Further needed definitions are added in the implementation section.

6.4.1. Multicomponent Models: Implementation

Consider a car as composed of M IFR components and one exponentially distributed component representing "the rest of the system". The "rest of the system" component can be thought of as composed of an array of M values, with each element representing the reliability of the system as a whole excluding one part, which is the item being considered at a given time. For each of the IFR parts, there are associated a set of costs for on-line and off-line replacements and failures, including C_{mF}, C_{mn}, C_{mx}, and C_S, corresponding to the costs of failure, out-of-pocket cost of a present replacement, expected cost replacement over the life of a new component, and the scheduled cost of replacing the m-th component. Similarly, there are a set of times associated with each component, such as T_m, and t_{mx} representing the scheduled replacement time and the expected time to first replacement for a new component given that survivors are replaced at T_m. Each component also has an associated failure distribution, $F_m(t)$. In addition to these, we can denote some variables which are common to the car as a whole, such as

166

the percent of time the car is off-line, P(O), the annual mileage rate, the switching costs, and the discount rate. In short, these are the same variables used in the two component model, subscripted appropriately for use in a multicomponent setting.

To these definitions, we can add the following:

J: an M x M symmetric matrix of the cost savings associated with joint repairs for each of the M IFR components. That is, $J(i,j)$ is the savings in costs associated with performing a replacement of parts i and j at the same time (such as wheelsets 1 and 2 described above). These savings can come from a number of sources, such as increased efficiencies in the repair process, more efficient inventory management, and so on. For most components, however, one might expect $J(i,j) = 0$, since the economies consist merely in switching the car to the repair track, and that has been separated out through the term S.

X: an M x 1 array which indicates whether or not a particular component is to be replaced at the present time. Each element is initially set to 0, indicating that no components are to be replaced.

The decision process for replacing components follows the pattern shown in Figure 6.5. Each component is first examined in terms of its suitability for replacing based on first order effects. After all the components have been examined for these, those which are to be replaced are compared to others with which they share joint economies.

The actual decision rules are as follows. For each component, k:

(1) If a component k has failed, then set $X(k) = 1$, and replace it;

(2) If a component k has been in service for a period longer than T_k, then replace k and set $X(k) = 1$;

(3) Calculate C_{kP} and C_{kL}. If $C_{kP} < C_{kL}$ and the "greedy" version is being followed, then set $X(k) = 1$ and replace k. If the "extended" version is being used, then calculate C_{kL}' and if $C_{kP} < C_{kL}'$ then set $X(k) = 1$ and replace k.

After each component has been examined for first order effects, if X includes any non-zero elements, the car is examined for additional economies. These second order effects are economies that proceed from the fact that some maintenance activity is being performed. This process is as follows:

(1) Create X', a copy of X after the first order analysis. For each non-zero

For k = 1 to Number of Parts DO.

Part k Failed ?

Yes

No

Part k "Overdue" ?

Yes

No

$C_{kp} < C_{kL}$?

Yes

No

K <= Number of Parts

Yes

No

k := k+1

Replace Part k

Test for Second Order Effects

$C_{kp} - J(1,k) < C_{kL}$

Yes

No

Replace Part k

Next K

Test for Higher Order Effects

Stop
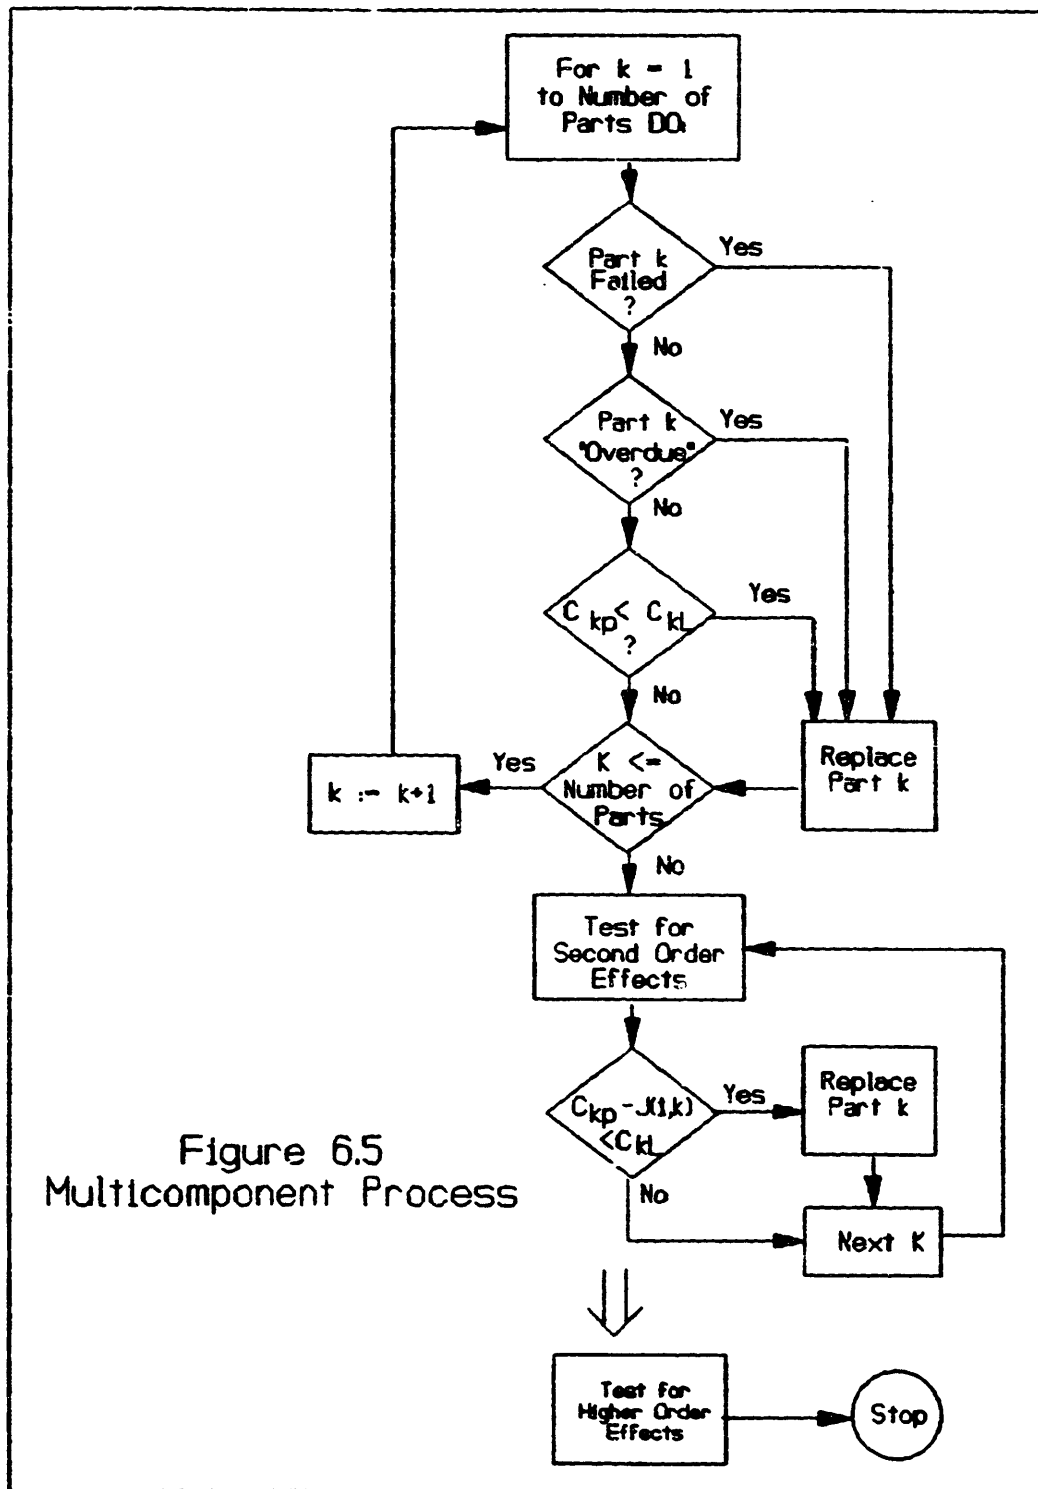
Figure 6.5
Multicomponent Process

element in X', denoted k, the corresponding elements in row $J(i,k)$, i = {1,2,3,...,M}, i $\neq$ k, are examined. If $J(i,k) > 0$, then if $X(i) = 0$, C_{iP} is reduced by the $J(i,k)$ and compared again to C_{iL} (or C_{iL}', in the "extended" version). If the present costs are now lower than the later costs, then $X(i) = 1$, and the part is replaced.

(2) If after examining all the non-zero elements in X', $X = X'$, then the process is completed. If not, then the process may be repeated to exhaust "third order" effects created by the previous pass. As a practical matter, if the number of non-zero elements in J is large, it may be desirable to limit the search to only second or third order effects. If, on the other hand, J is sparse, it may be possible to exhaust all the economies.

Several comments on the heuristic are in order. The first is that the use of part-by-part collecting of the economies means that some potential economies are bypassed. Consider the case where there are three parts which are subject to large economies if all three are replaced, but are subject to no special savings if any two alone are replaced. Under the heuristic these economies will not generally be realized. A second concern is that while the "extended" approach makes some use of the system reliability information, it does not make complete use of the available failure rate data. In particular, it may be that by waiting until a later failure of another component a substantial reduction in costs will become available. The probability of this event could be predicted with the available data and the cost effectiveness of waiting could be analyzed. The decision not to do so was based primarily on the need to limit the complexity of the system. This may be a fertile area for fu'ure research.

6.5. Incomplete Information

As has been indicated in earlier chapters, one of the factors which makes it difficult to manage the maintenance of railroad cars is incomplete information. In particular, when a car returns from an interline movement, the state of the car, and the car's maintenance history, are not completely known by the car's owner. The current rules regarding the reporting and billing of repairs require only that the bill be submitted within one year of the date of repair. Examination of Car Repair Billing data by the author found that most repairs are reported within 60-90 days, although some are not

reported for 180 or more days. The owner is thus faced with an additional source of uncertainty in planning car maintenance: the car which is being considered for either opportunistic or scheduled maintenance may have been repaired and the repair not yet reported.

To resolve the problem of incomplete information, it is necessary either to inspect the part, if possible, or to include in the cost relationships the possibility that the component has been replaced due to an as yet unreported in-service failure while off-line. For example, in the case of wheels, which are readily inspected, it is possible for repair personnel easily to detect that a wheelset which is being considered for preemptive maintenance is virtually new, and, at least estimate the time of the last replacement of the component. Unfortunately, many components cannot be inspected, or their true condition cannot be determined from an inexpensive inspection. The only solution in such cases is to add an additional term into the various cost equations to represent the probability that the component has been replaced and has not been reported.

What is needed is the probability that a part has failed in the previous year, given that it has not been reported to have failed, $P(F|NR)$. What is known (or can be readily estimated from most railroad car owners' information systems) is the probability that the component has not been reported given that it has failed, $P(NR|F)$, and the probability that it has failed in the previous 12 months, given that it had not failed prior to that time, $P(F)$. (We assume throughout that the car owner has essentially perfect information regarding the state of the component as of one year before, since repair bills must be submitted within a year.)

Since car repair bills are issued monthly through the A.A.R. to all participants in the CRB system, it is reasonable to use a discrete evaluation of $P(F|NR)$, so that

$$P(F|NR) = \sum_{i=1}^{12} P(F_i|NR_i) \qquad (6.32)$$

and

$$P(F_i \mid NR_i) = \frac{P(F_i \cap NR_i)}{P(NR_i)} \tag{6.33}$$

$$= \frac{P(NR_i \mid F_i)\, P(F_i)}{P(NR_i \mid F_i)\, P(F_i) + (1 - P(F_i))} \tag{6.34}$$

which are terms that can be evaluated with the available information. $P(NR_i \mid F_i)$ can be constructed from a table which links the typical reporting times by railroads on which the car owner's equipment is operated. The probability of a failure in each of the preceding months can be estimated using:

$$P(F_i) = \frac{F_0(t_a - A) - F_0(t_a - B)}{1 - F_0(t_a - ann)} \tag{6.35}$$

where A is the time (in miles) at the end of the month, B is the time (in miles) at the start of the month, and ann is the annual mileage of the car. When finally arrived at, $P(F_i)$ must be weigh·.d by the probability that the car was off-line.

Unreported repairs have no effect on C_P, the stream of costs associated with replacing a component at the present time, since the removed component has already been bought even if not yet paid for; i.e., it is a sunk cost. (If there were a market for used parts, this would not be the case, since the salvage value of the part would differ depending on the age or usage of the part). What is affected, under the heuristic described in this chapter, is the expected cost of replacing the part later, C_L, or C_L'. The possibility that the component has been replaced recently lowers the expected cost of a later replacement by potentially deferring that component's failure far into the future.

To assess the importance of the unreported repairs, the numerical example for the greedy heuristic was evaluated at 180,000 miles for the given failure distributions, annual mileage, and assuming a simple function to describe $P(NR \mid F)$. In particular, it was assumed that in the immediately prior month all repairs are still unreported, and that

P(NR|F) declines by half in each month. In other words, in the second previous month 50% of the repairs are unreported, 25% in the third previous, etc. This results in about 5% of all repairs reported in the period from 180 days to 1 year, and may be a reasonable approximation to the actual reportings. Applying the above formulas 6.12 - 6.15, P(F|NR) for the year prior to 180,000 miles is about 3.8%.

Applying this to the evaluation of C_L, there is approximately a 96% probability that the cost of a later replacement will be as estimated in the numerical example in section 6.3.1.a., and a 4% possibility that the costs will be less. Since most of the probability mass associated with P(NR|F) is concentrated in recent months, a reasonable assumption is that if an unreported replacement occurred in the past year it was most likely in those months. Assuming that the failure occurred three months ago, C_L evaluates to $3532 (a decrease of about .7%).

Inclusion of this information is not without its consequences. Notice that the previous decision to replace the part at 180,000 miles now becomes a decision to leave the part in service. This is, of course, something of an artifact of the particular time chosen. The point at which we are indifferent about replacing the parts appears to move from about 174,000 miles to about 182,000 miles, a change of about 4%. The extent to which this difference would affect the overall performance of the heuristic is left to future research.

6.6. Conclusions

In previous chapters, we determined that opportunistic maintenance seemed particularly appropriate for freight cars. Unfortunately, our review of such policies in Chapter 2 found that there were few practical implementations of such policies available, and none of them appeared to be appropriate to the unique cost structure faced by railroad car owners.

To resolve this, a new heuristic has been proposed in this chapter. The basic implementation, a "greedy" version, calculates whether it is economical to replace a component at the present time, given that you already have the car on the repair track, or wait for its scheduled time. The "extended" version of the heuristic considers the

possibility that the car may be available for another opportunity in the interval between the scheduled time and the present. Both versions use information that railroads generally have (or could have) available, and uses computational methods which are readily available to computer programmers. In the next chapter, these policies are compared with the present practices.

Chapter 7

Simulation of Freight Car Maintenance Policies:
Introduction and Overview

7.1. Introduction

In this and the following chapter, the maintenance policies employed by railroad car owners and some alternatives, including the heuristic presented in Chapter 6, are evaluated using a simulation model. This chapter explains the simulation model, the particular policies which were modelled, and an approach to evaluating the model's results and choosing among maintenance policies. The numerical results of the simulation are reported and discussed in Chapter 8.

Before examining the model itself, it may be useful to discuss why a simulation model is needed at all. If all the policies took the form of the "on condition" policies, a simulation would not be needed. In that case, one could simply calculate the expected number of failures (both on and off line) in an interval representing the life of the car, and compute the associated costs[1]. In the case of the opportunistic policies, such an approach is simply not feasible. Under these policies, the failure of any component (or the survival of a component to its schedule replacement time) creates an opportunity for each of the other components. This makes the calculation of the number of times any component will be replaced dependent on the life distributions of all the other components. An analytic solution to this highly complex probabilistic problem was not forthcoming. In this situation, the classic method available to the analyst, as a technique of last resort, is an event-structured probabilistic simulation. This chapter reports on such a model and the results obtained from it[2].

---

[1] Actually, even if all the policies lent themselves to closed form estimates of the number of replacements, a simulation might be necessary in order to evaluate the costs incurred, since policies differ not only in the number of parts replaced but in the timing of maintenance activities. With positive discount rates this can be a significant factor.

[2] Appendix B presents a technical description of the model for programmers familiar with object oriented Pascal.

The performance of the policies is measured quantitatively in terms of the three criteria presented in Chapter 5, namely. miles per maintenance event, miles per in service failure, and cost per mile. Particular emphasis is placed on the second and third of these, primarily as a reflection of the importance that most managers place on service reliability and cost control, and to a lesser degree because those measures highlight differences in the effects of the various policies. In addition to these indices, the sensitivity of the policies to various factors is explored. Clearly, a robust policy is generally preferable to one whose performance varies wildly depending on factors outside the control of the car owner or maintainer.

7.2. The Simulation Model

An event-based simulation model was developed which permits the analysis of car performance under a maintenance policy by monitoring the impact on the cost and service reliability of a large number of the car's components (including one which is an exponentially distributed "rest of the car" component[3]). The car (and its components) is operated for an extended period of time (2 million miles in the trials presented here), and parts are repaired and replaced in accordance with the specified maintenance policies.

The model requires information regarding failure distributions of the components to be modelled, the costs of repairs and maintenance activities (including any costs of in-service failures), and an explicit formulation of a policy. The last item, a formal statement of the decision rules employed by a policy, requires some judgement by the modeler. Unambiguous values must be attached to such items as the condition of a part and the likelihood of failure in an interval. In the railroad environment, these matters often turn on the expertise of maintenance workers and supervisors. In a model, a

---

[3] In principle, if all components are input into the model, the "rest of car" component becomes unnecessary. As a practical matter, however, this component is useful for representing all the parts which are not explicitly included in the maintenance program (for example because they are either too short-lived or inexpensive to be worth the cost of monitoring). The appropriateness of an exponential distribution to model these was shown in earlier chapters.

method for arriving at a numerical form must be developed. To address the many policy options available, a wide range of alternative formulations were examined. The simulation was designed and built with the express intention of comparing the effects of the policies, and attention was paid to insuring that each policy faced the same circumstances as other policies for a given trial or run. Each policy also faces the same set of "random" failure events so that no one policy is "luckier" than another.

The model allows the user to specify the number of components to be evaluated and their characteristics, and the nature of the car itself. Components are characterized by the following:

- a part name (wheel 1, for example);
- a failure distribution, which is represented by a 2-parameter Weibull distribution;
- a set of costs, including costs for on and off-line in-service failures, costs for scheduled replacements, and material and labor costs for scheduled or opportunistic replacements;
- an optimal replacement interval under a single component age replacement policy.

The optimal replacement interval is the "time" at which a part would be replaced if the component were a single, "stand alone" unit under the standard age replacement methods discussed in Chapter 2. (It corresponds to the variable T in the presentation of the heuristics in Chapter 6.) This was calculated using a Lotus 1-2-3 spreadsheet implementing the combined algebraic and graphical method described in Jorgenson, McCall, and Radner (1967).

As mentioned earlier, one of the components included, the "rest of car" component serves to indicate the overall reliability of the car, and is modelled with a Weibull shape parameter of 1. That is, failures are exponentially distributed with a given characteristic life. By changing the characteristic life, the car is subject to more or fewer "random" failures.

The car itself (and its usage) is described by the following fields:

- a car identifier (i.e., initial and number);
- car type;
- switching cost, i.e., the cost of sending the car to a repair track;

- annual mileage for the car;
- the discount rate.

The annual mileage and the discount rate are related to each other in several ways. The discount rate specifies the return available from competing uses of money, and is measured in terms of time (e.g., per annum). Since the model relates usage and wear in terms of mileage, the discount rate is converted to a per mile basis. Thus the higher the annual mileage, the lower the corresponding interest rate per mile. In the opportunistic maintenance heuristics, the interest rate per mile is used in determining the present value of the expected cost of leaving a part in service and of replacing the part now. In all the strategies, an increase in the annual mileage also has the effect of raising the cost per mile, since failures (or replacements under a planned maintenance program) will occur sooner in time than under a lower annual mileage. That is, the costs of a repair will not be discounted from as far a point *in time* under the high maintenance scenario since the mileage is reached sooner in time.

With these inputs, the sim ulation creates a vector of failure times for each of the components, using a pseudo-random number generator and the inverse of the failure distribution. These failure times represent the time to failure for each of the items in what is, in effect, an inventory of parts from which the model draws when a replacement is warranted. This vector of failure times represents an upper bound on the lifetime of components. A policy may, of course, result in a component being replaced sooner than the "predetermined" lifetime, but it cannot extend the life of a particular part beyond it. Because the random number generator can be explicitly "seeded", the same inventory of parts to be used can be generated for runs which apply various policies.

To clarify this approach, consider a simple example[4] of a single part and two policies, an on-condition one which allows the components to run until failure and another which replaces the part when it is 100 units old, or when it fails in use. If the predetermined lives of the part are given by the vector [120, 85, 125, 200], then the on condition policy will encounter an in-service failure at time 120, 205 (i.e., 120+85), 330,

----

[4] All the numbers in the following section are in units of thousands of miles.

177

and 530. The second policy will replace the part at 100, experience an in-service failure at 185, replace the part again at 285, and replace it again at 385. The effect is that a part which has a predetermined life that is shorter than the scheduled replacement interval will fail, and parts which have longer predetermined lives than the interval are replaced at the end of the scheduled interval. In this way each policy can be compared with the other policies without concern that one of the policies was somehow "favored" with better parts.

The simulation model was written in Turbo Pascal 5.5, an object-oriented implementation of that language, and runs on any IBM-compatible personal computer. Figure 7.1 is a flow chart showing the general structure of the simulation.

Results of the runs for each strategy can be compared using Wilcoxon's signed rank test[5]. This test estimates whether matched sets of data come from the same larger population; i.e., whether the results of sets of runs are different in a statistically significant way. Part of the appeal of the test is that sample sizes can be relatively small if the samples are well matched, as is the case in the simulation model. For each scenario, the model runs through 10 matched trials for each of the policies and then outputs the results. As will be seen in the results section below, this is a sufficient number of runs to discern statistically significantly different results at the 1% level.

The Wilcoxon signed rank test measures whether a set of items (such as the results of runs of a simulation) is larger or smaller than another set in a statistically significant way. While it incorporates information regarding the relative magnitude of differences, it does not measure the actual magnitude of the differences. This means that the results are compared in a pairwise manner, i.e., how two policies perform relative to one another. As the number of alternatives considered increases from n to n+1, the number of pairwise comparisons also increases by n+1. (I.e., an increase from 7 alternatives to 8 increases the number of pairwise comparisons by 8 as well.) Since we are interested in how a large number of policies perform over three different measures, a computer routine was developed which compares the results of each policy with those of all alternatives.

For the purposes of the simulation, these pairwise comparisons across a set of

---

[5] See, for example, Mosteller, F. and R.E.K. Rourke (1973).

FIGURE 7.1
Model General Structure

Input Data

Car Data
Parts Data
Cost Data

Create Part
Inventories

For Each Policy DO:

While Car Miles < 2,000,000
REPEAT:

Enqueue Next Event

(Either Random Failure,
Part Failure, or
Scheduled Event)

Get Next Event
And Process in
Accordance with Policy

Update Parts, and
Perform Accounting
of Events, etc.

Output Results
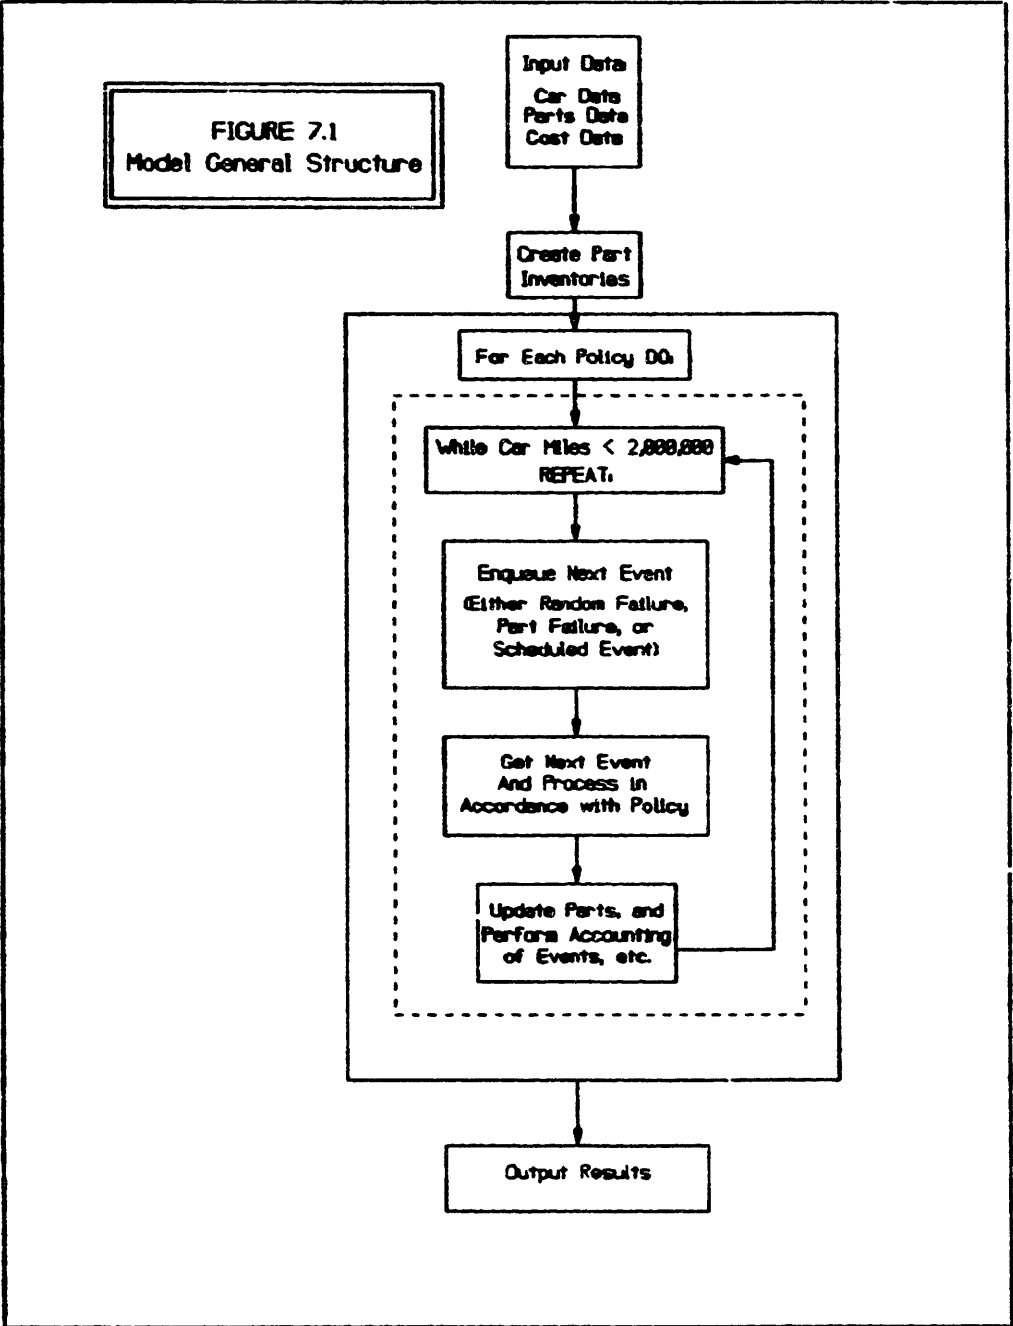
performance measures are quite acceptable, since we are primarily interested in the relative merit of policies. As a practical matter, however, railroad car maintenance managers are likely to be concerned with a number of additional factors, such as ease of implementation, or company philosophy. It is beyond the scope of this thesis to consider all the factors that might be faced in a particular situation.

In a simulation, the number and length of runs is an important decision. In this case, the decision of making 10 runs of 2 million car miles each was motivated by several factors. A run length of 2,000,000 miles represents more than double the "typical" lifetime of a car in the author's experience. For a car which is being used for 100,000 miles per year, this represents a lifetime of 20 years. (Note also that the average mileage of the Class I railroads' cars in 1988 was approximately 20,000 miles and the average age was slightly over 17 years [AAR(b, 1988)]). The use of a longer than normal lifetime has the effect of minimizing any "end effects", i.e., periods of the car's life where maintenance decisions would be dominated by the choice of whether or not to retire the equipment. Another concern is the "warm up" period, or initial conditions. In this case, the initial conditions are reasonably modeled by beginning all the parts at age 0, since cars are purchased as new systems with all new parts, and our interests are the typical costs and reliability over the life of the car.

On the matter of number of runs, the decision to limit the number of runs reflects both the desire to economize on time and the requirements for statistically significant results. As indicated, the primary intent of the model is to permit the comparison of one policy with another, in terms of cost and reliability. Given that no attempt has been made to include all the components which make up a freight car or to achieve complete accuracy in the costs at particular facilities, the particular means and variances of the runs for each policy are not critical outputs of the model (as opposed to the relative values). That is, the statistically significant results, using the Wilcoxon signed rank test, indicate that one policy is or is not higher (or lower) in some measure than another one, not necessarily the extent of the difference. The means and standard deviations for the three measures are presented, but most of our attention will focus on how the policies compare with each other in terms of the three previously discussed measures. The number of runs

was selected to insure that it was possible for a policy to outperform another in any measure at the 1% level of significance, a result which, in fact occurred a number of times.

As will be discussed below, an important criterion for evaluating alternatives is the sensitivity of a policy to a change in circumstances. In particular, if a policy performs very well with respect to some measure under one scenario, but performs poorly under another scenario, the policy may not be as desirable as one which performs well under both, especially if one is uncertain about which circumstances will be faced. The selection of a wide range of scenarios permits us to gain useful insights into which policies possess this important quality.

The limited number of runs (and the particular set of distributions used) have a practical implication for managers interested in applying the maintenance policies to their fleets. The actual magnitude of the differences (either savings or changes in reliability) which can be expected for a company's fleet will depend on the components in the cars and the actual cost structure faced by the car owner. They would likely be different than those presented here. If one wished to estimate accurately the mean values of the costs and reliability measures for the purpose of comparing alternative policies, more runs of the model using the appropriate parameters for the components of interest would be called for. The current input data, as explained in the following section, were chosen from a number of sources and do not reflect an actual car series.

