

IRTMM VALUE ANALYSIS MODULE

By
Shouh-Dauh Lin,
John Kunz, and
Paul Teicholz

**CIFE Technical Report #99
August, 1994**

STANFORD UNIVERSITY

Copyright © 1994 by
Center for Integrated Facility Engineering

If you would like to contact the authors please write to:

*c/o CIFE, Civil Engineering,
Stanford University,
Terman Engineering Center
Mail Code: 4020
Stanford, CA 94305-4020*

SUMMARY
CIFE TECHNICAL REPORT
Number 99

Title: IRTMM Value Analysis Module

Authors: Shouh-Dauh Lin, John Kunz, Paul Teicholz

Publication Date: August 1994

Funding Sources: Electric Power Research Institute, Pacific Gas and Electric, Southern California Edison, Shimizu Corporation of Japan

1. Abstract: We describe the Value Analysis Module (VA), which is part of the Intelligent Real Time Maintenance Management (IRTMM) project. The VA uses a decision-analytic approach to study the maintenance decision-making in a conventional steam power plant. We used two test cases from Southern California Edison (SCE) to validate the VA. In addition, a simulator was built to investigate the problem of scheduling multiple maintenance activities. The second half of this report studies three test cases from Pacific Gas and Electric (PG&E). The results show that the VA can handle individual maintenance activity very well. However, there are some problems remaining. A discussion of them and the implications for future research is presented at the end of this report.

2. Subject: This report focuses on building a decision support system, based on decision analysis, to study choice and timing of maintenance decisions.

3. Objectives/Benefits: The objectives of this research are to help the maintenance managers of a conventional steam power plant make informed and timely decisions regarding the timing and implementation of maintenance activities. The VA can facilitate clarifying the nature of the maintenance problems and choosing the best maintenance alternative available to the decision maker.

4. Methodology: A decision-analytic approach, which is represented as decision trees, is used to study the maintenance problems. The whole problem-solving process is developed on an expert system shell: ProKappa, which runs on a Sun workstation.

5. Results: The research shows that a decision-analytic approach can help people understand the tradeoff between running derated system and committing resources to maintenance. The problem formulation and analysis clarify the value judgments, choices, and chances faced by an operator. The computational results of the test cases are consistent with the practice of experts.

6. Research Status: Two SCE and three PG&E test cases have been used to validate the plausibility of the VA. Based on the results from the test cases, it is found that some issues, e.g., recurrent failure, are worth further study. An extension of the current VA is an ongoing research at CIFE.

ACKNOWLEDGEMENTS

The first author would like to thank Alex Brousilovsky and Andrew Velline for their help in computer programming.

CONTENTS

INTRODUCTION	1
1. BACKGROUND	2
2. METHODOLOGY	4
<i>Decision Analysis in the VA</i>	12
<i>Assumptions</i>	13
3. FORMULATION	15
<i>User Interface and Results</i>	19
4. SIMULATOR	24
<i>User Interface and Results</i>	29
5. PG&E TEST CASES	32
5.1. Test Case 1, HP Turbine First Stage Blade Replacement	32
<i>Problem Description</i>	32
<i>Formulation</i>	32
<i>VA Perspective</i>	33
<i>Baseline</i>	34
<i>Break-even Probability</i>	40
<i>The Results</i>	41
5.2. Test Case 3, FD Fan Bearing Maintenance	43
<i>Problem Description</i>	43
<i>Formulation</i>	44
<i>VA Perspective</i>	44
<i>Baseline</i>	44
<i>The Results</i>	47
5.3. Current Problems in the Stand-alone VA	49
6. TEST CASE TWO (RECURRENT FAILURE PROBLEM)	51
6.1. Test Case 2, Secondary Superheater Outlet Header Stub Tube Welds	
Level of Repair	51
<i>Problem Description</i>	51
<i>Formulation</i>	51
6.2. Simplified Approach to Analyze Recurrent Failure Problems	53
<i>Constant Failure Rate Over Time</i>	53
<i>Increasing Failure Rate Over Time (IFR)</i>	58
6.3. An Approach on Opportunistic Breakdown Maintenance	62

6.4. Two Step Preventive Maintenance	68
6.5. Sensitivity Analysis	71
<i>Effects of the Study Period Duration</i>	71
<i>Effects of the Annual Interest Rate</i>	72
7. DISCUSSIONS AND FUTURE WORK	76
7.1. Test Case Classification	76
7.2. Importance of Distinction	78
7.3. Statement of the Maintenance Problem	79
7.4. Future Work	81
<i>Criteria for Selecting Test Cases</i>	83
APPENDIX A. PG&E Three Test Case Statements	84
Test Case 1	84
Test Case 2	88
Test Case 3	91
APPENDIX B. Note on Minimal Repair	94
REFERENCES	98

INTRODUCTION

The Value Analysis Module (VA) was developed at the Center for Integrated Facility Engineering (CIFE) of the Stanford University. It is part of the Intelligent Real Time Maintenance Management (IRTMM) project, whose purpose is to address the issues encountered in maintenance activities of a conventional steam power plant.

For each maintenance plan given by the user, the VA attaches several timing strategies (repair as soon as possible, at the next period of low demand, at the next scheduled outage, and defer) to it. Then it uses a decision-theoretic approach, i.e., decision analysis, to analyze the problem and returns solutions to the user. We used two test cases provided by Southern California Edison (SCE) to validate the results. Each of them is a real case from SCE's maintenance history. The comparison shows that IRTMM's output is consistent with human judgment.

In addition to studying the two test cases from SCE, we built a simulator to investigate the problem of scheduling multiple maintenance activities. The second half of this report describes three test cases provided by Pacific Gas and Electric (PG&E) to show the capability and limitations of the stand-alone VA. All of them are more complicated than the previous two test cases from SCE. The results show that the VA can handle individual maintenance activity very well. However, as discussed below in Section 6, the VA fails to address the effects of recurrent failure and multiple maintenance activities on power production.

SECTION 1

BACKGROUND

The Value Analysis Module (VA) was originally developed together with two other modules: Situation Assessment (SA) and Planner, under the Intelligent Real Time Maintenance Management (IRTMM) project (Simpa, 1993; Simpa and Hayes-Roth, 1992; Jin and Levitt, 1993; Jin, Kunz, Levitt, and Winstanley, 1992). The project was implemented at the Center for Integrated Facility Engineering (CIFE) of the Stanford University, with funding from Electric Power Research Institute (EPRI), PG&E, SCE, and Shimizu Corporation of Japan.

The IRTMM project wants to address the issues encountered in maintenance activities of a conventional power plant. The three modules represent the three stages in a life cycle of maintenance. Figure 1.1 shows how the three modules communicate with each other.

First, the SA takes the monitoring data of the system operation as input. If it detects any abnormality in the data, then it reasons about the possible causes of the problem. At this stage the main purpose is to diagnose problems and make predictions about the failing system behavior in the future. An example would be a main bearing problem, which would be detected by an increase in shaft vibration.

Next, the diagnosis is passed on to the Planner. There is a physical model of the system built in the Planner. Using this knowledge and the diagnosis, the Planner can elaborate on activities necessary to fix the failing problem. Upon generating maintenance plans, the Planner also looks for mergeable plans and creates merged plans. At this stage, maintenance plans are generated based on the diagnosis and knowledge about the system.

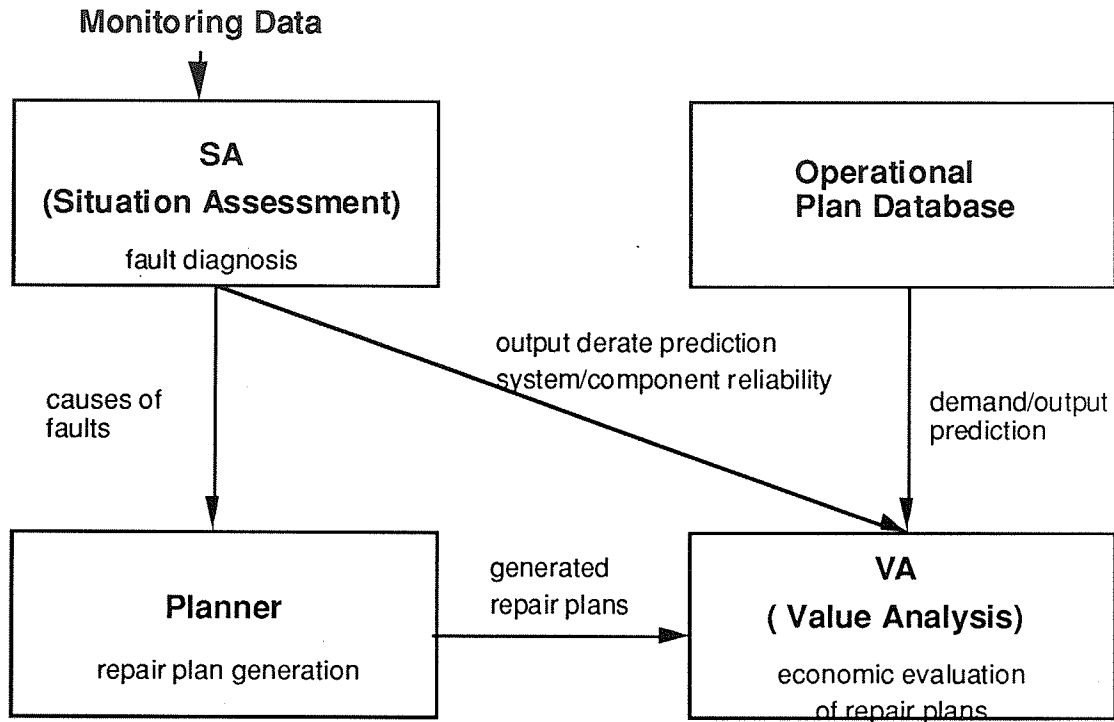


Figure 1.1 IRTMM Module Configuration.

Finally, the VA receives system behavior predictions from the SA and maintenance plans from the Planner. For each maintenance plan given by the Planner, the VA attaches several timing strategies (repair as soon as possible, at the low demand, at the next scheduled outage, and defer) to it. Based on each timing strategy, the VA generates a set of possible scenarios¹ and corresponding values, and then for that specific timing strategy it calculates the average value weighted by the failing component's life distribution. A life distribution is the cumulative probability of failure as a function of the failing component's operational life. At this final stage, the purpose is to evaluate the economic values associated with each maintenance plan and return them to the user.

1. A scenario is a combination of a possible outcome and the associated probability of occurrence.

SECTION 2

METHODOLOGY

The VA uses a decision-analytic approach to analyze maintenance problems. Given a maintenance activity from the user, the VA computes the cost and savings of performing it at different times. Therefore, the decision is about the timing of each maintenance activity. Given a choice, we take chances. The failing component might break at any time according to a prescribed life distribution. Thus, the decision we made cannot lead to a deterministic outcome but to a set of possibilities. To facilitate evaluating the effectiveness of each decision we use the standard Operations Research technique of defining an objective function on the outcomes. Usually it is either the cost or benefit. The impact of each decision is measured by the expected value of the objective function, weighted by the probability of occurrence associated with each outcome. The "best" decision is thus the one that has the optimal expected value.

A decision tree shows the choices and chances. A decision tree has two types of nodes: choice node and chance node, as shown in Figure 2.1.




Figure 2.1 Two Basic Nodes in a Decision Tree.

At the choice node we face a set of alternatives, out of which we can only choose one. In the maintenance problem, each alternative should be "Perform the maintenance activity at time ..." Suppose we have three timing alternatives. Then the choice node should be extended as Figure 2.2.




Figure 2.2 Available Alternatives of a Choice Node.

Suppose we choose the second alternative: Perform the maintenance activity at time T2. Then we have two chance outcomes: the failing component fails unexpectedly before the scheduled maintenance, or it survives to T2 and the maintenance is performed². We now incorporate this chance node and redraw the tree, as shown in Figure 2.3.

2. The failing component's life, T, should be a continuous random variable. However, the power demand/production prediction over time are usually discrete. For the convenience of calculation, the failure distribution has to be discretized. In this example, for illustrative purpose we only consider two possibilities: either failed or not failed before T2.

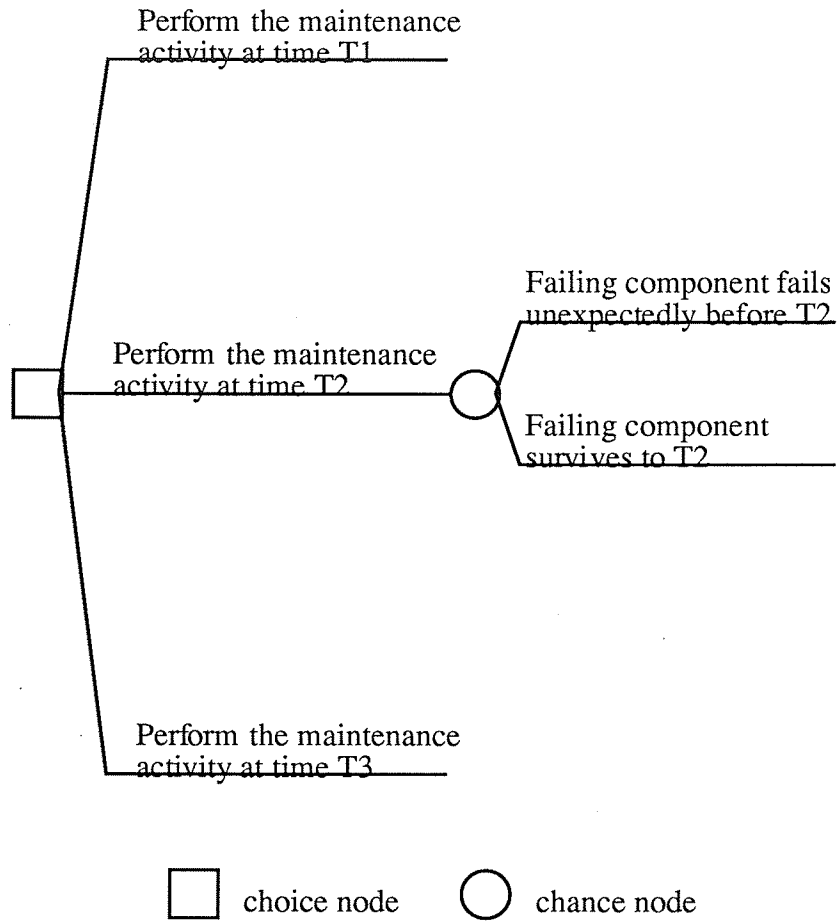


Figure 2.3 Chance Outcomes Given a Timing Alternative.

Now we assign the probability of occurrence to each chance outcomes. Define

$$P(T < T2) = P\{\text{Failing component fails unexpectedly before } T2\}$$

By the probability law,

$$P(T \geq T2) = P\{\text{Failing component survives to } T2\} = 1 - P(T < T2)$$

Now we redraw the decision tree with the probabilities of occurrence attached to it, as shown in Figure 2.4.

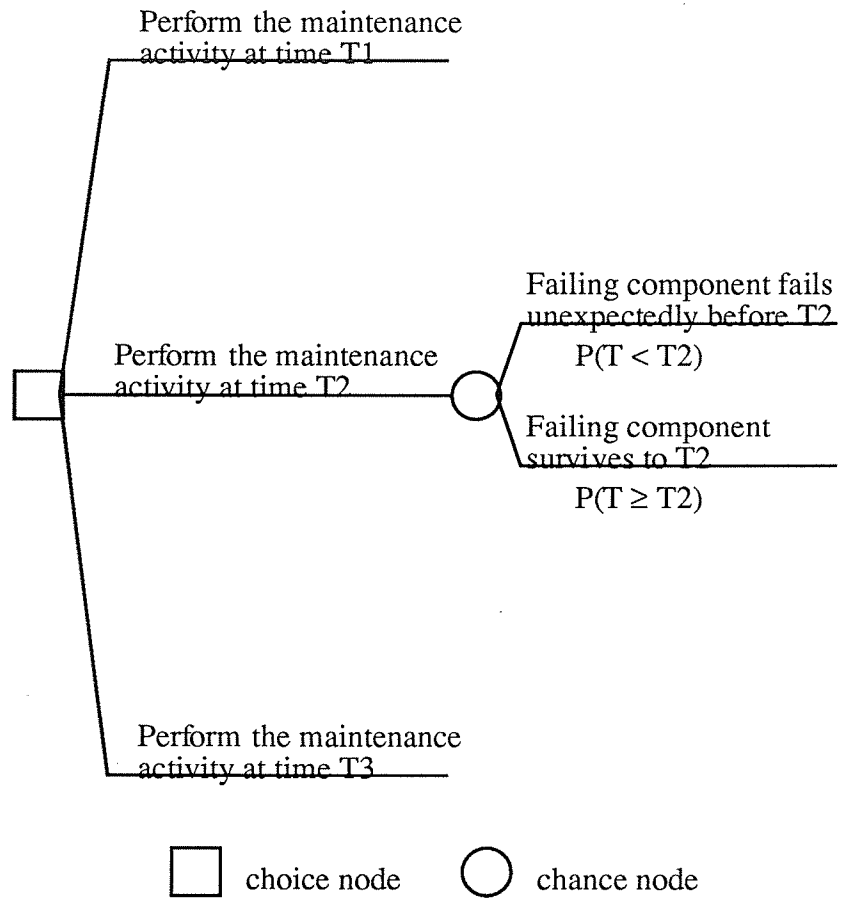


Figure 2.4 Scenarios Given a Timing Alternative.

Notice that a chance outcome with associated probability of occurrence is termed as a scenario.

To evaluate the value of each chance outcome, we need an objective function. Usually it is either cost or benefit. We will use benefit as our objective function.

To simplify the analysis, we compute the benefit of a choice relative to a reference value. We call the reference the *baseline*. The simplest baseline to choose is to assume performing no maintenance activity at any time and no unexpected failure is experienced. The time span, which is called the study period, can be either infinite or finite.

Now we define the benefit. A benefit, B , is a value associated with a scenario computed from the following three parts:

- cost savings, S , of operational mode relative to the baseline operational costs.
- cost of downtime, C_{downtime} , due to buying replacement power from other sources.
- cost of repair, C_{repair} , either planned or unplanned.

Symbolically, the benefit B can be represented as:

$$B = \sum_{\text{time from now} = 1}^{\text{end of study period}} (S - C_{\text{downtime}} - C_{\text{repair}}) \quad (2.1)$$

In the above example, let's assume the study period is composed of three periods: $T1$, $T2$, $T3$, with $T1 < T2 < T3$.

In the first chance outcome of the sample decision tree, the component fails unexpectedly before $T2$. This failure affects the power plant's electricity production and causes unplanned repair cost.

We assume the unexpected failure causes a forced shutdown, and the power demand has to be met by buying power from other sources. Denote the loss as $C(\text{loss}, T < T2)$. The unplanned repair usually costs more than a planned repair. Denote the cost for unplanned repair as $C(\text{unplanned repair})$. The benefit of the chance outcome, Failing component fails unexpectedly before $T2$, is thus the negative of the sum of $C(\text{loss}, T < T2)$ and $C(\text{unplanned repair})$. That is,

$$\begin{aligned} &\text{Benefit for Failing component fails unexpectedly before } T2 \\ &= - C(\text{loss}, T < T2) - C(\text{unplanned repair}) \end{aligned}$$

For the second chance outcome, Failing component survives to T2, we might still suffer from a plant shutdown to perform the scheduled maintenance activity. However, since it is planned in advance, the downtime duration should be shorter than that of a forced shutdown. Define the loss due to a planned shutdown as $C(\text{loss}, T \geq T2)$. Note that $C(\text{loss}, T \geq T2) < C(\text{loss}, T < T2)$. The cost of planned repair, $C(\text{planned repair})$, is usually smaller than $C(\text{unplanned repair})$. Therefore, we have

$$\begin{aligned} & \text{Benefit for Failing component survives to T2} \\ & = - C(\text{loss}, T \geq T2) - C(\text{planned repair}) \end{aligned}$$

Note that in the calculation, the component fails at most once because we assume that once the component is repaired (either by unplanned or planned repair), it never fails again within the study period by the same mechanism. Therefore, we have only one level of chance node in the decision tree. No additional failure is considered.

Now the decision tree is redrawn in Figure 2.5 to show the benefit for each chance outcome.

To evaluate the value of the second alternative, we compute the expected value of the benefit. That is,

$$\begin{aligned} & E[@T2] \\ & = \text{Expected value given decision to perform the maintenance activity at time T2} \\ & = P(T < T2) \times [- C(\text{loss}, T < T2) - C(\text{unplanned repair})] \\ & \quad + P(T \geq T2) \times [- C(\text{loss}, T \geq T2) - C(\text{planned repair})] \end{aligned}$$

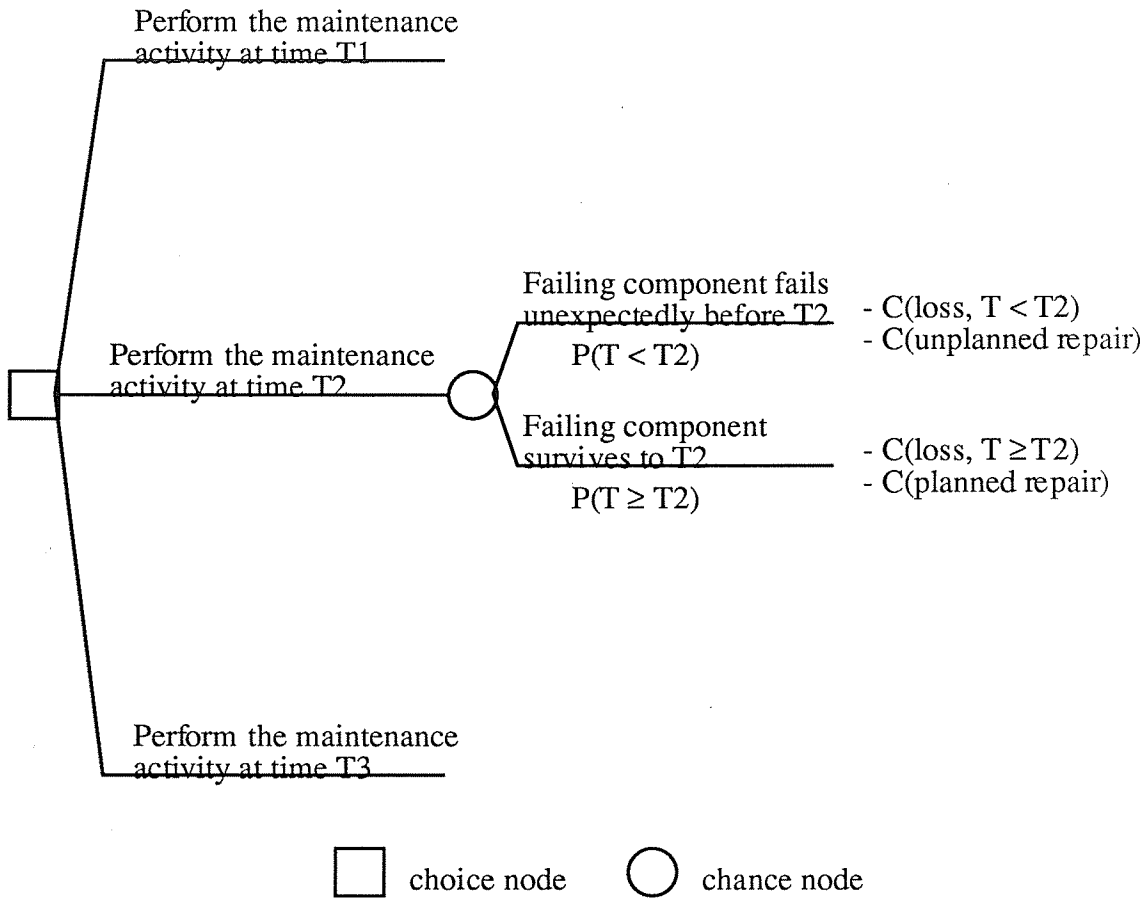


Figure 2.5 Scenarios with Objective Values Given a Timing Alternative.

The decision tree is drawn again with $E[@T2]$ attached to the second alternative, as shown in Figure 2.6.

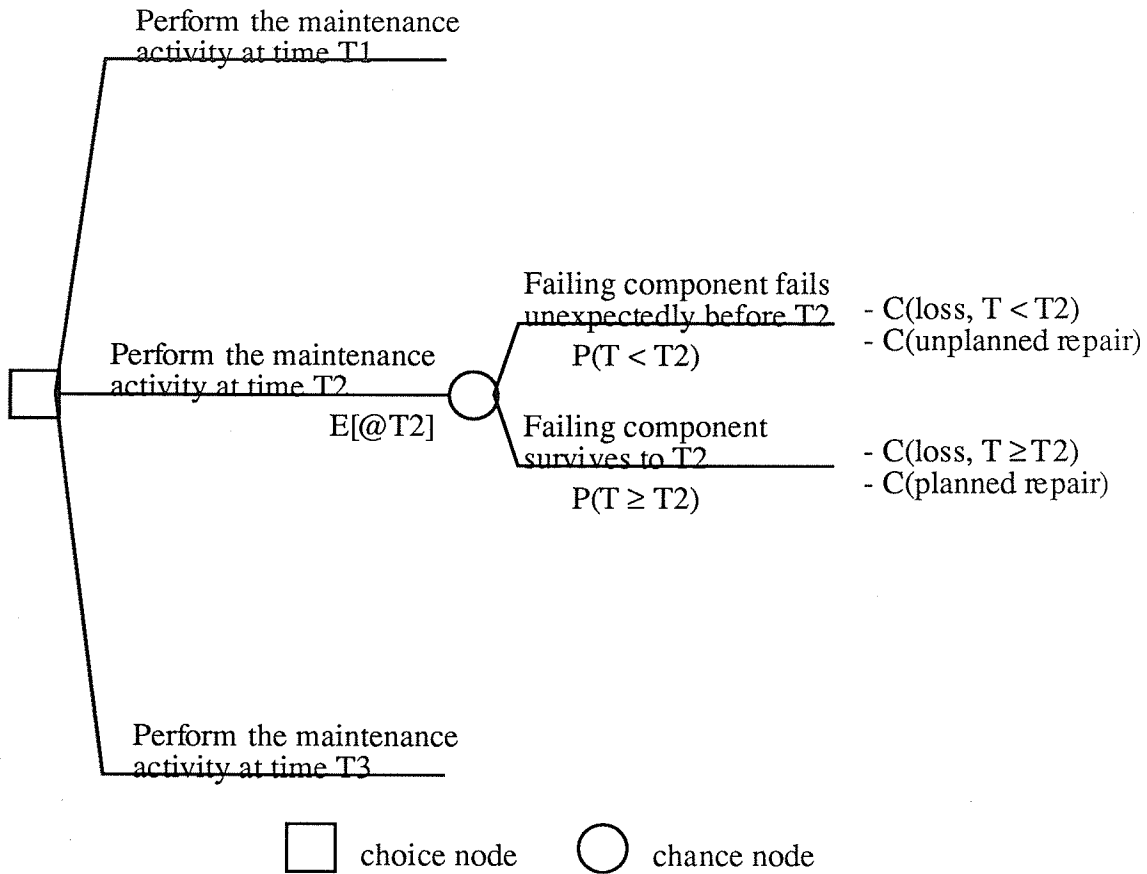


Figure 2.6 Expected Value for the Second Timing Alternative.

In the same way, we can complete the analysis of the whole decision tree and calculate the expected values for the all alternatives. This is shown in Figure 2.7.

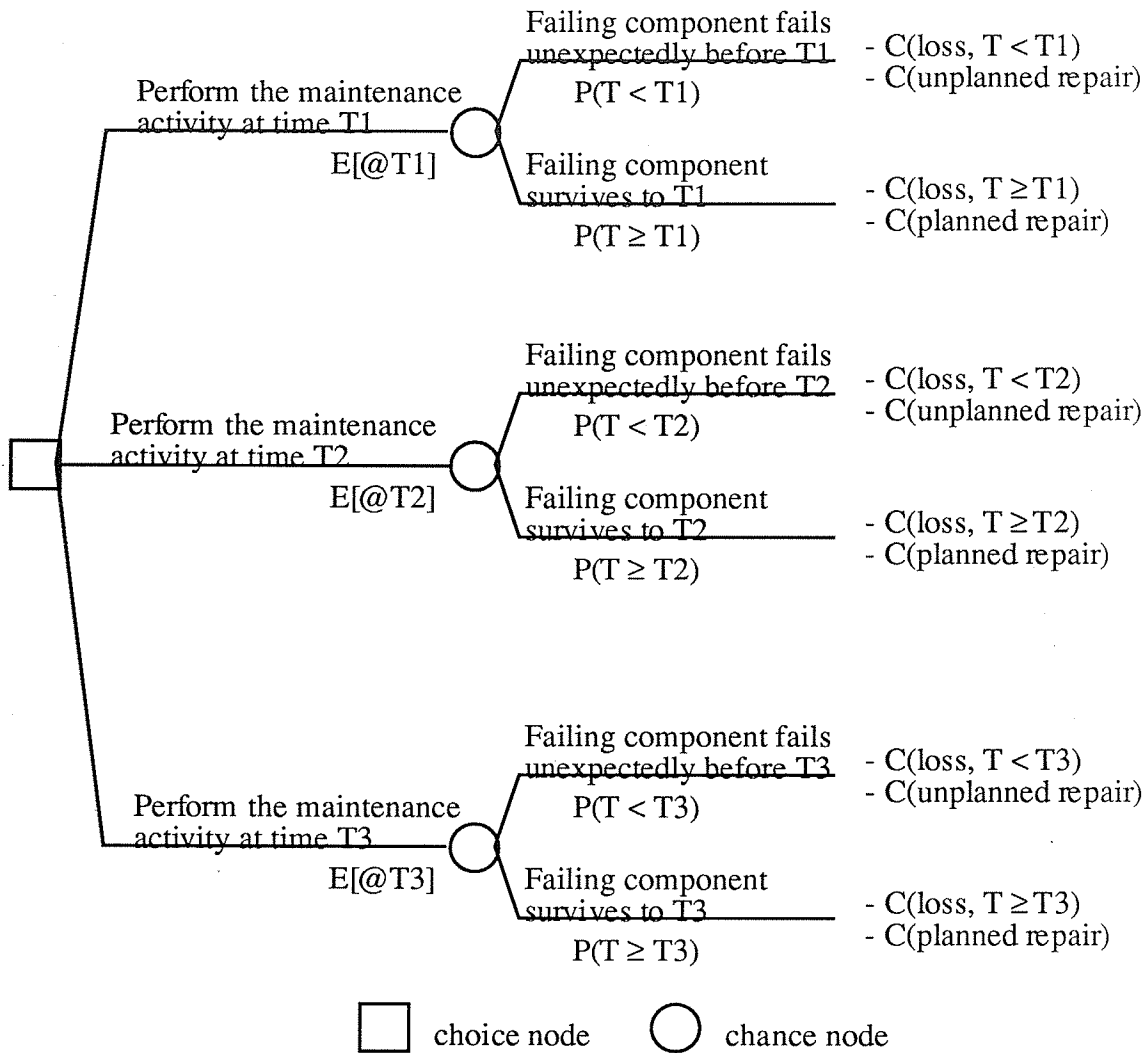


Figure 2.7 A Complete Decision Tree.

Suppose $E[@T3] > E[@T1] > E[@T2]$, then the optimal decision is to perform the maintenance activity at time T3 because it has the highest expected benefit.

Decision Analysis in the VA

The VA uses the decision trees. However, the tree structure is not shown in the user interface. Only the expected value for each timing alternative is returned to the user interface. For each alternative, the VA generates all possible scenarios based on the time unit defined by the input data. For example, if the power demand prediction is in units of

years and the study period is 14 years, we will have 4 chance outcomes if we schedule the maintenance activity in the fourth year. They are:

- Failing component fails unexpectedly in the first year;
- Failing component fails unexpectedly in the second year;
- Failing component fails unexpectedly in the third year;
- Failing component survives to the fourth year.

Once we have calculated the probabilities and benefits associated with all scenarios of a given timing alternative, we can easily calculate the expected benefit for that alternative and return it to the user. In the VA we do not suggest the optimal timing decision but show the expected benefit for all the timing alternatives input by the user.

Assumptions

Right now the VA considers only one failing component at a time, which means there is only one maintenance activity to be scheduled. The failing component will cause the power production reduction over time. The life distribution of this failing component is assumed to be Weibull because it is the most widely used parametric family of failure distributions (Barlow and Proschan, 1965: 13-16). The Weibull distribution has the form:

$$P\{T < t\} = F(t) = 1 - e^{-\lambda t^\alpha} \quad (2.2)$$

where T: failing component's life.