7.2.1. The Components Analyzed

The components chosen were relatively common parts for which failure data ' .J been previously estimated and presented in A.A.R. studies or developed in analyses by the author. Eight components were included in the model: each car had 4 wheelsets, 2 end-of-car cushioning devices, and 2 trucks. In addition, an exponentially distributed component ("the rest of the car") was included. The failure distributions for the wheelsets and trucks were those presented in Guins and Kyparisis (undated), and are actually components from a series of auto racks they studied. The distributions for the end-of-car cushioning devices are from a study privately sponsored by a railroad, consisting mostly of boxcars. The costs used are from the same sources, although the wheel prices also

181

reflect data gained while analyzing wheel failures in support of an A.A.R. research project [Little and Martland (1989)]. The base case distributions and cost data are presented in Table 7.1.

In order to test the sensitivity of the policies to the particular distributions, two alternative sets of distributions were also modelled. In particular, the characteristic life (i.e., Weibull scale parameter), or point at which 63% of the items fail, for wheels was shortened from 274,250 miles to 90,000 and 185,000 miles. Wheels represent half of the (non-random) components included in the simulation, so that a substantial altering of the distributions used for them reflects a fundamental change in the maintenance plan.

In the first instance, the use of a very short characteristic life (90,000 miles) represents a case where half the components are short-lived, and the other half long-lived (notice in Table 7.1 that the characteristic lives of end of car units and trucks are 350,000 and 564,000 miles respectively). One can think of the wheels in this case as a "proxy" for some set of other, shorter lived components (such as brake parts, or door fittings).

The second case, where the characteristic life of the wheels is 185,000 miles, reflects a concern expressed by some industry officials that some wheels are not so long lived as those used as the base case. This distribution examines how sensitive the policies are to modest changes in the failure distributions of the components. As will be seen in Chapter 8, some of the policies are quite sensitive to the failure distributions of the components.

For the base case, the following assumptions were made regarding the characteristics of the car:

- the cars were assumed to spend 50% of their time off-line;
- the cost of switching was assumed to be $150;
- the mean life of the exponentially distributed "random" component was assumed to be 50,000 miles;
- the annual mileage of the cars was assumed to be 66,000 miles;
- the discount rate used was 10%.

No second-order effects were calculated. That is, the only opportunities examined were those associated with switching costs. Second order effects, introduced in Chapter 6, are decisions to replace a component because another particular component is being

COMPONENTS	VALUES
WHEELS:	
Weibull Shape Parameter	3.5
Characteristic Life	274,250 miles
Costs:	
Offline Failure Cost	$1400
Online Failure Cost	$1100 + Switching Costs
Scheduled Replacement Cost	$700 + Switching Costs
Opportunistic Replacement Cost	$700
END OF CAR CUSHIONING UNITS	
Weibull Shape Parameter	4.0
Characteristic Life	350,000
Costs:	
Offline Failure Cost	$2000
Online Failure Cost	$1350 + Switching Costs
Scheduled Replacement Cost	$1100 + Switching Costs
Opportunistic Replacement Cost	$1100
TRUCKS:	
Weibull Shape Parameter	2.6
Characteristic Life	564,000 miles
Costs:	
Offline Failure Cost	$3000
Offline Failure Cost	$1850 + Switching Costs
Scheduled Replacement Cost	$1100 + Switching Costs
Opportunistic Replacement Cost	$1100
REST OF CAR:	
Mean Time Between Failures	50,000 miles
Costs	Not Applicable

Table 7.1
Distributions and Costs Used

183

replaced, usually because of some joint economies that these components share (e.g., replacing wheelset 1 because the car is being jacked up to replace wheelset 2). The calculation of second order effects depends on knowledge of the maintenance activities for each part, and creation of a matrix of joint economies. Because accurate cost data to estimate second order effects was not available, the simulation limits itself to first order effects, which accrue due to the presence of the car on the repair track independently of what is actually done to the car. This probably leads to an underestimation of the cost savings under opportunistic maintenance, since some "lower priced" replacement opportunities go undiscovered. These "overlooked" opportunities might also reduce the number of in-service failures. The effect of higher order effects on replacement decisions is a potentially important area for future research.

The selection of the particular values reflects the author's judgement and discussions with managers from railroads and the A.A.R. The switching cost used was conservative and does not include the loss of customer goodwill associated with the removal of a car from service. The annual mileage used was that of a group of cars in coal train service in the A.A.R. Car Maintenance Cost Data Base, which was used to estimate the parameters for "random failures" of the rest of the car. The percent of time spent off-line was chosen arbitrarily and was varied in subsequent runs.

Many of the values used in the base case were adjusted substantially in other scenarios, particularly the off-line percentages, the expected life of the "random component", and the annual mileage of the cars. The cost of switching and the discount rate were felt to be reasonable, and were left constant throughout[6].

The reader will note that the costs of in-service off-line failures in Table 7.1 are two to three times as high as those for scheduled or opportunistic maintenance activities. These costs reflect a number of factors, including the presumed greater efficiency of performing work at the desired location and the ability of on-line maintenance to

---

[6] The adjustment of the annual mileage has the same effect as adjusting the discount rate, since they are both used only in the calculations of the type $1/(1+i)^t$. (See equations 6.3 - 6.9 in the discussion of the heuristics.)

contribute toward fixed costs in maintenance plant and equipment. The costs were considered to be reasonable by several railroad officials, and are in keeping with the costs used by Guins and Kyparisis.

The selection of which values among the failure distributions and the car characteristics to vary merits further comment. One can think of the problem faced by the maintenance planner when deciding on a maintenance policy for a particular car series as consisting of two different types of uncertainties. The first of these concerns what to do about the things under his control. This includes whether or not to undertake a preventive maintenance program, and, if so, which components to include. The second source of uncertainty concerns the effects of matters outside his control on the car series in question. This includes a number of factors, such as the usage of the car (which is determined by markets and by car distributors), the quality of the car's design (which is usually a given for the maintenance planner, although a proper maintenance program will affect this over a period of time), and the firm's discount rate. In modelling several different failure distributions, we are asking about what happens if the maintenance planner chooses a different mix of parts to include in the maintenance program. If a policy is adversely affected by the parts mix, then the maintenance planner must be careful to select the "right" policy at the outset, especially if the costs of changing a policy are high. Choosing the "right" policy the first time may be very difficult, and later decisions to add components (as in the case of a "staged implementation" of a new policy) may undermine the policies effectiveness.

In examining the sensitivity of the policies to factors outside the control of the maintenance planner, such as usage rates or the discount rate, we are asking whether a seemingly "good" policy will become a victim of circumstances (and which circumstances are likely to have negative effects). Alternatively, knowing how a policy responds to various external circumstances may be helpful in deciding between two policies which are generally acceptable if the maintenance planner knows what circumstances the cars will face.

7.3. The Maintenance Policies Modelled

The maintenance policies simulated reflect both current practice and some reasonable alternatives to them. In Chapters 3 and 4 it was argued that railroad car owners generally follow two maintenance policies, "on condition" maintenance and "hard time" maintenance. The first, "on condition" maintenance, consists simply of allowing a car to remain in service until a component either wears to a level prescribed in a standard such as the A.A.R. interchange rules, or fails in service. The component is then replaced and the car is returned to service. Modelling the "on condition" policy is straightforward since the life of the components is known from the parts inventory, and parts are allowed to remain in service for that period of time. The "on condition" policy requires only that the accounting of events and failures be accurate.

The second policy, "hard time" maintenance, provides that a car is brought into a maintenance facility at fixed intervals and components which are believed to be near failure (however defined) are replaced. (The car is also shopped and repaired whenever an in-service failure occurs.) Both the intervals at which a car is brought in for scheduled maintenance and what components are to be replaced are generally considered matters of engineering and managerial judgement. To model them, however, a formal rule is needed. The rule used was that components which will reach their characteristic life during some specified part of the interval between the present time and the next scheduled maintenance activity are replaced. Two versions of this type of policy were modelled. The first, called "near sighted", calls for replacement of components which reach their characteristic life in the first half of the next interval. The other version, called "far sighted", mandates the replacement of a part which reaches its characteristic life at any point in the next interval.

To clarify the difference between the policies, consider again the simple example used to explain the "inventories" (Figure 7.2). Recall that the part inventory was given by the vector [120,85,125,200], and the part had a characteristic life of 110 units. Under a "near-sighted" policy, if the car is brought into the shop every 100 units, then on the first visit (at "time" 100) the part would be replaced, since it's characteristic life (110) plus the time of last replacement (0) is now within 50% of next replacement interval at

186

Figure 7.2
Near- vs. Far-Sighted Policies

Characteristic Life = 110

Characteristic Life = 110

Characteristic Life = 90

Characteristic Life = 110

Near-Sighted

Time

Failure at
185

Replace
(110 < 150)

No Replacement
(295 > 250)

Replace
(295 < 350)

100   200   300

Far-Sighted

Time

Failure at
185

Replace
(110 < 200)

Replace
(295 < 300)

Replace
(310 < 400)

100   200   300
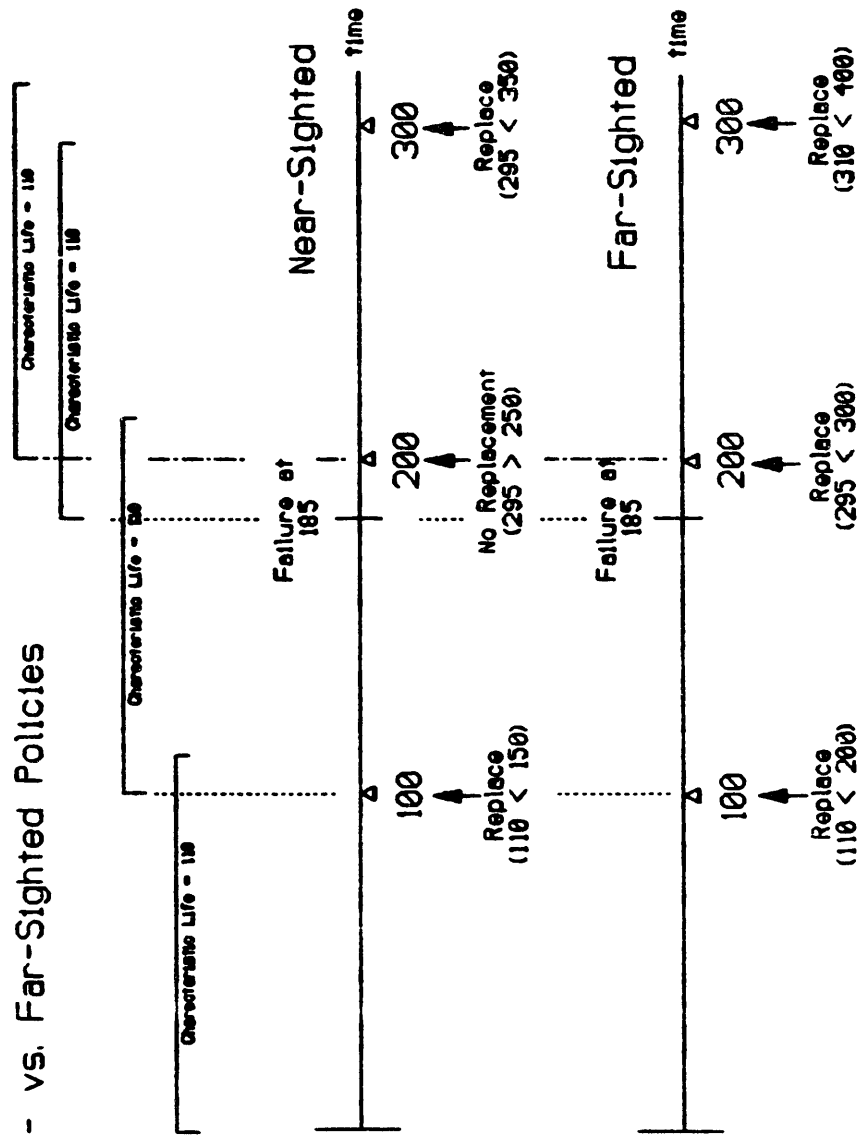
△ = Scheduled Servicing (000's of Miles)

187

time 200 (i.e., 150). In other words, the part is "expected" to fail in the next interval and so is replaced at this time. This part would then fail in service at 185 (because it only had 85 units of useful life). When the car is brought in at 200, the part would not be replaced, since the characteristic life would not be reached by 250, but it would be replaced during the shop visit at 300, since the sum (110 + 185) is less than 350, the midpoint of the next interval.

Under a "far-sighted" policy, the part will also be replaced at the first scheduled shop visit at time 100, and, as in the "near-sighted" case, that part will fail at 185. Unlike the "near-sighted" policy, however, at time 200 the part will be replaced, since it will reach its characteristic life at time 295, which is less than the next scheduled event at time 300. The part will be replaced again at time 300, for the same reasons.

The "near sighted" version appears to better reflect the current practice among railroad car owners than the "far sighted" version. In practice, car repair personnel generally estimate the remaining life in parts without using any formal statistical measures, so that relatively new components (such as that considered at the second scheduled shop visit in the above example) are left in service regardless of the characteristic life of the component. It is a subject for future research to develop more satisfying models for how "good" or "expert" repair personnel judge whether or not a part is "near" failure and the extent to which they are permitted to apply that judgement.

Since the intervals for bringing cars into shops for hard time policies seem to vary among car owners, both near- and far-sighted policies were evaluated for intervals of 50,000 miles, 100,000 miles and 200,000 miles between scheduled repairs.

In addition to the present maintenance practices of freight car owners, several alternatives were modelled. The simplest of these, known as "naive scheduling", treats each component as if it were the only part in the system. That is, for each part, the "single unit" replacement interval is calculated at the outset as if the part were a stand alone system, using any of the methods discussed in Chapter 2. Whenever any part has been in use for its individually calculated interval, the car is brought into the shop, the part is replaced, and the car is returned to service. Thus joint maintenance events occur only when two or more components happen to be "due" for replacement at precisely the

same time. This is likely to happen at the beginning of the car's life (corresponding to the beginning of a trial), when, for example, the wheels all have the same initial age and scheduled replacement time. As the car ages, however, and parts are placed in service with useful lives less than the scheduled interval, the scheduled replacement time for each of the car's parts will gradually become distinct. This policy attempts to explore the consequences of ignoring interactions among parts and maintenance events, and corresponds more or less to following a manufacturer's suggested replacement time blindly. The expected outcome of such a policy would be many trips to the repair track for scheduled replacements (corresponding to low miles per maintenance event), and an increase in miles per in-service failures relative to on-condition policies. The simulation makes no provision to penalize for excessive trips to the repair track. In practice, marketing and operating managers would probably consider a policy unacceptable if it required "too frequent" visits to the repair track (however defined). There could also be problems with repair track capacity[7].

The two versions of the heuristic proposed in Chapter 6 were also modelled. Recall that in both cases the decision regarding replacement of an unfailed part was based on a comparison of the cost of replacing the part at this time with the expected costs of allowing the part to remain in service up to the single part scheduled time. These costs and the expected costs are calculated. In the "greedy" case, if the costs now are less than or equal to the expected costs of waiting, the part is replaced. If not, the part is left in service. In the "extended " version, the expected costs are weighted using the probability that the random component will fail in the interval between now and the scheduled replacement of the component. As with all the policies, failed parts are always replaced immediately. The replacement interval for each part as a single unit or stand alone system is determined at the beginning, and parts which have been in use for that length of time are replaced. (This time acts as an upper bound on each part's life.)

---

[7] These concerns could be modelled in the simulation by assigning an increased switching cost for visits within some range, but since the penalty would depend on the particular circumstances of the car owner and users, this was not done.

Finally, as was indicated above, "far sighted" hard time policies were modelled. These policies call for the car to be sent to the repair track at fixed intervals (50, 100, or 200 thousand miles), and replacement of components which will reach their characteristic life before the next scheduled maintenance event.

7.4. Interpreting and Evaluating Competing Maintenance Policies

How one reports and interprets the results of a model can have a profound effect on what changes in behavior will be undertaken by practitioners in a field. It is therefore useful to preface any results of the simulation with a discussion of how the results can be understood and used. In this section, we first review the particular outputs of the model. We then focus on how these measures can be interpreted for a single scenario or set of circumstances. Finally, we look at the question of how to interpret results over a range of circumstances.

7.4.1. The Outputs Produced by the Simulation Model

The model presents results of simulated operations and maintenance activity of a freight car for an extended period of time. As has been indicated previously, these results take the form of three measures, miles per maintenance event, miles per in-service failure, and cost per mile. These measures are consistent with the concerns for service reliability and cost which were raised by managers in the case studies.

The first measure, miles per in-service failure, reflects concern with service reliability. In-service failures can have long-term consequences regarding the securing and retaining of business. An in-service failure was defined as any failure or exceeding of the AAR interchange standard by a component under study (including the exponentially distributed "rest of the car"). In-service failures do not include scheduled maintenance events or opportunistic replacements which are undertaken while the car is already in the shop.

The second measure, miles per maintenance event, measures the impact of random and scheduled events as well as in-service failures, since it is desirable to keep the equipment in use as much as possible. (Note, for example, that a car which is in storage suffers no in-service failures.) All maintenance events, including "random failures" and

190

scheduled activities are included in the computation of miles per maintenance event[8].

The third measure, cost per mile, reflects the cost of all maintenance events (repairs, replacements, switching costs, etc.) discounted back to the beginning of the simulation run (i.e., the "time" when miles equal 0) divided by the run length (2,000,000 miles). No switching or repair cost is assigned to failures of the "rest of the car" component, i.e. random failures, since each policy faces the same set of these events for a given run. Because of this, the cost per mile figures represent only the costs associated with the parts included in the maintenance program. If all costs were included, the differences between the policies relative to total maintenance cost would be reduced. Cost per mile is included because cost control is one of the primary concerns of maintenance managers, and is a critical element in railroad profitability.

More attention is focused on miles per in-service failure and cost per mile than miles per maintenance event. Miles per maintenance event essentially provides a check against achieving artificial service reliability by keeping the car in the shop at the expense of being in use. (The analogy might be the automobile owner who sends his car in for tuneups every 500 miles. He may rarely break down on the highway, but this approach makes it difficult to take a cross country trip.) Miles per in-service failure measures the effectiveness of the maintenance program in terms of providing vehicles which can be used to support revenue operations. It corresponds to a classical reliability measure, mean time between failures, and recognizes that the disruptions which correspond to in-service failures are to be avoided when possible. Cost per mile is an efficiency measure, not only of the maintenance plan, but in terms of measuring how much service-related use is purchased with the maintenance expenditure. Because of the importance attached to these two measures, they are presented in both numerical and graphical form, although the miles per maintenance event measure is reported in the tables for each policy and scenario, and discussed when of some importance.

---

[8] Although all the policies face the same set of "random events", they do not all use them the same way. The opportunistic policy exploits some of these "random events" to reduce the number of in-service failures.

7.4.2. Comparing Results for a Scenario

It is a simple matter to compare policies when one is clearly better in all measures. Such a superior policy can be referred to as "strictly dominant", and is to be preferred. Unfortunately, this is not always the case, and it becomes necessary to develop a means for assessing the performance of policies which are better in one measure and worse in another. In this section, we look at how to determine when a policy is strictly dominant, and propose a general approach for dealing with policies which are neither dominant nor dominated.

In reporting the results of simulation runs, we shall use a particular type of diagram which maps the reliability (measured in miles per in-service failure), and cost per mile. Figure 7.3 presents such a conceptual diagram for 5 alternative policies, labeled A through E. In considering a change from policy A, it is clear that one would never want to adopt policy C (or move anywhere in quadrant I), since this incurs higher costs and lower reliability. Policy B (or any policy in Quadrant IV), on the other hand is always preferable to A (i.e., a dominant solution), since it achieves higher reliability at a lower cost. Policies D and E (found in Quadrants II and III) require that some sort of tradeoff analysis be performed, in which the value of higher reliability is assessed relative to cost.

A caveat regarding the use of the diagrams is in order. Since they plot the means for the measures in question, it is possible for a policy to appear to be better than another in a measure when in fact the two are not statistically significantly different. In order to compare the measures properly, tables showing the results of all the pairwise comparisons are also presented, and should be used, particularly where the differences are modest.

In the next chapter, we will see that some policies perform quite poorly compared with others, making the decisions easy in comparing those alternatives. In other cases, however, the decision of whether to spend higher sums on maintenance in return for higher reliability requires a tradeoff analysis. These decisions will always depend on the situation and values of the managers making the decision, but it is possible to resolve them in an organized and analytical manner.

Other than in the most extreme cases, where the decision is either always to accept the low cost alternative (due, for example, to an impending bankruptcy), or to take the

192
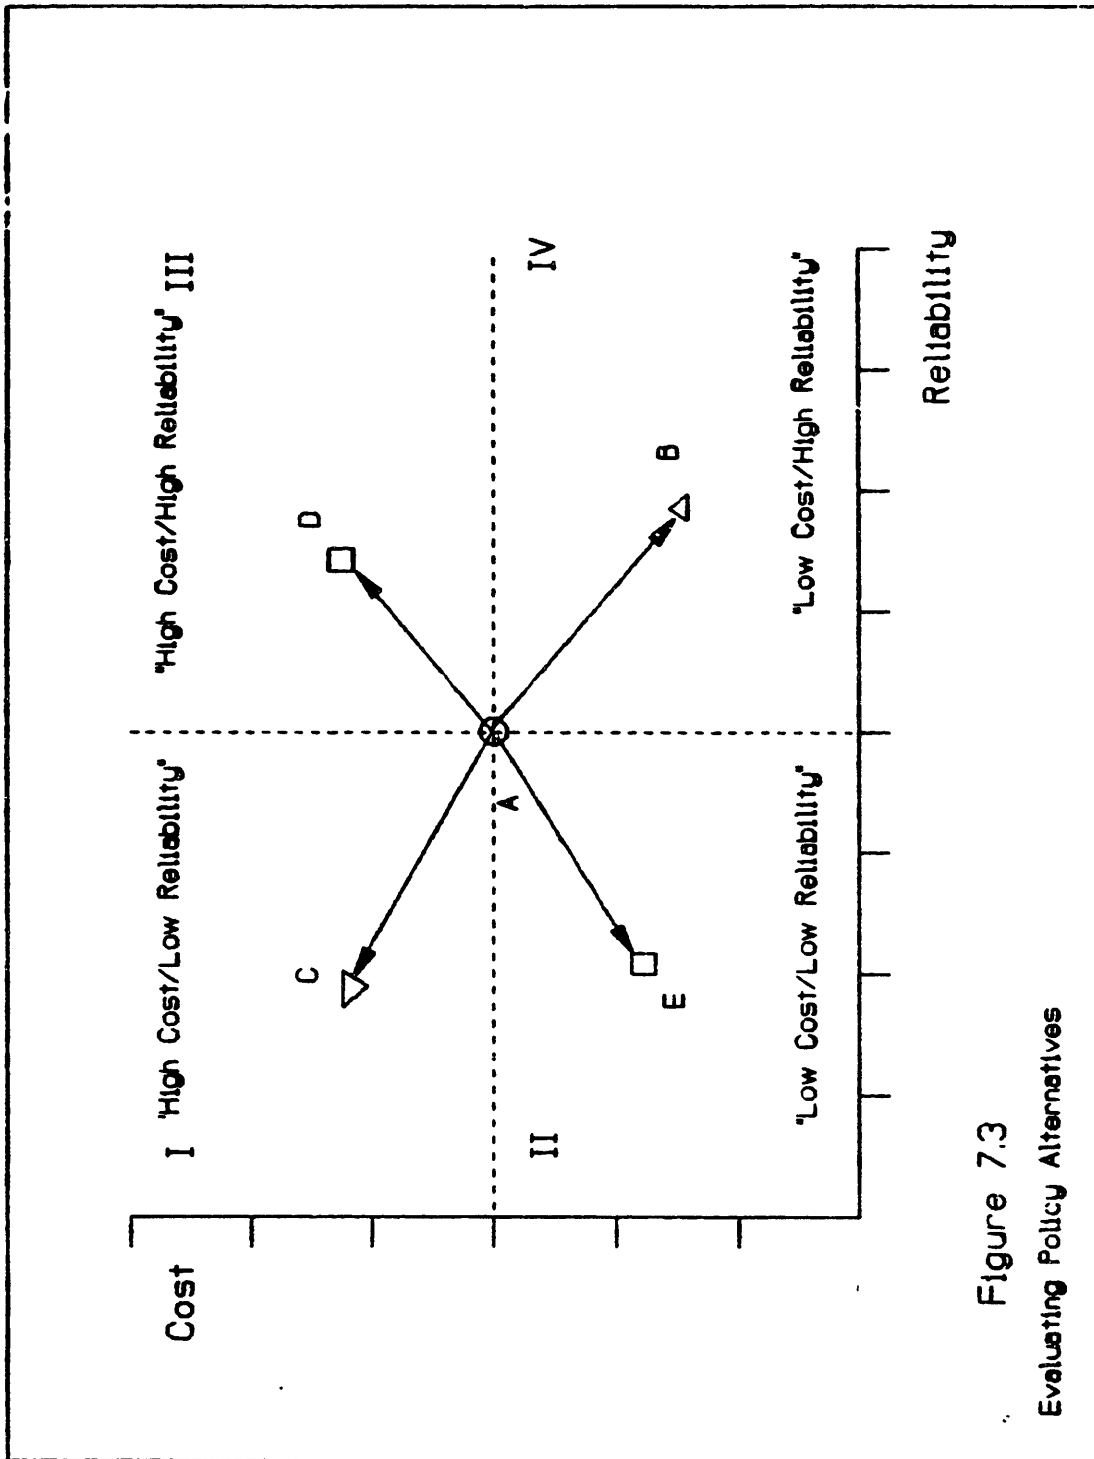
Figure 7.3

Evaluating Policy Alternatives

high reliability option (due to pending litigation or some corporate policy), the manager must find a means to translate the two axes into common units. The most direct translation would appear to be to attach an economic value to the reliability measure. This means one must ask what are the costs typically associated with more (or fewer) miles per in-service failure, or more simply, what is the true cost of an in-service failure.

While it is beyond the scope of this thesis to explore this matter fully, there would appear to be two sets of consequences to an in-service failure, the disruption of railroad operations and the disruption of the consignee who intends to use the goods carried in the car, i.e.

$$Cost_{In-Svc\ Failure} = Operations\ Costs + Shipper\ Costs \qquad (7.1)$$

The disruption of railroad operations can be further broken down into several components, including the capital costs associated with having assets idle (removing the car from service and repairing it) and the foregone net revenues the car could have earned had it been available. The costs in wasted assets associated with removing a car from service include the time value of all the equipment which is delayed while a car is removed from service (usually measured in terms of car hire or per diem), and the value of all equipment delays while the car is returned to service. In the case of removing the car from service, this delay may be half an hour or more, which in delaying a 100 car train may be on the order of $50 to $100 (either car hire due to others or foregone from a railroad's own cars), all incurred in the first half hour after the defect is discovered. In addition to setoff and pickup delays, the time while the car itself is unavailable can be considerable. Several railroad officials indicated that a car is typically out of service for 3 days while undergoing repairs for most defects. At car hire rates of $.50 to $1.00 per hour, this can add an additional $35-75. The revenue opportunity costs are the net revenues that would have been generated for the car owner had the car been available for service for the period of time it was in the shop (including, of course, any wait time for classification and return to service). Thus if a car is currently generating (on average) $200 per day of net revenue for its owner, the revenue opportunity cost of being out of service for 3 days would be $600. As a practical matter, one can calculate an equivalent

194

average revenue per mile operated, and, using the car's annual mileage, the foregone revenue mileage by being out of service.

On the shipper side, service unreliability causes the users of transportation services to hold greater inventories to prevent stockouts. One can calculate the optimal inventory based on the variance in transportation service times. If a shipper can be confident that the number of in service failures has been reduced due to an improved maintenance policy, that can be translated into smaller stockout inventories, and free up the cost of those inventories for other investments. In principle, the railroad car owner should be able to capture some of this shipper benefit in the form of either higher transportation rates, or higher equipment charges. In the case of shipper owned equipment, all the benefits accrue to the car owner, although in that case few, if any of the railroad operational benefits may be available (unless the car is part of a unit train consist, in which case both parts of equation 7.1 are fully captured by the car owner). There are a number of ways one can approximate the stockout inventory associated with in-service failures, but one of the simplest is to estimate the stockout inventory needed to protect against a single in service failure of an "average" length. This will depend on the consumption rate of the commodity and the value of the commodity itself.

These various terms can be formalized in the following equation:

$$
\begin{aligned}
C_I = &[(TL - 1) \times (SO + PU) \times CH_A] \times SHR_{C_{RR}} + \\
&[(SO + PU + RT_C) \times (CH_C + REV_C)] + \\
&[SHR_{CSR} \times C_{SR}]
\end{aligned}
\tag{7.2}
$$

with

C_I the cost of an in-service failure,

TL the average train length in which the car is used,

SO the average time (in hours) for setting off a car from a train due to an in-service failure,

PU the time to pick up a car which has been returned to service after a repair,

CH_A the average car hire rate for cars in the train,

$SHR_{C_{RR}}$ the share of the railroad operational savings that accrue to the owner of the car

RT_C the time a car is out of service due to an in-service failure,

CH_C the car hire rate for the car

195

REV_C the average hourly revenue the car earns when in service,

SHR_{CSH} the share of the shipper's savings that can be captured by the car owner by providing more reliable service,

C_{SH} the costs to the shipper due to mechanical unreliability (including excess stockout inventory, increased liability, and marketing effects).

The first line of equation 7.2 is the railroad operational savings, which are weighted by the extent to which the car owner can actually capture these savings. If the car is owned by a railroad and the car is on-line 100% of the time, then this share will be 100%. If, on the other hand, the car is owned by a shipper and is never used in conjunction with other cars owned by the same shipper, then SHR_{CRR} will be 0%, since all these costs are borne by other parties (car owners, other railroads, etc.). The second line of the equation is the capital and revenue costs which are incurred by the car itself while it is not usable due to an in-service failure. The third line is the cost to the shipper, and includes excess stockout inventory[9] held by the user of the car (generally the consignee) as protection against unreliable service due to in-service failures, any increased liability, and damaged customer relations. This is weighted by the share of these costs which are actually assumed by the car owner, SHR_{CSH}. If the car is owned by a railroad, this may be small, passed on in the form of altered rates reflecting service reliability. If the car is owned by the shipper, this may be as high as 100%, as in the case of case study Company C in Chapter 4. It is clear that the value of the commodity will have a direct effect on the value of the inventory, and we would thus expect to see companies who ship high value goods supplying their own cars and maintaining them to higher standards of reliability.

In Table 7.2, a numerical example is given, which contrasts a railroad owned car and a shipper owned car. In both cases, the car typically travels as part of 100 car trains,

---

[9] Calculating the value of the stockout inventory which a company chooses to hold is clearly afield from the thrust of this thesis, and for purposes of demonstration we will assume that a company carries an inventory equal to the expected number of in-service failures per year, up to a maximum of one in-service failure on hand. Thus a policy that reduces the number of expected in-service failures per year from .75 to .5 would reduce the inventory held specifically to protect against railroad mechanical failures by 33%.