λ : scale parameter.

α : shape parameter.

We assume the failing component is critical to the system, which means if the component fails unexpectedly or is under planned repair, the system is down. The only

thing that is random in the analysis is the failing component's life, T . The durations for planned and unplanned repair are assumed to be constant. The failing component is monitored continuously, thus we know exactly when it fails. Once it fails, the unplanned repair is implemented immediately. The time to perform unplanned repair is assumed to include both wait time to assemble parts and crew plus actual repair time. The deterioration in component life happens only when the component is operated. That is, when it is idle during a scheduled outage, it does not age with time.

Another important assumption is that once the component is repaired it is restored to an as-good condition and will not fail again within the study period.

SECTION 3

FORMULATION

In this project we are concerned about the condition-based maintenance. That is, the maintenance decision has to be made when the monitoring devices signal certain component deterioration. Therefore, we can take the time of the degradation signal as time zero, which corresponds to the age of the component at the time of this signal. Here we take the time unit to be one hour and the study period to be n hours.

Assume the component failure always occurs at the beginning of each hour. For example, the component failing at the k -th hour means that the breakdown happens at the beginning of the k -th hour (or the end of the $(k-1)$ -th hour).

Suppose the component fails at the k -th hour. Total cost of failure would be:

Total Cost if the component fails at the k -th hour
= costs due to deterioration from the first to the $(k-1)$ -th hour, e.g., cost of deration and marginal costs of operating an aged component, i.e., cost of production loss
+ costs due to shutdown (for repair) from the k -th to the $(k+t_r-1)$ -th hour
+ costs of unplanned repair labor and materials

where t_r : repair time.

Each of the first two terms can be further divided into two parts:

cost of buying replacement power to meet the demand - saved variable cost

The second term comes from the fact that when the derated output is less than demand, we spend more in buying replacement power and less in variable operating cost,

which means we save money in plant operation by buying replacement power from outside sources. Each of the two terms are defined as:

cost of buying replacement power

$$= \max [(\text{demand} - \text{production}), 0] \times \text{replacement power cost}$$

saved variable cost

$$= \max [(\text{demand} - \text{production}), 0] \times \text{variable operating cost of plant operation}$$

Taking the time value of money into account, we have

Total Cost if the component fails at the k-th hour

$$= \sum_{i=1}^{k-1} (\text{discount factor})^i \cdot (\text{replacement power cost} - \text{saved variable cost}) \\ + \sum_{i=k}^{k+t_r-1} (\text{discount factor})^i \cdot (\text{replacement power cost} - \text{saved variable cost}) \\ + (\text{discount factor})^k \cdot (\text{costs of unplanned repair labor and materials})$$

For the third term in the above equation, we assume that the cost of unplanned repair labor and materials is spent at the k-th period.

Now we will show how the probability of occurrence for each scenario is calculated.

As mentioned before, we use the Weibull distribution to describe the failing component's life. The two parameters, λ and α , that control the scale and shape of the curve have a specific relation with the component's mean time to failure μ and the standard deviation, σ , of life. Using method of moments (Rice, 1988: 230-234; Lloyd and Lipow, 1977: 137-138), we have

$$(1 + r^2)\Gamma^2\left(\frac{1}{\alpha} + 1\right) - \Gamma\left(\frac{2}{\alpha} + 1\right) = 0 \quad (3.1)$$

$$\lambda = \left[\frac{\Gamma\left(\frac{1}{\alpha} + 1\right)}{\mu} \right]^\alpha \quad (3.2)$$

where μ : failing component's mean time to failure;

σ : standard deviation of failing component's life;

λ : scale parameter of the Weibull distribution;

α : shape parameter of the Weibull distribution;

$r = \frac{\sigma}{\mu}$: coefficient of variation;

$\Gamma(\bullet)$: gamma function.

Given μ and σ we can calculate r . Since α is just a function of r (see equation (3.1)), we can calculate α by numerical methods. From the second relation, λ is obtained once α is known. Notice in equation (3.1) the coefficient of variation r uniquely determines α . When r is less than 1, α is greater than 1, which in term causes the failure rate to increase over time. That is,

$$r < 1 \Leftrightarrow \alpha > 1 \Leftrightarrow \text{increasing failure rate}$$

If we plug α and λ into equation (2.2), we can easily obtain the probability of occurrence for each scenario. Through this transformation, once the user specifies the component's mean time to failure μ and standard deviation of component's life σ , the VA can quickly calculate the probabilities of occurrence for all scenarios.

If we choose a alternative of repairing the component at the m -th hour, then we face m possible scenarios: the component fails before the scheduled repair at the j -th hour with probability of P_j , $1 \leq j \leq m - 1$; and the component survives to the m -th hour with probability P_s . Obviously, by probability law,

$$\sum_{j=1}^{m-1} P_j + P_s = 1$$

where $P_j = P\{j-1 \leq T < j\}$;

$$P_s = P\{m-1 \leq T\}$$

T: failing component's life.

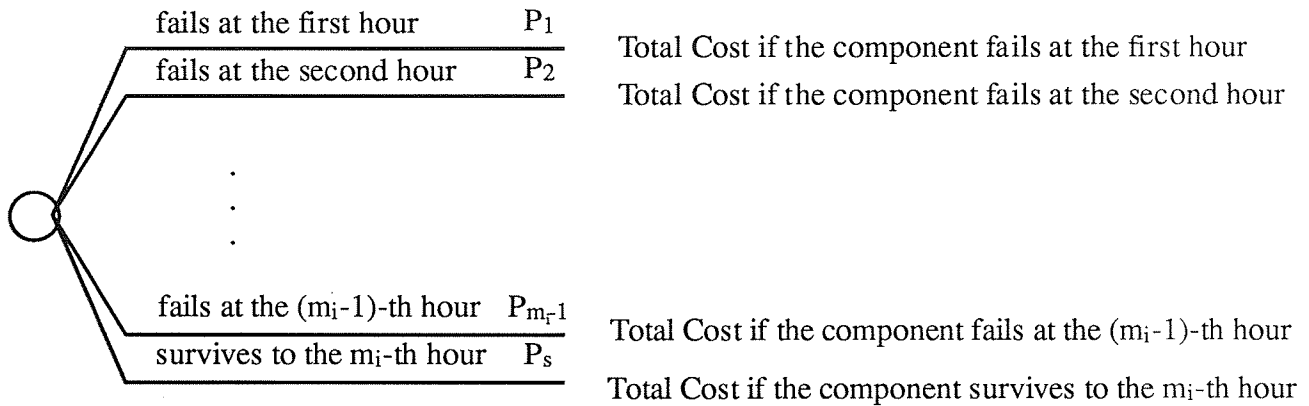
Therefore,

$$\begin{aligned} & \text{Expected Total Cost for the alternative to repair the component at the } m\text{-th hour} \\ &= \sum_{j=1}^{m-1} P_j \times \text{Total Cost if the component fails at the } j\text{-th hour} \\ & \quad + P_s \times \text{Total Cost if the component survives to the } m\text{-th hour} \end{aligned}$$

The expected cost of deferring the repair is slightly different from that of the above. The component may survive to the end the study period without any repair. Thus the Total Cost if the component survives to the end of the study period is just the cost caused by degraded performance within the study period. There is no cost of repair in it. Therefore, the expected cost of deferring the repair would be

$$\begin{aligned} & \text{Expected Total Cost for the defer alternative} \\ &= \sum_{j=1}^n P_j \times \text{Total Cost if the component fails at the } j\text{-th hour} \\ & \quad + P_s \times \text{Total Cost if the component survives to the end of the study period} \end{aligned}$$

Suppose that there are $q+1$ alternatives: repair at the m_i -th hour, $i = 1, \dots, q$; defer the repair. There are $n+1$ scenarios for the last alternative and the number of scenarios of other alternatives is fewer. For example consider the alternative of repairing at the m_i -th hour. The possible scenarios for this alternative can be represented as the following tree:



Note that the first $(m_i - 1)$ scenarios are the same as those of the defer alternative. Therefore, we can just calculate the scenarios for the defer alternative and use part of them to calculate the expected total costs for other alternatives.

If we want to address the decision maker's risk attitude, we should replace the cost (or benefit) for each scenario with the decision maker's utility about the outcome. However, determining the decision maker's risk attitude and preferences is outside the scope of the present study. Thus we only calculate the expected cost (or benefit) for each timing alternative. Once the decision maker's utility function is given, we can easily incorporate it into the VA. The calculation of the expected utility is the same as that of expected cost (or benefit).

User Interface and Results

We have used the demand prediction and cost data from SCE to study the timing of one maintenance activity. Figure 3.1 shows a screen dump of the user interface of the VA.

The dialog box shows that the current problem is the boiler tube leak in Unit 3. The prediction about the failing component's life distribution is specified by the mean time to failure, 1000 hours, standard deviation, 250 hours, and time it has been operated (i.e., its current age), 350 hours. The information about the cost structure and power production prediction can be accessed from the **Parameters** button at the bottom. Figure 3.2 shows the parameter dialog box.



Repair Plan Evaluation

Components with Fault: **UNIT_3 SUBFP BEARING WEAR_EXTENT**
UNIT_3 BOILER_TUBE LEAK

Mergeables:

Predecessors: UNIT_3_SUBFP_BEARING_WEAR_EXTENT

Values of UNIT_3_BOILER_TUBE_LEAK

Mean Time To Failure (Hrs):	1000	Standard Deviation (Hrs):	250	Time Component Operated (Hrs):	350
Repair Time:	ASAP	Off Peak	Next Outage	Defer	User Defined
Start Time (Hrs):	16	n/a	n/a	669	
Duration (Hrs):	55	n/a	n/a	55	
Probability of Survival:	1.0	n/a	n/a	0.49	n/a
E(Operating Cost) (K\$):	4225.0	n/a	n/a	4147.0	n/a
E(Replacement Power Cost)(K\$):	438.0	n/a	n/a	556.0	n/a
E(Repair Cost) (K\$):	68.0	n/a	n/a	35.0	n/a
Revenue (K\$):	5563.0	n/a	n/a	5563.0	n/a
E(Net Profit) (K\$):	832.0	n/a	n/a	826.0	n/a

Close Parameters Evaluate Explain VA List Plan Gantt Merge Help

TopP IdleShe Simulated SFPPanel XTerm

Figure 3.1 The User Interface of the VA.

Console Login
 I
 gantt% sd
 gantt% sd

IRTTMM

- MAIN MENU**
- Open/Close P&ID
 - Comm Monitor
 - Work Requests
 - Load Plant Model
 - Start Simulation
 - Exit IRTMM
 - Machine Mode

VA Parameters

Power Replacement Cost (cents/KwHr) 6

Sale Price (cents/KwHr) 5

Production Variable Cost (cents/KwHr) 4

Production Fixed Cost (\$ per Hr) 100

Degradation offset (Mw) 300

Degradation slope (Mw/hr) -0.2

OK Cancel

Repair Plan Eva

UNIT_3_SUBFP_BEARING_WEAR_EXTENT

Components with fault:
UNIT_3_BOILER_TUBE_LEAK

Mergeables:
 UNIT_3_SUBFP_BEARING_WEAR_EXTENT


Protections:
 UNIT_3_SUBFP_BEARING_WEAR_EXTENT

Values of UNIT_3_BOILER_TUBE_LEAK

Mean Time To Failure (Hrs):	1000	Standard Deviation (Hrs):	250	Time Component Operated (Hrs):	350
Repair Time:	ASAP	Off Peak	Next Outage	Defer	User Defined
Start Time (Hrs):	16	n/a	n/a	669	n/a
Duration (Hrs):	55	n/a	n/a	55	n/a
Probability of Survival:	1.0	n/a	n/a	0.49	n/a
E(Operating Cost) (K\$):	4225.0	n/a	n/a	4147.0	n/a
E(Replacement Power Cost) (K\$):	439.0	n/a	n/a	555.0	n/a
E(Repair Cost) (K\$):	66.0	n/a	n/a	35.0	n/a
Revenue (K\$):	5383.0	n/a	n/a	5553.0	n/a
E(Net Profit) (K\$):	832.0	n/a	n/a	825.0	n/a

Close Parameters Evaluate Explain VA List Plan Cont Merge Help

Figure 3.2 Parameters Dialog Box.



The first two entry boxes specify the price to purchase electricity from another source and the price to sell the electricity to customers, respectively. The two entry boxes in the middle are for the fixed and variable operating costs of electricity generation. The last two entry boxes show the future power production prediction. The Degradation offset shows current production level, 300 MW in Figure 3.2, and the Degradation slope shows how the production degrades with time, the production decreases at a rate of 0.2 MW per hour in Figure 3.2. In Figure 3.2, the power production degradation is resulted from the boiler tube leak, which causes steam escape and less heat to the turbine. Right now the production is assumed to degrade linearly with time.

There are only two timing strategies considered in the above example. Namely, repair as soon as possible (ASAP) and defer the repair until after the study period (Defer). The study period duration is 4 weeks (672 hours) in this example. Figure 3.1 shows that ASAP is better than Defer in this case, because the expected net profit (see the row of **E[Net Profit]**) for ASAP is \$832K, which is higher than that of Defer, \$826K. Note that all the rows with **E[•]** represent the *expected* value. For example, **E[Repair Cost]** is the expected value of repair cost (both planned and unplanned). Therefore, the value for **E[Repair Cost]** of ASAP is different from that of Defer, because the probabilities of failure for the two timing alternatives are different.

The user can click on the **Explain VA** button to see how each timing alternative affects the failure distribution, the demand/production prediction, and the revenue/cost curve over time. Figure 3.3 shows those curves for deferring the maintenance activity. The horizontal axis in each graph represents the time from now. Since the curves are colored on the monitor screen, it is not so distinguishable on the black-and-white screen dump. However, it is shown here to emphasize that the user can use this feature to understand the effects of each timing alternative on the future outcomes.

Console Login

gantt@sd
gantt@sd
gantt@sd


IRTMM

MAIN MENU


- Open/Close P&ID
- Comm Monitor
- Work Requests
- Load Plant Model
- Start Simulation
- Exit IRTMM
- Machine Mode

Parameter Images Panel


Option: UNIT_3_BOILER_TUBE_LEAK_Defer_op, Probabilities



Option: UNIT_3_BOILER_TUBE_LEAK_Defer_op, Production and Demand



Option: UNIT_3_BOILER_TUBE_LEAK_Defer_op, Costs and Revenue



Close

Figure 3.3 ExplainVA Dialog Box.

TopPIDShe Simulated

SECTION 4

SIMULATOR

The VA only considers one single maintenance activity when calculating the expected value of each timing alternative. If more than one maintenance activities are to be scheduled simultaneously, we can follow the method of Section 2 to draw a corresponding decision tree. However, the tree will grow very big such that it will take a huge amount of computation time to do the calculation. Figure 4.1 is just a part of a decision tree regarding the scheduling of two maintenance activities with a time span of 12 periods. In a power plant, there may be dozens of maintenance activities to schedule over a year. If the granularity of the demand/production period is hour, we cannot easily generate the whole tree and do the calculation.

An alternative approach to study the multiple-activity scheduling problem is to use Monte Carlo simulation to sample some of the possible scenarios within the decision tree and then use a statistical method to estimate the true expected value associated with each scheduling decision. Figure 4.2 shows how the simulation is implemented.

As shown in the figure, the class `Components` lists all the failing components that affect the system's output. The class `System Configuration` controls how those components are structured in the system and thus decides the system's behavior. These two classes comprise the system behavior model. The `Simulator` retrieves each failing component's life distribution from `Components` and maintenance information from the class `Maintenance Activities`. Then it generates a set of `Events`, which occur within the study period. Each event stands for performing a planned (or unplanned) maintenance activity on a failing (or failed) component at a scheduled (or failure) time. Once the events are generated, the `Simulator` combines the system behavior and events to calculate the value for this specific simulation.

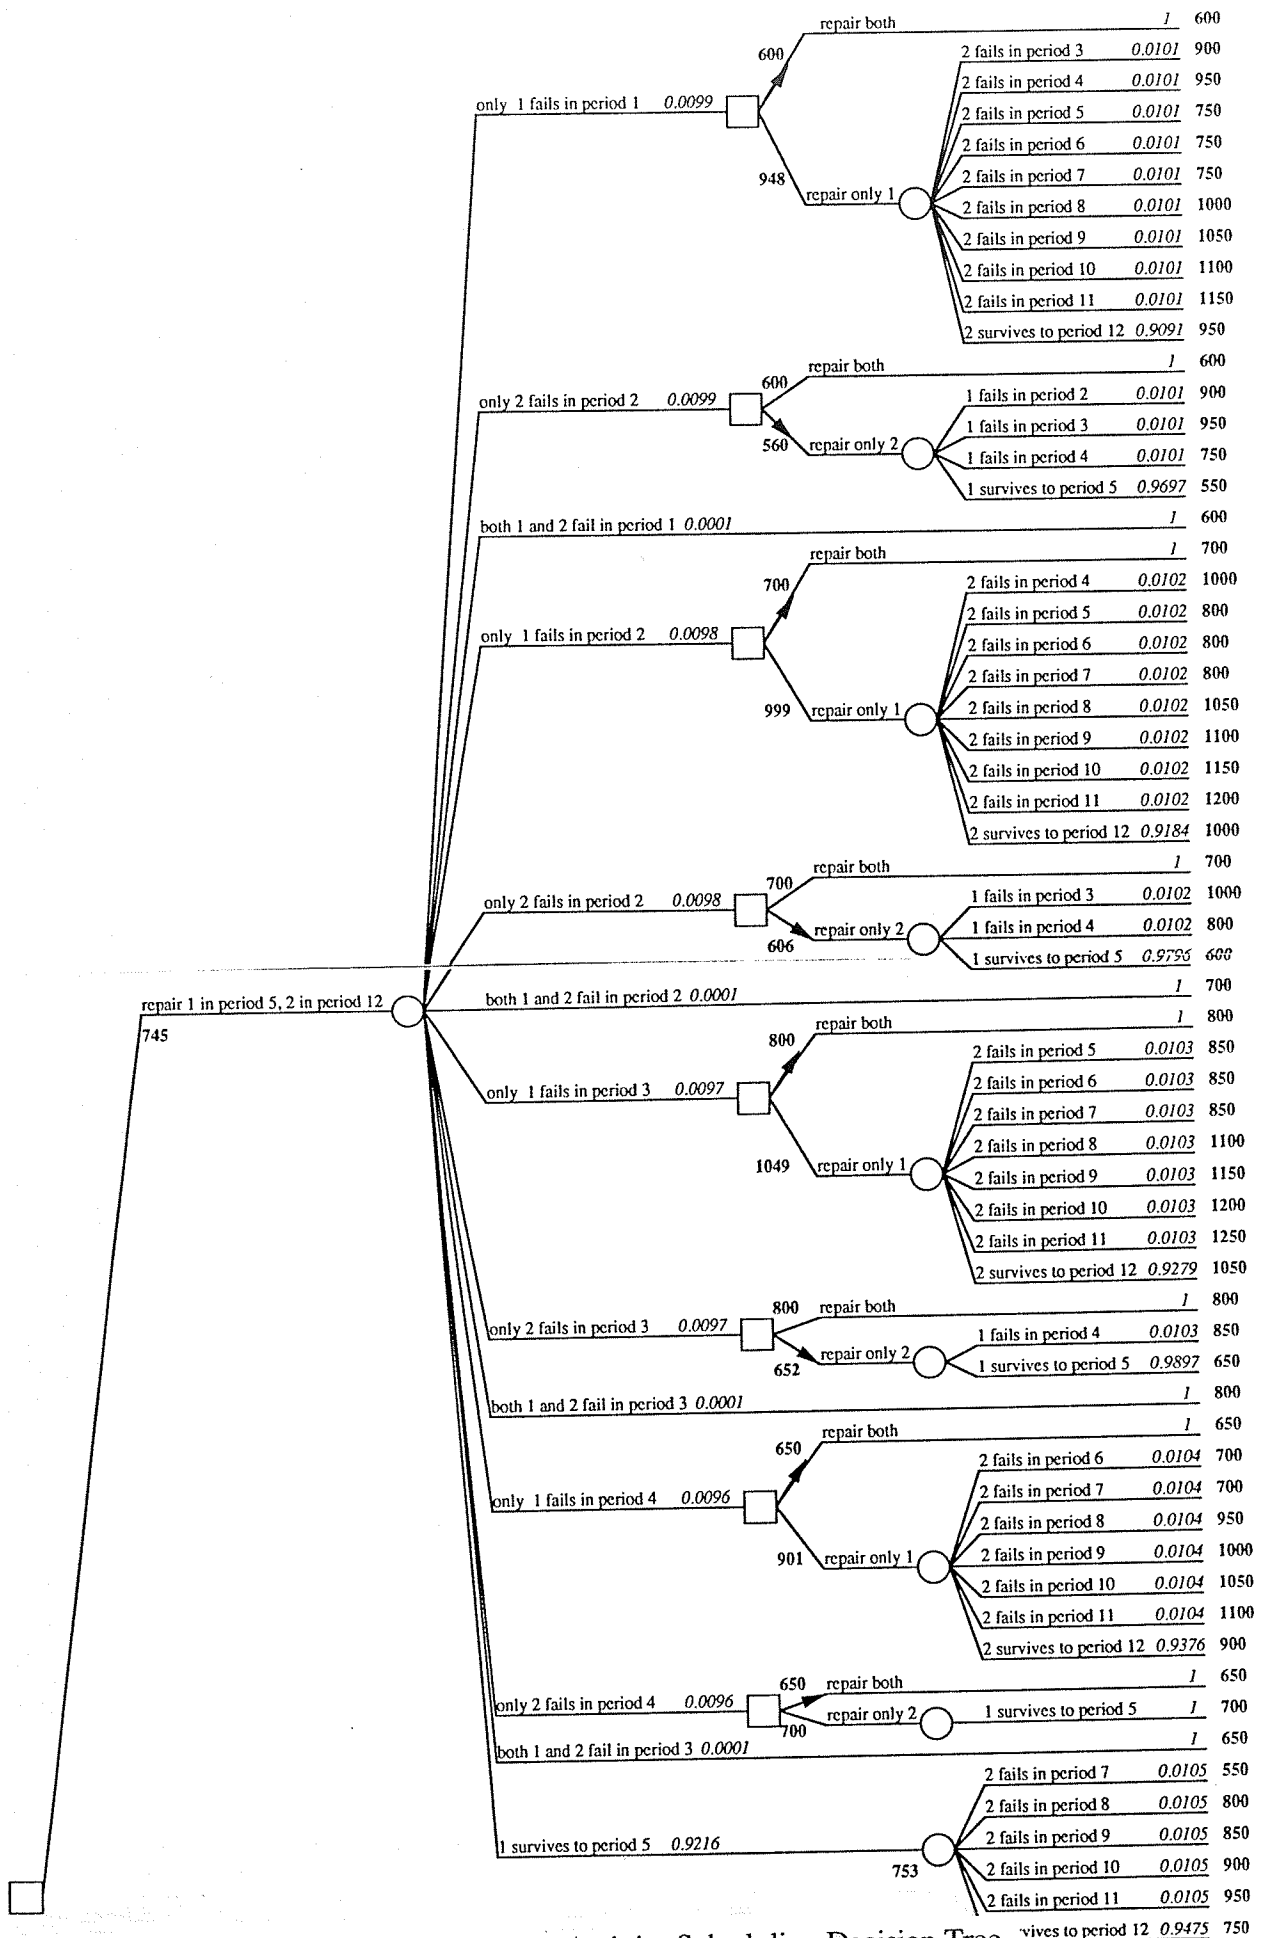


Figure 4.1 Part of Two Maintenance Activity Scheduling Decision Tree.

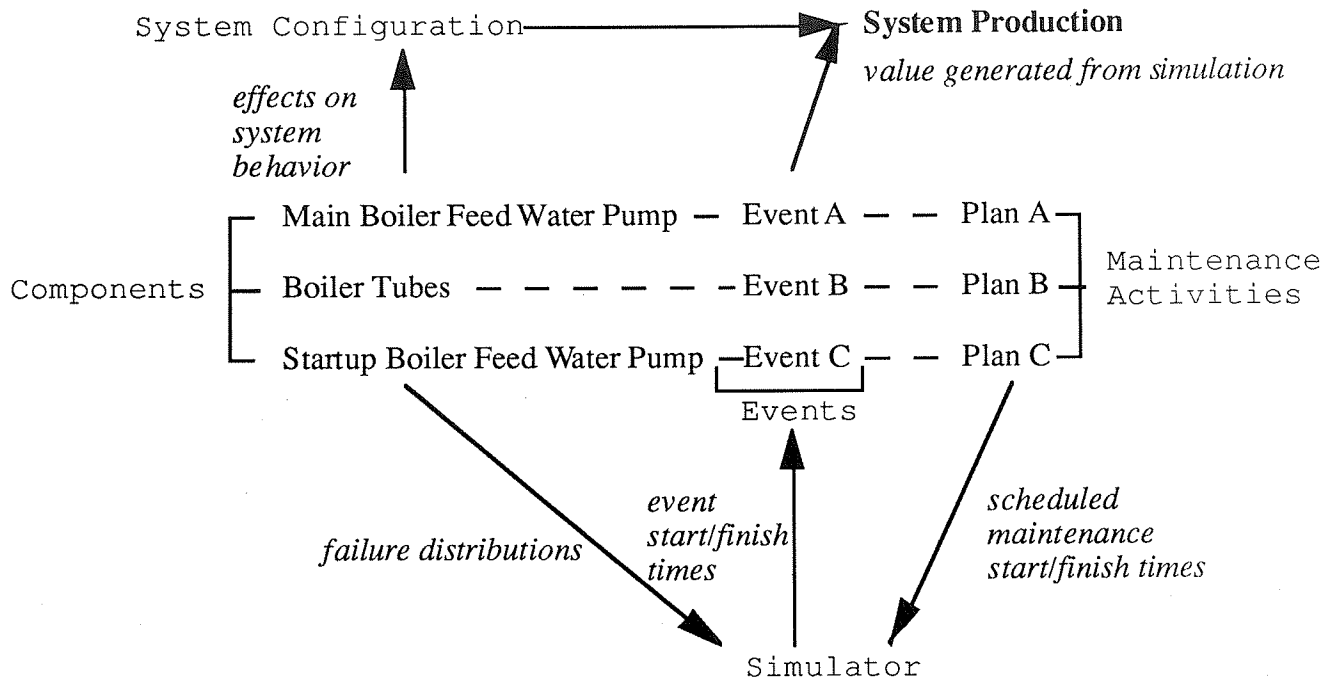


Figure 4.2 The Configuration of the Simulator.

Below is a brief description of how the user runs the simulator.

Before the simulation is implemented, the user has to specify what failing components are to be repaired within the study period. Also included are the failing component's operating characteristics: the effects of its deterioration on the power production and the time to failure prediction. A maintenance activity is then assigned to repair each component. The user has to specify the costs, durations of repair (either planned or unplanned), and timing for all maintenance activities. The last step for the data input is the number of simulations to be implemented. We will discuss how many simulations are needed later.

After the data input for the simulation is done, the simulator generates a set of random numbers to designate the failure times of all failing components. If the simulated failure time of a failing component is earlier than the scheduled repair time intended for it, then we have an unexpected failure, otherwise we have a planned repair. Each unexpected failure or planned repair represents an event in the simulation. Upon determining the start

time of each event, the simulator begins to simulate the event evolution over time and calculate the value associated with it. Each simulation is a repetition of the above process. Notice it is not necessary that we have identical results from any of the two simulations, because the failure times are randomly generated from the random number generator.

Using simple statistics we can calculate the estimate for the true expected value of each scheduling decision. By law of large numbers, the estimate will approach to the true expected value as the number of simulations grows.

In reality, we cannot run infinite number of simulations. Even if it is possible, it should be easier to generate the whole decision tree and calculate the true expected value for each scheduling decision. Thus we face the problem of deciding the appropriate number of simulations.

If we run n simulations then the simulation result contains n values, each corresponds to one simulation. Denote X_i as the value generated from the i -th simulation, $i = 1, \dots, n$. We can construct an unbiased estimate $\bar{X}(n)$ for the true expected value $E[X]$ and confidence interval based on the t-statistics. That is,

$$\bar{X}(n) = \frac{\sum_{i=1}^n X_i}{n} \quad (4.1)$$

confidence interval with $100(1 - \alpha)$ percent confidence

$$= [\bar{X}(n) - \delta(\alpha, n), \bar{X}(n) + \delta(\alpha, n)] \quad (4.2)$$

where $\delta(\alpha, n) = t_{n-1, 1-\alpha/2} \sqrt{\frac{S^2(n)}{n}}$

$$S^2(n) = \frac{\sum_{i=1}^n (\bar{X}(n) - X_i)^2}{n - 1}$$

Given a $100(1 - \alpha)$ percent confidence, we can choose the number of simulations n such that the corresponding confidence interval is less than a prespecified range. The problem with this approach is that the random number X_i is not normally distributed and thus the confidence interval based on t-statistics is only an approximate one. To overcome this problem, we use a method proposed by (Law and Kelton, 1991: 532-536) to determine the appropriate number of simulations.

First we run 500 experiments of simulation. Each experiment constitutes n simulations. After each experiment we perform the simple statistics as above to get the estimate $\bar{X}^j(n)$ and $\delta^j(10, n)$, $j = 1, \dots, 500$. Note that we construct 90% confidence interval for each experiment and thus α is equal to 10. Then for each experiment we compare the 90% confidence interval $[\bar{X}^j(n) - \delta^j(\alpha, n), \bar{X}^j(n) + \delta^j(\alpha, n)]$ with the true expected value $E[X]$, which is known, to see if $E[X]$ lies in that interval. If it does, we say it is covered by the confidence interval generated from the j -th experiment. Upon finishing 500 experiments, we calculate the estimator \hat{p} and corresponding 90% confidence interval by

$$\hat{p} = \frac{\text{number of coverages}}{500} \quad (4.3)$$

$$\text{Confidence interval with 90\% confidence} = \hat{p} \pm z_{0.95} \sqrt{\frac{\hat{p}(1 - \hat{p})}{500}} \quad (4.4)$$

Equation (4.4) is based on the fact that $(\hat{p} - p)/\sqrt{\hat{p}(1 - \hat{p})/500}$ is approximately distributed as a standard normal random variable (Hogg and Craig, 1970: 187). In this case, p is equal to 0.9. Running the 500 experiments repeatedly for different n , choose an n such that 90% (the percentage of confidence for $E[X]$) lies within the above 90% confidence interval for p .

After running several sets of the 500 experiments for different n , we conclude that 20 simulations are acceptable to make the simple statistics approximately correct.

Right now the simulator can simulate the maintenance outcome with a finite number of maintenance activities. However, it is still assumed that each failing component can have only one failure within the study period. That is, once the failing (or failed) component is repaired, we assume that it will not fail again within the study period. This single-failure assumption is plausible when the repaired component has an expected life much longer than the remaining time of the study period. Nevertheless, if the repaired component still has high failure rate (i.e., short expected life as compared with the study period duration), then the above assumption fails. This will be discussed in more detail for the three PG&E test cases. The simulator can be easily modified to address the recurrent failure problem.

User Interface and Results

The user interface of the simulator is different from that of the VA. We use a Gantt chart to show the time lines of all the scheduled maintenance activities. The user can easily see when each activity is scheduled. After each simulation the Gantt chart also shows when the activities are actually performed. Some of the maintenance activities are implemented earlier than scheduled because of the unexpected failures. Viewing both the planned and unplanned results gives the user a feeling about how frequently an unexpected failure will occur and when it is mostly likely to happen. An experienced user can compare these dynamics with past experiences to see if the assumed life distributions of the failing components are appropriate. If they are not, the user can easily adjust the parameters that control the life distributions or change to other types of distributions. Although the user can make changes by looking at the curves of the life distributions, the dynamics give a more concrete feeling about the failing components' behavior. Secondly, looking at the demand prediction, a user can decide when to schedule the maintenance activities and later use simulator to compute expected benefits of the alternative.

Figure 4.3 shows the simulator's user interface. There are three scheduled maintenance activities, RepairPlanA, RepairPlanB, and RepairPlanC. The user can click on any of the time lines to pop up a parameter dialog box, which contains the information

regarding the activity and the designated component. Here the dialog box for RepairPlanB is shown. The plan is designed to fix the problems associated with the boiler tubes, whose current efficiency is 85% and is decreasing with a rate of 0.15% per hour.