BASIC DATA

Conversion Factors

TrainLength:	100 Cars	
SetOff:	30 Minutes	0.500 Hours
PickUp:	5 Minutes	0.083 Hours
CarHire_A:	$0.50 per Hour	
RepairTime:	72 Hours to Repair	542 Miles
CarHire_Car:	$1.00 per Hour	
Rev_C:	$0.50 per Mile	
C_SH/Day	$500 per Day (Stockout Inventory only)	
Annual Miles:	66000 Miles	

Railroad Owned Car

Share_RR:	100%	
Share_SH:	0%	

Cost/ISF:		$29	(Line 1, equation 7.2)
	+	$344	(Line 2, equation 7.2)
	+	$0	(Line 3, equation 7.2)

		$373	

Cost/Mile: $0.0056

Shipper Owned Car

Share_RR:	0%	
Share_SH:	100%	

Cost/ISF:		$0	(Line 1, equation 7.2)
	+	$344	(Line 2, equation 7.2)
	+	$1,500	(Line 3, equation 7.2)

		$1,844	

Cost/Mile: $0.0279

TABLE 7.2
Example of Calculations: Mid Value Commodity

and is subject to efficient set off, pick up and repair times. The shipper costs are the stockout inventory per day that the car is typically out of service due to an in-service failure. In the railroad owned case, all the railroad-related costs are borne by the owner and none of the shipper costs. The shipper owned car represents the opposite extreme, assuming none of the railroad costs and all the shipper costs accrue to him. The calculations are such that in the railroad owned case, line 1 of equation 7.2 is $29; in both cases the costs associated with the car being under repair are $344; in the shipper owned case the shipper costs (due to increased inventories) are $1500. The costs per mile of an in-service failure are $0.0056 per mile for the railroad owned car and $0.0279 per mile for the shipper owned car. What this means as a practical matter is that the shipper is much more concerned with the reliability of this car than the railroad, since the costs of in-service failure which befall him are much higher. In terms of tradeoffs, if an alternative maintenance policy cost $0.00279 per mile in additional direct maintenance costs (an increase of 10%), it must increase the car's miles per in-service failure by 10% to be attractive to the shipper. In the railroad owned case, that same increase in direct costs must be accompanied by an almost 50% increase in miles per in-service failure to be attractive. When examining the performance of the various policies we will return to this sort of calculation.

In some circumstances, a policy will perform dramatically better than another in one measure, and only slightly worse in another measure. If, for example, a policy results in an improvement of several hundred percent in miles per in-service failure at an increase in cost of only .01 percent, it is hard to imagine circumstances in which managers would not wish to "invest" in the more costly policy. In such cases, the obviously better policy will be referred to as "virtually dominant". Naturally, in practice, managers must decide for themselves at what point a policy is so overwhelmingly superior in some measure that it no longer requires more formal analysis.

7.4.3. Comparing Results Across Scenarios

An important question to ask about any maintenance policy is how well it will perform if circumstances change. An ideal policy would be one which is strictly dominant under virtually all events which the car owner might face. In the next chapter,

we will see that none of the alternatives appear to be such an optimal policy. In the absence of a policy which is always "best", a reasonable goal would be a policy which performs well (although not necessarily best) under a large variety of circumstances. Such a policy can be said to be "robust". It is useful to explore this notion of robustness in some detail, and be careful not to confuse it with another property, "sensitivity".

A policy will be said to be robust if it meets some criteria for performing well under all circumstances. This corresponds to the notions of robustness applied to a person's health. A person is said to be robust if they are found to be physically well over a long period of time and circumstances. Consider the policies in figure 7.4. If the minimum standard for a policy is .8, then Policies B and C are robust, since they perform above the minimum standard in all cases. Policy A is not robust, since it never performs well; Policy D is not robust either, since it performs poorly in circumstances C2, C3, C6, and C7. Notice that Policy D is actually the best policy in terms of the measure of effective.iess in circumstance C5. If one is certain that that is the case that will always be faced, robustness will not matter. If, on the other hand, one is not certain what circumstances will be encountered, one of the robust policies is probably to be preferred.

A related concept is that of "sensitivity". This measures the effect on a policy of a change in circumstances. If a policy is robust under a wide range of circumstances, it may still exhibit some variation within the acceptable range due to external factors. If, for example, the changes in the performance of Policy B can be related to some external factor (such as the intensity of use of a freight car), then that policy can be referred to as both robust and sensitive. Knowledge about what a policy is sensitive to can help the maintenance planner in deciding among policies when some of the conditions the car will face are known. Indeed, one could argue that it is preferable that a policy be responsive to certain factors (such as use or overall reliability), if that sensitivity reflects an incorporation of information into the maintenance process in a way which improves maintenance performance.

Policies which are not robust (or which are very sensitive to certain external factors) may still be desirable under those circumstances where the maintenance planner can readily predict the conditions in which the car will be used. This notion of "tuning"

199
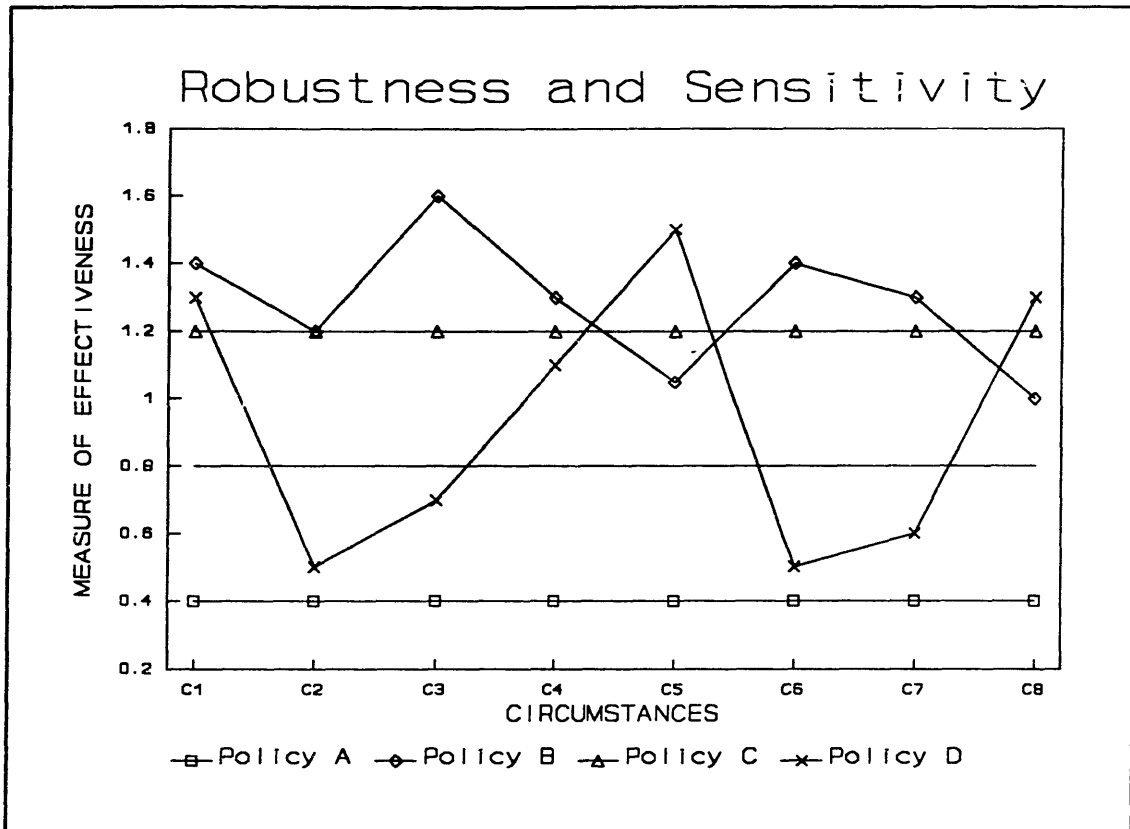
Figure 7.4

a policy to particular circumstances is likely to be most attractive in those cases where the cars are of high importance to the owner (for example, used for very high revenue shipments), or where maintenance activities can be easily adjusted. Policy "tuning" is likely to be undesirable when the maintenance facilities are at capacity (or inventories are limited), or when the fleet is large and diverse. In these cases a policy which exhibits robustness is likely to be preferable, since the maintenance manager can set the policy and not fear that an unnoticed change in conditions will cause high costs or low reliability.

7.5. Conclusion

In this chapter we have introduced the model, and its inputs and outputs. The various policy alternatives have been discussed, and a framework for interpreting the results has been proposed. In the next chapter the results of the simulation are presented and evaluated.

Chapter 8

Simulation of Freight Car Maintenance Policies:
Results

8.1. Introduction

In this chapter, the results of the simulation model described in Chapter 7 are presented. The results are organized into three sections. First, the base case results are presented and discussed. The next section examines the effect of changing the failure distributions of the parts included in the maintenance program. The sensitivity of the various maintenance policies to usage and other factors is then explored. The organization of the results along these lines represents a distinction between factors which are under the control of the maintenance manager and ones which are not. In particular, while the maintenance manager cannot control the failure rate of individual parts, he can decide which parts to include as part of a maintenance program, and which parts to maintain on an "on-condition" basis. He also can generally select the quality of parts to put on cars (for example, parts which use premium steels). These are the sorts of factors represented in the alternate distributions section. Other factors, including the annual usage of the car, the routing of the car on other railroads, and even the overall quality of the fleet already purchased are outside the maintenance manager's control. The second set of sensitivity analyses examine these issues.

For each scenario or alternative, a basic description of the relevant inputs and circumstances is given, along with a discussion of the purpose or intent of the scenario. In general, the various scenarios were intended to test the sensitivity of one or more of the policies to a single parameter. Following the description, the numerical results are given, and, where appropriate, graphs are presented. The numerical results in the tables are the mean values (and associated standard deviations) for ten matched runs of the simulation (trials), which are subject to the caveats expressed in the previous chapter. A reliability/cost tradeoff diagram using the mean values is given for each scenario, using miles per in service failure as the reliability measure. (The base case also includes a tradeoff diagram for cost and miles per maintenance event.) For each case a table of the

pairwise comparisons for the three measures of interest is also given. As previously indicated, the pairwise comparisons are based on the results of a Wilcoxon signed rank test. (Only policies which outperform another in a measure at the 10% level of significance are reported as preferred.) Each of the pairwise comparison tables are organized so that the upper left compares the currently followed policies to each other, the upper right compares the current policies to the various alternatives, and the lower right compares the alternatives to each other. Following the results for each of the scenarios, some general observations are made.

8.2. The Base Case

The base case, as noted above, is characterized by the failure distributions and cost data given in Table 7.1, annual mileage of 66,000 miles, random failures occurring with a mean frequency of 50,000 miles, and an off-line mileage percentage of 50% (henceforth called the 66/50/.5 scenario). The base case examines the following policies:

- **on condition maintenance (OC)**, which is the most widely used policy by railroads;

- **near sighted hard time policies**, in which the car is brought to the shop at fixed intervals, and parts which reach their characteristic lifetime in the first 50% of the next interval are preventively replaced, tested using intervals of 50, 100, and 200 thousand miles (N50, N100, and N200, respectively).

- **naive scheduling (NS)**, which treats replacement decisions for each component as if it were the only item;

- the **greedy (GO)** and **extended (XO) versions of the opportunistic heuristic** described in Chapter 6;

- **far sighted hard time policies**, in which the car is brought in for preventive maintenance at fixed intervals, and all components which will reach their characteristic life in the next interval are replaced upon visits to the shop, with trips to the shops scheduled at intervals of 50, 100, and 200 thousand miles (F50, F100, and F200, respectively);

On condition and near sighted hard time policies best describe the maintenance strategies currently in use by the major North American railroads and private car owners. The results for the base case are given in Tables 8.1 and 8.2, and Figures 8.1 and 8.2. Table 8.1 presents the means and standard deviations for the 10 runs for each of the policies. Table 8.2 gives the results of the pairwise Wilcoxon signed rank tests among all the

	Current Policies				Alternative Policies					
	On Cnd.	Near Sighted Hard Time			Naive Sched.	Opportunistic		Far Sighted Hard Time		
		50K	100K	200K		Greedy	Xtnded	50K	100K	200K
Miles per In Service Failure (std. dev.)	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	31114 (3436)	30235 (3178)	28497 (3246)	31818 (4376)	33474 (5006)
Miles per Maintenance Event (std. dev.)	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	29912 (3075)	27932 (2476)	18103 (1311)	24063 (2490)	28605 (3635)
Cost per Mile (std. dev.)	.01059 (.00075)	.01068 (.00085)	.01030 (.00096)	.01033 (.00086)	.01010 (.00076)	.01112 (.00051)	.01038 (.00045)	.01050 (.00075)	.01018 (.00073)	.01092 (.00086)

Base Case
Table 8.1

	Current Policies				Alternative Policies					
	On Con OC	Near Sighted Hard Time N(n)			Naive Sched. (NS)	Opportunistic		Far Sighted Hard Time F(n)		
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	· · ·	N50* OC* OC#	N100* OC* N100*	N200* OC+ N200*	NS* OC* NS*	GO* GO* OC*	XO* XO* ---	F50* OC* F50#	F100* F100* F100*	F200* F200* OC#
N 50		· · ·	N50* N100* N100*	N50* N200* N200*	NS* NS* NS*	GO* GO* N50+	XO* XO* XO+	F50* F50* F50+	F100* F100* F100*	F200* F200* ---
N 100			· · ·	N100* N200* ---	NS* --- NS#	GO* GO* N100*	XO* XO* ---	F50* N100* N100#	F100* F100* ---	F200* F200* N100*
N 200				· · ·	NS* N200* NS+	GO* GO* N200*	XO* XO* ---	F50* N200* N200+	F100* F100* ---	F200* F200* N200+
NS					· · ·	GO* GO* NS*	XO* XO* NS+	F50# NS* NS*	F100* F100* ---	F200* F200* NS*
GO	KEY Each Cell Contains Preferred Policy for:					· · ·	GO* GO* XO*	GO* GO* F50*	--- GO* F100*	F200* GO* ---
XO	Miles / In Svc. Failure Miles / Maint. Event Cost / Mile						· · ·	XO* XO* ---	F100* XO* F100#	F200* --- XO*
F 50	(--- indicates none preferred) Level of Significance:							· · ·	F100* F100* F100*	F200* F200* F50+
F 100	* : 1% + : 5% # : 10%								· · ·	F200* F200* F100*
F 200										· · ·

Base Case
Table 8.2

Figure 8.1
Base Case

KEY

OC: On Condition
NN: Near-Sighted
Hard Time, Interval
of N-thousand Miles
NS: Naive Scheduling
G0: Greedy Heuristic
X0: Extended Heuristic
F(N): Far-Sighted
Hard Time, Interval
of N-thousand Miles

Cost/Mile

.0110
.0108
.0106
.0104
.0102
.0100

20 22 24 26 28 30 32 34 36

Miles/In Service Failure
(Thousands)

OC
N200
N50
N100
F50
NS
G0
F200
X0
F100

205

Figure 8.2
Base Case

KEY

OC: On Condition
N(N): Near-Sighted
Hard Time, Interval
of N-thousand Miles
NS: Naive Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic
F(N): Far-Sighted
Hard Time, Interval
of N-thousand Miles

Cost/Mile

Miles/Maintenance Event
(Thousands)

policies for the three measures. Figures 8.1 and 8.2 are tradeoff diagrams comparing cost per mile with miles per in-service failure and miles per maintenance event.

Several points immediately leap out in examining the base case results. The first of these is that the currently followed set of policies are virtually always outperformed by the alternatives. Notice, for example that the on condition policy (OC), which is the most widely followed maintenance policy among the U.S. railroads, is outperformed in miles per in-service failure by every other policy. The other currently followed policies, the "near sighted" hard time policies (N50, N100, and N200) are also outperformed by the alternatives in miles per in service failure and frequently in miles per maintenance event. This is shown graphically in Figure 8.1, where the currently followed policies are all clustered on the left side of the figure. It is noteworthy that these policies sometimes do better in cost per mile (for example, than GO and F200), which reflects the traditional concern with cost management among railroad managers. Even in this regard, however, suitable alternatives are available. Both the extended version of the opportunistic heuristic and the F100 "far sighted" hard time policy are superior in both cost and reliability (i.e., strictly dominant) to the current practices. Even where a tradeoff is required, it is difficult to construct a circumstance where managers would not seek a 40 percent increase in reliability in return for a 5 percent increase in cost, as in the case of moving from on condition to the greedy opportunistic policy. The first conclusion from the base case is that maintenance managers should be seriously examining alternatives to the policies they are following at the present time.

The second point to note in examining the results is how well the greedy and extended versions of the opportunistic heuristic performed in miles per maintenance event. The opportunistic policies should be expected to perform well in this measure, since they make efficient use of on-line maintenance events to improve the reliability of other components. As might have been expected, the greedy version is superior to the extended version in service reliability but achieves this at a higher cost. This is to be expected since the greedy version generally leads to parts being replaced earlier, and the earlier replacement of IFR parts leads to higher survival probabilities, as noted in chapter 2. The extended version is able to achieve lower costs by virtue of two different phenomena.

By delaying replacements, the costs incurred are discounted more deeply, and by allowing parts to remain in service longer the total number of parts replaced is generally lower for a fixed number of miles (2 million miles in this case). This suggests that if one can choose between the greedy and opportunistic versions, the decision will generally turn on the sort of tradeoff analysis discussed in Chapter 7. In this case, the increase in miles per in-service failure is 2.8%, which is achieved at an increase in costs of 7.1%. Following the tradeoff example from chapter 7, this would mean that for both the shipper-owned and railroad-owned car the XO policy would be more attractive, since the increase in reliability is less than the increase in cost. The relevant tradeoff can also be thought of in terms of the annual expenditures involved. Recall that in this case the annual mileage of the cars was 66,000 miles. The expected increase in annual maintenance cost due to a switch from XO to GO would be $48.80 per car (66,000 miles times $0.01112 minus $0.01038). In return for this investment, the car owner might expect to go about 880 miles further between in service failures. This is equivalent to a reduction of only 0.06 in-service failures per year.

The naive scheduling policy was an impressive performer in cost per mile, more or less average in miles per in service failure, and fared poorly in miles per maintenance event. Only the F50, N50, and N100 policies did worse in this measure. This suggests that the policy is not easily extended to many more parts than the eight components considered in the base case. The frequent maintenance events and low costs also suggest that the model may tend to underestimate the costs of frequent visits to the shop. In particular, by assigning the same "switching" or fixed cost to all visits to the repair track, the impact of such a policy on shop capacity is ignored. In practice, it is likely that such costs would quickly lead managers to seek an alternative. If, for example, the repair shops are heavily utilized, then managers will seek ways to get more out of each visit to the repair track than can be achieved under the NS policy. Similarly, if the car must be removed from service frequently for scheduled maintenance events, it may be necessary to own a larger fleet, resulting in higher capital costs. The model has no method for estimating these sorts of effects.

A most important point to note is how well some of the "far sighted" versions of

the hard time policies performed. The F200 policy achieved the highest miles per in service failure of all the policies, albeit at the highest cost, and the F100 policy was strictly dominant over many of the other policies. This seems to contradict the experience of railroad car owners regarding the absence of a "best" single time to assign cars to the shop for planned maintenance. If all cars can be productively sent to the shop at fixed intervals, maintenance management could be made much simpler. Unfortunately, the next section shows that the effectiveness of the "far sighted" policies depends to a high degree on the failure distribution of the parts included in the maintenance program.

8.3. Alternative Distributions

An important test of the various policies, and particularly the hard time policies, is what effect a different failure distribution would have on the performance. In particular, since the best policy in the base case in terms of miles per in service failure was the 200,000 mile far-sighted hard time policy (F200), and the F100 policy dominated many other policies, alternative distributions were selected to test the sensitivity of those policies to the parts modelled. All the various policies were simulated with the characteristic life for the "wheels" part reduced from 274,000 miles to 90,000 miles and to 185,000 miles. These particular distributions examine the impacts of two different circumstances. In the case of the very short characteristic life, what is really being modelled is the effect of a different "mix" of parts in the maintenance program. In particular, this case serves as a test of the consequences of undertaking a program in which one half of the parts have very short lives, while the other half have long lives. (Recall that end of car units and trucks have characteristic lives of 350,000 and 564,000 miles.)

The second alternate distribution simply tests the effect of one of the distributions changing to a somewhat lower level. In both cases, all other inputs were as in the base case scenario (Table 7.1).

The case with the very short lived parts included was referred to as alternate distribution 1 and the 185,000 mile wheels as alternate distribution 2. First the results of the alternate distributions are presented, followed by some comments on the policies as a whole. In particular, after presenting the numerical results of the two alternate

distributions, we can address the issue of "robustness", that is, how well a policy performs under a variety of circumstances.

8.3.1. Alternate Distribution 1

In this case, we are essentially asking what happens when half of the eight components modelled are short-lived and the other half are long-lived. The decision of which parts to include in a planned maintenance program is one which maintenance managers cannot avoid, but can actually be quite difficult. Ideally, managers would like to include all the parts which can have a significant impact on costs or reliability, and exclude all the others. To appreciate the difficulty managers and planners face, consider some of the alternate rules that might be applied to reach a decision. If, for example, the determination is made to include all the parts with costs in excess of some amount (say $100), then an inexpensive part which is of great importance to the overall reliability may be excluded. If, on the other hand, all the parts which can lead to an in-service failure are included, it is possible that the number of parts may be very great, making it difficult to implement the program. One of the ways that managers typically attempt to deal with uncertainty of this type is to attempt to implement programs gradually, concentrating first on aspects with the greatest potential (i.e, starting with parts which exhibit both high cost and high importance for reliability). Then, if that is successful, the program is expanded (in our case, more parts are added to the maintenance program). The use of such "staged implementation" requires that the policy be sufficiently robust that the effectiveness of the maintenance policy is not diminished as the composition of the parts included under the program changes. If the results of a policy depend greatly on which parts are included and which are excluded, then managers are likely to find such staged implementation difficult.

The results for the alternative distribution are given in Tables 8.3 and 8.4 and Figure 8.3. Table 8.3 presents the means and standard deviations for each of the policies for the 10 runs. Table 8.4 is the pairwise Wilcoxon results for all the policies for the three measures. Figure 8.3 is a tradeoff diagram for cost per mile and miles per in-service failure.

| | Current Policies | | | | Naive Sched. | Alternative Policies | | | | |
| | On Cnd. | Near Sighted Hard Time | | | | Opportunistic | | Far Sighted Hard Time | | |
		50K	100K	200K		Greedy	Xtnded	50K	100K	200K
Miles per In Service Failure (std. dev.)	12853 (683)	15259 (1401)	14489 (1012)	14273 (881)	17857 (1653)	21063 (2226)	20065 (1830)	28554 (3950)	17689 (1512)	15868 (1119)
Miles per Maintenance Event (std. dev.)	12853 (683)	11675 (827)	12650 (765)	13319 (766)	11731 (427)	18303 (1600)	17116 (1253)	18102 (1614)	15018 (1088)	14697 (960)
Cost per Mile (std. dev.)	.02445 (.00130)	.02381 (.00156)	.02378 (.00148)	.02400 (.00149)	.02310 (.00133)	.02490 (.00089)	.02407 (.00091)	.02431 (.00106)	.02623 (.00155)	.02578 (.00159)

Alternative Distribution 1 ($\alpha = 90,000$)
Table 8.3

| | Current Policies | | | | Alternative Policies | | | | | |
| | On Con OC | Near Sighted Hard Time N(n) | | | Naive Sched. (NS) | Opportunistic | | Far Sighted Hard Time F(n) | | |
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	-	N50*	N100*	N200*	NS*	GO*	XO*	F50*	F100*	F200*
	-	OC*	OC+	N200*	OC*	GO*	XO*	F50*	F100*	F200*
	-	N50*	N100*	N200*	NS*	OC#	XO#	---	OC*	OC*
N 50		-	N50*	N50*	NS*	GO*	XO*	F50*	F100*	F200+
		-	N100*	N200*	---	GO*	XO*	F50*	F100*	F200*
		-	---	---	NS*	N50*	---	N50#	N50*	N50*
N 100			-	N100+	NS*	GO*	XO*	F50*	F100*	F200*
			-	N200*	N100*	GO*	XO*	F50*	F100*	F200*
			-	N100*	NS*	N100*	---	N100*	N100*	N100*
N 200				-	NS*	GO*	XO*	F50*	F100*	F200*
				-	N200*	GO*	XO*	F50*	F100*	F200*
				-	NS*	N200+	---	---.	N200*	N200*
NS					-	GO*	XO*	F50*	---	NS*
					-	GO*	XO*	F50*	F100*	F200*
					-	NS*	NS*	NS*	NS*	NS*
GO						-	GO*	F50*	GO*	GO*
						-	GO*	---	GO*	GO*
						-	XO*	F50+	GO*	GO+
XO							-	F50*	XO*	XO*
							-	F50*	XO*	XO*
							-	---	XO*	XO*
F 50								-	F50*	F50*
								-	F50*	F50*
								-	F50*	F50*
F 100									-	F100*
									-	F100*
									-	F200#
F 200										-
										-
										-

KEY

Each Cell Contains Preferred Policy for:

Miles / In Svc. Failure
Miles / Maint. Event
Cost / Mile

(--- indicates none preferred)

Level of Significance:

* : 1%
+ : 5%
# : 10%

Alternate Distribution - 1 ($\alpha = 90,000$)
Table 8.4

Figure 8.3

Alternate Distribution 1

The most striking result under this scenario is the change in the performance of the various "far sighted" hard time policies. The F100 and F200 perform quite poorly, while the F50 now becomes the best of all the policies modelled in terms of miles per in-service failure. This dramatic change in the performance of the F100 and F200 policies raises questions about whether or not these policies are stable across different distributions. In particular, it suggests that one must select the maintenance interval carefully to match the parts mix. If the mix is all long lived, one can follow a longer set of intervals (as in the base case), but if the parts include short-lived components, one must select a shorter interval. This means that managers who wish to change the parts in the program (using, for example, a "staged implementation" approach) must also change the interval as the mix changes.

The second thing to note is that the two versions of the opportunistic heuristic seem to do quite well even under these circumstances. Looking at Table 8.4, notice that the greedy version (GO) performs better in miles per in-service failure and miles per maintenance event than any other policy except the F50, and as well as the F50 in miles per maintenance event. It does so, however, at a high cost relative to the currently followed policies. The extended version is outperformed only by F50 and GO in miles per in service failure, and only by the naive scheduling in cost per mile. In other words, the opportunistic heuristic, while not the single best policy in this case, is once again one of the best, and is bettered by a different policy than in the base case. This quality of "robustness" may be quite attractive to maintenance planners who are uncertain as to which components to include in the maintenance program from the outset, or who intend to modify the maintenance program over a period of time.

It is appropriate to comment briefly on a trade off analysis at this time with respect to the XO and the F50 policy. The F50 achieves a 42 percent increase in miles per in-service failure over the XO, at an increase of only one percent in cost per mile. Without any formal analysis such as that in Section 7.4, it seems apparent that the increased reliability is a bargain. In return for an annual increase in costs per car of $16, and increase of 8489 miles per in service failure is achieved. (This is equivalent to a reduction of one in service failure each year.) Only a car owner who is unable to capture

any of the benefits of increased service reliability would forego such an increase at such a cost. This suggests that if one knew for certain that the circumstances were as described in this scenario, F50 is the policy to follow.

8.3.2. Alternative Distribution 2

Like alternate distribution 1, this scenario serves to demonstrate the sensitivity of the "far sighted" hard time policies to the components included in the maintenance program, and the robustness of the opportunistic policy. In this case, however, the focus is less on the "mix" of included components and more on the effect of a modest change in the failure distribution of one of the components. The results are given in Tables 8.5 and 8.6 and Figure 8.4. Table 8.5 presents the means and standard deviations for the policies for the 10 runs. Table 8.6 is the pairwise Wilcoxon results. Figure 8.4 is a tradeoff diagram for cost per mile and miles per in-service failure.

Once again the extended version of the opportunistic policy was found to be quite a good performer. Only the F100 and GO policies were better in miles per in service failure, and the N200 and NS in cost per mile. Although the GO is superior to XO in miles per in service failure, it is also higher in cost per mile. An increase of 3.5% in miles per in service failure is achieved at a cost per mile increase of 4.2%. Returning to our example from Chapter 7, we find this might be attractive when the benefits of in-service reliability accrue principally to the car owner (the shipper case), but not where the benefits are shared with another (the railroad case). (An annual increase of $38 per car results in an increase of 933 miles between in service failures.) In any event, the seeming robustness of the extended version further reinforces the idea that if maintenance planners are beginning a planned maintenance program, or want to alter the composition of the program in terms of the car types or parts included, the XO policy may be particularly attractive.