Figure 4.3 The Simulator User Interface.

SECTION 5
PG&E TEST CASES

5.1. Test Case 1, HP Turbine First Stage Blade Replacement

Problem Description

See Appendix A.

Formulation

Study Period: 15 years (1993 - 2007).

Available Choices:

Governor valve strategies

1. 2-valve operation;
2. 4-valve operation until a new rotor is available, then 2-valve operation;
3. 4-valve operation;
4. 5-valve operation.

Turbine rotor strategies

1. Refit rotor and install as soon as possible (in 1995);
2. Refit rotor and install in 1997 scheduled outage;
3. Refit rotor and install in 2007 scheduled outage.

- Chances:
1. Minor failure;
 2. Major failure;
 3. Survival.

Assumptions:

1. After installing the new rotor, there is no risk of failure. Therefore, the operation is switched to 2-valve mode.
2. Cost of repair for any unexpected failure (minor or major) that occurs within the study period is not affected by how long the turbine is operated. It is a function of failure mode only (i.e., minor or major).
3. 5-valve operation is the baseline and without risk of failure.
4. After repairing the failed turbine, operation should be restored to 5-valve until the rotor installation.
5. The conditional probabilities that the failure mode is minor (or major) given a failure occurs are constant over time.
6. If an unexpected failure occurs and the new rotor is available, then the rotor is installed and the problem is fixed. That is, under this circumstance we don't pay for the repair cost of fixing the failure, but only pay the rotor installation cost.
7. Study period finishes at the end of the 2007 scheduled outage. No matter what turbine rotor strategy is employed, after this outage the new rotor is installed anyway and is risk-free for all strategies. Therefore, we don't have to consider chances after this outage.
8. The duration of the rotor installation is less than either of the two scheduled outage durations.
9. Any unexpected failure (minor or major) or scheduled outage will occur in the middle of the time period (the time unit is year in the present case).

VA Perspective

The HP Turbine problem can be fit into the VA's framework, which is mentioned in Section 3. However, this decision problem is more complicated than a simple maintenance timing problem because we have to consider the mode of operation as well.

Besides, we have two failure modes: minor and major, which means one level of chance node is not enough.

Baseline

The baseline in this test case is: Run 5-valve operation within the study period without installing the new rotor. Obviously the benefit B for a baseline is zero. Therefore,

$$B_{\text{baseline}} = 0 \tag{5.1}$$

Since we are supplied with annual data, the time unit for this test case is year. Therefore, the number of periods within the study period is 15.

Consider the deterministic case. Cost savings S_2 of 2-valve (or S_4 of 4-valve) operation from fuel efficiency contribute to relative benefit of this choice. The other three parts that contribute to the benefit are costs:

C_{minor} (or C_{major}): repair cost for minor (or major) failure,

C_{rotor} : overhaul cost for rotor installation,

C_{downtime} : downtime cost (including the power replacement cost).

Notice that C_{downtime} is calculated from the Replacement Power Value (RPV) cost.

$$C_{\text{downtime}} \text{ in the } i\text{-th year} = \text{RPV}(\$/\text{MW-hr}) \times \text{downtime duration (hr)} \times D(i)(\text{MW}) \tag{5.2}$$

where $D(i)$: average power demand (MW) in the i -th year.

Consider one possible chance outcome from the 2-valve operation. We have an unexpected minor failure in 1994 and the new rotor, available in 1995, is to be installed during the 1997 scheduled outage. The benefit can be broken down into the four parts

mentioned above. The table below shows the cost breakdown for this assumed chance outcome.

	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007
Event	-	minor failure	rotor available	-	rotor installation	-	-	-	-	-	-	-	-	-	-
S_2	>	> 0	-	-	> 0	> 0	> 0	> 0	> 0	> 0	> 0	> 0	> 0	> 0	> 0
C_{rotor}	-	-	-	-	> 0	-	-	-	-	-	-	-	-	-	-
C_{minor}	-	> 0	-	-	-	-	-	-	-	-	-	-	-	-	-
$C_{downtime}$	-	> 0	-	-	-	-	-	-	-	-	-	-	-	-	-

By the ninth assumption, the unexpected minor failure occurs in the middle of 1994. Therefore, in the first half year of 1994 the operational mode is 2-valve. However, the new rotor is not installed until 1997. By assumption 4, in the second half of 1994, all of 1995 and 1996, and the first half of 1997 the mode of operation is 5-valve. Therefore, there are no cost savings during this interval. By assumption 1, upon installing the new rotor, we should switch to 2-valve operation, which means we have cost savings again. The years when C_{rotor} and C_{minor} incur are obvious. By assumption 8, the time to install the new rotor is less than the duration of a scheduled outage. Therefore, we don't have $C_{downtime}$ in 1997.

Before calculating the benefit for this particular scenario, we need to know the cost and demand data. The table below is the average demand $D(i)(MW)$ and $RPV(\$/MW-hr)$ prediction over the next 15 years.

	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007
D(i)	125.46	124.1	118.32	112.2	137.7	128.18	131.58	117.98	112.54	104.04	114.92	107.78	113.56	99.9	116.62
RPV	3.1	3.38	3.62	3.82	3.82	4.24	4.68	5.18	5.74	5.84	6.46	7.13	7.77	8.31	8.48

The other assumed data are shown below.

$$S_2 = \$250\text{K/year}$$

$$C_{\text{minor}} = \$851\text{K}$$

Repair duration for minor failure = 10 weeks = 1680 hrs

$$C_{\text{rotor}} = \$1,200\text{K}$$

Annual discount rate = 10%

$$\text{Discount factor } \beta = \frac{1}{1 + \text{annual discount rate}} = \frac{1}{1 + 0.1} = 0.9091$$

The total cost savings for this scenario is calculated as follows.

total cost savings

$$\begin{aligned}
&= S_2, '93 + S_2, '94 + S_2, '97 + S_2, '98 - '07 \\
&= 250 + \beta \times 250 \times \frac{1}{2} \times \frac{365 \times 24 - 1680}{365 \times 24} + \beta^4 \times 250 \times \frac{1}{2} \times \frac{365 \times 24 - 1680}{365 \times 24} + \sum_{i=5}^{14} \beta^i \times 250 \\
&= \$1,460\text{K}
\end{aligned}$$

The first term is for the cost savings from 2-valve operation in 1994. In 1995, we have only about half of annual cost savings S_2 , which occurs in the first half of year 1995. The third term accounts for the cost savings which occurs in the second half of 1997. Notice that in 1996 the operational mode is 5-valve. Therefore, there is no cost savings for it. The last term calculates the cost savings for the remaining ten years.

Let's now calculate C_{downtime} , which incurs in 1995. Using equation (5.2), we have

$$\begin{aligned}C_{\text{downtime}} &= \beta \times \text{RPV@1994} \times \text{Repair Duration for Minor Failure} \times D(1994) \\ &= 0.9091 \times 3.38 \times 1680 \times 124.1 / 1000 \\ &= \$641\text{K}\end{aligned}$$

Now we can calculate the benefit for this scenario.

$$\begin{aligned}B(\text{minor failure @1994 with rotor installed @1997}) & \\ &= \text{total cost savings} - \beta C_{\text{minor}} - \beta^4 C_{\text{rotor}} - C_{\text{downtime}} \quad (5.3) \\ &= 1,460 - 0.9091 \times 851 - 0.9091^4 \times 1,200 - 641 \\ &= - \$774\text{K}\end{aligned}$$

Once we know how to calculate the benefit for each scenario, we can proceed to draw the whole decision tree and calculate the expected benefit for each alternative.

Part of the decision tree for this problem is shown in Figure 5.1. The remaining parts of the tree are similar.

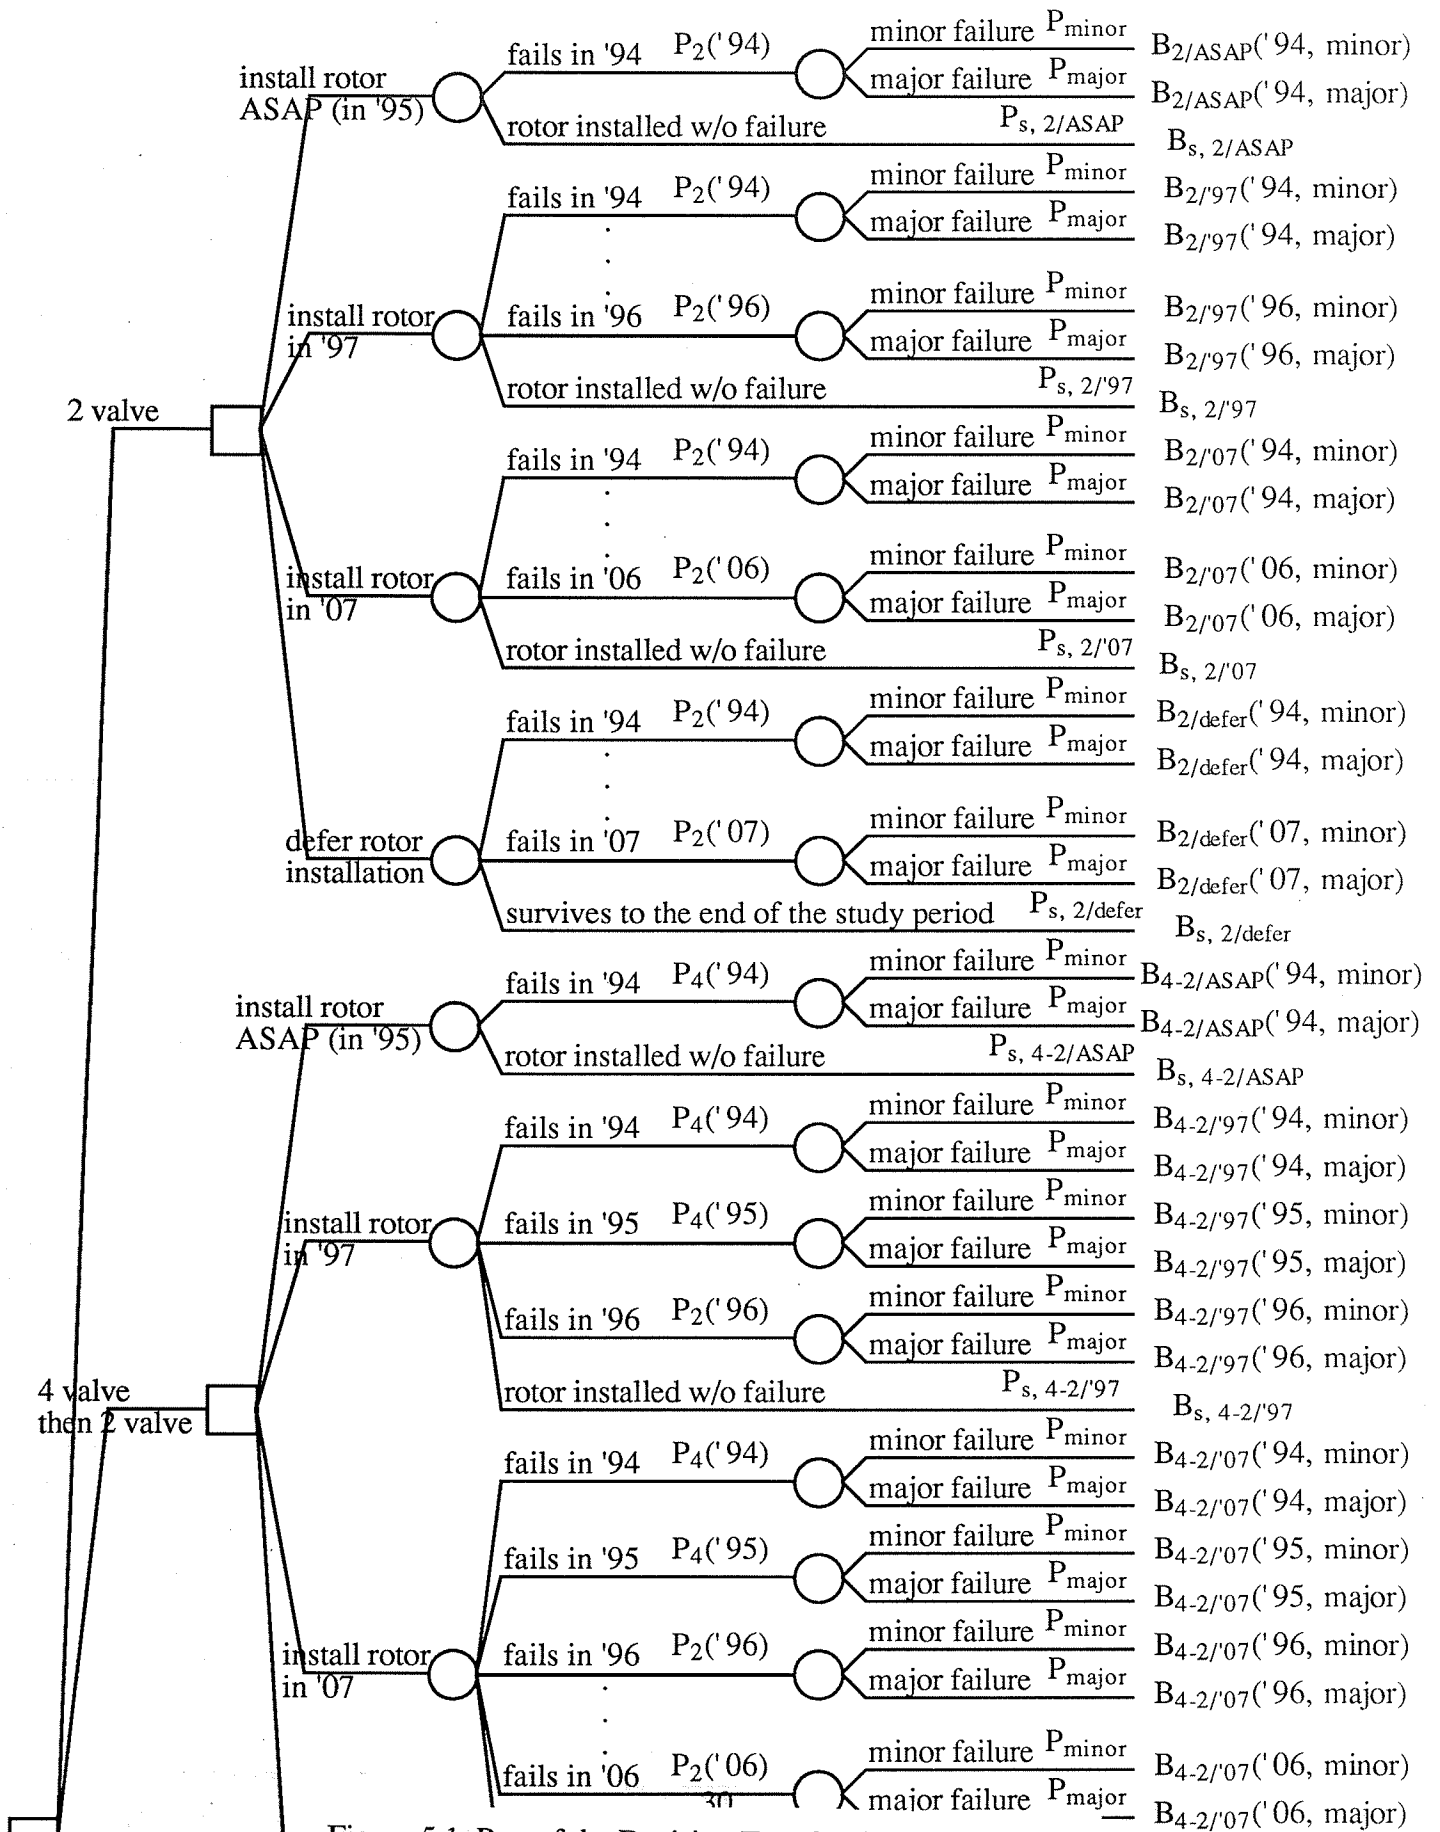


Figure 5.1. Part of the Decision Tree for the First Test Case

The conditional probability of failure mode is defined as

$$P_{\text{minor}} = P\{\text{failure is minor} \mid \text{failure occurs in the } i\text{-th year}\};$$

$$P_{\text{major}} = P\{\text{failure is major} \mid \text{failure occurs in the } i\text{-th year}\}.$$

By assumption 5, P_{minor} and P_{major} are constant over time, as can be seen on the second level of chance outcomes. Also notice that the sum of P_{minor} and P_{major} should be one. That is,

$$P_{\text{minor}} + P_{\text{major}} = 1 \tag{5.4}$$

As mentioned before, in the VA, the life of a failing component follows a Weibull distribution. Therefore, the probability that the failing component fails unexpectedly in 1994 under 2-valve operation, $P_2('94)$, is calculated as

$$\begin{aligned} P_2('94) &= P\{1/1/1994 \leq T < 12/31/1994\} \\ &= F(1994 - 1993) - F(1993 - 1993) \\ &= (1 - e^{-\lambda 1^\alpha}) - (1 - e^{-\lambda 0^\alpha}) \\ &= 1 - e^{-\lambda} \end{aligned}$$

Notice that we choose the beginning of 1994 as time zero.