In this case, the best of the "far sighted" hard time policies is clearly the F100 policy, in which the car is brought in at 100,000 mile intervals. The performance of the F200 deteriorates greatly from the base case, so that it is now quite expensive on a cost per mile basis, and is outperformed by most of the alternatives considered. The F100 is, however, slightly more expensive than some of the alternatives, and performs at the same

215

| | Current Policies | | | | | Alternative Policies | | | | |
| | On Cnd. | Near Sighted Hard Time | | | Naive Sched. | Opportunistic | | Far Sighted Hard Time | | |
		50K	100K	200K		Greedy	Xtnded	50K	100K	200K
Miles per In Service Failure (std. dev.)	18935 (1353)	22632 (2601)	21714 (2291)	21164 (2081)	25025 (2789)	27964 (2651)	27031 (2994)	26246 (3047)	34697 (4772)	25692 (2659)
Miles per Maintenance Event (std. dev.)	18935 (1353)	15541 (1221)	17814 (1534)	19124 (1695)	17338 (841)	25819 (2087)	24187 (2203)	17165 (1292)	25675 (2663)	22745 (2086)
Cost per Mile (std. dev.)	.01388 (.00091)	.01373 (.00114)	.01340 (.00115)	.01336 (.00113)	.01318 (.00087)	.01430 (.00072)	.01373 (.00072)	.01338 (.00093)	.01395 (.00071)	.01544 (.00102)

Alternative Distribution 2 ($\alpha = 185,000$)
Table 8.5

| | Current Policies | | | | Alternative Policies | | | | | |
| | On Con OC | Near Sighted Hard Time N(n) | | | Naive Sched. (NS) | Opportunistic | | Far Sighted Hard Time F(n) | | |
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	- - -	N50* OC* N50#	N100* OC* N100*	N200* --- N200*	NS* OC* NS*	GO* GO* OC+	XO* XO* ---	F50* OC* F50*	F100* F100* ---	F200* F200* OC*
N 50		- - -	N50* N100* N100*	N50* N200* N200*	NS* NS* NS*	GO* GO* N50*	XO* XO* ---	F50* F50* F50#	F100* F100* ---	F200* F200* N50*
N 100			- - -	N100+ N200* ---	NS* N100# NS+	GO* GO* N100*	XO* XO* ---	F50* N100* ---	F100* F100* N100+	F200* F200* N100*
N 200				- - -	NS* N200* ---	GO* GO* N200*	XO* XO* N200#	F50* N200* ---	F100* F100* N200+	F200* F200* N200*
NS					- - -	GO* GO* NS*	XO* XO* NS*	F50* --- NS#	F100* F100* NS*	F200# F200* NS*
GO	KEY Each Cell Contains Preferred Policy for:					- - -	GO# GO+ XO*	GO* GO* F50*	F100* --- F100+	GO* GO* GO*
XO	Miles / In Svc. Failure ·Miles / Maint. Event Cost / Mile						- - -	XO* XO* XO#	F100* F100* XO#	XO* XO* XO*
F 50	(--- indicates none preferred) Level of Significance:							- - -	F100* F100* F50*	F50# F200* F50*
F 100	* : 1% + : 5% # : 10%								- - -	F100* F100* F100*
F 200										- - -

KEY
Each Cell Contains
Preferred Policy for:

Miles / In Svc. Failure
·Miles / Maint. Event
Cost / Mile

(--- indicates none preferred)

Level of Significance:

    * : 1%
    + : 5%
    # : 10%

Alternate Distribution - 2 ($\alpha = 185,000$)
Table 8.6

KEY

OC: On Condition
N(N): Near-Sighted
        Hard Time, Interval
        of N-thousand Miles
NS: Naive Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic

F(N): Far-Sighted
        Hard Time, Interval
        of N-thousand Miles

Cost/Mile

.0155
.0150
.0145
.0140
.0135
.0130

18   20   22   24   26   28   30   32   34   36

Miles/In Service Failure
(Thousands)

F200

GO

XO

F50

NS

N50

N100

N200

OC

F100

Figure 8.4

Alternate Distribution 2

218

level as the GO in miles per maintenance event.

Once again, a trade off analysis between XO and the "best" of the hard time policies (F100 in this case) suggests that the F100 is to be preferred *if the conditions are known with certainty.* (An increase of more than 28% in miles per in service failure is obtained for an increase of about 2% in costs).

Together with the results of alternative distribution 1 and the base case, it is interesting to conjecture about a rule for selecting the optimal interval for such policies. It would appear that the best maintenance interval for hard time policies is less than the lowest characteristic life of any of the parts included, but not "too much lower". This may be a fruitful area for future researchers.

8.3.3. Comments on the Alternative Distribution Cases

In Section 7.4 of the preceding chapter, it was indicated that one can compare policies within a scenario and across several scenarios. In the preceding sections, the focus has been on how the policies performed within a scenario. In this section, we look at how the policies performed across the distributions, focussing on robustness.

In the base case and in each of the alternate distribution cases, we found that one of the hard time policies was the best in each case, in terms of miles per in service failure, and sometimes in cost per mile. In each case, however, it was a *different* policy that was best in miles per in service failure. The versions of the opportunistic heuristic, on the other hand, were consistently near the top in miles per in service failure.

To formalize our inquiry into robustness, it is necessary to find a way to permit comparison over different trials. Since we are interested in the relative performance of the policies, one method is to "normalize" the results of the trials about the mean value for all the policies within that scenario. In other words, for the base case, the mean for each of the three measures for all the policies is computed, and then the measure for each policy is divided by that mean. In the case of miles per in service failure and miles per maintenance event, a higher than 100% rating indicates a better than average performance. For cost per mile, the lower the rating, the better the performance of the policy. The results of this for the base case and the alternate distributions are given in Table 8.7.

Looking at Table 8.7, one is struck most immediately by how poorly the currently

		Current Policies				NS	Alternative Policies					AVG
		OC	N50	N100	N200	NS	GO	XO	F50	F100	F200	AVG
Mi / ISF	Base Cse	79%	94%	90%	87%	100%	110%	107%	101%	113%	119%	100%
	Alt - 1	72%	86%	81%	81%	100%	118%	113%	160%	100%	89%	100%
	Alt - 2	75%	90%	87%	87%	100%	110%	108%	105%	138%	102%	100%
Mi / ME	Base Cse	96%	75%	88%	88%	89%	130%	121%	78%	104%	124%	100%
	Alt - 1	88%	80%	87%	87%	81%	126%	118%	124%	103%	101%	100%
	Alt - 2	93%	76%	87%	87%	85%	126%	118%	84%	126%	111%	100%
Cost/Mi	Base Cse	101%	102%	98%	98%	96%	106%	99%	100%	97%	104%	100%
	Alt - 1	100%	97%	97%	97%	95%	102%	98%	99%	107%	105%	100%
	Alt - 2	100%	99%	97%	97%	95%	103%	99%	97%	101%	112%	100%

Table 8.7

Alternative Distributions

Performance of Individual Policies v. Average for All Policies

220

followed policies performed over all three of the distributions. None of the currently followed policies performed above average in either miles per in-service failure or miles per maintenance event. Only in cost per mile did the current policies achieve better than average performance, and only the N100 and N200 were better than average in this measure in each case. The currently followed policies (OC, N50, N100 and N200) are certainly not robust, since they never perform well (although they are consistent in their level of performance).

Similarly, NS was not particularly robust in the reliability measures, as it was well below average in miles per maintenance event for all three distributions. It is, however, quite robust in cost per mile, and average in miles per in service failure, suggesting that tradeoffs may be appropriate for those car owners who are concerned primarily with costs. The relatively poor performance in miles per maintenance event may tend to limit the appeal of NS to the few cases where costs are virtually the only consideration. Even in those cases, there is the possibility that congestion-related costs have been systematically underestimated.

To facilitate discussions of the opportunistic heuristic and far sighted hard time policies, the result given in Table 8.7 are presented in graphical form in Figures 8.5, 8.6, and 8.7.

Turning first to miles per in service failure, it is apparent from Figure 8.5 just how consistently well the GO and XO policies perform. These policies are above average in all the circumstances, and do not appear to be affected very much by changes in the components included in the maintenance program. This is strong evidence that these policies are robust enough to be used under a wide range of component mixes, and are not adversely affected by changes in the failure distributions of components assigned to the maintenance program.

This is strongly contrasted by the far sighted policies. Each of the policies performs well in one or two of the situations, but is at or below average in one or two others. Indeed, the F200 policy, when applied to a mix of parts including many short-lived components does quite badly. The F100 ranges from an average performance with alternative distribution 1 to almost 40% better than average with the second alternative

221

Miles per In Service Failure
Robustness

Figure 8.5

Miles per Maintenance Event
Robustness
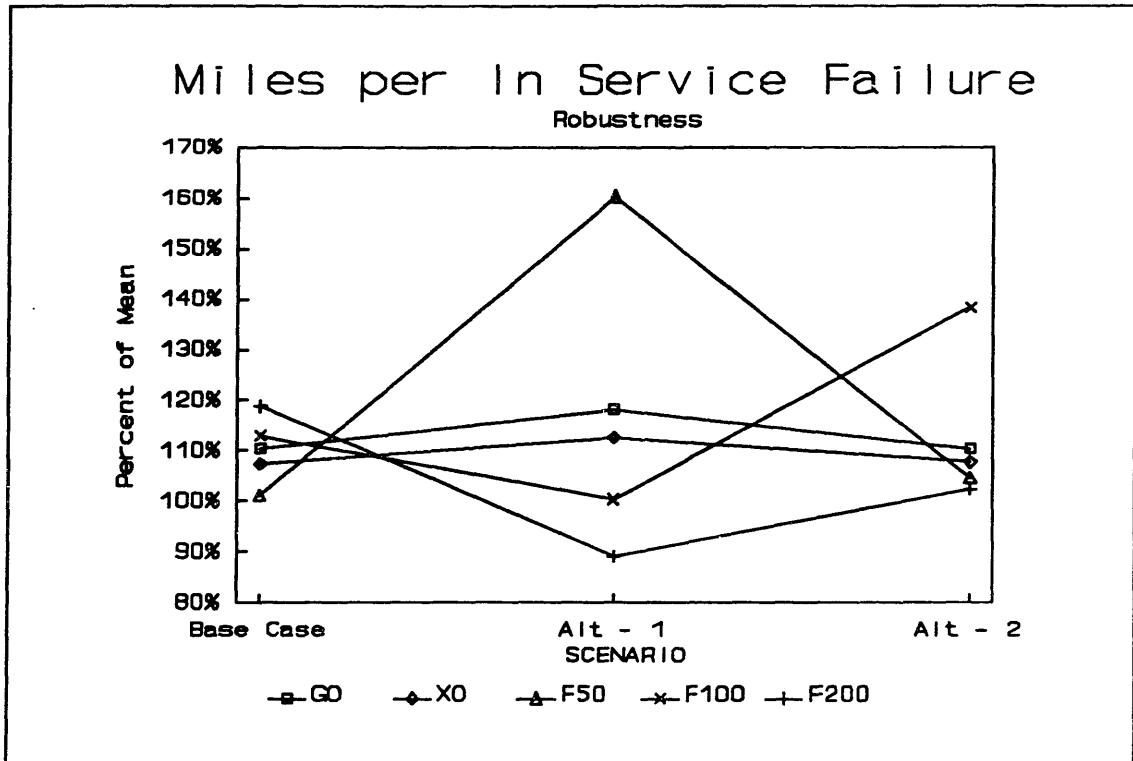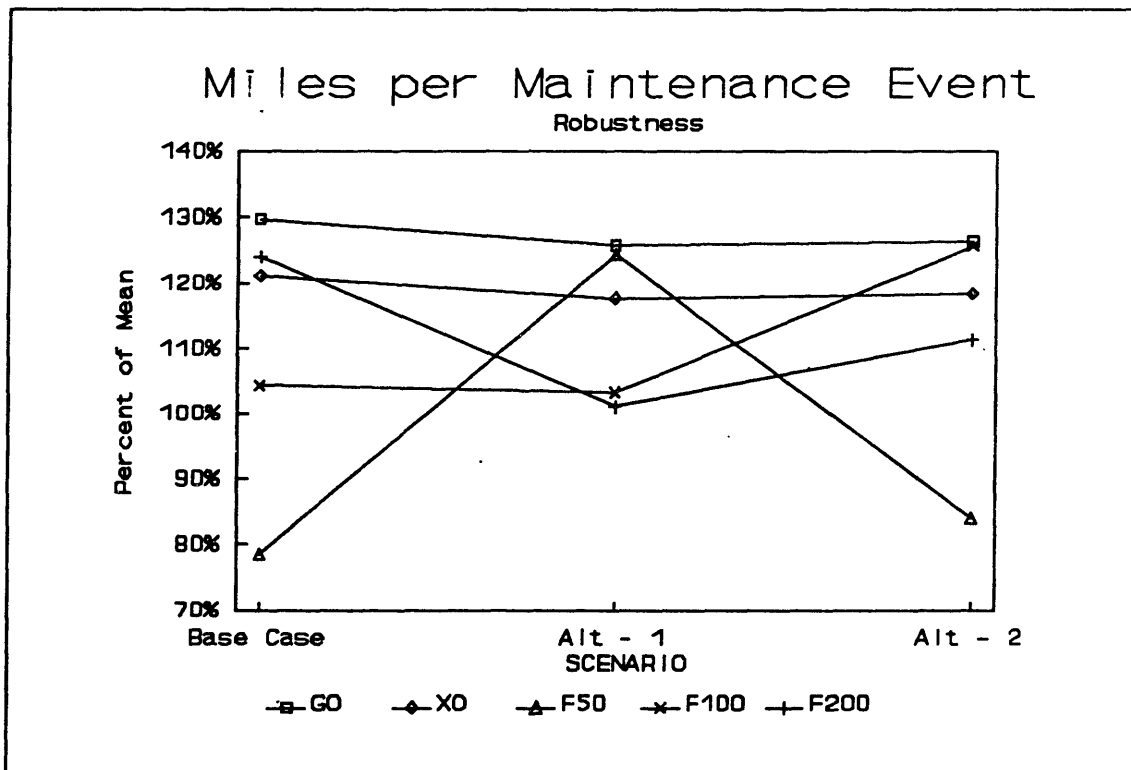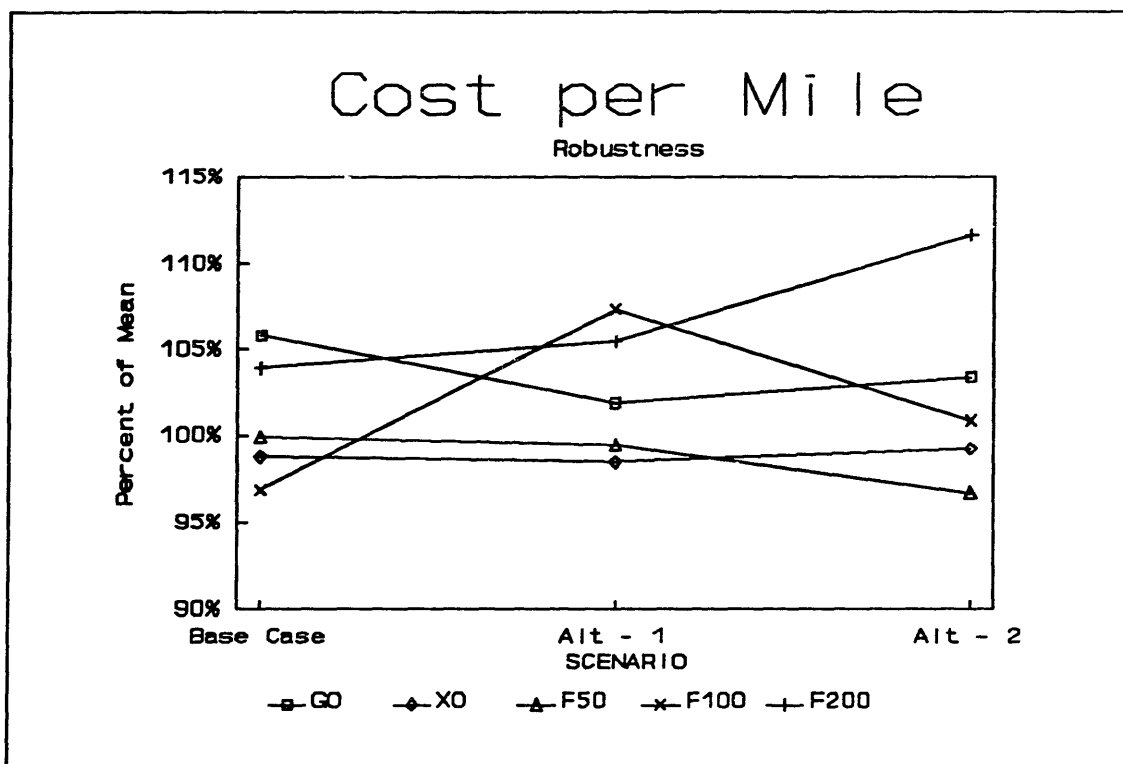
Figure 8.6

Figure 8.7

distribution.

The miles per maintenance data is also revealing (Figure 8.6). Because the opportunistic heuristics utilize failures of one component to reduce future problems with other components, they achieve high values in this measure in all the cases examined. The far-sighted hard time policies perform well in some cases, and only average in others. The F50 policy, because of its frequent scheduled visits to the shop (every 50,000 miles), is usually below average in miles per maintenance event.

The cost per mile data undercuts the performance of three of the policies. The F200 and GO policies both are higher than average (i.e., worse than average) in all three of the circumstances. The F100 experiences high costs under alternative distribution 1 (short and long-lived parts). Only the XO and F50 achieve average or better cost per mile performance under all the circumstances.

Putting all the measures together in the absence of information about the circumstances likely to face the maintenance planner is not possible, but some general

223

conclusions can be drawn. The first is that the XO policy is the most robust policy over all the measures and all the circumstances. While it was never the very best among the alternatives, it is always among the best, consistently better than average, and never performed poorly. None of the other policies studied exhibited this behavior. Maintenance managers who face an uncertain environment or who wish to stage the implementation of their maintenance program may wish to consider the XO both for the overall performance and for robustness.

The second conclusion is that the far sighted hard time policies must be "tuned" to the particular distributions of the parts included in the program. If there are institutional reasons for limiting the intervals to long times, or for refusing to change the intervals once set, the maintenance planner may want to adopt "hybrid" strategies, in which one brings the car into the shop at relatively long intervals, and then follows an alternative rule for the shorter-lived components in the interim. The alternative rule could take a number of possible forms, including use of opportunistic maintenance in the interim. These "hybrid" approaches should be studied more thoroughly. Research into this area should also examine what are appropriate criteria for deciding which parts to include under each of the policies.

8.4. Sensitivity Analysis of the Policies

In the previous section, we examined the effects of some of the factors which are under the control of the maintenance manager, such as the composition of the maintenance program and the quality of the parts included in the program. In this section the focus shifts to factors which are outside the control of the maintenance managers, but which may affect the effectiveness of the program. The factors examined are the overall quality of the cars being maintained, the level of usage of the car, and the share of time (or miles) the car is used on and off line; the effect of discounting is also discussed briefly, although that was not examined using the simulation model. That the first factor, the overall quality of the fleet, is exogenous to the maintenance manager may seem strange at first, since one of the purposes of a maintenance program is to enhance overall reliability. The quality of a car will, however, be dependent on other exogenous factors

such as the car's design, age, previous maintenance activities, and complexity. Over a period of time one would expect that a well maintained car will experience fewer "random" failures, i.e., will become more reliable; in the interim, the reliability of the car may be treated as exogenous.

The exogenous factors which were examined are not explicitly included in the decision rules of most of the maintenance policies. Indeed, only the opportunistic heuristic makes use of this information as part of the determination of whether or not to replace an unfailed part. Because the non-opportunistic policies do not directly use these factors in the decision rule, the results are usually unchanged from the base case, or changed only in minor ways. The fact that the non-opportunistic policies do not use such factors as the mileage or the overall reliability of the cars suggests that those policies overlook potentially valuable information which the maintenance planner might wish to consider.

8.4.1. The Effect of Overall Reliability

In these scenarios, the number of random failures (i.e., the reliability of the car excluding the modelled components) was varied. Recall that in the base case random failures occurred using an exponential distribution with mean of 50,000 miles. To test the effect of this parameter, the model was run with random failures occurring with means of 25,000 miles, a "low reliability" scenario, and 75,000 miles, a "high reliability" scenario. Only the opportunistic heuristics are affected by these parameters in terms of the costs[1], since the other policies schedule maintenance events independently of the reliability of the car as a system. All the policies are affected in terms of miles per in service failure and miles per maintenance event, since the random failures are

_____

[1] A qualification is in order here. The cost per mile figures which remain unchanged are the costs associated with the repair or replacement of the components being modelled. Random failures are treated as being costless in themselves, since each policy faces that same set of random events on a given trial. In reality, as the number of random failures increases, the total cost per mile would be expected to increase, although it would presumably increase the same amount for any and all policies.

incorporated into these measures. These scenarios, then, are tests of two things, the sensitivity of all the policies to the number of random failures (i.e., the number of opportunities to perform preventive or preemptive maintenance), and the degree to which the opportunistic heuristics are able to mitigate unreliability in the car as a whole.

The results of the reliability scenarios are found in Tables 8.8 - 8.10, and Figures 8.8 - 8.9. Table 8.8 presents the means and standard deviations for the 10 trials for the low and high reliability scenarios (and the base case). Tables 8.9 and 8.10 give the results of the pairwise comparisons for each scenario. Figures 8.8 and 8.9 are the tradeoff diagrams for each of the two scenarios.

To highlight how the policies are affected, a table of "normalized" values was created for the reliability scenarios like that used for the alternate distributions. For each scenario, the mean of each measure over all policies and runs is computed. This value is set to 100%, and the value for each individual policy is then compared to this. A higher value is better for miles per in service failure and miles per maintenance event. A lower value is better for cost per mile. These values are given in Table 8.11.

As expected, the overall reliability of the car affects all policies in terms of miles per in service failure. As the overall reliability decreases (i.e., the miles between "random" failures declines), the miles per in service failure also declines. More noteworthy is that the opportunistic policies do not seem to be as adversely effected by the decline as the best of the far sighted hard time policies. Notice in Figure 8.10 that when the overall reliability is high, then a properly "tuned" hard time policy outperforms both GO and XO. When reliability is low, however, GO becomes more attractive. The differences between GO and F200 become quite small; F100 and GO are not statistically significantly different. This has important implications for maintenance planners. There are times in the life of most cars when the overall reliability may be quite low, such as when a car is new and undergoing a period of "breaking in", or when a car is old and many of its parts are wearing out. Changes in the usage characteristics, such as more intensive loadings or carrying of certain commodities, may also cause a car to experience low reliability. These results suggest that in those cases, GO may be desirable. The obverse is also true. When a car is of very high reliability, the use of a properly tuned

Standard Deviations in Parentheses		On Cond.	Near Sighted Hard Time 50K	Near Sighted Hard Time 100K	Near Sighted Hard Time 200K	Naive Sched.	Opportunistic GO	Opportunistic XO	Far Sighted Hard Time 50K	Far Sighted Hard Time 100K	Far Sighted Hard Time 200K
			Current Policies				Alternative Policies				
Low	Mi/ISF	15762 (1160)	17735 (1810)	17342 (1872)	16870 (1538)	18518 (1899)	20211 (2233)	19291 (2087)	18687 (2023)	20056 (2484)	20730 (2837)
	Mi/ME	15762 (1160)	13068 (977)	14760 (1350)	15549 (1302)	14880 (952)	19336 (2331)	18430 (1795)	13575 (1053)	16674 (1695)	18756 (2305)
	Cost/Mi	.01059 (.00075)	.01068 (.00085)	.01030 (.00096)	.01033 (.00086)	.01010 (.00076)	.01107 (.00052)	.01049 (.00054)	.01050 (.00075)	.01018 (.00073)	.01092 (.00086)
Base	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	31114 (3436)	30235 (3178)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	29912 (3075)	27932 (2476)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.01059 (.00075)	.01068 (.00085)	.01030 (.00096)	.01033 (.00086)	.01010 (.00076)	.01112 (.00051)	.01038 (.00045)	.01050 (.00075)	.01018 (.00073)	.01092 (.00086)
High	Mi/ISF	25955 (1812)	31803 (3795)	30502 (3688)	29108 (2865)	34354 (3891)	38272 (5421)	37120 (4505)	35018 (4552)	40253 (6681)	43095 (8377)
	Mi/ME	25955 (1812)	19380 (1416)	23319 (2140)	25385 (2175)	23666 (1434)	35185 (4234)	33168 (3550)	20519 (1579)	28557 (3352)	35283 (5556)
	Cost/Mi	.01059 (.00075)	.01068 (.00085)	.01030 (.00096)	.01033 (.00086)	.01010 (.00076)	.01068 (.00070)	.01036 (.00056)	.01050 (.00075)	.01018 (.00073)	.01092 (.00086)

Table 8.8
Effect of Overall Reliability

227

	On Con OC	Near Sighted Hard Time N(n) 50K	100K	200K	Naive Sched. (NS)	Opportunistic Grdy (GO)	Xtnd (XO)	Far Sighted Hard Time F(n) 50K	100K	200K
OC	• • •	N50* OC* OC#	N100* OC* N100*	N200* OC+ N200*	NS* OC* NS*	GO* GO* OC+	XO* XO* ---	F50* OC* F50#	F100* F100* F100*	F200* F200* OC#
N 50		• • •	N50* N100* N100*	N50* N200* N200*	NS* NS* NS*	GO* GO* N50+	XO* XO* XO#	F50* F50* F50+	F100* F100* F100*	F200* F200* ---
N 100			• • •	N100* N200* ---	NS* --- NS#	GO* GO* N100+	XO* XO* ---	F50* N100* N100+	F100* F100* ---	F200* F200* N100*
N 200				• • •	NS* N200* NS+	GO* GO* N200*	XO* XO* ---	F50* N200* N200+	F100* F100* ---	F200* F200* N200+
NS					• • •	GO* GO* NS*	XO* XO* NS*	F50# NS* NS*	F100* F100* ---	F200* F200* NS*
GO		**KEY** Each Cell Contains Preferred Policy for: Miles / In Svc. Failure Miles / Maint. Event Cost / Mile (--- indicates none preferred) Level of Significance: * : 1% + : 5% # : 10%				• • •	GO* GO* XO*	GO* GO* F50+	--- GO* F100*	F200+ GO* ---
XO							• • •	XO* XO* ---	F100* XO* F100#	F200* --- XO+
F 50								• • •	F100* F100* F100*	F200* F200* F50+
F 100									• • •	F200* F200* F100*
F 200										• • •

Low Reliability
Table 8.9

	Current Policies				Alternative Policies					
	On Con OC	Near Sighted Hard Time N(n)			Naive Sched. (NS)	Opportunistic		Far Sighted Hard Time F(n)		
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	· · ·	N50* OC* OC#	N100* OC* N100*	N200* OC+ N200*	NS* OC* NS*	GO* GO* ---	XO* XO* ---	F50* OC* F50#	F100* F100* F100*	F200* F200* OC#
N 50		· · ·	N50* N100* N100*	N50* N200* N200*	NS* NS* NS*	GO* GO* ---	XO* XO* XO#	F50* F50* F50+	F100* F100* F100*	F200* F200* ---
N 100			· · ·	N100* N200* ---	NS* N100* NS#	GO* GO* N100+	XO* XO* ---	F50* N100* N100+	F100* F100* ---	F200* F200* N100*
N 200				· · ·	NS* N200* NS+	GO* GO* N200#	XO* XO* ---	F50* N200* N200+	F100* F100* ---	F200* F200* N200+
NS					· · ·	GO* GO* NS*	XO* XO* NS#	F50+ NS* NS*	F100* F100* ---	F200* F200* NS*
GO						· · ·	--- GO* XO*	GO* GO* ---	F100# GO* F100*	F200* --- GO*
XO							· · ·	XO* XO* ---	F100* XO* F100+	F200* F200+ XO*
F 50								· · ·	F100* F100* F100*	F200* F200* F50+
F 100									· · ·	F200* F200* F100*
F 200										· · ·

KEY
Each Cell Contains
Preferred Policy for:

Miles / In Service Failure
Miles / Maint. Event
Cost / Mile

(--- indicates none preferred)

Level of Significance:

* : 1%
+ : 5%
# : 10%

High Reliability
Table 8.10

Figure 8.8
Low Reliability

KEY

OC: On Condition
NN: Near-Sighted
   Hard Time, Interval
   of N-thousand Miles
NS: Naïve Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic
FN: Far-Sighted
   Hard Time, Interval
   of N-thousand Miles

Miles/In Service Failure
(Thousands)

Cost/Mile

KEY

OC: On Condition
NN: Near-Sighted
    Hard Time, Interval
    of N-thousand Miles
NS: Naive Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic

F(N): Far-Sighted
    Hard Time, Interval
    of N-thousand Miles

Figure 8.9

High Reliability

Cost/Mile

Miles/In Service Failure
(Thousands)

		Current Policies				Alternative Policies						AVG
		OC	N50	N100	N200	NS	GO	XO	F50	F100	F200	
Mi / ISF	Low Rel	85%	96%	94%	91%	100%	109%	104%	101%	108%	112%	100%
	Base Cse	79%	94%	90%	87%	100%	110%	107%	101%	113%	119%	100%
	High Rel	75%	92%	88%	84%	99%	111%	107%	101%	117%	125%	100%
Mi / ME	Low Rel	98%	81%	92%	97%	93%	120%	115%	84%	104%	117%	100%
	Base Cse	96%	75%	88%	95%	89%.	130%	121%	78%	104%	124%	100%
	High Rel	96%	72%	86%	94%	88%	130%	123%	76%	106%	130%	100%
Cost/Mi	Low Rel	101%	102%	98%	98%	96%	105%	100%	100%	97%	104%	100%
	Base Cse	101%	102%	98%	98%	96%	106%	99%	100%	97%	104%	100%
	High Rel	101%	102%	98%	99%	97%	102%	99%	100%	97%	104%	100%

Table 8.11

Effect of Overall Reliability

Performance of Individual Policies v. Average for All Policies

232

far sighted hard time policy may achieve the best overall results.