By the probability law, the sum of probabilities of occurrence for all scenarios is 1. Therefore, the probability of survival to scheduled rotor installation for each alternative, taking $P_{s, 2/'97}$ as an example, is calculated by

$$P_{s, 2/'97} = 1 - P_2('94) - P_2('95) - P_2('96)$$

Notice that the subscript s in $P_{s, \text{ alternative}}$ and $B_{s, \text{ alternative}}$ stands for survival. Obviously, the component's life distribution need not be confined to the Weibull form. It can be normally distributed or any form that the user deems appropriate. Since the VA is modularly coded, we can easily change the component's life distribution within the VA.

As mentioned before, the new rotor is available in 1996. Therefore, the scenarios for the *4-valve then 2-valve* alternative have a more complicated life distribution than those of other alternatives. Before 1996, the scenarios follow the life distribution of 4-valve operation, $P_4(\bullet)$, while the life distribution changes to that of 2-valve operation, $P_2(\bullet)$, in and after 1996 because the alternative tells us to switch to 2-valve when the new rotor is available in 1996.

Once we have the probability of occurrence and benefit associated with each scenario, we can easily calculate the expected benefit for each alternative by taking the sum of benefits over all scenarios weighted by the probabilities of occurrence. For example, the expected benefit of performing *2-valve and install rotor in '97* alternative is

$$\begin{aligned}
 & E[\text{benefit of } 2\text{-valve and } \textit{install rotor in '97}] \\
 &= P_2('94)[P_{\text{minor}}B_{2/'97}('94, \text{ minor}) + P_{\text{major}}B_{2/'97}('94, \text{ major})] \\
 &+ P_2('95)[P_{\text{minor}}B_{2/'97}('95, \text{ minor}) + P_{\text{major}}B_{2/'97}('95, \text{ major})] \\
 &+ P_2('96)[P_{\text{minor}}B_{2/'97}('96, \text{ minor}) + P_{\text{major}}B_{2/'97}('96, \text{ major})] \\
 &+ P_{s, 2/'97}B_{s, 2/'97}
 \end{aligned}$$

Upon calculating the expected benefits for all alternatives, the VA returns the values to the user and let the user make the final decision.

Break-even Probability

Usually the decision-maker is not sure about the failing component's mean time to failure μ . What he has is a "gut feeling" about the possible range of μ , based on past experiences. To overcome the uncertainty about the failing component's life distribution,

we can analyze the sensitivity of the value over the possible range of μ . This approach is embodied in the break-even analysis.

The idea of break-even analysis is to compare the expected benefit for a given alternative with a deterministic baseline alternative and to see what mean time to failure μ makes these two alternatives have identical benefits. This specific μ is called break-even mean time to failure for that alternative, denoted as μ_{BE} , alternative. Now we can use μ_{BE} , alternative as a reference against our belief about the possible range of μ for that alternative. If the range of μ is greater than μ_{BE} , alternative, then we expect this alternative is more attractive than the baseline because the former is less risky in that it has larger mean time to failure μ .

Since the Weibull distribution is determined by two parameters, α and λ , we cannot uniquely define the component's life distribution by assigning μ only. Standard deviation σ should be considered as well. To simplify the break-even analysis process, we set the coefficient of variation r to be a constant, which makes σ linearly proportional to μ . Once μ is specified, σ is determined accordingly.

The Results

Figure 5.2 shows the expected benefits of all alternatives as a function of component's mean time to failure μ , whose time unit is year. σ is assumed to be 0.9μ . Notice that the expected benefit increases with the mean time to failure.

CC7 Tri-Pin Blade Problem




Figure 5.2 The Expected Benefits of Alternatives in Test Case One.

The highest expected benefit alternative is to run at 5-valve and replace the rotor at the first outage, scheduled in 1997, if failure is expected in less than 7 years. The highest expected benefit alternative is to run at 2-valve and replace the rotor at the first outage if failure is expected between 7 and 16 years. Otherwise, run at 2-valve and defer the rotor installation is the highest expected benefit alternative. Based on Figure 5.2, Table 5.1 shows the dominating alternatives for different ranges of μ .

Table 5.1 Ranges of Mean Time to Failure μ Favoring Each Alternative.

repair timing operating mode	install rotor ASAP	install rotor in 1997	install rotor in 2007	defer rotor installation
2-valve	-	[7, 16]	-	> 16
4-valve then 2-valve	-	-	-	-
4-valve	-	-	-	-
5-valve	-	[0, 7]	-	baseline

* - indicates no interval of μ favors this alternative

For example, if we believe the failing rotor has an expected life of 10 to 14 years, then from Table 5.1 we know that we should run 2-valve mode and prepare to install the new rotor in the 1997 scheduled outage. Notice that if we expect the rotor to fail very soon, then we should choose the 5-valve mode, which is assumed to be risk-free, as the best alternative. On the other hand, if we believe the rotor is not likely to fail in the foreseeable future, then we should choose riskier alternative, run 2-valve and defer the rotor installation.

5.2. Test Case 3, FD Fan Bearing Maintenance

Problem Description

See Appendix A.

Formulation

Study Period: 1 week.

- Available Choices:
1. Change oil off-line ASAP;
 2. Change oil on-line ASAP;
 3. Defer to the end of the study period.

- Chances:
1. Unexpected bearing failure;
 2. Oil spill;
 3. Survival.

Assumptions:

1. Oil spill and unexpected bearing failure are independent.
2. Oil spill can only happen at the beginning of on-line oil change, which is assumed to occur in the first period, with a probability of P_{spill}
3. If the oil spill occurs, we have to shut down the system and change oil completely. The time to clean up the spilled oil is long enough such that oil can be changed off-line simultaneously.
4. After oil is changed, no failure occurs within the study period.
5. The costs for changing oil off-line and changing oil on-line are the same.
6. The time to change oil on-line is short enough such that no bearing failure can possibly occur during the work.

VA Perspective

The FD fan bearing problem can also be fitted into the VA's framework.

The appropriate time unit for this test case is hour. With a study period of one week we have 168 periods.

Baseline

The baseline is defined as the one that oil change is implemented right at the end of the study period. Notice that no downtime cost is considered in the baseline. Only the cost

of off-line oil change, which is equal to the cost of on-line oil change (see assumption 5), is spent.

Obviously, the benefit for the baseline is zero. That is,

$$B_{\text{baseline}} = 0 \tag{5.1}$$

The table below gives the data needed for calculating the expected benefits for the three alternatives.

	Oil Spill	Oil Change	Bearing Failure
Cost (\$K)	5	-	12
Repair Duration (hrs)	48	48	72

Notice the cost for oil change is not shown in the table because we do not need it in calculating the benefit. The reason is that the oil change cost is included in the baseline and oil change appears in all the scenarios. Therefore, in calculating the benefit against the baseline it is canceled out.

The decision tree for this test case is shown in Figure 5.3.

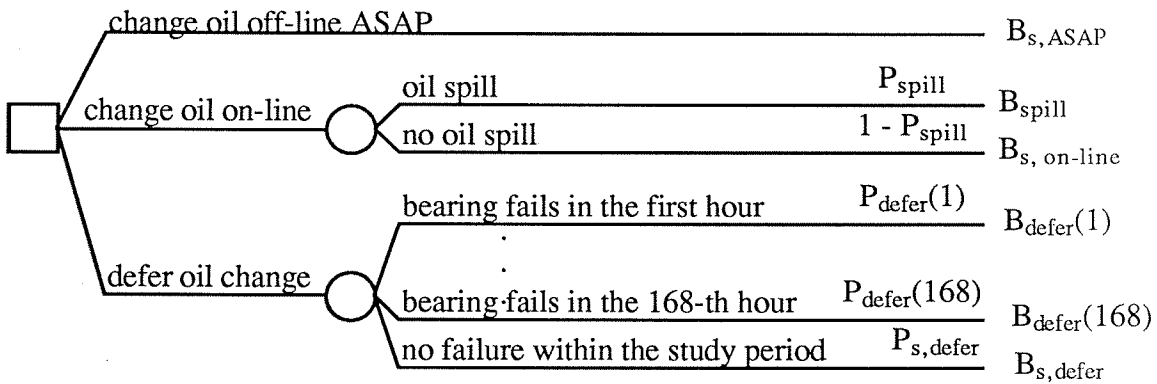


Figure 5.3 Decision Tree for the Third Test Case.

Notice for the ASAP alternative, there is only one scenario: system is shut down immediately and oil is changed off-line. Therefore,

$$B_{s, \text{ASAP}} = - C_{\text{downtime}} = - \sum_{i=1}^{48} D(i) \times \text{RPV}(i)$$

Not like the first test case, in which we have to decide the mode of operation and calculate the corresponding cost savings S , in this case no modes of operation are involved. Therefore, we do not have to consider S . The calculation for C_{downtime} is same as that for the first test case. We will not give numerical example here because the calculation involves 48 hours of downtime, which means 48 terms in C_{downtime} .

For the *change oil on-line* alternative, the calculation is as follows.

$$B_{\text{spill}} = - C_{\text{downtime}} - C_{\text{spill}} = - \sum_{i=1}^{48} D(i) \times \text{RPV}(i) - C_{\text{spill}}$$

By assumptions 2 and 6, if there is no oil spill in the first hour, we assume that we can finish changing the oil on-line without having an unexpected bearing failure. Therefore, if we successfully change oil on-line, the only cost we have to pay is the cost of oil change. In comparing with the baseline, we conclude that

$$B_{s, \text{on-line}} = 0$$

For the *defer oil change* alternative, if the system fails in the k -th hour, the benefit is:

$$B_{\text{defer}}(k) = - \sum_{i=k}^{k+72-1} D(i) \times \text{RPV}(i) - C_{\text{bearing failure}} \quad 1 \leq k \leq 168$$

If there is no unexpected bearing failure, then it is equivalent to the baseline. Therefore, the benefit would be:

$$B_{s, \text{defer}} = 0$$

The Results

Figure 5.4 shows the best alternative and associated expected benefit, relative to the *ASAP/Off-line* alternative, as a function of the probability of bearing failure and probability of oil spill. For example, when the probability of oil spill is high and the probability of bearing failure is low, we should defer the off-line oil change. When both the probability of oil spill and probability of bearing failure are high, the riskiest situation, we should immediately shut the system down and change the oil off-line. Otherwise, we can just change oil on-line.

Table 5.2 shows the optimal decision in terms of the bearing failure and oil spill probabilities. Also shown are relative benefit of each optimal decision to that of the second best.

PPP7 FD Fan Bearing Problem




Figure 5.4 The Relative Expected Benefit of Best Alternative in Test Case Three.

Table 5.2 Decision Table in Terms of $P_{\text{bearing failure}}$ and P_{spill} .

benefit relative to the second best (K\$)	$P_{\text{spill}} = 0.1$	$P_{\text{spill}} = 0.2$	$P_{\text{spill}} = 0.3$	$P_{\text{spill}} = 0.4$	$P_{\text{spill}} = 0.5$	$P_{\text{spill}} = 0.6$	$P_{\text{spill}} = 0.7$	$P_{\text{spill}} = 0.8$	$P_{\text{spill}} = 0.9$	$P_{\text{spill}} = 1.0$
$P_{\text{bearing failure}} = 0.1$	0.1	11.7	23.5	35.3	47.1	58.9	70.7	82.5	94.3	101.1
$P_{\text{bearing failure}} = 0.2$	12.4	0.6	11.2	23	34.8	46.6	58.4	70.2	82	88.8
$P_{\text{bearing failure}} = 0.3$	25.1	13.3	1.5	10.3	22.1	33.9	45.7	57.5	69.3	76.1
$P_{\text{bearing failure}} = 0.4$	38.4	26.5	14.74	2.9	8.9	20.7	32.5	44.3	56.1	62.9
$P_{\text{bearing failure}} = 0.5$	52.3	40.5	28.7	16.9	5.1	6.8	18.6	30.4	42.2	49
$P_{\text{bearing failure}} = 0.6$	67.1	55.3	43.5	31.7	19.9	8.1	3.7	15.5	27.3	34.1
$P_{\text{bearing failure}} = 0.7$	83.1	71.3	59.5	47.7	35.9	24.1	12.3	0.5	11.3	18.1
$P_{\text{bearing failure}} = 0.8$	101	89.2	77.4	65.6	53.8	42	30.2	18.4	6.6	0.2
$P_{\text{bearing failure}} = 0.9$	101.2	89.4	77.6	65.8	54	42.2	30.4	18.6	6.8	5
$P_{\text{bearing failure}} = 1.0$	101.2	89.4	77.6	65.8	54	42.2	30.4	18.6	6.8	5



defer



on-line



ASAP

The patterns indicate three preferred repair alternatives. Numbers in boxes show relative benefit of preferred alternative over second-best alternative.

5.3. Current Problems in the Stand-alone VA

As evidenced by the above two test cases provided by PG&E, the stand-alone VA can handle some power plant maintenance problems very well. Given the power demand prediction, monetary information, and component failure modes, the VA can return the user with maintenance strategies regarding timing, probability of survival, and the break-even analysis.

However, the calculation is based on a simplified assumption: the maintenance recovers the failing component to a brand new condition such that it will never fail again within the study period by the same mechanism. That is, the maintenance is assumed to be perfect. This assumption is plausible only if the study period is short compared with the

expected life of the repaired component. If it is not, the VA will return an over-optimistic result because it will underestimate the effects of recurrent failure³.

Take the Contra Costa Turbine Blade Problem (Test Case One) and Pittsburgh Heater Tube Leaks (Test Case Two, which will be addressed later) as examples. In the former problem, once a new rotor is installed we can safely assume the repair is perfect because the study period is in the order of ten years while the new rotor's expected life is about fifty years. However, this is not the case for the tube leak problem. A tube's expected life is about half a year, but the study period is much longer (seven years in this case). Under this circumstance, assuming that after the repair the component can never fail again is a very serious mistake. There is another situation in which the perfect repair assumption does not work. We might want to partially repair the component (either it is failing or has already failed) because we want to put the system back on line immediately, and the resource to perform a complete repair is not available at the moment or the partial repair is relatively short or inexpensive. This is the place where the current VA's analysis capability ends and a new research begins.

3. Because the VA considers the failure only once, the subsequent failures due to imperfect maintenance are truncated. This is the cause of underestimating failure effects.

SECTION 6

TEST CASE TWO (RECURRENT FAILURE PROBLEM)

6.1. Test Case 2, Secondary Superheater Outlet Header Stub Tube Welds

Level of Repair

Problem Description

See Appendix A.

Formulation

Study Period: 1994 - 2001 (7 years).

Available Choices:

Repair strategies

1. Do nothing (*Do Nothing*);
2. Grind out and replace all stub welds (*All Welds*);
3. Grind out and replace 60 stub welds (*60 Welds*);
4. Replace all stub tubes with standard stub tubes (*All Standard*);
5. Replace 60 stub tubes with standard stub tubes (*60 Standard*);
6. Replace all stub tubes with new B&W sleeves and insert (*All B&W*);
7. Replace 60 stub tubes with new B&W sleeves and insert (*60 B&W*);
8. Ultrasonic test welds and replace only cracked welds (*UT/Reweld*).

Timing strategies

1. Repair during the scheduled outage in 1994;
2. Repair during the scheduled outage in 1996.

- Chances: 1. Failure;
2. Survival.

Assumptions:

1. The cost to fix an unexpected tube leak⁴ after the repair is the same for alternatives 1, 2, 3, 5, 7, and 8, which is \$50K per unexpected tube leak.
2. There is not too much information provided for the UT/Reweld Cracks alternative. It is assumed this repair can reduce the failure rate and can be finished within the scheduled outage.