The relative performance of the policies in miles per maintenance event serves to further reinforce this point. When the overall reliability is low, GO is the best policy in this measure. As the overall reliability increase, F200 improves in relative terms until, in the high reliability case, F200 is as good in this measure as GO. Not to be lost in this is the good performance of the XO policy. The good performance of the GO and XO policies, particularly at lower levels of overall reliability is consistent with the design of the opportunistic heuristics. The opportunistic policies incorporate the overall reliability of the car into the decision framework. The greedy version does so in an indirect way, since a more reliable car will receive fewer opportunities for preemptive replacement, and so the expected time of replacement will be later resulting in lower miles per in service failure and lower costs. The extended version incorporates the reliability of the car in an explicit way, and so causes the "point of indifference" to adjust. Recall from Chapter 6 that this is the time at which one is indifferent between replacing a part and leaving it in

Figure 8.10

233
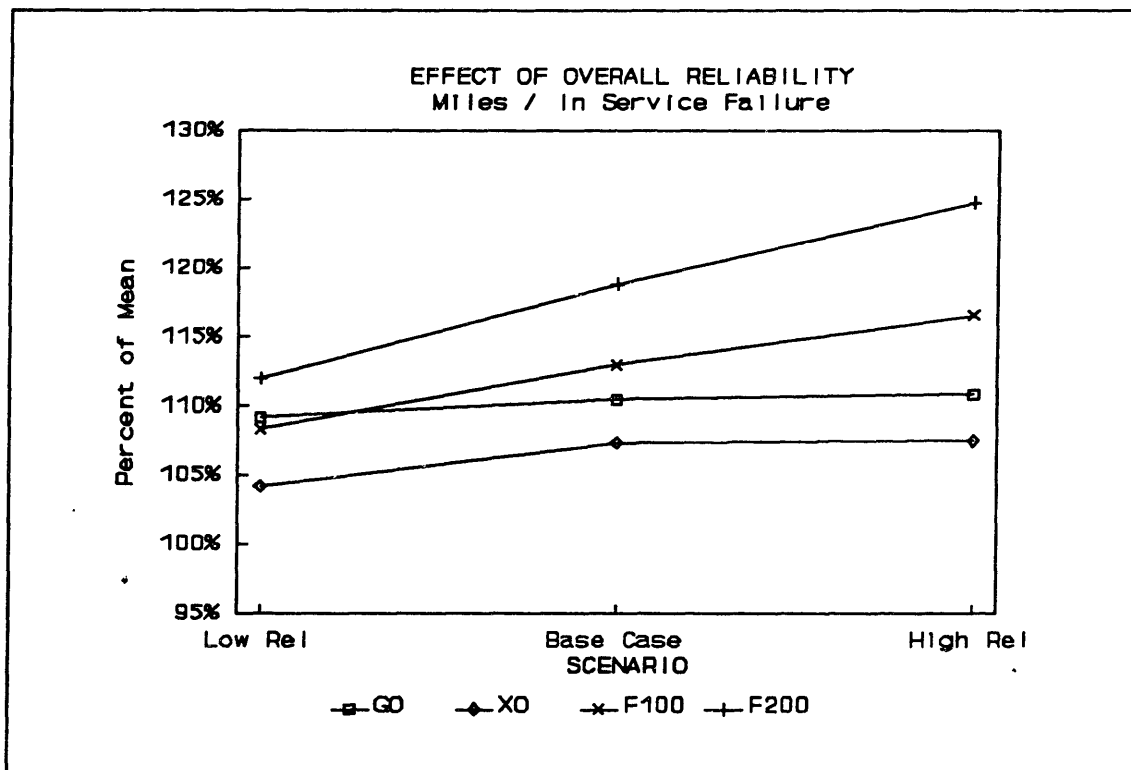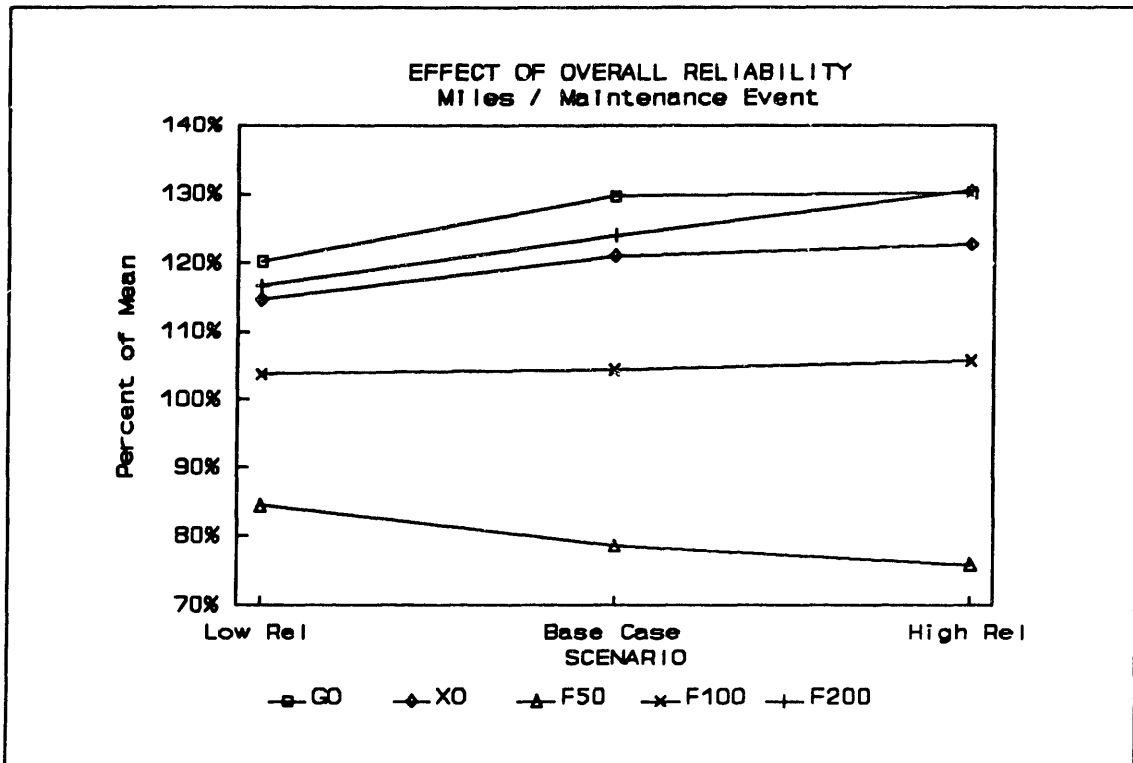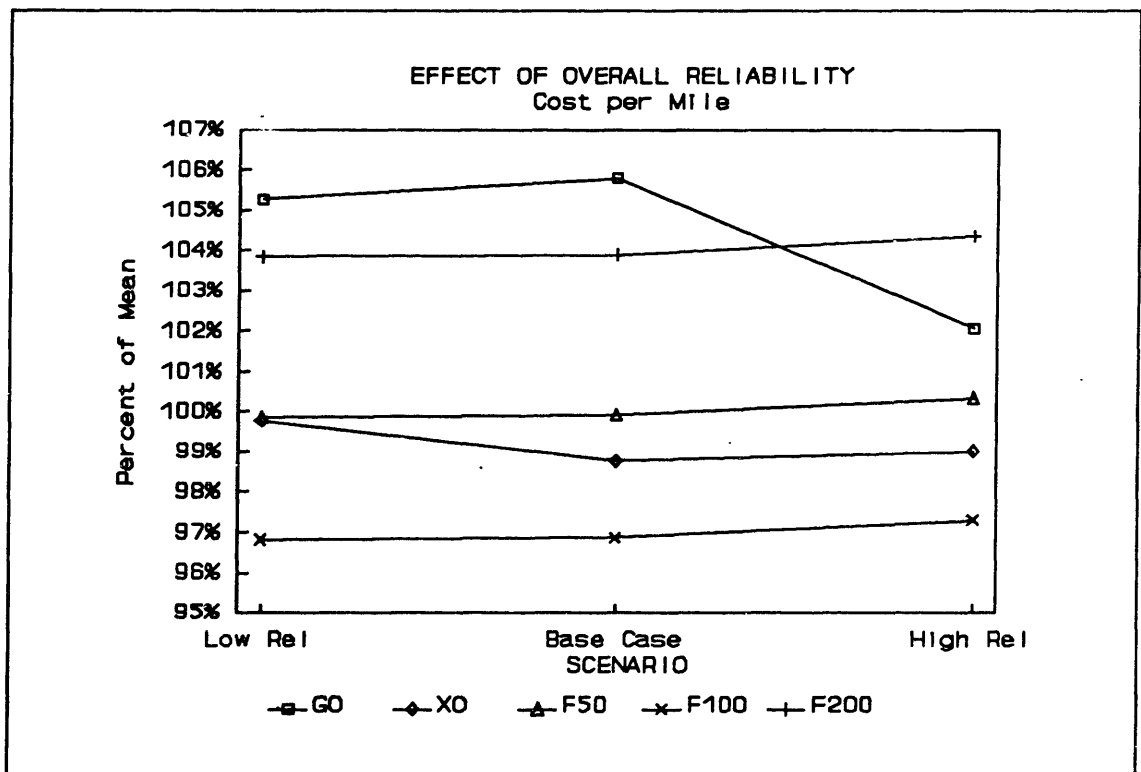
Figure 8.11

Figure 8.12

service. As the car becomes more prone to random failures (i.e., less reliable as a whole system), the point of indifference increases, since a later, "better" opportunity becomes more likely.

The impact of overall reliability on cost per mile is shown in Figure 8.12. The relative performance of most of the policies is not greatly affected in this measure by the overall reliability of the car. (Only the GO policy changes greatly among the far sighted and opportunistic policies. This is because as fewer opportunities arise in the high reliability case, GO becomes more and more similar to XO.) Figure 8.12 also shows that only the F50, F100, and XO policies achieve lower than average costs per mile. This again supports the notion that the XO policy achieves "good" performance under a wide range of circumstance, with low expected costs, and relatively high reliability measures.

8.4.2. The Effect of Annual Mileage

In this scenario, we examine the effects of high and low annual mileage, relative to the base case. Only the opportunistic heuristics incorporate the annual mileage of the car directly into the replacement decision, but all the policies are affected in terms of cost per mile, since higher annual mileage leads to earlier replacements (either as part of the maintenance plan or due to failures). Since costs per mile are discounted back to the beginning of the simulation (i.e., when the car was new), higher (lower) annual mileage will lead to higher (lower) costs per mile. The other effect, which occurs in the case of the opportunistic maintenance policies, is that higher annual mileage leads to a lower discount rate per mile, thus making the expected costs of an early replacement lower than under a lower mileage scenario. (Recall that in equation 6.4 - 6.9, the expected costs of an early replacement includes the present discounted value of future replacements, with the discounting done using the discount rate adjusted for mileage.)

Recall that the base case annual mileage was 66,000 miles per year. The policies were tested for both lower and higher annual mileage. In these scenarios, the annual mileage was changed to 33,000 miles per year and to 99,000 miles per year. All other variables were as in the base case. The results are given in Tables 8.12 - 8.15, and Figures 8.13 and 8.14. Table 8.12 presents the mean values (and standard deviations) for

235

Standard Deviations in Parentheses

| | | Current Policies | | | | Naive Sched. | Alternative Policies | | | | |
| | | On Cond. | Near Sighted Hard Time | | | | Opportunistic | | Far Sighted Hard Time | | |
			50K	100K	200K		GO	XO	50K	100K	200K
Low	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	31053 (3736)	29526 (3346)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	29229 (3110)	26514 (2232)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.00458 (.00053)	.00476 (.00057)	.00454 (.00061)	.00453 (.00058)	.00443 (.00053)	.00480 (.00034)	.00458 (.00042)	.00470 (.00051)	.00459 (.00044)	.00492 (.00045)
Base	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	31114 (3436)	30235 (3178)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	29912 (3075)	27932 (2476)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.01059 (.00075)	.01068 (.00085)	.01030 (.00096)	.01033 (.00086)	.01010 (.00076)	.01112 (.00051)	.01038 (.00045)	.01050 (.00075)	.01018 (.00073)	.01092 (.00086)
High	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	31411 (3737)	30295 (3500)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	30175 (3292)	28328 (2776)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.01531 (.00095)	.01531 (.00108)	.01483 (.00127)	.01488 (.00111)	.01455 (.00096)	.01591 (.00079)	.01501 (.00068)	.01506 (.00097)	.01456 (.00101)	.01568 (.00124)

Table 8.12
Effect of Annual Mileage

| | Current Policies | | | | Alternative Policies | | | | | |
	On Con OC	Near Sighted Hard Time N(n) 50K	N(n) 100K	N(n) 200K	Naive Sched. (NS)	Opportunistic Grdy (GO)	Xtnd (XO)	Far Sighted Hard Time F(n) 50K	F(n) 100K	F(n) 200K
OC	• • •	N50* OC* OC*	N100* OC* ---	N200* OC+ ---	NS* OC* NS*	GO* GO* OC+	XO* XO* ---	F50* OC* OC*	F100* F100* ---	F200* F200* OC*
N 50		• • •	N50* N100* N100*	N50* N200* N200*	NS* NS* NS*	GO* GO* ---	XO* XO* XO*	F50* F50* ---	F100* F100* F100#	F200* F200* ---
N 100			• • •	N100* N200* ---	NS* --- NS+	GO* GO* N100+	XO* XO+ ---	F50* N100* N100*	F100* F100* ---	F200* F200* N100*
N 200				• • •	NS* N200* NS+	GO* GO* N200+	XO* XO* ---	F50* F50* N200*	F100* F100* ---	F200* F200* N200*
NS					• • •	GO* GO* NS*	XO* XO* NS+	F50# NS* NS*	F100* F100* NS+	F200* F200* NS*
GO						• • •	GO* GO* XO*	GO* GO* ---	--- GO* F100*	F200* GO* GO#
XO							• • •	XO+ XO* XO+	F100* XO* ---	F200* F200* XO*
F 50								• • •	F100* F100* F100#	F200* F200* F50+
F 100									• • •	F200* F200* F100*
F 200										• • •

KEY

Each Cell Contains
Preferred Policy for:

Miles / In Svc. Failure
Miles / Maint. Event
Cost / Miles

(--- indicates none preferred)

Level of Significance:

* : 1%
+ : 5%
# : 10%

Low Annual Mileage
Table 8.13

	Current Policies				Alternative Policies					
	On Con OC	Near Sighted Hard Time N(n)			Naive Sched. (NS)	Opportunistic		Far Sighted Hard Time F(n)		
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	- - -	N50* OC* ---	N100* OC* N100*	N200* OC+ N200*	NS* OC* NS*	GO* GO* OC+	XO* XO* XO#	F50* OC* F50+	F100* F100* F100*	F200* F200* ---
N 50		· · ·	N50* N100* N100*	N50* N200* N200*	NS* NS* NS*	GO* GO* N50+	XO* XO* ---	F50* F50* F50+	F100* F100* F100*	F200* F200* N50#
N 100			· · ·	N100* N200* ---	NS* --- NS+	GO* GO* N100*	XO* XO* ---	F50* N100* N100#	F100* F100# F100#	F200* F200* N100*
N 200				· · ·	NS* N200* NS+	GO* GO* N200*	XO* XO* ---	F50* N200* N200+	F100* F100* F100#	F200* F200* N200+
NS					- - •-	GO* GO* NS*	XO* XO* NS#	F50# NS* NS*	F100* F100* ---	F200* F200* NS*
GO						- - -	GO* GO* XO*	GO* GO* F50*	--- GO* F100*	F200* ·GO* ---
XO							- - -	XO* XO* ---	F100* XO* F100+	F200* --- XO*
F 50								- - -	F100* F100* F100*	F200* F200* F50+
F 100									- - -	F200* F200* F100*
F 200										- - -

KEY

Each Cell Contains
Preferred Policy for:

Miles / In Svc. Fail
Miles / Maint. Event
Cost / Mile

(--- indicates none preferred

Level of Significance:

* : 1%
+ : 5%
# : 10%

High Annual Mileage
Table 8.14

KEY

OC: On Condition
NN: Near-Sighted
   Hard Time, Interval
   of N-thousand Miles
NS: Naive Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic
FN: Far-Sighted
   Hard Time, Interval
   of N-thousand Miles

F200

F100

GO

F50

XO

NS

N50

N100

N200

OC

Miles/In Service Failure
(Thousands)

Cost/Mile

.0049
.0048
.0047
.0046
.0045
.0044

22   24   26   28   30   32   34

Figure 8.13
Low Annual Mileage
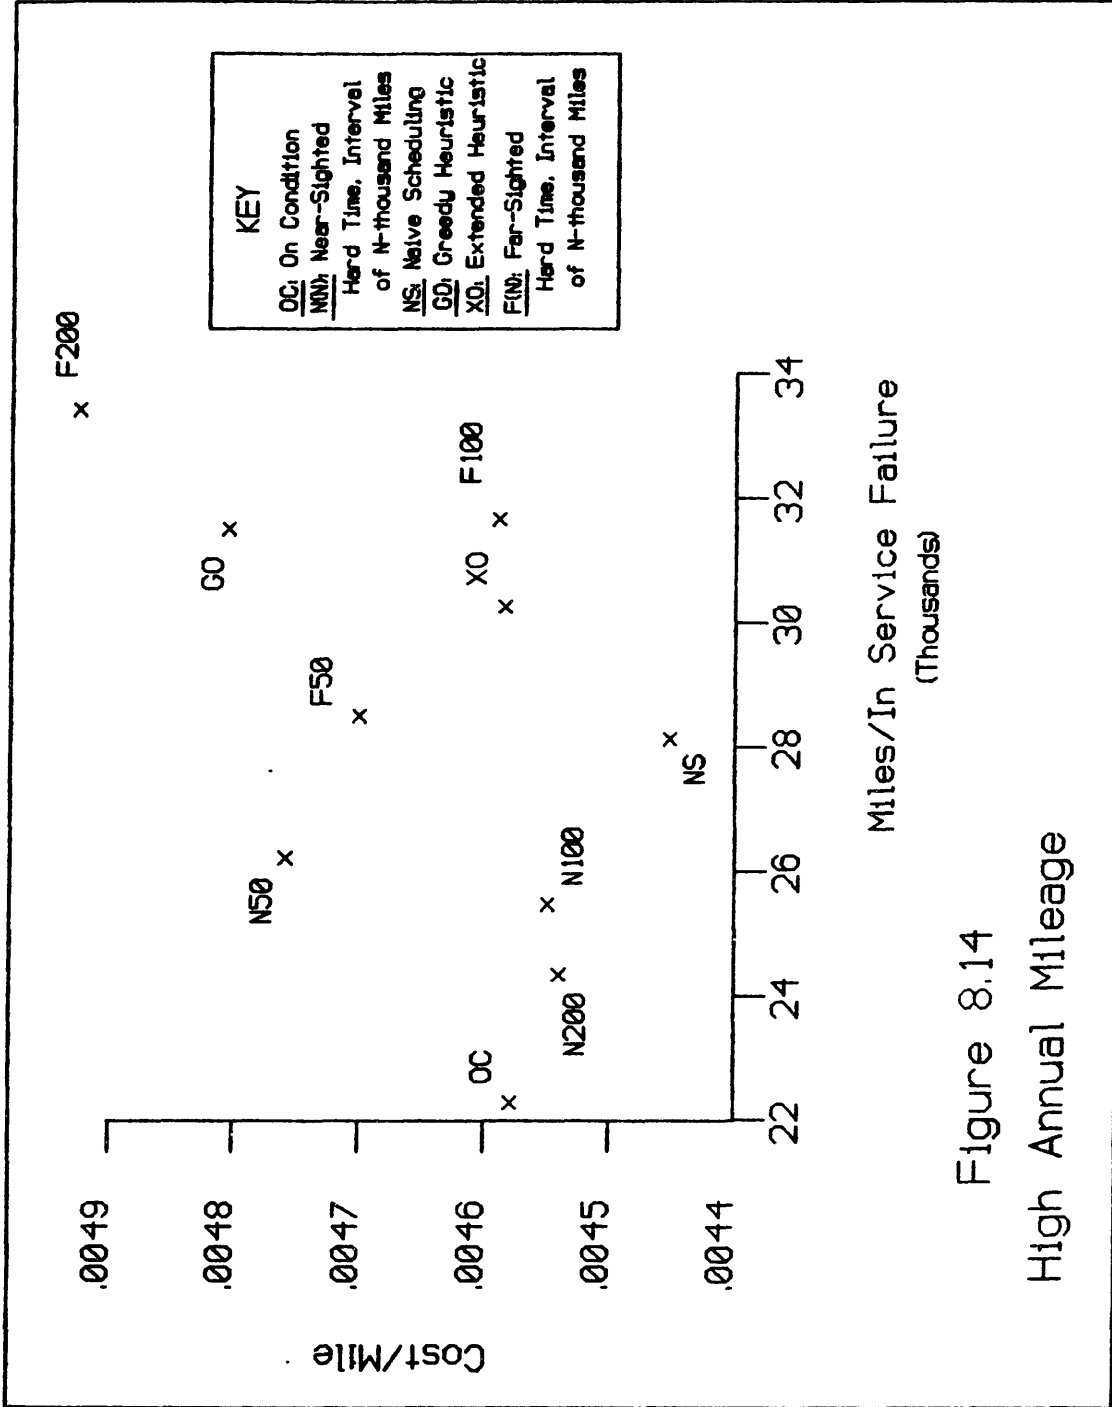
Figure 8.14

High Annual Mileage

KEY

OC: On Condition
NN: Near-Sighted
    Hard Time, Interval
    of N-thousand Miles
NS: Naive Scheduling
G: Greedy Heuristic
XO: Extended Heuristic

F(N): Far-Sighted
    Hard Time, Interval
    of N-thousand Miles

Cost/Mile

.0049
.0048
.0047
.0046
.0045
.0044

Miles/In Service Failure
(Thousands)

22 24 26 28 30 32 34

F200
GO
F50
N50
N100
N200
OC
NS
XO
F100

		Current Policies				Alternative Policies						
		OC	N50	N100	N200	NS	GO	XO	F50	F100	F200	AVG
Mi / ISF	Low Ann	79%	94%	91%	87%	100%	111%	105%	101%	113%	119%	100%
	Base Cse	79%	94%	90%	87%	100%	110%	107%	101%	113%	119%	100%
	High Ann	79%	94%	90%	87%	100%	111%	107%	101%	113%	119%	100%
Mi / ME	Low Ann	97%	75%	89%	95%	90%	128%	116%	79%	105%	125%	100%
	Base Cse	96%	75%	88%	95%	89%	130%	121%	78%	104%	124%	100%
	High Ann	96%	74%	88%	94%	89%	130%	122%	78%	104%	124%	100%
Cost/Mi	Low Ann	99%	103%	98%	98%	95%	103%	99%	101%	99%	106%	100%
	Base Cse	101%	102%	98%	98%	96%	106%	99%	100%	97%	104%	100%
	High Ann	101%	101%	98%	98%	96%	105%	99%	100%	96%	104%	100%

Table 8.15
Effect of Annual Mileage
Performance of Individual Policies v. Average for All Policies

241

the 10 trials for each of the three measures for both scenarios and the base case. Tables 8.13 and 8.14 are the pairwise comparisons of the policies under the Wilcoxon signed rank test. Figures 8.13 and 8.14 are tradeoff diagrams between miles per in service failure and cost per mile. Table 8.15 presents the "normalized" values for each policy about the overall mean for that set of trials.

As has been indicated, only the versions of the opportunistic policy (GO and XO) use the annual mileage in their decision making rules, so all the other policies are unchanged from the base case in miles per in service failure and miles per maintenance event. All the policies are changed in cost per mile because of the impact on the effective discount rate. The greedy and extended versions of the heuristic are changed in all the measures. The changes in both the absolute and normalized values of the policies are quite modest, as can be seen by examining Tables 8.12 and 8.15. The GO and XO policies show slow but steady increase in miles per in service failure. There is a similar small increase in miles per maintenance event for the GO version. The XO policy shows a noticeable increase in miles per maintenance event, but this is actually not very large (about 6.8%).

The key conclusion is that although both versions of the opportunistic heuristic are slightly sensitive to the annual mileage of the car, the annual mileage does not appear to be a significant factor in deciding among competing maintenance policies.

8.4.3. The Effect of On-line v. Off-line Usage

One of the most important issues in railroad car maintenance management is the problem with the car being maintained to an "on condition" standard when the car is off-line, i.e., on another railroad. Only the opportunistic heuristics explicitly take this into account in decision making, although all the policies are subject to the effects of potentially higher costs incurred when the car fails (or exceeds A.A.R. standards) and is repaired off-line. These scenarios attempt to examine the effects of higher and lower percentages of off-line mileage. In the low off-line case, the car is assumed to accumulate only 10% of its miles off-line, while in the high scenario, the car is off-line for 90% of the miles in service. While the high off-line case is potentially realistic, the low off-line mileage case suffers from a problem. If a car were under the owner's control

for 90% of the miles or time it is in use, the appropriate tradeoff would not be between being repaired at on-line or off-line facilities so much as between being repaired at the owner's efficient or inefficient shops. In this regard, the assumption that all on-line repair points have the same cost structure falls down somewhat.

The results of the high and low off-line mileage are reported in Tables 8.16 - 8.19 and Figures 8.15 and 8.16. Table 8.16 presents the means (and standard deviations) for the 10 trials for the two scenarios and the base case. Tables 8.17 and 8.18 are the pairwise comparisons from the Wilcoxon signed rank tests. Table 8.19 is the "normalized" comparison of the means of the various measures over the set of alternatives. Figures 8.15 and 8.16 are tradeoff diagrams.

As in the annual mileage scenarios, only the opportunistic heuristics use the share of miles spent off line in decision making. In the low off line mileage case, the greedy version obtains many opportunities for replacing unfailed components and as a result it achieves very high miles per in service failure and miles per maintenance event. As the share of miles spent off line increases, the performance of GO in these measures steadily decreases. The XO, on the other hand, appears to be only slightly affected in miles per in service failure, although it also declines in miles per maintenance event. These declines in the reliability measures as the off line share declines is probably related to the decreasing number of *usable* maintenance events for opportunistic maintenance. In particular, as the car spends less time on line, component failures are more likely to occur off line and not create opportunities for preemptive maintenance. As can be seen in Figure 8.17, as the share of miles off line increases, XO becomes more attractive relative to GO, particularly since XO is also lower in costs per mile. (Although the mean miles per in service failure of XO is higher than GO in the high off line case, the difference between the two are not statistically significant.)

In terms of costs per mile, the on condition (OC) policy is the most profoundly influenced by the share of off line miles. As can be seen in Table 8.19, OC varies from being one of the lowest in cost per mile to one of the highest, as the off line share increases. Lest managers of fleets with low off line shares take too much comfort in this news, it should be noted again that in practice, managers may still wish to address the

Standard Deviations in Parentheses		On Cond.	Current Policies			Naive Sched.	Alternative Policies				
			Near Sighted Hard Time				Opportunistic		Far Sighted Hard Time		
			50K	100K	200K		GO	XO	50K	100K	200K
Low	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	32039 (3765)	30243 (3902)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	31186 (3354)	28546 (3154)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.00963 (.00061)	.00996 (.00071)	.00955 (.00077)	.00963 (.00072)	.00943 (.00062)	.01027 (.00042)	.00974 (.00050)	.00984 (.00062)	.00960 (.00056)	.01050 (.00067)
Base	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	31114 (3436)	30235 (3178)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	29912 (3075)	27932 (2476)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.01059 (.00075)	.01068 (.00085)	.01030 (.00096)	.01033 (.00086)	.01010 (.00076)	.01112 (.00051)	.01038 (.00045)	.01050 (.00075)	.01018 (.00073)	.01092 (.00086)
High	Mi/ISF	22219 (1548)	26389 (3104)	25466 (2897)	24481 (2282)	28099 (2997)	29859 (3051)	30128 (3552)	28497 (3246)	31818 (4376)	33474 (5006)
	Mi/ME	22219 (1548)	17225 (1309)	20259 (1814)	21795 (1798)	20505 (1121)	26082 (2316)	25962 (2352)	18103 (1311)	24063 (2490)	28605 (3635)
	Cost/Mi	.01144 (.00085)	.01130 (.00098)	.01095 (.00107)	.01102 (.00098)	01062 (.00091)	.01096 (.00054)	.01073 (.00068)	.01101 (.00090)	.01060 (.00081)	.01124 (.00091)

Table 8.16
Effect of Share of Miles Offline

| | Current Policies | | | | Alternative Policies | | | | | |
| | On Con OC | Near Sighted Hard Time N(n) | | | Naive Sched. (NS) | Opportunistic | | Far Sighted Hard Time F(n) | | |
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	•	N50*	N100*	N200*	NS*	GO*	XO*	F50*	F100*	F200*
	•	OC*	OC*	OC+	OC*	GO*	XO*	OC*	F100*	F200*
	•	OC*	---	---	NS*	OC*	---	OC*	---	OC*
N 50		•	N50*	N50*	NS*	GO*	XO*	F50*	F100*	F200*
		•	N100*	N200*	NS*	GO*	XO*	F50*	F100*	F200*
		•	N100*	N200*	NS*	N50+	XO#	F50+	F100*	N200*
N 100			•	N100*	NS*	GO*	XO*	F50*	F100*	F200*
			•	N200*	---	GO*	XO*	N100*	F100*	F200*
			•	---	---	N100*	---	N100*	---	N100*
N 200				•	NS*	GO*	XO*	F50*	F100*	F200*
				.	N200*	GO*	XO*	N200*	F100*	F200*
				•	---	N200*	N200#	N200*	---	N200*
NS					•	GO*	XO*	F50#	F100*	F200*
					•	GO*	XO*	OC*	F100*	F200*
					•	NS*	NS+	NS*	---	NS*
GO						•	GO*	GO*	---	F200+
						•	GO*	GO*	GO*	GO*
						•	XO*	F50+	F100*	---
XO							•	XO*	F100*	F200*
							•	XO*	XO*	---
							•	---	---	XO*
F 50								•	F100*	F200*
								•	F100*	F200*
								•	F100+	F50*
F 100									•	F200*
									•	F200*
									•	F100*
F 200										•
										•
										•

KEY
Each Cell Contains
Preferred Policy for:

Miles/In Svc Fail
Miles/Maint Event
Cost/Mile

(--- indicates none preferred)

Level of Significance:

* : 1%
+ : 5%
# : 10%

Low Offline Mileage
Table 8.17

	Current Policies				Alternative Policies					
	On Con OC	Near Sighted Hard Time N(n)			Naive Sched. (NS)	Opportunistic		Far Sighted Hard Time F(n)		
		50K	100K	200K		Grdy (GO)	Xtnd (XO)	50K	100K	200K
OC	-	N50*	N100*	N200*	NS*	GO*	XO*	F50*	F100*	F200*
	-	OC*	OC*	OC+	OC*	GO*	XO*	OC*	F100*	F200*
	-	N50#	N100*	N200*	NS*	GO+	XO*	F50*	F100*	---
N 50		-	N50*	N50*	NS*	GO*	XO*	F50*	F100*	F200*
		-	N100*	N200*	NS*	GO*	XO*	F50*	F100*	F200*
		-	N100*	N200*	NS*	GO+	XO*	F50*	F100*	---
N 100			-	N100*	NS*	GO*	XO*	F50*	F100*	F200*
			-	N200*	---	GO*	XO*	N100*	F100*	F200*
			-	---	NS*	---	XO#	---	F100+	---
N 200				-	NS*	GO*	XO*	F50*	F100*	F200*
				-	N200*	GO*	XO*	N200*	F100*	F200*
				-	NS*	---	XO+	---	F100*	---
NS					-	GO*	XO*	F50#	F100*	F200*
					-	GO*	XO*	NS*	F100*	F200*
					-	NS#	---	NS*	---	NS+
GO						-	---	GO*	F100*	F200*
						-	---	GO*	GO*	F200*
						-	XO*	---	F100#	---
XO							-	XO*	F100*	F200*
							-	XO*	XO*	F200*
							-	XO#	---	XO+
F 50								-	F100*	F200*
								-	F100*	F200*
								-	F100*	---
F 100									-	F200*
									-	F200*
									-	F100*
F 200										-
										-
										-

KEY
Each Cell Contains
Preferred Policy for:

Miles/In Svc Fail
Miles/Maint Event
Cost/Mile

(--- indicates none preferred)

Level of Significance:

* : 1%
+ : 5%
# : 10%

High Offline Mileage
Table 8.18

246

Figure 8.15
Low Offline Mileage

KEY

OC: On Condition
NN: Near-Sighted
Hard Time, Interval
of N-thousand Miles
NS: Naive Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic
FN: Far-Sighted
Hard Time, Interval
of N-thousand Miles