The estimated preventive maintenance costs are shown in Table 6.1, along with other parameters.

The third column represents the belief about the boiler tubes' aggregated initial failure rate after the repair. That is, the belief about the range of the initial failure rate after repair represents the effectiveness of each repair alternative. Both the alternatives *All Standard* and *All B&W* are assumed to be risk-free in operation for 7 years, either one of them can be used as a baseline case. We assume that they take the same extra downtime (2 weeks for each, assuming the scheduled outage lasts only 6 weeks)⁵. Therefore, the only difference between them is the preventive maintenance cost. Since the preventive maintenance cost for *All Standard* is less than that of *All B&W*, it is chosen as the baseline case.

4. There are about 280 tubes in the boiler.

5. The extra downtime cost is \$70.9K/week.

Table 6.1 Case 2 Input Parameters.

Alternative	Preventive Maintenance Cost (\$K)	Belief about the Initial Failure Rate γ after Maintenance ⁶ (faults/year)	Scheduled Repair Duration (weeks)
1. Do Nothing	0	> 2	0
2. All Welds	164	0 for 4 years, then $1.8 < \gamma < 2$	6
3. 60 Welds	50	$1.3 < \gamma < 1.5$	< 6
4. All Standard	211	0 for 7 years	8
5. 60 Standard	102	$1.2 < \gamma < 1.5$	< 6
6. All B&W	244	0 for 7 years	8
7. 60 B&W	117	$0.9 < \gamma < 1.3$	< 6
8. UT/Reweld Cracks	60	$1.4 < \gamma < 1.5$	6

- Each alternative is performed right at the beginning of the study period. No other preventive maintenance is followed up. Each unexpected fault is fixed upon failure, which costs \$50K.
- The duration of a scheduled outage is 6 weeks. Any repair that takes more than 6 weeks to finish requires extra downtime. Extra downtime cost is \$70.9K/week.

6.2. Simplified Approaches to Analyze Recurrent Failure Problems

Constant Failure Rate Over Time

As mentioned before, the perfect repair assumption, used in the current VA, cannot directly handle a recurrent failure problem like the Heater Tube Leaks. Therefore, we have to find a new way to solve the problem. The simplest solution is to assume the failure rate is constant over time and calculate the expected total cost.⁷ The calculation breaks the cost

6. The numbers are used for illustrative purpose, which might not be true.

into two parts: cost for planned maintenance activity and cost to repair unexpected failures after the maintenance.

The total cost for *All Standard*, which is now the baseline case, is

$$\text{preventive maintenance cost} + \text{extra downtime cost} = 211 + 70.9 \times 2 = \$353\text{K}$$

If we assume the initial failure rate after repair γ is constant over time, we can easily calculate the expected total cost. Take one of the alternatives, Replace 60 tubes with standard tubes (*60 Standard*), as an example. If we perform this maintenance activity during the scheduled outage in 1994, it costs us \$102K. Since the outage lasts 6 weeks and the maintenance activity takes time less than that, we don't have cost caused by extra downtime. Suppose after the maintenance the initial failure rate⁸ of tube leaks is reduced to 1.4 faults per year and each faults costs \$50 K. Assume the annual discount rate is 10% and the study period is 7 years, then the expected total cost of the repair alternative over the study period would be

$$\begin{aligned} & 102 + 0 \\ & + 1.4 \times \left[\frac{50}{(1+10/100)^1} + \frac{50}{(1+10/100)^2} + \frac{50}{(1+10/100)^3} + \frac{50}{(1+10/100)^4} + \frac{50}{(1+10/100)^5} \right. \\ & \quad \left. + \frac{50}{(1+10/100)^6} + \frac{50}{(1+10/100)^7} \right] \\ & = \$443\text{K} \end{aligned}$$

Symbolically,

-
7. This is the method used to solve the Tube Leak Problem for the presentation at the Pittsburg Power Plant on Dec. 16, 1993.
 8. Since the failure rate is assumed to be constant over time, the failure rate at any time is equal to the initial failure rate after repair γ .

$$\begin{aligned} \text{expected total cost} &= \text{preventive maintenance cost} + \text{extra downtime cost} \\ &+ \text{failure rate} \times \sum_{\text{year from now} = 1}^{\text{end of study period}} \frac{\text{cost per fault}}{(1 + \text{discount rate})^{\text{year from now}}} \end{aligned} \quad (6.1)$$

We can use equation (6.1) to calculate the expected total cost for each repair alternative. The results are shown in Figure 6.1.

Using the above approach we can calculate the expected total cost of each alternative as a function of our belief about its effectiveness, i.e., the initial failure rate γ of the boiler tubes upon the implementation of the repair alternative. For example, if we implement the seventh alternative, *60 B&W*, and subsequently the tubes are believed to have an initial failure rate of 0.97 faults/year, the expected total cost for this alternative is \$353K, which is equal to the cost of implementing the baseline alternative, *All Standard*. If we believe that the initial failure rate γ for *60 B&W* is less than 0.97 faults/year then it should be superior to the baseline alternative, because the smaller the failure rate is, the lower is the expected total cost. Table 6.2 shows the results.




Figure 6.1 Expected Total Cost (Net Present Value (NPV) with 10% annual discount rate) over 7 year study period, assuming a constant failure rate over time.

Cost is shown for each alternative for the given range of belief about the initial failure rate after repair γ (see second column of Table 6.2), which is represented by cost range enclosed by dash symbols. The cost range of *Do Nothing* is unbounded in the upper limit because only lower limit of the initial failure rate after repair is specified. The failure rate is assumed to be constant over time. For this case, the *All Welds* and *60 B&W* options can have costs less than the baseline alternative *All Standard*, so these three alternatives are shown in Table 6.2 as candidate repair alternatives.

The Candidacy column indicates whether the alternative is ever the highest benefit alternative. The third column in Table 6.2, Expected total cost (\$K) Given the Belief about γ , shows the range of possible expected total cost based on the belief about the failure rate γ given each alternative, which is shown in the second column. For example, the belief about

the failure rate of tube leaks upon performing *60 Welds* lies between 1.3 and 1.5 faults/year. The corresponding expected total cost is thus in the range between \$366K and \$415K, using equation (6.1).

Table 6. 2 Case 2 with Constant Failure Rate Assumption.

Alternative	Belief about the Initial Failure Rate γ after Maintenance (faults/year)	Expected Total Cost (\$K) Given the Belief about γ	Candidacy
1. Do Nothing	> 2	> 487	No
2. All Welds	0 for 4 years, then $1.8 < \gamma < 2$	[317, 334]	Yes
3. 60 Welds	$1.3 < \gamma < 1.5$	[366, 415]	No
4. All Standard	0 for 7 years	353	Yes
5. 60 Standard	$1.2 < \gamma < 1.5$	[394, 467]	No
6. All B&W	0 for 7 years	386	No
7. 60 B&W	$0.9 < \gamma < 1.3$	[336, 433]	Yes
8. UT/Reweld Cracks	$1.4 < \gamma < 1.5$	[401, 425]	No

Assumptions:

1. Each alternative is performed right at the beginning of the study period. No preventive maintenance is followed up.
2. The study period is 7 years.
3. The annual discount rate is 10%.
4. The calculation is based on equation (6.1).

Comparing the above expected total costs of individual alternatives with that of baseline alternative *All Standard*, we conclude that both *All Welds* and *60 B&W* are possible candidates because their expected total costs are less than or equal to that of *All Standard*.

If we look at the formulation more carefully, there are some drawbacks with this approach. First of all, in the calculation the failure rate is constant over time. In reality, very often it increases with time. Second, the approach assumes breakdown maintenance because the repair is only performed upon unexpected failure. The cost of breakdown repair is higher than that of replacing a failing one in advance, so breakdown maintenance is optimal when the frequency of failure is low and the cost of unplanned shutdown is not high. Therefore, assuming there are no follow-up preventive maintenance activities within the study period is not appropriate except when the study period is short in comparison with the component's expected life.

Increasing Failure Rate Over Time (IFR)

Now we discuss how the decision would be different if the constant failure rate over time assumption is relaxed.

Suppose in the above example the failure rate increases with time, using the same data we would obtain different conclusion. Assume the failure rate increases 0.5 faults for each year. Without a second maintenance activity the expected total cost for 60 Standard is

$$\begin{aligned}
 &72 + 0 + 30 \\
 &+ \frac{50 \times 1.4}{(1+10/100)^1} + \frac{50 \times (1.4+0.5)}{(1+10/100)^2} + \frac{50 \times (1.4+1)}{(1+10/100)^3} + \frac{50 \times (1.4+1.5)}{(1+10/100)^4} + \frac{50 \times (1.4+2)}{(1+10/100)^5} \\
 &+ \frac{50 \times (1.4+2.5)}{(1+10/100)^6} + \frac{50 \times (1.4+3)}{(1+10/100)^7} \\
 &= \$762K
 \end{aligned}$$

Symbolically, the calculation can be expressed as:

$$\begin{aligned}
&\text{expected total cost} = \text{preventive maintenance cost} + \text{extra downtime cost} \\
&+ \sum_{\text{year from now} = 1}^{\text{end of study period}} \frac{\text{cost per fault} \times (\text{initial failure rate} + \text{increment per year} \times \text{year from now})}{(1 + \text{discount rate})^{\text{year from now}}}
\end{aligned}
\tag{6.2}$$

Using equation (6.2) we can calculate the expected total costs for individual alternatives, which are shown in Figure 6.2.

As we can see from Figure 6.2., no alternative has an expected total cost lower than that of the baseline alternative *All Standard*. Using this approach, we obtain the expected total cost for each alternative, which is shown in Table 6.3. The shaded cells are those whose candidacy or values have changed in comparison with Table 6.2, because of the assumption about the failure rate (constant failure rate vs. increasing failure rate).

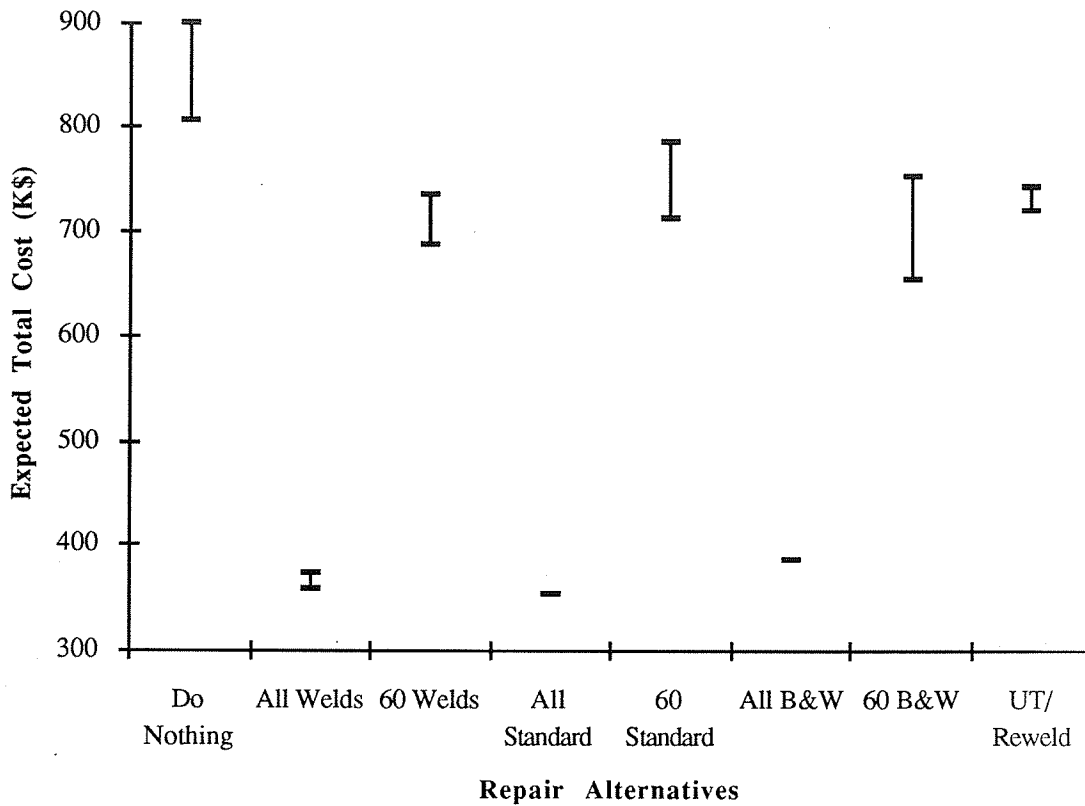


Figure 6.2 Expected Total Cost (Net Present Value (NPV) with 10% annual discount rate) over 7 year study period, assuming a linearly increasing failure rate over time (IFR).

Cost is shown for each alternative for a given range of belief about the initial failure rate after repair γ (see second column of Table 6.3), which is represented by cost range enclosed by dash symbols. The cost range of *Do Nothing* is unbounded in the upper limit because only lower limit of the initial failure rate after repair is specified. For this case, no alternative has cost less than the baseline alternative *All Standard*, so only this baseline alternative is shown in Table 6.3 as candidate repair alternative.

Obviously, the expected total cost for each alternative with IFR is much higher than the counterpart with constant failure rate. Using the belief about the initial failure rate after repair γ in the second column, we conclude we should choose *All Standard*. This choice contrasts with the decisions made with constant failure rate assumption, in which *All Welds* and *60 B&W* are possible candidates.

Table 6.3 Case 2 with IFR Assumption.

Alternative	Belief about the Initial Failure Rate γ after Maintenance (faults/year)	Expected Total Cost (\$K) Given the Belief about γ	Candidacy
1. Do Nothing	> 2	> 806	No
2. All Welds	0 for 4 years, then $1.8 < \gamma < 2$	[357, 374]	No
3. 60 Welds	$1.3 < \gamma < 1.5$	[686, 734]	No
4. All Standard	0 for 7 years	353	Yes
5. 60 Standard	$1.2 < \gamma < 1.5$	[713, 786]	No
6. All B&W	0 for 7 years	386	No
7. 60 B&W	$0.9 < \gamma < 1.3$	[655, 753]	No
8. UT/Reweld Cracks	$1.4 < \gamma < 1.5$	[720, 744]	No

Assumptions:

1. Each alternative is performed right at the beginning of the study period. No preventive maintenance is followed up.
2. The study period is 7 years.
3. The annual discount rate is 10%.
4. The calculation is based on equation (6.2).
5. The failure rate increases at a rate of 0.5 fault per year.

We used a simplified calculation in equation (6.2) to analyze a problem with increasing failure rate over time. This approach is only appropriate when a minimal repair policy is employed. A minimal repair policy is one that any breakdown maintenance only recovers the system's functionality, but does not disturb the system's age (Barlow and Proschan, 1965: 96-98). The minimal repair policy is only plausible when a failing system is comprised of a large collection of identical components. This is because when one of the N components (N is very big) breaks unexpectedly, the breakdown maintenance only

repairs that specific component. It is either fixed or replaced and is assumed to be as good as new after repair. However, this zero-age component has very little effect on the system's age because it is just one of the N components. The other N-1 components are not fixed opportunistically. Equation (6.2) actually is an approximation of calculating the expected cost of breakdown maintenance by using a minimal repair policy (see Appendix B).

6.3. An Approach on Opportunistic Breakdown Maintenance

Suppose the breakdown maintenance is not planned only for failed components. It is designed such that when an unexpected failure occurs, some components other than those failed are fixed opportunistically. Obviously, the opportunistic breakdown maintenance can affect the failing system's age, and thus its failure rate. Therefore, the minimal repair does not apply in this case. However, the system's age now depends on how extensive the opportunistic repair is. The following calculation assumes that the opportunistic breakdown maintenance recovers the failing system back to the condition as that after the initial preventive maintenance at time zero. Figure 6.3 shows a possible realization over time with implementing *60 Welds* at time zero and opportunistic breakdown maintenance in the future. Note that every unexpected tube leak is fixed by an opportunistic breakdown maintenance, which restores the boiler's initial failure rate back to 1.4 faults/year. This figure is based on the repair effectiveness of *60 Welds implemented* at time zero.