Cost/Mile

Miles/In Service Failure
(Thousands)
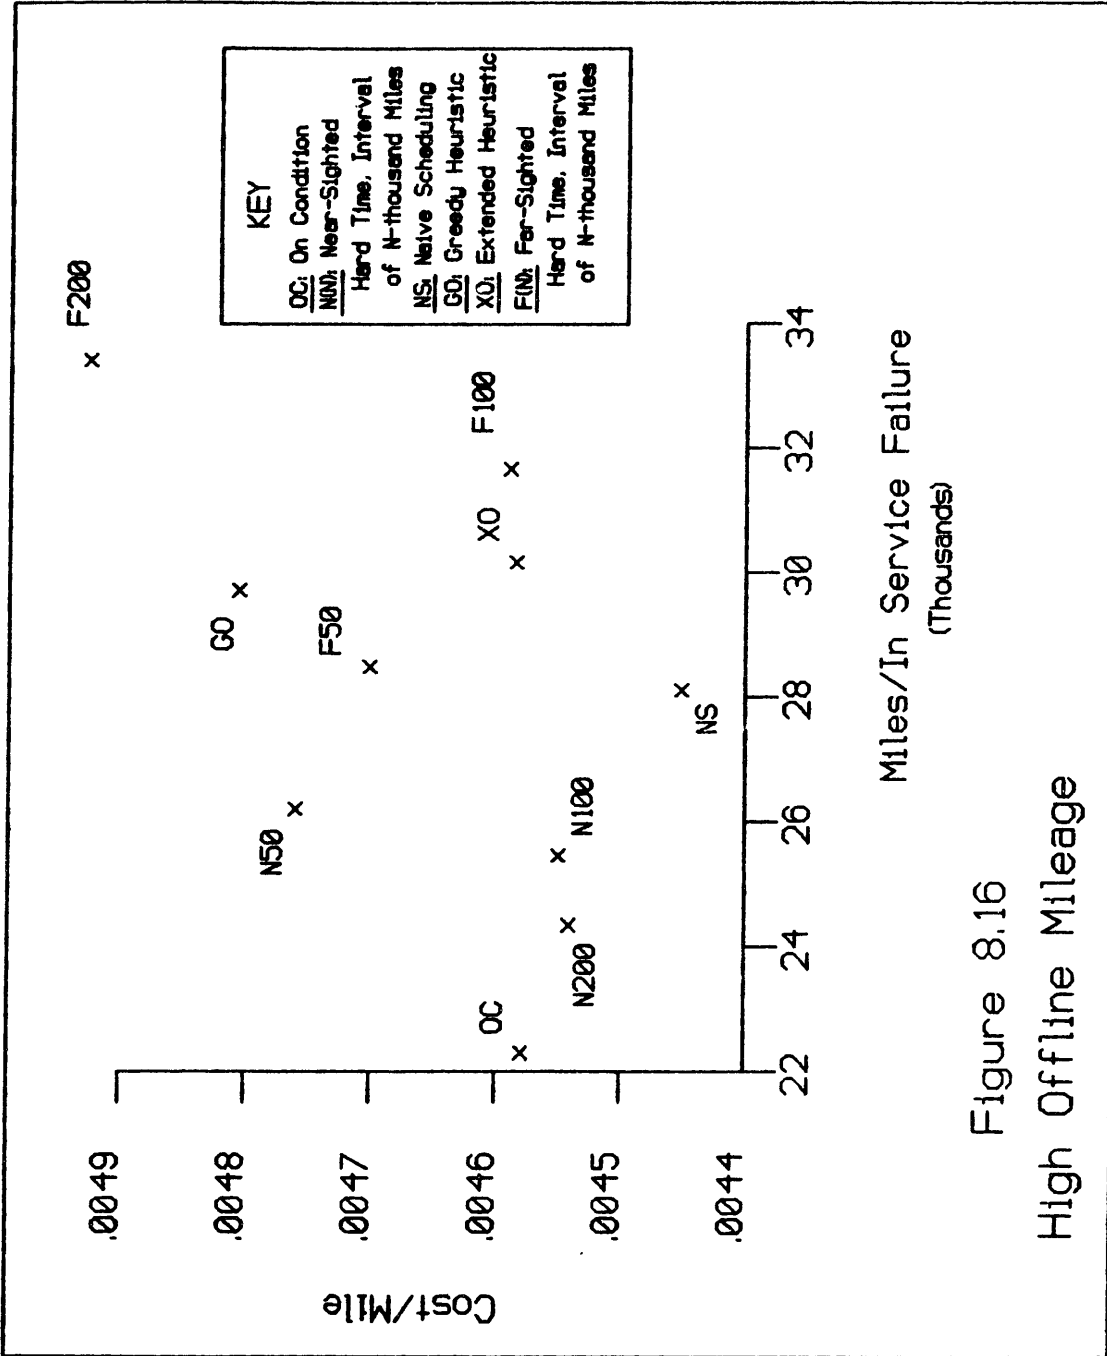
Figure 8.16

High Offline Mileage

KEY

OC: On Condition
NN: Near-Sighted
Hard Time, Interval
of N-thousand Miles
NS: Naive Scheduling
GO: Greedy Heuristic
XO: Extended Heuristic
FN: Far-Sighted
Hard Time, Interval
of N-thousand Miles

| | | Current Policies | | | | Alternative Policies | | | | | | |
		OC	N50	N100	N200	NS	GO	XO	F50	F100	F200	AVG
Mi / ISF	Low Off	79%	93%	90%	87%	99%	113%	107%	101%	113%	118%	100%
	Base Cse	79%	94%	90%	87%	100%	110%	107%	101%	113%	119%	100%
	High Off	79%	94%	91%	87%	100%	106%	107%	102%	113%	119%	100%
Mi / ME	Low Off	96%	74%	87%	94%	88%	134%	123%	78%	103%	123%	100%
	Base Cse	96%	75%	88%	95%	89%	130%	121%	78%	104%	124%	100%
	High Off	99%	77%	90%	97%	91%	116%	115%	81%	107%	127%	100%
Cost/Mi	Low Off	98%	101%	97%	98%	96%	105%	99%	100%	98%	107%	100%
	Base Cse	101%	102%	98%	98%	96%	106%	99%	100%	97%	104%	100%
	High Off	104%	103%	100%	100%	97%	100%	98%	100%	96%	102%	100%

Table 8.19
Effect of Share of Miles Used Offline
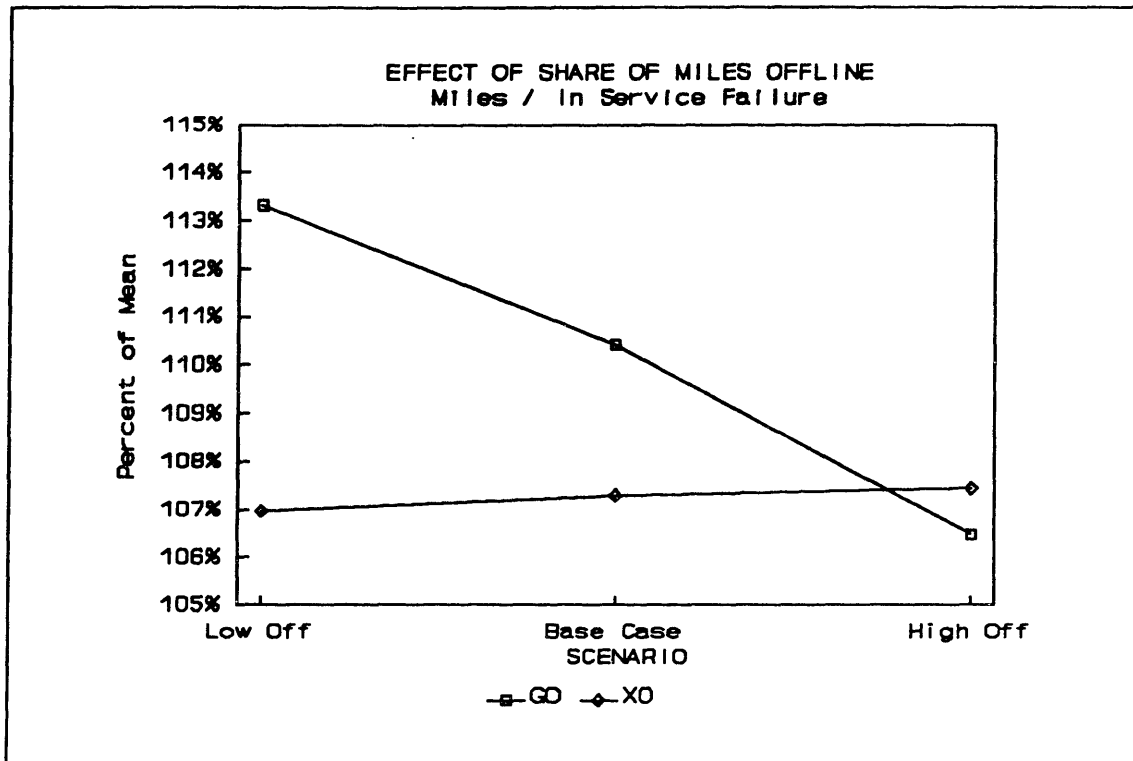Performance of Individual Policies v. Average for All Policies

249

Figure 8.17

choice between more and less efficient repair points under their control. The poor performance under the high off line share suggests that companies such as Company A in Chapter 4 (a regional carrier) should look very seriously at finding an alternative to on condition maintenance, since that policy achieves poor reliability and high costs.

Under the high off-line mileage scenario, all the various policies experience increases in costs per mile relative to the base case, except for the greedy version of the heuristic (GO), which exhibits a slight decline. The reason for the overall increase is that when in service failures occur under this scenario, the repairs are billed at the off-line (A.A.R.) rate, which can be two to three times higher than the best on-line rates. The reason why GO does not experience this increase is because it has so many preemptive replacements occur on-line that it avoids some off-line in service failures. This is because the opportunistic heuristics use the percent of time the car is off-line in the decision process, and causing the "point of indifference" between replacing and leaving in service to shift to an earlier time or mileage. The extended version exhibits this same phenomena

250

of "adapting" to the changed usage; the extended version exhibits the second smallest increase in cost per mile among the remaining policies.

8.4.4. A Comment on Discounting

One of the weaknesses noted in previous studies of reliability and maintenance is the failure to include discounting in making replacement decisions and evaluating their consequences. The simulation model discounts costs back to the present time (time 0), and the opportunistic heuristics incorporate the discount rate into the decision rule. The effect of discount rates has been captured, more or less, in the annual mileage scenarios, and is not believed to be a crucial factor in selecting a maintenance policy. There is, however, a point which bears mentioning in this regard. SInce most studies do not discount future replacement costs, those studies have a bias toward preventive maintenance programs, since there is no penalty for early replacement (i.e., opportunitiy cost), and there are rewards in the form of reduced out of pocket costs and higher reliability. This bias has been avoided in this thesis by discounting. One of the effects of this is that the heuristic does not dominate all other alternatives. Several runs of the model were made (using an earlier version of the heuristic) in which replacement costs were not discounted. The effect was that the opportunistic heuristics achieved the lowest cost per mile of any policy. The conclusion is that research results which do not discount replacement costs

8.5. Observations and Commentary on the Policies

Having presented the numerical results of the simulations of the policies, it is useful to extract some insight regarding the policies themselves and the circumstances in which they are most likely to be effective. To accomplish this, we discuss each policy separately, and then draw some general conclusions.

8.5.1. On Condition Policies

The on condition policies, which are among the most commonly followed by railroads in the U.S. and Canada, did not generally perform very well. Under every scenario and set of distributions, the on condition policies were the worst in terms of

miles per in service failure. In miles per maintenance event, while certainly not the worst policy, on condition was always below average in this measure. Even in costs, which are presumably what has lead to continued use of this policy, the on condition approach was generally not a particularly good performer except in the low annual mileage and low off line mileage cases. Even in those cases where the costs are lower than most of the alternatives, it is hard to imagine circumstances where the car owner would be unwilling to pay a 1-2% increase in costs for a 36% increase in miles per in service failure (comparing OC and XO in the low off line mileage case). The only conclusion that one can draw from this is that maintenance managers in a wide variety of circumstances should be considering the use of alternate policies rather than on condition maintenance.

8.5.2. Hard Time Policies

These policies are divided into two groups, the "near sighted" and "far sighted" versions. From discussions with the car owners reported in Chapter 4 (and subsequent interviews with other railroad officials) it is apparent that the "near sighted" policies reflect the current practice, since all the companies rely on the experience and judgement of their maintenance personnel, and all acknowledge that they have no formal standards for determining which components are worn enough to merit replacement as part of planned maintenance. In the simulations, both the currently followed "near sighted" policies and a more aggressive "far sighted" approach were modelled. The results suggest that for a given schedule of visits to the repair track, the "far sighted" policies are to be preferred, particularly in terms of miles per in service failure and miles per maintenance event. When the maintenance interval can be "tuned" to the failure distributions of the parts included in the maintenance program, it would appear that the "far sighted" policies are also preferable to the "near sighted" in terms of cost per mile. This would suggest that if maintenance managers intend to stay with hard time policies as their primary means of performing planned maintenance, they should develop aggressive standards for parts replacement (such as those used in the simulation which depend upon the failure distributions of the parts), and encourage their maintenance workers to apply them aggressively. This flies in the face of one of the concerns of most railroad managers,

which is to avoid "gold-plating" things, i.e., to be very cautious about spending money or replacing parts which are not broken. Recent reports in the railroad trade press indicate that a shift toward aggressive preventive maintenance is underway [Shedd (1990)].

Of more concern is the apparent lack of robustness of these policies to a variety of failure distributions and component mixes. When the failure distributions in the base case were varied and alternate distributions used, the performance of the "far sighted" policies varied wildly. Some intervals, such as the F200 went from being one of the best policies to one of the worst, depending solely on which parts were included in the maintenance program. This strongly suggests that maintenance managers should expect some difficulties and inefficiencies in implementing hard time policies as the appropriate time interval and parts mix is developed. Consider the difficulty if the decision is made to bring a series of cars in every 100,000 miles and then later some parts are included which have characteristic lives of 75,000 miles. If the hard time interval is not reduced to less than 75,000 miles, the policy may become inefficient; changing the time interval for the cars may, however, be problematic, particularly if the shop is at or near capacity and many different car series must be planned for.

One option would be to limit hard time programs to very long lived parts, and to use some other policy (such as the opportunistic heuristics) for the other parts. Such "hybrid" policies would appear to be a fruitful area for research.

8.5.3. The "Naive Scheduling" Approach

These policies were originally proposed as a sort of "straw man" to show why one cannot simply treat each of the parts as individual systems. In one of the ironies that make research interesting, the naive scheduling (NS) policy proved to have interesting characteristics of its own. The first interesting characteristic was that it tended to outperform the currently followed policies in both miles per in service failure and cost per mile over virtually all scenarios. Indeed, NS was generally superior in cost per mile to all other policies modelled. The good performance in cost may reflect a flaw in the model by assuming that increases in the number of maintenance events have no additional

253

cost penalty beyond higher switching costs; this failure to consider potential problems in capacity and disruptions to customer service by revisiting the "rip" track makes NS seem more attractive in cost than it would really be. The real drawback, however, is that NS was generally a poor performer in miles per maintenance event, which suggests strongly that this is not a robust approach to planned maintenance. In particular, if the number of components was increased from 8 to 16, one would expect that the miles per maintenance event would continue to drop, until it became apparent that there is a need to combine maintenance events according to some other rule.

Notwithstanding these difficulties, naive scheduling, using, for example, the manufacturers' recommended lifetimes for a small number of parts may be an acceptable method of implementing a planned maintenance program for small railroads (such as short lines) which are currently using on condition policies and are not experiencing shop capacity problems.

8.5.4. The Opportunistic Heuristics

An important objective of this chapter was to evaluate the two versions of the opportunistic heuristics developed in Chapter 6. One must conclude that these approaches offer robust and significant improvements over the currently followed maintenance policies. The extended version, in particular, is generally better in both reliability and cost per mile than any of the current approaches, and is able to retain its effectiveness under a wide variety of circumstances. This quality of robustness under a variety of circumstances should make it very attractive to maintenance planners, who can be confident that reliability goals established under this policy may be achieved even if circumstances change somewhat. The greedy version is much better in miles per in service failure than all but the very best of the "far sighted" hard time distributions. Given that the "far sighted" policies can do quite poorly if the "wrong" parts or failure distributions are included, one can argue that the opportunistic policies are the best overall performers under a wide range of circumstances.

It is interesting to compare the greedy and the extended versions of the heuristics. The first conclusion one can draw is that the greedy version of the heuristic is virtually

always better than the extended version in miles per in service failure and miles per maintenance event. The extended version, on the other hand, is almost always lower in costs per mile. These differences are to be expected given the nature of the two approaches. The greedy version calls for earlier replacement of parts, which should lead to higher survival probabilities for those parts. The extended version bases the decision on the expectation of lower future costs due to better opportunities, and so should achieve lower costs per mile than the greedy version. Because of this, most of the scenarios would call for a tradeoff analysis to determine whether the increased reliability under the greedy version is justified by the higher cost per mile.

Even without performing the trade off analysis, however, there would appear to be some circumstances under which one or the other of the policies would appear to be preferable. The greedy policy seems best suited to the cases of low reliability, and low off line mileage. In each of these cases, the greedy policy is able to take advantage of circumstances to achieve quite high reliability. In these cases the high reliability follows from the many opportunities for preemptive replacements.

The extended version appears to be particularly well suited to the high off line mileage case and the base case, where it achieves high miles per in service failure and retains control over costs. Its most noteworthy feature may be how stable it is over the full range of scenarios. The extended version virtually always varied from the base case mean values by less than the greedy version in miles per in service failure and miles per maintenance event. In those cases where the base case failure distributions were used, XO was also generally more stable in cost per mile than the greedy version. In other words, the extended version appears to be less sensitive to changes in circumstances than is the greedy version.

If the circumstances which the maintenance planner faces are not well known, or the plan to be followed is subject to changes, the extended version of the opportunistic heuristic may be preferable to all other policies because of its demonstrated robustness. Put simply, one need not fear that a change in the mix of parts included in the program (or their quality), or changes in the usage or overall characteristics of the car will result in a dramatic shift in the expected outcome of the maintenance policy.

This attractiveness of the heuristics in cases of management uncertainty is somewhat moderated by the need for reasonably accurate information regarding the failure distributions of components. Any of the policies which depend upon statistical estimates of failure rates demand that the estimates be well made and properly used. Which policies perform best when inaccurate estimates are used is a subject for further research.

8.6. Conclusions

In this chapter we have examined the current approaches to railroad car maintenance and some alternatives by using an event based simulation model. The policies tested included the on condition and "near sighted" hard time policies currently followed by railroad car owners, a naive scheduling approach, a more aggressive hard time approach and the two versions of the heuristics presented in Chapter 6. The policies were tested under a number of failure distributions and operating scenarios.

The most important conclusions are:

- railroad car owners could achieve considerable performance improvements at lower costs by replacing **on condition** policies with a planned maintenance strategy.

- Among the alternatives, the **"far sighted" hard time policies** obtained the best reliability, but were very sensitive to the failure distribution and mix of components being considered. Failure to match the hard time interval with the failure distributions can result in very poor performance.

- Among the **hard time** policies, the **"far sighted"** policies greatly outperformed the currently followed "near sighted" ones, which should encourage maintenance managers to be very aggressive in conducting planned maintenance programs.

- The **opportunistic heuristics** outperformed all the currently followed policies in miles per in service failure and in miles per maintenance event. The **extended** version of the heuristic achieved both high reliability and low cost per mile. Both versions were very robust under a variety of circumstances, which may be important to maintenance planners. Indeed, only the

opportunistic heuristics were able to achieve high reliability under a broad mix of failure distributions.

The purpose of the chapter was to evaluate maintenance policies followed by the owners of freight cars, and some feasible alternatives. What we have learned is that car owners can realize substantial improvements in both cost and reliability by adopting alternative maintenance strategies. Among the alternatives, the extended version of the opportunistic heuristic dramatically outperforms the current practices in all measures, and the greedy version does so in the reliability measures. These heuristics do not only improve upon the performance of the current policies. They do so while exhibiting a robustness across all scenarios which is not matched by any of the alternatives. Only when the maintenance interval can be properly selected are the "far sighted" hard time policies to be preferred. Even in those cases, "hybrid" approaches may be appropriate to realize the benefits of both hard time and opportunistic approaches. In the next chapter, we turn to the problem of how to manage the information needed to implement an effective maintenance policy.

Chapter 9

Improving Railroad Car Maintenance:
Implementation Issues

9.1. Introduction

It was concluded in Chapter 5 that the effective maintenance of railroad freight cars is constrained by three serious problems:

- the policies followed,

- the measures used to monitor car maintenance activities, and

- the information systems used to store and manage the data needed by maintenance planners and managers to implement effective programs.

The first two of these problems have been addressed in earlier chapters. To the current group of maintenance policies has been added a set of opportunistic maintenance heuristics which perform as well as, or better than the current approaches; specific measures for vehicle maintenance effectiveness and efficiency have been introduced and used to evaluate alternative maintenance policies. What remains is to provide a means of implementing these new approaches in the actual environment faced by the railcar owner. That is the focus of this chapter. In particular, this chapter introduces a construct which permits managers to overcome the problems associated with the current information systems and permit them to implement and monitor improved policies.

The information systems currently used by most freight car owners were developed to meet accounting, auditing, and regulatory needs, and are now widely used by many departments and functions. As such, these information systems represent long term investments that car owners have made on behalf of a wide range of interests within and outside the company. The effect of this is that information which is needed by freight car maintenance managers is often found in many very large data sets containing use, repair, and economic information. Even within these groupings, the information is often in several data sets.

What is called for is a unified data source which allows maintenance managers access to the information they need while permitting the company to continue to use its

current information systems and the application programs they support. In order to meet both these objectives, the approach must accurately reflect the way that maintenance managers and planners think about cars rather than the way that data processors store data. It must provide all the information needed to plan and implement the appropriate maintenance policies examined in Chapters 7 and 8. It must also be sufficiently inclusive to permit the development of new techniques and technologies (including better maintenance policies) such as expert systems. Indeed, some of the new technologies have the potential to permit managers to undertake approaches to maintenance management that would have been unthinkable several years ago.

The chapter is organized into three sections. The first section outlines in more detail the information and data problems currently faced. These problems are clarified by looking at the information needs required to support an opportunistic maintenance program or to examine the effects of a "hard time" policy. The second section presents a remedy for these problems, the "structured car history". The third section speculates briefly on ways that these structured car histories might be applied to address other aspects of vehicle maintenance.

9.2. Information/Data Problems

Railroads were among the first industries to use computerized information systems, partly in response to extensive regulatory reporting requirements and partly from a desire to reduce clerical costs. While the result of this was certainly beneficial in meeting the companies' needs at the time, one residual effect is that the data which is collected and stored today often reflects historical needs and departmental interests rather than current managerial directions. The wealth of data in old data sets, user programs, and supporting systems provide a powerful force against reorganizing the data and information systems to meet new needs[1]. At the same time, the complexity of the information systems can serve to make it an expensive proposition to gather data from several sources within a company to support systematic approaches to solving problems.

---

[1] See, for example, Lott (1971).

Consider, for example, the information needed to perform the calculations for the greedy opportunistic heuristic from Chapter 6 and the data sources from the railroad car owners studied in Chapter 4. The heuristic calls for the following information:

- *Failure Distributions*, which should be derived from the historical maintenance data of the company, such as that used to support the company's participation in the A.A.R. Car Repair Billing system, and a usage measure, such as

- *Annual Mileage*, which would generally be derived by linking the car's trips (reflected in a waybill master file, which may not even be keyed on particular car initial and number data) and a mileage table.

- *Repair Costs*, which in the case of off-line repairs would be found in A.A.R. billing data tapes or the **Office Manual**, are also needed. For on-line repairs, if the costs are available at all, they would consist of some measure of labor inputs (either the A.A.R. labor requirements or a company's own standard) which must be matched to the company's wage rates (and overheads), and costs for materials, either from purchasing or inventory data sets.

- *The Discount Rate*, which ought to reflect the company's financial decisions about the available return it can get by investing in alternative projects, is needed to adjust future expected costs, but it may well be found only in a few LOTUS 1-2-3 spreadsheets in financial and marketing department offices[2].

As can be seen, to implement a relatively straightforward heuristic for opportunistic repairs, the car owner must gather data from a number of diverse sources, some of which may not be structured to easily support such analysis (if available at all), and link them together.

---

[2] It is somewhat artificial to imagine that there is a single number used throughout most companies to represent the "proper" discount rate used for evaluating investments, particularly in what are more or less operational matters such as maintenance. In practice it may be more appropriate to estimate the effects of investments over a number of possible values for the discount rate in order to determine at what values projects are or are not attractive. Larger companies generally provide guidelines for discounting.

The hard time policies are just as demanding of appropriate information. These policies, if they are to be implemented effectively, require information regarding the failure distributions (including usage data) and costs of replacing components, and additionally require some sort of guidance to field personnel on the handling of short-lived components. Recall, for example, how sensitive the F(200) policy (i.e., far-sighted hard time maintenance at 200,000 mile intervals) was to the failure distributions in the simulations in Chapter 8. If the parts being replaced exhibited characteristic lives less than 200,000 miles, the policy performed quite poorly. In practice, this means that maintenance planners would want to examine maintenance data carefully to select which components to replace during scheduled visits to the shop.

For both opportunistic and hard time policies, the data must come from a number of diverse sources. Consider the data sets of Company A, a regional railroad presented in Chapter 4. The maintenance records are stored as part of a car repair data base which is also used for generating and auditing car repair billing system transactions. These records within the data base are therefore organized in terms of the individual components and maintenance actions which often comprise a larger repair. Thus, what a manager considers a wheelset replacement is entered into the car repair database as two wheels, two roller bearings, an axle, a labor charge, and other miscellaneous parts. The creation of failure distributions thus entails a redefinition of what is being analyzed into the relevant parts of the system as stored in the data base. Mileage data for individual cars is only available for off-line car moves (because this data is provided by the A.A.R. as part of it's TRAIN II car movement reporting system). Within the information systems of Company A, this information is collected and stored only for series or groups of cars, not for individual cars. The off-line mileage, of course, could be stored for each car, if there were an appropriate repository for the data. In order to generate on-line miles for each car, it would be necessary to inspect the waybill file, which records the car movement authority for all on-line car moves, and match this to a table of mileage between points. Cost data for off-line repairs is provided through the CRB system, and is stored for each repair on each car. Company A also "charges" itself the A.A.R. rate for on-line repairs. To generate the actual on-line labor costs, the company would have

to either match the A.A.R. labor requirements for each (component-level) job code (which it stores as part of its car repair data base) or some substitute standard to their own wage rates (which are considerably lower than the industry average). For on-line material costs, Company A can refer to purchasing records, which are stored in a separate file.

It is important to recognize that the barriers that restrict a car owner from gathering this data from the various sources each time a repair decision is to be made are not technological. Rather the barriers are institutional, and it seems clear that an integrated approach to the management of the data is needed to improve maintenance programs. This is especially true where the costs of developing and maintaining information systems are high, or where there is uncertainty regarding the final form of the maintenance policy. In the next section, an integrated approach is presented.

9.3. Structured Car History

As part of a demonstration project at the A.A.R.'s Affiliated Laboratory at M.I.T., a study was made to analyze the usefulness of knowledge based systems (KBS), also known as expert systems, for railroad freight car maintenance [Little and Martland, 1989]. The focus of the project was primarily on using KBS to search through large data sets such as that generated by the Car Repair Billing System to find cars which might be experiencing excess component replacements, and diagnose the cars for appropriate maintenance actions. One of the important outcomes of that research was the recognition that the data sets which store relevant information for car maintenance do not reflect the way car maintenance managers think about and analyze car repairs. As a result, maintenance managers had to invest considerable effort to gain access to data which they needed to manage and monitor car performance. Even when they obtained the data, it was organized along lines which reflected the concerns of other users of the information, or of the information systems professionals who manage the data. The desire to find a way to structure the data to correspond to the needs of maintenance managers led to the notion of the "structured car history".

The idea behind the structured car history was to put together into a single data set all the various information about a freight car which a manager might reasonably need

when making maintenance decisions. Toward this end, a structured car history is a single data record for a car which includes information about the car itself, the usage of the car, maintenance data, and performance measures. Gathering all the relevant data about the fleet into a file of structured history records makes it relatively easy for managers to access the information needed to plan and monitor maintenance activities. This is particularly true when compared with the current situation, in which the necessary information may be in several different files each with its own unique format. More importantly, the data can be organized in the way that managers think about freight cars rather than the way the originators of the company's data sets thought about accounting, or car hire payments, or car movement systems. The most important element is deciding what managers really want and need to know about a car. This process of determining how problem solvers organize knowledge to think about problems is frequently performed as part of building expert systems, and is known in that field as *knowledge engineering*[3]. The fundamental step in the process is to spend time with experts in the field, and watch them analyze situations or problems, asking them questions, and noting the way that they organize all the information available to them.

It was found by Little and Martland (1989) that there were four basic areas that the experts seemed consistently to refer to in analyzing cars. The structured history was formed around these four basic areas:

1. *Car characteristics*: These include basic data such as car initial and number, car type and age, and engineering information regarding the design and structure of the car such as brake mounting type, control valve type, etc. Much of this data is available in standard sources, such as the Uniform Machine Language Equipment Register (UMLER), and in other cases could be derived by analysis of the repairs made on a car. In general, the owner of the car has ready access to this information.

2. *Usage characteristics*: These include annual mileage, the type of service the

---

[3] For more information on knowledge engineering, the reader is referred to Waterman (1986) or Harmon and King (1985).

car is used for (unit train vs. free running), the commodities the car typically carries, and whether the car is subject to any sort of specialized handling, or is part of any assigned pools. Revenue information should be available, or at least some indication of the value attached to the car as an operating asset; expected future demand for the car should be included if available. Once again, this data is readily available to the car's owner, and requires updating only on a periodic basis.

3. *Maintenance activity:* This section summarizes the type, date, and cause of various repairs performed on the car. Unlike the CRB data, which lists each component separately, the structured history should summarize the repairs in ways that reflect the way that managers analyze cars; for example, the 6-10 items associated with replacing a car's wheels are structured as a single item, the wheelset. In addition, any information regarding unusual maintenance or repair activities should be included, such as derailments, rebuilds, or visits to the shop for special maintenance activities. It is important that this section not be limited only to "big ticket items" or even to items for which preventive maintenance is being considered. Little and Martland (1989) found, for example, that an inexpensive item can often be an indicator of potentially much more expensive part failures (e.g., brake shoes as precursors to wheel failures). If failure distributions for components are known, or scheduled activities are to be performed at the component level, they also belong under this category; that is, maintenance activity refers to future acts as well as past ones.

4. *Performance measures:* These should include the measures proposed in Chapter 5, i.e., cost per mile, miles per maintenance event, and miles per in-service failure. In addition, however, one might include total labor and material costs for repairs, the number of failure events which required removal of the car from service, and the months a car has gone without experiencing a failure of a "critical component". Indeed, any measures which managers believe might prove useful should be included.

	FIELD	EXAMPLE
Car Information	Car Number	101
	Type	B414
	Built Date	10/1979
	Brake Mounting	Body
	Control Valve	ABDW
Usage Information	Mileage	Unknown
	COTS Date	<nil>
	Derailed Date	05/1984
Summarized Maintenance Information	Brake Shoes Used	34
	L1 Thermal Failures	1
	L2 Thermal Failures	2
	L3 Thermal Failures	2
	L4 Thermal Failures	1
	...	...
	R4 Thermal Failures	0
	L1 Flange Failures	1
	L2 Flange Failures	2
	L3 Flange Failures	2
	L4 Flange Failures	1
	...	...
	R4 Flange Failures	0
	Brakebeam 1 Date	<nil>
	Brakebeam 2 Date	11/84
	Brakebeam 3 Date	<nil>
	Brakebeam 4 Date	<nil>
	Brakerod Date	05/84
	Slack Adjuster Date	<nil>
Performance Information	Problem Start Date	06/84
	Most Recent Problem	12/84
	Failure Months	8
	Labor Cost	3153.19
	Material Cost	8043.88

Table 9.1
Structured Car History

Table 9.1 presents an example of the structured histories developed in Little and Martland.

A related issue is the overall maintenance of the structured history file. It seems clear that this data should be kept as up-to-date as is possible. Several times when updates would be particularly important are

- when information is received from the Car Repair Billing system regarding off line repairs;
- when repair actions are being considered for a particular car;
- when maintenance performance is being evaluated.

To accomplish this, it would seem that the car owner would want to be able to update individual cars (as when actions are being considered for that car), and the fleet as a whole (as when performance is being assessed.) The balancing of the costs of developing software to provide timely updates and the value of that information is an important matter that should be carefully considered by the car owner.

By collecting together the data from many diverse sources and organizing it so that it corresponds to the approaches of the manager, it should be possible not only to provide the material on which to base a planned maintenance program (whether opportunistic or not), but also should permit a number of new applications to be developed. In the next section, some of the possible new approaches are speculated upon, and a particular example of how the structured history has already been used is presented.

9.4. Uses for Structured Car Histories

In the previous section, we saw in general terms how a freight car owner might organize a structured car history, and motivated much of that discussion by providing the necessary information for a planned maintenance program, whether of the hard time or opportunistic sort. In this section, we look at a few examples of related but distinct aspects of car maintenance management to which the structured history is also applicable here. The intent of this section is not to advocate these particular applications, but rather to show that the use of an integrated approach to freight car information can lead into other areas as well as the design and support of planned maintenance programs. No

pretense is made that the list is exhaustive, and some of the proposed applications may not be relevant to all freight car owners.

9.4.1. Railroad Car Diagnostic System (RCDS)

Not surprisingly, the first example is drawn from the work which led to the concept of the structured history, the diagnosis of freight cars with excessive and repeated consumption of critical components. Little and Martland (1989) found that some railroad cars, running on what are generally considered to be well managed fleets were experiencing extraordinarily high consumption of wheels and brake shoes. (The study was limited to cars with brake and wheel-related problems because it was intended as a demonstration of expert systems technology rather than a production system.) That study, which examined the complete fleet of a small railroad and a sample from a Class I carrier found that in both cases some 3-5% of the cars were responsible for 12-20% of the maintenance expenses, even after correcting for the usage patterns[4]. These "problem cars", as they were characterized, were generating thousands of dollars in repair bills for the same components. They also found that the cars could be identified, and, in many cases diagnosed, leading to specific repair recommendations by using an expert system, Railroad Car Diagnostic System (RCDS).

RCDS operates by analyzing cars by a four step process:

1. The available data for a car is processed to create a structured history.

2. The structured history is evaluated using a set of rules and thresholds to determine if there is evidence that the car is experiencing excess parts consumption; that is, is it a "problem car". The decision regarding the condition of the car can result in several possible outcomes:

    - the car is OK;

---

[4] Because no mileage data was available for individual cars, the "problem cars" were separated into groups with the same car types and ages. The data was then examined to determine whether the problem cars were clustered into particular groups. They were not; that is, the problem cars were "outliers" with respect to the separated groups as well as the fleet as a whole. In interviews with the cars owners no indications were reported that the excess component consumption was due to any usage patterns. In other words, it appears that the problems were due to defects in the cars themselves.

- the car is experiencing excess brake related wheel failures;
- the car is experiencing excess wear-related wheel failures;
- the car is consuming an excessive number of brake shoes.

3. If the car is not OK, then a set of diagnostic rules appropriate to its condition are invoked and possible reasons for the car's problem are determined. For example, if a car was involved in a derailment and began having problems shortly thereafter, the problem may be that the car's brake system was damaged or wrongly repaired at the time of the derailment.

4. After the car is diagnosed, possible actions to correct the problem are recommended. These recommendations can be of either an engineering or a managerial nature. If, for example, it appears that a car is having problems due to a mechanical defect, then a corrective action is recommended. If, on the other hand, the use of the car may be inappropriate (e.g., using a car with an older air brake design in a consist with newer cars), then alternative management options are given.

RCDS used a simplified version of the structured car history as one of its basic elements. The reason for the use of the structured history was twofold:

1. The structured history permitted the volume of data to be reduced to a size that would permit the use of personal computers for the expert system (reducing a 30 Mbyte repair data base to 179 Kbytes); and, more importantly,

2. the structured history conformed closely to the way that their panel of freight car engineering and repair experts "thought about" the problem, so that they could construct rules and diagnostic approaches which captured the experts' knowledge.

The conclusion of the study was that the Class I railroads could save over $60 million annually in excess wheel repairs alone by early detection, diagnosis, and repair of so-called "problem cars".

9.4.2. A Repair Track Scheduling System

A second, more speculative, use of the structured car history might be in the area of repair track scheduling, i.e., deciding which cars to bring into the shop on a given day.

As was discussed in Chapter 3, this decision is based on a number of factors, including the availability of parts, and the priority which is attached to a car. Automation of this decision process either by an expert system, or through a conventional program, would be greatly simplified by the use of a structured car history. Consider just a few of the items which would likely be required to build an effective decision-support system for this problem:

-   Basic information about the car, such as owner, the car hire rate, and the car's current loaded or empty status;

-   the nature of the defect, including what labor and materials might be required to repair the car;

-   the value of the car as an economic asset (i.e., should the car be fixed at all?

-   the expected future demand for the car, and any seasonalities which affect demand (e.g., fixing a grain car during planting season is probably a lower priority than just before the harvest);

-   is the car due for other repairs or maintenance?

Other than the car's current loaded status and the reason why the car is currently shopped, each of the above items is part of the structured history as proposed.[5] In other words, a railroad using a structured history could build such an expert system with significantly less difficulty than a railroad that had each of these items stored in separate data bases. The operation of the system using a structured history would probably be much faster as well, as the system could be optimized in searching one file rather than having to access many different data sources. In other words, the use of the structured history for planning and monitoring maintenance would also have the benefit of making operational decisions more efficient.

---

[5] It might prove to be the case that the demand information could be more efficiently kept in another file, which would record the demand for general classes of cars, and the structured history would indicate to which general class the car belonged. That is an implementation issue that would depend on the structure of the actual system,

9.5. Conclusions

In this chapter, we have addressed the third major barrier to effective freight car maintenance first presented in Chapter 5, namely, that the information systems used by most railroad car owners are difficult to use for effective maintenance management. While it is true that the information systems have evolved over an extended period of time for purposed unrelated to car maintenance, and that these systems now have a committed base of users who would be resistant to changes, the problem is not insurmountable. What is called for is to gather the data which is required by maintenance managers into a single file, the structured car history, and allow that to be the primary information focus of information users. This not only frees the maintenance managers from accessing multiple data sets to gain information, it permits that data to be restructured in ways that are more compatible with the methods of analysis used by car repair managers and planners.

The structured car history has already been shown to be a useful way to organize car repair and related data in another area of the car repair arena. In that case, it was used as to support the development of an expert system for diagnosing problem cars. It is easy to speculate on other areas of car maintenance management that might also benefit from the use of structured histories.

The primary focus remains, however, the development of more productive maintenance policies and the tools with which to monitor them. In that context, the use of the structured history promises to be particularly useful.

Chapter 10

Conclusions and Future Research

10.1 Introduction

This thesis can be summarized in a simple phrase, "the maintenance of railroad freight cars can be improved". That simple statement is, in a sense, both a commentary on the current practices and policies and an expression of optimism, since a review of car owners shows them willing to adopt the better ways of going about their business proposed in this thesis. In this chapter, we briefly review the current state of affairs in car maintenance and repair, which were shown to suffer from at least three shortcomings. The solutions to these problems are then highlighted. Finally, some future directions for research are proposed.

10.2 The Current State of Freight Car Maintenance

In the U.S., railroads and other companies own and maintain a fleet of more than 1.2 million freight cars, at an annual maintenance expense of approximately $2 billion. This fleet is becoming progressively smaller, older, and more intensively used as railroads and shippers alter their transportation patterns in response to an evolving business climate. In this thesis, we have looked at the maintenance policies and practices that support that car fleet, by reviewing general trends, examining in detail the behavior of several companies, and by simulating the consequences of the currently followed approaches to maintenance.

In Chapters 3,4, and 5 the current practices and policies of freight car owners were examined, both in general terms, and in three detailed case studies. The case studies presented a spectrum of companies, including a small railroad, a large railroad, and a private car owner whose fleet is used to support the company's production and marketing of chemicals and petroleum products. This review focussed on the maintenance policies and practices, the performance measures applied to maintenance activities, and the information systems used to support car maintenance. These companies each had

different maintenance plans and policies in terms of details and timing, but their basic strategies could be summarized as falling into two categories:

- *on condition maintenance*, in which a car is operated until a component either fails or exceeds a wear standard (such as the AAR's interchange rules), in which case the failed or worn component is replaced and the car is returned to service; and

- *hard time maintenance policies*, in which a car is sent to the repair shop at fixed intervals for replacement of worn components. Between scheduled visits to the shop the car is operated using an on condition policy.

Both policies were found to have advantages, primarily ease of implementation and uniformity across carriers. The first advantage can be important, since railroads are geographically disparate, and subject to contentious labor-management relations. The second advantage is also potentially important, since the cars are maintained according to an on condition policy while they are being used on another railroad, i.e., in interchange service.

On the other hand, both policies are also subject to problems and inefficiencies. On condition maintenance fails to exploit any of the potential economies of scale in maintenance, and leaves the car subject to a high level of unreliability (since all component replacements are, in effect, in-service failures). Hard time policies are sensitive to the interval of time chosen and the failure distributions of the components which are serviced. Indeed, there are good theoretical reasons for believing that if many components are included, there is no single best time for scheduling maintenance activities. One of the case study companies seemed to bear this out with a continuing increase in its intervals for performing maintenance activities.

In addition to the case studies, the overall state of the industry was reviewed by examining articles on car maintenance in the trade press, usually a particularly optimistic forum. Interviews with maintenance managers at the case study companies and comments in the trade press all suggested that railroad car owners would be receptive to better policies if they could be developed and were practical.

The second general problem which faces freight car owners is that the measures and standards used to monitor maintenance effectiveness and efficiency are often

inappropriate, and may even lead to undesirable maintenance activities. The measures being used include percent of the fleet bad ordered, number of cars set out from enroute trains, and cost per loaded mile. Only the last measure, cost per loaded mile, can be considered an appropriate measure based on standards for measures of maintenance effectiveness in the standard reliability literature. Alternative measures were developed and used, as discussed below.

The third problem which was uncovered in the research was in the information systems used by freight car owners, and most particularly in the case study railroads. The information most needed by maintenance managers to plan and monitor car maintenance activities is generally inaccessible to them, either because of the complexity of the data structures or because of the sheer volume of data. In many cases the information systems and data bases used are monuments to the history of railroading, built to meet regulatory requirements which have since passed, or to support management and accounting functions without regard for their potential as managerial assets. The result is that many maintenance management decisions are undertaken without the full information systems support that could be made available. Implementing more effective maintenance policies is particularly dependent on organizing data regarding car design and engineering, repair histories, car usage, and expected future demand for the car. Cost control also requires that the data be organized in a way that permits managers to determine what the low cost facilities are, and then route the cars to them. In short, an integrated approach to maintenance information management is needed.

10.3 Solutions and Conclusions

Although the current state of policies and practice are subject to criticisms, one of the most positive outcomes of the research is the awareness that the industry is willing to listen to new ideas and approaches to maintenance, and is willing to invest the necessary resources to improve. Because of this generally positive attitude toward change, it was possible to press forward with constructive methods and techniques for addressing the problems listed above. The solutions and some further conclusions are discussed in the next three sections.

10.3.1. Maintenance Effectiveness Measures

The effectiveness of railroad freight car maintenance could be better measured by the adoption of three measures:

1. *Miles per maintenance event*

2. *Miles per in-service failure*

3. *Cost per mile.*

The first of these measures serves to indicate both how effectively the maintenance program is grouping together maintenance activities, and the extent to which the maintenance program itself is "interfering" with revenue operations. It is desirable that this measure be as high as possible, all other measures being equal. There can, however, be a tradeoff between this measure and miles per in-service failure for some maintenance programs (for example, the so-called "naive scheduling" and on condition policies examined in Chapters 7 and 8). In this case, miles per maintenance event is serving a "checks and balances" function against policies which are achieving seemingly high service reliability at the cost of spending too much scheduled time in the shop.

The second measure, miles per in-service failure is the direct analogue of the standard reliability measure mean time between failures, with an adjustment for the units of usage for a freight car. One can argue that trips, i.e., mileage, is the product that is sold by freight railroads, and so it is more desirable to measure the usage of cars in terms of mileage than time. Miles per in-service failure is a measure of service reliability as provided by the maintenance program. As such, it is an effectiveness measure.

The third measure, cost per mile, is an efficiency measure, which indicates how much money is spent on maintenance in order to provide a unit of transportation service (miles). This measure is based directly on the need to provide a cost-effective program of maintenance to compete in the marketplace.

These three measures were applied to the various maintenance policies developed and tested in Chapters 6, 7 and 8.

10.3.2. Maintenance Policies

It was shown that the maintenance policies currently being used by freight car owners are subject to some practical and theoretical deficiencies. One of the most

important contributions of the research was the demonstration that practical opportunistic maintenance policies can be developed and applied in railroad car maintenance. Opportunistic maintenance consists of treating the failure of one component as an opportunity to perform other maintenance actions when economies of scale in maintenance exist. Previous research suggests that such policies, while clearly desirable in theory, do not lend themselves to straightforward implementations. Indeed, most of the published research in this area has been to find ways of extending the boundaries of optimal solutions to more than a handful of components and failure distributions. Unfortunately, there are also few published heuristics for suboptimal solutions to the problem.

Freight car maintenance does not appear at first glance to be a very welcoming environment in which to apply the tools of the theory of reliability. Both the control over the car and the resulting cost structure are complex because of the need for cars to move freely over a network belonging to many companies. The industry's labor management relationships further complicate matters in many cases by effectively restricting the degree of discretion that can be allowed at the shop level.

At the same time, the railroad environment is an almost perfect case for opportunistic maintenance. Virtually every maintenance action is subject to at least some economies of scale, since the costs of removing a car from a train and switching it to the repair track are quite high. Many of the components are related in the maintenance activities required to access them. For example, changing one wheelset entails many of the same activities (such as jacking the car) as changing the other wheelset on that end of the car. Because of this high potential for returns on opportunistic maintenance, an effort was made to develop a practical heuristic which is consistent with the theory of reliability.

Two versions of an opportunistic maintenance heuristic were developed. In both cases the essence of the heuristic is to calculate the cost of replacing a component at the present time (given that the car is in the shop for some other reason or failure), and calculate the expected costs of allowing the component to remain in service until a later date (which is calculated based on standard single component replacement theory). If the

costs of replacing the component now (including the effects of a shortened component life such as foregone wear life) is less than the expected costs of replacing the component at the future "scheduled" time, then the component is replaced opportunistically. If not, the component is allowed to remain in service. Two versions of the heuristic were developed and tested:

1. *A "Greedy" version*, in which components are replaced whenever the car is on a repair track and the costs of replacing now are less than the expected costs of waiting; and

2. *An "Extended" version*, in which the future expected costs associated with leaving the component in service are weighted by the probability of a random failure in the interval between the present time and the "scheduled" time for that component.

These heuristics were tested against a number of alternatives representing variations on the current policies using simulation. The other policies tested included on condition maintenance, several types of hard time maintenance policies, and a naive scheduling approach. The policies were tested for eight components over a wide range of operating circumstances. The model developed can be used to test other approaches as they are developed.

A number of important conclusions were reached:

- The **on condition policies** currently being followed performed worse than almost any of the planned maintenance strategies in both cost and reliability. This suggests that railroad car owners who are following this sort of policy could benefit greatly by adopting a planned maintenance program.

- While the **hard time policies** which are used by some car owners are very sensitive to the way in which they are implemented and to the failure distributions of the components included in the maintenance program, a properly "tuned" hard time policy can perform very well. In particular, the hard time polices were most effective in achieving high service reliability when components were aggressively replaced during each scheduled visit to the shop and when the maintenance interval was matched to the failure

276

distributions of the components included in the maintenance program. The policies were especially ineffective in achieving service reliability when the components being maintained are shorter-lived than the maintenance interval. In other words, maintenance managers who adopt hard time policies should carefully study the components included, and, for appropriate components, should be willing to perform many preemptive maintenance actions during visits to the shops.

The hard time policies were among the most expensive policies, but, when appropriately matched to the lives of the components, achieved the highest service reliability, measured in miles per in-service failure.

• **The opportunistic heuristics** achieved very good results over a wide range of circumstances, a quality known as "robustness". Both versions were consistently among the most reliable, and the **"extended" version** was among the lowest in costs. These policies proved to be robust under the alternative distributions, and are not subject to the potential problems with short-lived components found in hard time policies. Indeed, the addition of a few short-lived components may be beneficial to these policies, based on the simulations of low versus high reliability cars. That result found that the opportunistic policies achieved higher miles per in-service failure with lower reliability cars (i.e., more random failures) than they did with more reliable cars, since those failures created more opportunities to replace other components preventively.

The **"greedy" and the "extended" versions** both performed well under a wide range of circumstances; the "greedy" version achieves higher reliability, but at a higher cost. The "extended" version is more robust. The "greedy" version is the particularly attractive when the cars are very new, very old, or otherwise subject to low reliability, which is a common situation among U.S. railroads. There may be other circumstances in which the greedy heuristic is preferable, since it does not depend on accurate estimation of the overall failure distribution of the car.

- Car owners who adopt a long-term hard time maintenance policy may also want to consider adopting an opportunistic maintenance policy for components which are relatively short lived. Such a **hybrid** policy was not studied, however, and so its full effects cannot be ascertained from the results of this research.

- A method for conducting **tradeoff analysis** between the benefits of fewer in-service failures and the higher maintenance costs of these policies was developed.

As can be seen from the above results, there is considerable room for improvement by freight car owners over the current policy of on condition maintenance. Each of the alternatives depends on a careful analysis of the failure distributions of the components and of the repair costs. That analysis would be greatly aided by easier access to all the relevant information about the cars and their use, as discussed in the following section.

10.3.3. Information Integration

The current situation of storing the data needed for management decision making in a number of very large files which are each uniquely organized is a considerable barrier to change. Because the current data bases are used for a number of important accounting and other management functions, it appears that the best course would be to establish a separate, management-oriented data file of all the relevant information needed by maintenance planners and managers. Such a file could be drawn from the current information systems in a way that is essentially transparent to the users of the data. There is a precedent for such an approach, which was developed in a study applying expert systems technologies to the diagnosis of "problem cars" [Little and Martland (1989)]. The approach, the **structured car history**, used knowledge engineering techniques to define four basic groups of management information about freight cars:

- *car characteristics*, including the basic "facts" about the car such as type, age, and design features;

- *usage characteristics*, including the mileage, typical commodities, and the type of service the car is in;

- *maintenance activity*, summarizing repair histories in the way that

278

maintenance mangers think about cars and components rather than the way the CRB system organizes billing information;

- *performance measures*, including the three given in section 9.3.1.

This structured car history should not only prove useful in easing the implementation of planned maintenance programs such as the opportunistic heuristic or the hard time policies, but can be used to support new maintenance management systems such as diagnostic systems and repair track scheduling systems. In short, this approach should make it much easier for maintenance professionals to use the information in the company's data bases as a management asset.

10.3.4. Contributions of the Research

In this thesis, three specific contributions have been achieved:

- documenting the state of current practice in the area of railroad car maintenance,

- developing a better approach to planned maintenance by applying the tools of reliability theory to a complex environment, and

- developing a practical set of tools to guide the maintenance manager in the decision making process (including the structured history, the simulation model, and the method for trade off analysis among policy alternatives.)

It is to be hoped that these will prove to be useful to other researchers in the field, and, particularly, to the companies who are attempting to improve their maintenance activities.

10.4. Future Research Needs

This research has highlighted many other areas that could fruitfully be explored, ranging from theoretical issues to the practical details of implementation on railroad properties. In this section a few of the most important of those issues are presented and briefly discussed.

10.4.1. Additional Maintenance Policies

There are many possible variations on the opportunistic maintenance heuristics proposed and tested in Chapters 6, 7 and 8. At least several of these should be carefully

formulated and evaluated. These include:

Hybrid approaches: It was suggested in Chapter 8 that there are circumstances when both the hard time policies and the opportunistic heuristics together might combine some of the best characteristics of each. In particular, if the hard time policy calls for the car to be brought in only at long intervals, it may be desirable to maintain shorter-lived components using an opportunistic heuristic. Similarly, even if the decision is made to follow an opportunistic policy, if the car is very reliable as an overall system, it may be beneficial to bring the car in at fixed intervals simply to apply the opportunistic heuristic and thus reduce the number of in-service failures. Such hybrid approaches should be modelled and evaluated.

"Fuller" heuristics: The "extended" version of the heuristic which was tested only weights the expected future costs associated with leaving a component in service by the probability of "random failures". An important alternative might be to develop the joint probability of a future maintenance opportunity based on all the other components individually. In particular, since the components being modelled include a number of other potential opportunities, to exclude them is to disregard valuable information. Development of such a heuristic should, however, proceed with an eye toward actual application.

10.4.2. Extensions to the Simulation Model

One of the shortcomings of the maintenance policies presented and analyzed using the simulation models is that no information is presented regarding shop inventories and material management. In effect, it was assumed that the repair tracks had sufficient capacity to undertake maintenance actions to support any of the policies. (The opportunistic policies dealt with this by assigning a cost for more frequent labor usage to support preemptive maintenance in the calculation of the costs at the present time.) We know that in practice this is not the case. Some shops are severely constrained in their capacity, and are forced to undertake very limited maintenance activities in order to avoid the formation of queues. Along with the capacity problem is the problem of inventories. The simulation in effect assumes that the parts for a repair are always

available. Enhancing the model to capture both matters would increase its usefulness.

Another assumption that the model makes is that there is only one repair facility (or that all the on-line facilities share the same cost structure). It would certainly be useful to car owners to be able to distinguish between their shops based on costs and efficiency. Improving the model to simulate operations over a network would, however, greatly increase the model's complexity. It is not entirely clear whether a general purpose model could be constructed for this or whether each freight car owner would, in effect, have to create it's own custom version.

10.4.3. Data Processing/Information Systems Issues

The issue of integrating the various information and data systems associated with car maintenance has been addressed in the context of the structured history. It would, however, serve a useful purpose to undertake a more extensive knowledge engineering effort to develop a standard format for a structured history which could be used by all car owners. This would not only make data exchanges easier, but would permit car owners to work cooperatively to develop information systems and expert systems. This sort of cooperation could both reduce the development costs of such systems and permit the sharing of expertise across company lines. It does not appear that the car owners gain any competitive advantage from the unreliability of other companies' cars. A common approach to the structured history developed under the aegis of the AAR might serve to benefit all car owners.

10.4.4. Other Issues

One of the most poorly understood areas of maintenance planning is the effect of maintenance policies on inventories. Yet the management of inventories to reduce the volume and value of goods held on hand has been one of the primary areas of interest in the field of logistics. Because of the size and complexity of the railroad operating environment, an understanding of the impact of maintenance policies on these stocks of goods has the potential to lead to significant financial contributions to the car owners and equally significant contributions to the field of logistics.

As a final note, the use of maintenance as a strategic tool in Chapter 2 was a conjecture. It is in the nature of such decisions that data is rarely available, and then is

subject to many interpretations. It would be a significant contribution in the area of economics known as industrial organization if it could be shown that service industries such as transportation use maintenance (or even service levels) in the same way that manufacturing companies use capacity, i.e., to define markets and control their competitive environment.

10.5. Final Remarks

In this research an important problem to practitioners has been shown to be amenable to an application of reliability theory. The ability to use the theory has only been possible by looking at the problem from the perspective of seeking practical solutions first, and advances to the theory as a secondary matter. It seems quite reasonable that more gifted theoreticians (and there are many) could make much greater contributions by simply listening to the problems of transportation managers.

REFERENCES

"44 Utilities and Their Unit Train Costs" (1988), **Progressive Railroading**, June 1988, pp. 47-48.

AAR Freight Car Utilization Program (1980), **The Impact of Freight Car Design and Serviceability on Freight Car Utilization**, AAR Report No. R-438, November, 1980.

Armitage, W. (1970), "Maintenance Effectiveness", in A.K.S. Jardine (ed.), **Operational Research in Maintenance**, (Manchester University Press (Barnes and Noble): New York), 1970.

Association of American Railroads (a), **Railroad Facts**, (AAR: Washington, D.C.), published annually.

Association of American Railroads (b), **Analysis of Class I Railroads**, (AAR: Washington, D.C.), published annually.

Association of American Railroads (c), **Field Manual of the A.A.R. Interchange Rules**, (AAR: Washington, D.C.), published annually.

Association of American Railroads (d), **Office Manual of the A.A.R. Interchange Rules**, (AAR: Washington, D.C.), published annually.

Association of American Railroads (e), **Manual of Standards and Practices**, (AAR: Washington, D.C.), published annually.

Barlow, R.E., "Maintenance and Replacement Policies", in Zelen, Marvin (ed.), **Statistical Theory of Reliability**, (University of Wisconsin Press: Madison), 1963.

Barlow, R.E. and F. Proschan (1965), **The Mathematical Theory of Reliability**, (Wiley: New York), 1965.

Barlow, R.E., "Analysis of Retrospective Failure Data Using Computer Graphics", **Proceedings 1978 Annual Reliability and Maintainability Symposium**, (IEEE: New York), 1978.

Bauer, Edward J. (1989), "BN Explores Car, Locomotive Needs", **Modern Railroads**, November, 1989, pp. 39-40.

Belovaric K., and J.T. Kneafsey (1972), **Determinants of Line Haul Reliability**, SROE, Vol 3., Massachusetts Institute of Technology, 1972.

Blaine, D.G., (1979), **Modern Freight Car Air Brakes**, (Simmons-Boardman, Omaha), 1979.

Blundell, James K and K.W. Beard (1985), " Maintenance Strategies for Aero Engines", **Proceedings 1985 Annual Reliability and Maintainability Symposium**, (IEEE: New York), 1985.

Brownlee, James H. (1985), "Preparing for the Future", **Modern Railroads**, June, 1985, pp. 37-40.

Carter, A.D.S. (1986), **Mechanical Reliability**, (Macmillan: London), 1986.

Corder, A.S. (1978), **Maintenenance Management Techniques**, (McGraw Hill: London), 1978.

Day, J. and L. George (1982), "Opportunistic Replacement of Fusion Power System Parts", **Proceedings 1982 Annual Reliability and Maintainability Symposium**, (IEEE: New York), 1982.

Dingle, A.D. (1977), **Case Studies on Car Rejection, Inspection, and Conditioning**, AAR Report No. 307, October 1, 1977.

Dixit, A. (1980), "The Role of Investment in Entry Deterrence", **The Economic Journal**, March, 1980, Vol. 90, pp. 95-106.

Dixit, A. (1982), "Recent Developments in Oligopoly Theory", **American Economic Association Papers and Proceedings**, May, 1982, Vol. 72, No.2, pp. 12-17.

Drenick, R.F. (1960), "Mathematical Aspects of the Reliability Problem", **J. Soc. Industrial Appl. Math.**, Vol. 8, No. 1, March 1960, pp. 125-149.

Ferney, James R. (1978), **The Interface of Federal Railroad Administration Regulations with the U.S. Railroad Industry**, unpublished S.M. Thesis, Sloan School of Management, Massachusetts Institute of Technology, 1978.

"Freight Car Statistics" (1989), **Progressive Railroading**, April, 1989, pp. 33-50.

Flehinger, B.J. (1962), "A General Model for the Reliability Analysis of Systems Under Various Preventive Maintenance Policies", **Ann. Math. Statistics**, Vol. 33 (1962), No. 1, pp. 137-156.

George, L.L., H. Mahlooji, and P.W. Hu (1979), " Opportunistic Replacement and Build Policies", **Proceedings 1979 Annual Reliability and Maintainability Symposium**, (IEEE: New York), 1979.

Ghemewat, Pankaj (1984), "Capacity Expansion in the Titanium Dioxide Industry", **Journal of Industrial Economics**, December, 1984, Vol. 23, No. 2, pp. 145-163.

Gregory, E. (1973), "Reliability in Engineering. Aircraft Maintenance", **Conference on Improvement of Reliability in Engineering,** (Institution of Mechanical Engineers: London), 1973.

Guins, Thomas S., and Michael B. Hargrove (1980), **Status Report: Car Maintenance Cost Data Base**, AAR Report R-429, April 1980.

Guins, T.S. (1983), "Analyzing Freight Car Repairs by Computer, Part III", Presentation to the Car Department Officers Association Annual Meeting, September 19, 1983.

Guins, T.S. (1984), "The Braking System and Freight Car Maintenance", Presentation to to Car Department Officers Association 1984 Annual Meeting, September, 1984.

Guins, T.S. (1987), "A Statistical Analysis of Roller Bearings", Paper presented to ASME Conference, Boston, 1987.

Guins, T.S., M.B. Hargrove, J. Kyparisis, N. Singpurwalla, R. Soyer (1984), **Reliability Analysis of Freight Car Data with Implications for Scheduled Maintenance,** Unpublished AAR Report, 1984.

Guins, T.S. and J. Kyparesis (undated), "Equipment Reliability Analysis Systems", unpublished A.A.R. working paper, undated.

Harmon, P. and D. King (1985), **Expert Systems**, (John Wiley and Sons: New York), 1985.

Haven, Paul J. (1980), **Transit Vehicle Maintenance: A Framework for the Development of More Productive Programs**, Unpublished S.M. Thesis, Massachusetts Institute of Technology, 1980.

Jorgenson, D.W., J.J. McCall, and R. Radner (1967), **Optimal Replacement Policy**, (North-Holland: Amsterdam), 1967.

L'Ecuyer, P. and A. Haurie (1983), "Preventive Replacement for Multicomponent Systems: An Opportunistic Discrete Time Dynamic Programming Model", **IEEE Trans. Reliability**, R-32 (April, 1983), No. 1, pp. 117-118.

Lewis, E.E. (1987), **Introduction to Reliability Engineering**, (Wiley: New York), 1987.

Liang, T.Y. (1985), "Optimum Piggyback Preventive Maintenance Policies", **IEEE Trans. Reliability**, R-34 (December, 1985), No. 5, pp. 529-538.

Little, Patrick and Carl D. Martland (1989), **Improving Railroad Car Maintenance By Using Knowledge Based Systems**, M.I.T. Studies in Railroad Operations and Economics, Volume 39, December 15, 1989.

Lott, Richard W. (1971), **Basic Systems Analysis**, (Canfield Press: San Franscisco), 1971.

Manheim, Marvin, L. (1979), **Fundamentals of Transportation Systems Analysis. Vol. 1: Basic Concepts**, (MIT Press: Cambridge, MA), 1979.

Mann, Nancy R., R.E. Schafer, N.D. Singpurwalla (1974), **Methods for Statistical Analysis of Reliability and Life Data**, (John Wiley and Sons: New York), 1974.

Martland, Carl D. (series editor) (SROE), **Studies in Railroad Operations and Economics**, series of monographs and papers published jointly by the Association of American Railroads and the M.I.T. Rail Group.

Mao, Chi-Kuo (1982), **Integrating Organizational and Technological Perspectives - An Approach to Improve Rail Motive Power Management**, unpublished Ph.D. Dissertation, Massachusetts Institute of Technology, 1982.

Mao, C.K. (1988), "Transportation Operations Management: Integrating Organizational and Technological Perspectives", M.I.T. Rail Group Working Paper 88-3, January 1988.

McCall, John J. (1963), "Operating Characteristics of Opportunistic Replacement and Inspection Policies", **Management Science**, Vol. 10 (1963), No. 1., pp. 85-97.

McCall, John J. (1965), "Maintenance Policies for Stochastically Failing Equipment: A Survey", **Management Science**, Vol 11, No. 5, March 1965, pp. 493-524.

Messner, M. G. (1980), **A System for Controlling a Railroad's On-Line Freight Car Cycle**, unpublished S.M. thesis, Massachusetts Institute of Technology, Center for Transportation Studies, 1980.

Mishkoff, H.C. (1986), **Understanding Artificial Intelligence**, (Radio Shack: Fort Worth, TX), 1986.

Mosteller, F. and R.E.K.Rourke (1973), **Sturdy Statistics**, (Addison-Wesley: Reading, MA), 1973.

Nakagawa, T. (1981), "Generalized Models for Determining the Optimal Number of Minimal Repairs Before Replacement", **J. Oper. Res. Soc. Japan**, 24 (1981), pp. 213-227.

Nakagawa, T. (1987), "Optimum Replacement Policies for Systems With Two Types of Units", in Osaki and Cao (1987).

Nelson, Wayne (1982), **Applied Life Data Analysis**, (John Wiley and Sons: New York), 1982.

Niño, L.E. (1974), "Predict Failures and Minimize Downtime", **Hydrocarbon Processing**, January 1974, pp 108-110.

Nowlan, Frank S. (1964), "Relationship Between Reliability and the Periodicity of Scheduled Maintenance", **Reliability in American Industry**, (Western Periodicals Co.: No. Hollywood, CA), 1964.

O'Conner, Patrick D.T. (1985), **Practical Reliability Engineering**, (Wiley: Chichester), 1985.

Ohashi, M. and S. Miyamoto (1987), "Monotone Optimal Maintenance Policy for Markovian Deteriorating Systems and Minimal Repair", in Osaki and Cao (1987).

Osaki, S. and J. Cao (1987), **Reliability Theory and Applications**, (World Scientific: Singapore), 1987.

Özekici, S. (1988), "Optimal Replacement of Multicomponent Reliability Systems", **Operations Research**, Vol. 36 (1988), No. 4, pp. 542-552.

Philip, C.E. (1980), **Improving Freight Car Distribution Organization Support Systems: A Planned Change Approach**, unpublished Ph.D. Dissertation, Massachusetts Institute of Technology, 1980.

Radner, Roy and D.W. Jorgenson (1963), "Opportunistic Replacement of a Single Part in the Presence of Several Parts", **Management Science**, Vol. 10 (1963), No. 1, pp.70-84.

Ramaswamy, R. (1983), **Design of an Optimal Preventive Maintenance Program for Transit Vehicles**, unpublished S.M. Thesis, Massachusetts Institute of Technology, 1983.

REMSA (1966): need cite for simulation studies book from D. Snyder.

Rennicke, William J. and Anthony A. Davis (1989), "New Cars: Boom or Blip", **Railway Age**, April, 1989, pp. 35-37.

Riggs, J.L. and T.M. West (1986), **Engineering Economics**, (McGraw Hill: New York), 1986.

Salop, Steven C. and David T. Scheffman (1983), "Raising Rivals Costs", **American Economic Association Papers and Proceedings**, May, 1983, Vol. 73, No. 2, pp. 267-271.

Sasieni, M.W. (1956), "A Markov Chain Process in Industrial Replacement", **Operational Research Quarterly**, Vol 18. (1956), pp. 148-154.

Sethi, Davinder P. (1977), "Opportunistic Replacement Policies", in Tsokos, C.P., and I.N. Shimi (eds.), **The Theory and Applications of Reliability**, Vol. 1, (Academic Press: New York), 1977., pp.433-447.

Sheaffer, Warren W. and G. L. Stern (1986), "The Influence of Locomotive Reliability Upon Locomotive Acquisition and Maintenance Policies", **Proceedings of the Twenty-seventh Annual Meeting of the Transportation Research Forum**, Seattle, Wa, 1986.

Shaw, Robert B. (1961), **Down Brakes, A History of Railroad Accidents, Safety Procedures, and Operating Practices in the United States of America**, (Macmillan: London), 1961.

Shedd, Tom (1990), "If It Ain't Broke...Fix It Anyway", **Modern Railroads**, September, 1990, pp. 33-44.

"Smaller Fleet, Greater Utility" (1989), **Modern Railroads**, October, 1989, pp. 27-31.

Stigler, G. (1971), "The Theory of Economic Regulation", **Bell Journal of Economics**, Spring 1971, 2, pp. 3-21.

Symonds, A.J. (1986), "An Introduction to IBM's Knowledge System Products", **IBM Systems Journal**, 25:2, (pp. 134-146), 1986.

"The AAR and the Freight Car" (1989), **The Progressive Railroading Freight Car Yearbook, 1988-1989**, (Murray-Richter: Chicago, IL), 1989, pp. 25-34.

THomas, L.C. (1986), "A Survey of Maintenance and Replacement Models for Maintainability and Reliability of Multi-Item Systems", **Reliability Engineering** 16 (1986), (pp. 297-309).

"TT Beefs Up Rack-Car Fleet" (1989), **Railway Age**, October, 1989, pp. 41-43.

Ulen, Thomas. S. (1980), "The Market for Regulation: The ICC for 1987 to 1920", **American Economic Association Papers and Proceedings**, May, 1980, Vol. 70, pp. 306-310.

United States Congress, Office of Technology Assessment (1978), **An Evaluation of Railroad Safety,** Report No. OTA-T-61, (U.S. Government Printing Office: Washington, D.C.), May, 1978.

Vergin, R.C. (1968), "Optimal Renewal Policies for Complex Systems", **Naval Research Logistics Quarterly,** Vol 18., 1968, pp. 523-534.

Waterman, D.A. (1986), **A Guide to Expert Systems,** (Addison Wesley: Reading, MA), 1986.

Watson, C. (1970), "Is Preventive Maintenance Worthwhile?", in A.K.S. Jardine (ed.), **Operational Research in Maintenance,** (Manchester University Press (Barnes and Noble): New York), 1970.

Widell, G.W. (1986), **Freight Car System Maintenance,** unpublished AAR Working Paper, June 24, 1986.

Williamson, Oliver (1968), "Wage Rates as a Barrier to Entry: The Pennington Case", **Quarterly Journal of Economics,** February, 1968, Vol. 85, pp. 85-116.

Winston, P.H. (1988), Keynote Address, **Proceedings for the Artificial Intelligence Applications for Military Logistics Symposium and Workshop,** (American Defense Preparedness Association: Arlington, VA), 1988.

Appendix A

A Look at Freight Car Maintenance Costs and An Estimate of the Average Maintenance Cost Per Mile

In this appendix, the freight car maintenance expenditures of the Class I railroads are examined and used as a basis for estimating the average maintenance cost per mile for all freight cars operated in the U.S. This figure must be estimated (as opposed to simply calculated) because only the Class I railroads are required to report their expenditure and mileage data, and these carriers won only about 60% of the overall fleet of cars operated in this country. The other cars, owned by smaller railroads and private car owners, are not required to report on either their level of expenditure or their level of usage.

To accomplish the estimate, some of the facts and figures of the Class I railroads are reviewed, and these are used as a basis for estimating the expenditures of the other car owners.

It is shown in Chapter 3 that the age and composition of the car fleet has changed in recent years. The work force of the Class I railroads assigned to the maintenance function has also changed dramatically. In 1980, the number of employees in the category "Maintenance of Equipment and Stores" was 99,487. By 1988, that had fallen to less than half, 45,209. The average annual compensation for those workers rose from $22,160 to $33,591 in the same years[1].

The various changes have been translated into changes in the amount of maintenance dollars spent, and on the way those dollars are spent. In 1980, direct labor and fringe benefits expenses by the Class I railroads for the maintenance of freight cars was $856 million ($693 million in direct labor and $163 million in fringe benefits). By 1988, that fell to $649 million ($453 million in direct labor and $196 million in fringe benefits). Material and supplies also fell, from $640 million to $518 million. Some of this reduction was undoubtedly caused by the shift of work on privately owned cars to contract shops, but the majority was caused by the decision to retire older, high maintenance cars, particularly during the recession of 1982-1983 when such cars were simply not needed. Taking direct labor, fringe benefits, and materials as an approximation of the out-of-pocket costs of car maintenance by the Class I railroads, one can form some interesting statistics. In 1980, the out-of-pocket cost by the Class I carriers per car in the total U.S. fleet was $874. In 1988 this had risen to $941, which is less than the rate of inflation over the period[2]. That is, in spite of an increased aging

---

[1] Unless otherwise indicated, the data used in this section is from **Analysis of Class I Railroads** [AAR (b)], for the respective year.

[2] Need a formal measure of this, such as the wholesale price index.

and higher utilization, the railroads' expenditure on maintenance per car declined in real terms.

If, on the other hand, one assumes that the Class I railroads are primarily engaged in performing maintenance on their own cars, leaving the private car owners to contract shops, a different result emerges. The out-of-pocket cost per Class I owned car rose from $1280 in 1980 to $1795 in 1988. It seems clear that the truth about the level of maintenance expenditures lies in some middle ground. There has been a shift away from railroad performed maintenance by private car owners, and there have been attempts to reduce costs of maintenance by the railroads to gain efficiencies in the workplace and to deal with competitive pressures from other modes.

In order to estimate the costs, we first assess the total costs by the Class I railroads. This includes both the direct expenditures as reported to the ICC and an estimate of the share of general and administrative expenses which can be allocated to the freight car maintenance function. These expenses by the Class I railroads are then used to extrapolate the overall expenses by all the freight car owners.

The first thing that must be done, then, is to estimate the total costs of car repair and maintenance for all cars which operate on the Class I railroads. Ideally, this should include the costs reported by the Class I's, any costs at contract shops incurred by private cars, and on line repairs be railroads smaller than Class I. Since little or no data is available on expenditures by Non-Class I car owners an facilities other than those owned by Class I railroads, it is necessary to make some assumptions.

The first assumption that must be made is regarding which expenses to include from among the reported costs. The "known" costs those reported in the Analysis [AAR(b)][3]. It seems clear, for example, that the equipment lease/rental expenses and depreciation are not properly part of the car repair and maintenance expenses. If these are excluded, the 1988 remaining portion of the freight car account is $1.196 billion. To this, however, must be added the car maintenance program's "share" of the general administrative and executive expenses (Line 249). In 1988, maintenance of equipment (including locomotives) constituted 21.9% of the total operating expenses (Line 298). Of Freight cars accounted for approximately 55% (based on Lines 182, 190, and 198), and locomotives and other equipment for 45%. It is conservative, then to assign 10% (i.e., approximately one half of the 21.9% of the operating expenses to freight cars. Using this 10% figure, we can assess a share of the General and Administrative expenses (Line 249) to car maintenance. In 1988, this category was $3.466 billion, so that the car maintenance share is approximately $350 million.

Thus the total Class I car repair expense is approximately $1,550,000,000. To this must be added the expenses which private car companies and non-Class I railroads spend at contract shops and their own facilities. Since no data is available on this, one must again make an estimate.

---

[3] These include primarily lines 183-190 "Freight Service Expense - Freight Cars", and an allocation of Line 249, "General and Administrative".

The private and non-Class I railroads own approximately 40% of the car fleet. We will assume that they perform one-third of their cars' maintenance at Class I facilities and two-thirds at contract and other shops. This seems to be a conservative figure, in light of data presented regarding Company C in Chapter 4, and looking at reports of car repair practices of coal car owners reported in the trade journal **Progressive Railroading** ["44 Utilities... "]. In many cases, the Class I railroads may be performing as little as 10% of the maintenance on private cars. We assume also that the cost to maintain a private car is the same as for a railroad-owned car. There are reasons to reject this assumption, but they seem to weigh just as heavily on believing the private car is more expensive or less expensive than a railroad-owned car. Some private cars are specially equipped with parts which are expensive to maintain (such as tanks), while others are quite simple cars (e.g., coal hoppers). Most contract shops offer lower labor rates than the AAR rates, but most contract shop users seem to base their decisions about maintenance more on reliability than on costs alone, so may tend to maintain their cars to a higher level of reliability than the Class I railroads. The point is, for the entire fleet of private cars, there is no simple rule for assuming a higher or lower cost structure relative to that of the Class I railroads. Using these assumptions, we would then assume that the Class I railroads' expenses are made up of all the costs of maintaining the 60% of the fleet they own, plus 33% of the costs of maintaining the other 40% of the fleet. That is, the $1.5 billion the Class I railroads spend is about 75% of the total amount spent on car repair in the U.S. Thus, the approximate total amount spent in 1988 was $2 billion.

The other part of estimating average costs per mile is to determine the total freight car miles operated in the U.S. The total freight car miles in 1988 on Class I railroads was 26.3 billion car miles. The Class I railroads operate approximately 90% of all the railroad trackage in the U.S. This is somewhat misleading, however, since these lines are used to a much higher degree of intensity than the other 10%, and account for more than 90% of the total car miles in the U.S. If we disregard the mileage accrued on non-Class I owned track, the total cost per mile for all cars in the U.S. is

$$\frac{\$2,000,000,000}{26,300,000,000\ miles} = \$0.076.$$

That is, the average maintenance cost per mile for all car owners in the U.S. is approximately $0.07-$0.08 per mile.

These numbers are higher than those reported by unit coal train operators in a recent survey ["44 Utilities... "], but are lower than those for found for Company C in Chapter 4; they are in the same range as those for the two case study railroads.

Appendix B

CARSIM: A Simulation Model for Evaluating Freight Car Maintenance Policies

B.1 Introduction

CARSIM is an event-based simulation for estimating the effects of alternative maintenance policies on the reliability and cost of railroad freight cars. The simulation is programmed in Turbo Pascal 5.5, an object-oriented implementation of that language, and runs on IBM-compatible personal computers. The simulation tracks the state of a set of components and the overall state of a freight car over a period specified by the user (2 million car miles, in the cases studied). The "states" of the components take on the value of failed or unfailed at any given time, depending on the failure distribution of the component and the particular time at which the state is being assessed. Given the state of the components and a set of cost information, the model applies a maintenance policy to the car and determines what, if anything, is to be done to each component. Events in the simulation correspond to the car being sent to a repair facility, either due to the failure of a component or as a scheduled event. The model is capable of analyzing a number of maintenance policies, including:

- *On condition maintenance*, in which a components are replaced when their condition no longer meets a predetermined standard;

- *Far-sighted hard time maintenance*, in which a car is brought into the shop at fixed intervals of time (or miles), and all components which are "expected" to fail between the present time and the next scheduled interval are replaced. The maintenance intervals are selected by the user.

- *Near sighted hard time maintenance*, in which the car is brought into the shop at fixed intervals, and components that are "expected" to fail in some percentage of the next interval are replaced. The percent of the interval can be set in the range from 0% to 100%.

- *"Naive scheduling"*, which calls for each of the car's components to be scheduled using a single component age replacement approach [Barlow and Proschan (1965)], and the car brought into the shop whenever a component is "due" based on its schedule. If a component fails before its scheduled date, the failed part is replaced immediately and the scheduled interval for that component is reset to begin from the time of failure.

- A *"greedy" opportunistic policy*, in which each of the components are evaluated as candidates for "preemptive" maintenance whenever the car is brought into a shop under the control of the car owner. The car is brought into the shop whenever a component fails or a component reaches the scheduled age as in the "naive scheduling" approach. The rule for deciding which components to replace, given that the car is in the shop, is to replace any component which has a lower total replacement cost now than the

293

expected replacement cost at the single part scheduled age. The calculation of the costs used in the decision rule are explained in Chapter 6, and the implementation is given below.

- A *"full" opportunistic policy*, which, like the greedy policy, treats any trip to an on-line repair facility as an opportunity to decide whether or not to replace other, unfailed components; in the "full" case, however, the expected future costs are weighted by the probability that another opportunity will arise in the interval between the present time and the scheduled time for the component in question.

The output of the model is a group of files reporting on the number and type of failures for the car and its components, which can then be analyzed statistically. The inputs, outputs and processes of the simulation model are presented in the following sections. A basic knowledge of the language PASCAL is presumed, although the readers who are familiar with any structured programming language should be able to grasp most, if not all, of the model's structure.

B.2. The Decision Structure and Processes of the Model

The central notion in CARSIM is to follow and record all the relevant events that "happen" to a railroad freight car while it is operated for an extended period of time. To accomplish this, the car is structured around the programming concept known as an "object"[1].

B.2.1. The "Objects" Modelled in the Simulation

By structuring the car and its attributes as an object, it is possible to organize data and information about the car in a way that facilitates thinking about the operations and maintenance of the car. In particular, objects are data structures which are capable of having procedures and functions specifically attached to them. While this may seem unimportant to many readers, this property makes it possible for the program to direct an object to "update itself", for example. Thus by organizing both cars and their components parts as objects, the programming task was greatly simplified.

The two most important objects in CARSIM are the part and the car. The part serves as the basic or parent object for describing a component of a railroad car. It is characterized by a number of data fields and methods. Data fields are the aspects of the object that take on particular values like a field in a conventional record. Methods are the procedures and functions associated with the object which update the values in the fields. This tight coupling of fields and the operations which set the values is known as *encapsulation*. A part is basically a listing of the failure distributions and failure costs

---

[1] Object oriented programming has been widely discussed in a number of journals and books in recent months. An excellent introduction for programmers is the **Turbo Pascal Object Oriented Programming Guide**, (Borland: Scotts Valley, CA), 1989, which accompanies version 5.5 of that language.

for a particular component. All the parts are modelled using Weibull failure distributions (see Chapter 2). The costs for the failure of a part include the cost of an off line in service failure, an in-service on line failure, and a scheduled replacement at the most efficient facility on line. These costs are stored in an array.

The actual structure of the object "Part" is:

```
Part = Object
    Name  : String[10];
    Shape: Real;
    Scale,
    Optimal : Longint;
    NumFails : Array[0..FailClass] of Integer;
    FailCosts : Array[0..FailClass-1] of Integer;
    MaterialShare,
    LaborShare: Integer;
    LastFailTime,
    NextFailTime : Longint;
    NextFailType : Shortint;
    Inventory : Array[1..InvSize,1..2] of LongInt;

    Function GetName: String;
    Function GetShape: Real;
    Function GetScale: LongInt;
    Function GetFailCosts(K: Integer) : Integer;
    Function GetMaterialShare: Integer;
    Function GetLaborShare: Integer;
    Function GetOptimal: Longint;
    Function GetNumFails(K: Integer) : Integer;
    Function GetLastFailTime: LongInt;
    Function GetNextFailTime: LongInt;
    Function GetNextFailType: ShortInt;
    Procedure Create(OffPct: Real);
    Procedure SetName;
    Procedure SetShape;
    Procedure SetScale;
    Procedure MakeInventory(Threshold : Real);
    Procedure SetCosts;
    Procedure SetShares;
    Procedure SetOptimal;
    Procedure InitNumFails;
    Procedure SetFailTimes;
    Procedure AddaFail(K: Integer);
```

As can be seen, the data fields of a part or component are composed primarily of cost and failure distribution information, and the time and number of failures. An exception to this is the array known as an inventory. An inventory is a vector of mileages which represent the maximum possible lifetimes of a group of components, and whether those components will fail on or off line if they survive for their full potential lifetime. It is from this inventory that replacement parts are drawn when a component is removed either as a result of a failure or a preventive maintenance action. The methods consist primarily of means to set or change the values. Many of this are simply

295

input and output routines which allows the user to set initial values (such as SetScale, which is used to set the scale parameter of the failure distribution for a part).

The manner in which the particular values are set is discussed below.

The object "Car" is composed of all the parts, i.e., components, and a set of data fields (and corresponding methods) that are used either by the parts to determine values or by the simulation as a whole to determine costs and events. The Car object includes:

- *Identifier fields* (and corresponding methods), such as Init (name), Number, and TypeCar, which establish which car is being modelled.
- *General Usage fields* (and methods), such as ArnualMiles and OffLinePct which provide information about the use of the car.
- *Part fields*, which include the objects used to model components, and the necessary supporting fields, such as NumParts.
- *Reporting and Accounting fields and methods*, which keep track of when and where events occur (or will occur), the interval for any hard time maintenance events, and printing out the results of the simulation to a file.

The actual structure of the object "Car" is:

```
Car = Object
    Init : String[4];
    Num  : String[6];
    TypeCar : String[4];
    OffLinePct : Real;
    AnnualMiles : LongInt;
    NumParts : Integer;
    Parts: Array[0..MaxParts] of Part;   { Part 0 is Random Failures }
    SwitchCost : Integer;
    Switched : Boolean;
    RIPTrackEvents: Array[0..2] of Integer;
    LastEvent,
    NextEvent,
    HardInterval,
    HardDueDate : LongInt;
    NextSite : ShortInt;
    Procedure Initialize;
    Procedure SetInit;
    Function GetInit: String;
    Procedure SetNumber;
    Function GetNumber : String;
    Procedure SetAnnualMiles;
    Function GetAnnualMiles: LongInt;
    Procedure SetHardInterval;
    Function GetHardInterval: LongInt;
    Procedure InitHardDueDate;
    Procedure SetHardDueDate;
    Function GetHardDueDate: LongInt;
    Procedure HardDueDateIsIrrelevant;
    Procedure SetType;
    Procedure SetNumParts;
    Procedure SetPctOffLine;
    Function GetPctOffLine: Real;
```

Function GetNumParts: Integer;
Procedure SetSwitchCost;
Procedure SetSwitched(State : Boolean);
Function GetSwitched : Boolean;
Function GetSwitchCost: Integer;
Procedure SetRIPTrackEvents;
Procedure AddRIPTrackEvent(OnOrOff : Integer);
Function GetRIPTrackEvents(OnOrOff : Integer): Integer;
Procedure SetNextEvent;
Function GetNextEvent: LongInt;
Function GetNextSite: ShortInt;
Procedure InitParts;
Procedure PrintResults;

The reader may notice that many of the methods are designed to set or get the values within the data fields of the objects. It would be faster to perform these functions directly, i.e., to simply assign values to fields, but this not only violates the "spirit" of object oriented programming, it actually has the effect of making the code much more difficult to maintain and change over time. If, for example, it is desired to use a new method to enter the values of fields associated with Part.Inventories, for example, all that needs to be done under the current structure is to change the method Part.MakeInventory. If, on the other hand, the inventories were directly set in some other procedure in the program, then it would be necessary to "track down" the code in question, modify it, and ascertain that no other fields or values were affected by the changes.

B.2.2. The Structure of the Simulation Process

In this section, the overall structure of the simulation and the corresponding program steps are presented. In general an overview of the various processes are presented, by in some cases the actual details of procedures are given.

The model begins with an initialization procedure (InitializeRun) which resets the clock (CurrentTime) to 0, allows the user to attach a name to the series of runs, and seeds the random number generator. (The same random number seed is for each policy in a trial, which limits the extent of model induced variation within the same set of trials.) InitializeRun also selects the discount rate to be used for the series of runs.

Following this, the car itself is initialized, using the method Car.Initialize. This procedure sets the initial, number, type, percent of miles off line, annual mileage, switching costs, hard time maintenance interval to be followed, and the number of parts the car is composed of (i.e., the components to be modelled). This procedure also resets the various counters associated with the car to zero. The parts themselves are then initialized, using the method Car.InitParts.

Car.InitParts is basically a loop over the number of parts in the car (including the "rest of the car" part, which is considered part number 0. This calls the parts method, Part.Create.

Part.Create sets the name, shape parameter, scale parameter, and costs associated with each part. If the part is subject to an IFR distribution then it also asks the user to input the optimal replacement interval. (A future enhancement of the model would be to calculate this.) After setting all the part related counters to zero, this procedure creates

297

an inventory of part lives. The inventory is created by using the random number generator and an inverse of the Weibull cumulative distribution function to solve for t, the failure time. The equation used is:

$$t = \alpha \, [-\ln(1 - P)]^{\frac{1}{\beta}} \qquad \text{(B.1)}$$

where t is the time, P is the random number, α and β are the scale and shape parameters respectively[2]. Based on the percent of off line and on line miles that the car is used for, each failure time is assigned to occur either on or off line. Part.Create also sets the next failure time to the appropriate time.

The program then begins a series of loops, running through each of the maintenance policies for each car and inventory. Each policy has assigned to it a number, from 1 to the number of policies, and the loop simply passes through each process from CurrentTime of 0 to RunLength (2 million miles).

The process for each policy is described in the following sections.

B.2.2.1. On Condition Maintenance

This is the easiest of the policies to model. The process simply consists of always finding the lowest failure time among the variable Part.NextFailTime associated with each of the parts. The CurrentTime is then incremented to this time, the counters are incremented accordingly, and the costs are increased. This part's LastFailTime is set to CurrentTime, and the NextFailTime for this part is set to CurrentTime plus the lifetime of the next part in the inventory for that part. The simulation then looks for the next failure event among the parts.

B.2.2.2. Naive Scheduling

This process uses the optimal replacement interval for each part as an upper bound on the life of the component. If the part has a lifetime greater than the optimal replacement interval, then the part's lifetime is restricted to the replacement time. Naturally, in that case the replacement always occurs on line. Thus under this policy, the simulation still looks for the next failure among the components. Upon finding one, however, the simulation looks to see if other components are also due for replacement at this time. If so, then those parts are replaced without incurring additional switching costs or trips to the repair track.

B.2.2.3. Hard Time Maintenance

In this case, the Car.HardTimeInterval data field is exploited. When the CurrentTime is equal to the value Car.HardDueDate, then the car is sent to an on line repair facility. Depending on the strategy being followed, all the components which are

---

[2] Nelson (1982) gives the inverse for all the commonly used distributions in reliability under the name 100Pth percentile.

expected to fail in the next hard time interval (or some user specified fraction) are replaced. The determination of which components are expected to fail is based on the component's scale parameter. If the scale parameter plus the last fail time of the component is less than the sum of CurrentTime plus the HardTimeInterval (or some fraction thereof), then the part is replaced, and the appropriate counters are incremented. If not, the part is allowed to remain in service.

In the interval between hard time events, the car is operated under an on condition policy.

B.2.2.4. Greedy Opportunistic Maintenance

Under the greedy policy, each time the car undergoes a trip to a repair track, i.e., each time CurrentTime is incremented to the lowest NextFailTime, then if the failure is occurring off line, the replacement proceeds as a simple in service failure as under the on condition maintenance. If, on the other hand, the visit to the repair track occurs on line (either due to an on line in service failure or because a component is due for maintenance as under the naive scheduling), then all the other components are tested for early replacement. This test consists of calculating the cost of replacing the part now versus the expected costs of leaving the part in service. The cost of replacing the part now is given by the sum of the labor and materials associated with the replacement, plus the share of the remaining life (i.e., the interval between the present time and the replacement time divided by the total optimal replacement interval), plus the associated labor with more frequent replacements. These actual equations are given in chapter 6, and are not repeated here.

The expected cost of failure is the sum of the cost of replacing the part at the scheduled time weighted by the survival probability to that time (discounted from that time) plus the probability of failure in the interval times the cost of an in-service failure (which is an average of the off and on line costs weighted by the off and on line mileage percentages), discounted from the middle of the interval.

If the cost now is less than the expected cost later, the part is replaced at this time; if not, the part is left in service. If the part is replaced, the relevant counters are incremented. Note that no switching cost is assigned to an opportunistic replacement, since that cost has already been assigned to the failed (or scheduled part).

B.2.2.5. "Extended" Opportunistic Maintenance

This policy proceeds exactly as the greedy policy, except that in this case the expected cost is weighted by the probability that another opportunity will present itself in the interval. The opportunity is assumed to occur in the middle of the interval, which simplifies the calculation of the costs of a random opportunity. The cost of materials and labor remains the same as the cost now, the foregone part life becomes half its current value, as does the excess labor cost required to perform more frequent maintenance activity. The calculation of the value of the probabilities proceeds as in Chapter 6.

B.2.2.6. Adding New Policies

What is required to add a new policy is to create a "process procedure", similar

to the ones which exist for each member of the current set of policies. This is relatively straightforward if the policy is at all similar to any of the existing policies. When, for example, the hard time policies were added, the fact that the poicy between HardDueDates was on condition meant that all the processes for those times could be re-used. The object orientation of the program also makes coding new policies simple for the experienced programmer because all the adjustments to the car and its parts are performed by using the methods associated with that object. (By way of example, coding the hard time policy took less than 1 person week of programming and debugging becasue it could build on the existing structures.)

B2.3. The Outputs of the Model

The outputs of the model are simply a set of files. Each trial results in a small file with reporting the number of replacements by component and by type (i.e., off line, on line, scheduled, and opportunistic), and the number of random failures that occurred on and off line. The processing of these files is currently performed manually, although it would be a straightforward matter to enhance to program to perform tabulations and statistical tests among runs. The data can also be fed into other programs, such LOTUS 1-2-3 spreadsheets using those programs' import facilities, and the results processed using macros.