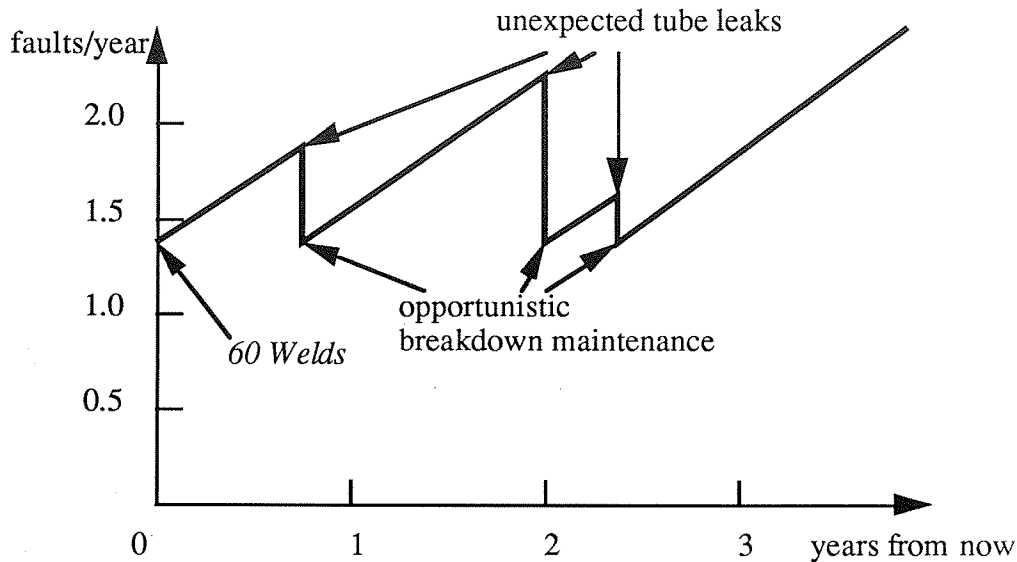



Figure 6.3 A Possible Realization of Boiler's Failure Rate over Time.

60 Welds is implemented at time zero, with initial failure rate after repair $\gamma = 1.4$ faults/year.

Consider the above example again. Suppose *60 Welds* is implemented in time zero and it reduces the system's initial failure rate g down to 1.4 faults/year. Let the probability of tube leaks in the first year be P . If we have tube leaks in the first year, then upon the opportunistic breakdown maintenance, the recovered system will have a failure rate of 1.4 faults/year, as stated in the previous paragraph, instead of 1.9 faults/year in the second year. This can be represented in the following tree.



That is, we have 1.4 faults/year in the second year with probability P and 1.9 faults/year with probability $1-P$. Thus the expected failure rate in the second year should be less than 1.9 faults/year. The calculation in Section 6.2 assumes that the failure rate in the second year is 1.9 faults/year with probability one, which is based on the minimal repair

assumption. Obviously, it overestimates the failure rate and consequently has higher expected total cost. This is because it does not take the reduction in the system's failure rate into account if the opportunistic breakdown maintenance is implemented. Therefore, the previous calculation cannot handle the increasing failure rate problem correctly when an opportunistic breakdown maintenance is assumed. A more rigorous approach is needed to analyze the problem with opportunistic breakdown maintenance.

Assuming that failure rate increases linearly by 0.5 fault/year, the linear IFR can be represented by the following function.

$$\gamma(t) = \gamma_0 + 0.5t \quad (6.3)$$

where γ_0 : initial failure rate;
 t : years from now.

By the reliability theory (Barlow and Proschan, 1981: 53), the life distribution is

$$F(t) = 1 - \exp\left(-\int_0^t \gamma(u)du\right) = 1 - \exp(-\gamma_0 t - 0.25t^2) \quad (6.4)$$

A rigorous way to calculate the expected total cost for each alternative is

$$E[\text{total cost}] = C_1 + C_d + C(m) \quad (6.5)$$

where C_1 : preventive maintenance cost;
 C_d : extra downtime cost;
 m : duration of the study period;
 $C(m)$: expected cost caused by unexpected failures within a duration of m .

$C(m)$ can be calculated from the following recursive equation.

$$C(m) = \int_0^m \beta^t [C_f + C_b + C(m-t)] dF(t) \quad (6.6)$$

where β : discount factor;

C_f : cost of fixing the unexpected fault⁹;

C_b : opportunistic repair cost¹⁰.

For example, the expected total cost for *60 Standard*, with initial failure rate $\gamma = 1.4$ faults/year, is

$$\begin{aligned} & E[\text{total cost for } 60 \text{ Standard}] \\ &= 102 + 0 + C(7) \\ &= 102 + \int_0^7 \frac{1}{(1+0.1)^t} [50 + 7 + C(7-t)] d[1 - \exp(-1.4t - 0.25t^2)] \\ &= 102 + \int_0^7 \frac{1}{1.1^t} [57 + C(7-t)] (1.4 + 0.5t) \exp(-1.4t - 0.25t^2) dt \\ &= 102 + 456 \\ &= \$558K \end{aligned}$$

Compare the above result with that from the simplified calculation with minimal breakdown maintenance, \$762K. We learn that in this case a opportunistic breakdown maintenance is much cheaper than a minimal breakdown maintenance.

-
9. Repairing the unexpected fault is assumed to be \$50K. See the note in Table 6.1.
 10. Opportunistic repair cost is assumed to be \$7K in the following calculation. The figure is much less than the cost of fixing the unexpected fault because the fixed cost, e.g., equipment cost, is attributed to the latter.

Figure 6.4 shows the calculation results using the rigorous approach with opportunistic breakdown maintenance. The expected total cost for each alternative obtained from this approach is shown in Table 6.4. The shaded cells are those whose candidacy or values have been changed in comparison with Table 6.2, because of the assumption about the failure rate and breakdown maintenance.




Figure 6.4 Expected Total Cost (Net Present Value (NPV) with 10% annual discount rate) over 7 year study period, assuming a linearly increasing failure rate (IFR) and considering the reduction in failure rate following opportunistic breakdown maintenance.

Cost is shown for each alternative for a given range of belief about the initial failure rate after repair γ (see second column of Table 6.4), which is represented by cost range enclosed by dash symbols. The cost range of *Do Nothing* is unbounded in the upper limit because only lower limit of the initial failure rate after repair is specified. For this case, no alternative has cost less than the baseline alternative *All Standard*, so only this baseline alternative is shown in Table 6.4 as candidate repair alternative.

Table 6.4 Case 2 with IFR Assumption, Considering Effect of Opportunistic Breakdown Maintenance on the Initial Failure Rate.

Alternative	Belief about the Initial Failure Rate γ after Maintenance (faults/year)	Expected Total Cost (\$K) Given the Belief about γ	Candidacy
1. Do Nothing	> 2	> 586	No
2. All Welds	0 for 4 years, then $1.8 < \gamma < 2$	[358, 375]	No
3. 60 Welds	$1.3 < \gamma < 1.5$	[483, 528]	No
4. All Standard	0 for 7 years	353	Yes
5. 60 Standard	$1.2 < \gamma < 1.5$	[512, 583]	No
6. All B&W	0 for 7 years	386	No
7. 60 B&W	$0.9 < \gamma < 1.3$	[459, 550]	No
8. UT/Reweld Cracks	$1.4 < \gamma < 1.5$	[516, 538]	No

Assumptions:

1. Each alternative is performed right at the beginning of the study period. No preventive maintenance is followed up.
2. The study period is 7 years.
3. The annual discount rate is 10%.
4. The calculation is based on equation (6.5).
5. The failure rate increases at a rate of 0.5 fault per year.

Compare Table 6.4 with Table 6.2, we conclude that the decision options are now fewer. *All Standard* is the only one choice. Also the expected total cost is different, which is higher for using a rigorous approach because the failure rate is increasing linearly over time. Notice that using the rigorous approach with opportunistic breakdown maintenance we have the same choice (*All Standard*) as the decision made by the simplified approach with IFR and minimal breakdown maintenance (see Candidacy column in Table 6.3). However, the expected total costs are different.

6.4. Two Step Preventive Maintenance

The rigorous approach with opportunistic breakdown maintenance suggests that the baseline alternative, *All Standard*, is the optimal among the 8 alternatives. However, all of the alternatives are one-shot preventive maintenance in that there is no follow-up preventive maintenance within the study period. Intuitively, after the one-shot preventive maintenance the failure rate drops immediately and then increases linearly over time. As we approach the end of the study period, the failure rate becomes very high, which causes high expected total cost. If we can schedule a follow-up preventive maintenance sometime within the study period we might be able to reduce the failure rate and in term lower the expected total cost. We will use the above example with opportunistic breakdown maintenance to show the feasibility of two-step preventive maintenance.

Consider there is a second preventive maintenance activity, *60 Welds*, following an initial one, *All Welds*. The initial failure rate upon the first preventive maintenance activity is 1.8 faults/year. If we perform the second preventive maintenance activity some time in the future, it will reduce the initial failure rate to 1.3 faults/year. They are the lower bounds of our belief about the effects of *All Welds* and *60 Welds* on the failing boiler tubes' initial failure rates after repair, respectively (see column 3 of Table 6.1). The question is: Is it beneficial to schedule it at all? If the answer is yes, then when should it be scheduled?

Suppose we schedule the second maintenance activity s years from now. The expected total cost is

$$E[\text{total cost}] = C_1 + C(s) + \beta^s [C_2 + C(m - s)] \quad (6.7)$$

where C_1 : cost for the first preventive maintenance activity, which is *All Welds*;

C_2 : cost for the second preventive maintenance activity, which is *All Welds*;

m : length of the study period;

β : discount factor;

$C(t)$: expected cost caused by unexpected failures within a length of t , see equation (6.6).

Plug in the numbers, we have

$$\begin{aligned}
 & E[\text{total cost of implementing } All\ Welds \text{ now and } 60\ B\&W \text{ } s \text{ years from now}] \\
 &= 164 + \int_4^s \frac{1}{1.1^t} [50 + 7 + C(s - t)] d[1 - \exp(-1.8t - 0.25t^2)] \\
 &\quad + \frac{1}{1.1^s} \times \left\{ 50 + \int_0^{m-s} \frac{1}{1.1^t} [50 + 7 + C(m - s - t)] d[1 - \exp(-1.3t - 0.25t^2)] \right\}
 \end{aligned}$$

Notice that we assume that the first four years are risk-free upon performing *All Welds*. Therefore, the lower bound of the first integral is 4 instead of 0. Solving the equation numerically, we can obtain a graph as Figure 6.5.

If we schedule the second preventive maintenance activity, *60 Welds*, 4 years from now, the expected total cost is \$350K. This is the minimum point along the curve, which happens to be right at the beginning of the fifth year. This is the time when the boiler tubes begin to be prone to failure. Also notice that the expected total cost of this two-step preventive maintenance is lower than that of the baseline case, *All Standard*, which costs \$353K.

It is interesting to note that if we implement one-shot preventive maintenance, both *All Welds* and *60 Welds* are inferior to the optimal preventive maintenance activity *All Standard*. However, when combining them together we enhance the effectiveness of preventive maintenance and have lower expected total cost. Although the second preventive maintenance activity causes additional preventive maintenance cost, the risk reduction, as embodied in lower expected cost due to unexpected failures, justifies the decision.




Figure 6.5 The Expected Total Cost, (Net Present Value (NPV) with 10% annual discount rate) over 7 year study period, of a two-step preventive maintenance with opportunistic breakdown maintenance as a function of the timing of the second preventive maintenance activity *60 Welds*.

Table 6.5 shows the results of this simple two-step preventive maintenance together with other eight one-shot preventive maintenance alternatives, all based on opportunistic breakdown maintenance. The shaded cells are those whose candidacy or values have been changed in comparison with Table 6.2, because of the assumption about the increasing failure rate, opportunistic breakdown maintenance, and use of multiple repair alternatives.

Table 6.5 Case 2 with IFR Assumption, Considering Effect of Opportunistic Breakdown Maintenance on the Initial Failure Rate after Repair.

Alternative	Belief about the Initial Failure Rate γ after Maintenance (faults/year)	Expected Total Cost (\$K) Given the Belief about γ	Candidacy
1. Do Nothing	> 2	> 586	No
2. All Welds	0 for 4 years, then $1.8 < \gamma < 2$	[358, 375]	No
3. 60 Welds	$1.3 < \gamma < 1.5$	[483, 528]	No
4. All Standard	0 for 7 years	353	Yes
5. 60 Standard	$1.2 < \gamma < 1.5$	[512, 583]	No
6. All B&W	0 for 7 years	386	No
7. 60 B&W	$0.9 < \gamma < 1.3$	[459, 550]	No
8. UT/Reweld Cracks	$1.4 < \gamma < 1.5$	[516, 538]	No
9. All Welds now and 60 Welds 4 years later	$\gamma = 1.8$ upon implementing All Welds, $\gamma = 1.3$ upon implementing 60 Welds	350	Yes

Assumptions:

1. Each alternative is performed right at the beginning of the study period. No preventive maintenance is followed up except the ninth alternative: *All Welds* now and *60 Welds* four years later.
2. The study period is 7 years.
3. The annual discount rate is 10%.
4. The calculation is based on equation (6.7).
5. The failure rate increases at a rate of 0.5 fault per year.

6.5. Sensitivity Analysis

Effects of the Study Period Duration

All the above computations calculate the cost accumulated over a study period of 7 years. Obviously, the length of study period can affect the choice of repair alternatives. Since the utility industry is being deregulated, the utility companies are shifting their focus

from power generation to electric grid service. As a result, they may sell some power plants to other companies. Before the sale, the plants still need maintenance. However, if the study period duration is short, *Do Nothing* might be the best alternative because any unexpected failure is unlikely to happen before the plant is sold.

Though the previous calculation does not take the salvage price into account, we can still see the effects of study period duration on the maintenance decision. Figure 6.6 shows the expected total cost for each repair alternative as a function of the study period duration, which varies from 1 to 7 years, with opportunistic breakdown maintenance. If the length of study period is one year, *Do Nothing* could be the best choice, even though it causes the initial failure rate to be higher than 2 faults/year. As the length of study period increases, the optimal alternative is shifted from the myopic one, *Do Nothing*, to *All Welds*. As the trend shows, eventually *All Standard*, which is risk-free for 7 years, becomes the best choice.

Effects of the Annual Interest Rate

Figure 6.7 shows how the annual interest rate affects the maintenance decision. As we might have expected, the higher the annual interest rate is, the greater discounting is the cost of unexpected failure, which is fixed by the breakdown maintenance accompanied by an opportunistic repair on other non-failed tubes. This implies that the higher is the annual interest rate, the less we are concerned about future failures. Therefore, when the annual interest rate is 5% or 10%, we choose *All Standard*, which is risk-free for 7 years, as the best alternative. When the annual interest rate is increased to 15%, however, the best alternative, is switched to *All Welds*, which is risk-free only in the first 4 years.

From the above discussion we see that the assumption about the failure rate and breakdown maintenance policies, minimal or opportunistic, have striking effects on the choice and outcomes, which cannot be gained from the simplified calculation with the constant failure rate assumption.




Figure 6.6 Expected Total Cost (Net Present Value (NPV) with 10% annual discount rate) as a Function of Study Period Duration for each Repair Alternative, assuming (like Figure 6.4) linear IFR and considering the reduction in initial failure rate after opportunistic breakdown maintenance.

The All Welds has fixed cost, i.e., the preventive maintenance cost, for 4 years, while failure rate is assumed to be 0. The cost interval between identical symbols represents the range of expected total cost for an alternative based on the study period duration and belief about the initial failure rate repair.




Figure 6.7 Expected Total Cost (Net Present Value) as a Function of Annual Interest Rate for each Repair Alternative.

The failure rate is linearly increasing over time and the opportunistic breakdown maintenance is used. The study period duration is 7 years.

The VA uses a decision-analytic approach to analyze the possible outcomes and calculate the expected total cost, weighted by probabilities, for each alternative. For the recurrent failure problem, it is still possible to obtain the solution by the same approach. However, the decision tree becomes so complicated that it eventually makes the computation intractable.

Based on the above observations we conclude that the simplified calculation with constant failure rate assumption oversimplifies the problem while the computational inefficiency prevents the current VA from solving the problem quickly. Recurrent failure and partial repair occur very often in plant maintenance. If we can analyze the problem accurately, it will help people to make better maintenance policies.

SECTION 7

DISCUSSIONS AND FUTURE WORK

7.1. Test Case Classification

We have studied three PG&E test cases. It is found that test case 1 and test case 3 can be solved by the present Value Analysis Module. However, the VA cannot solve the second test case because of the recurrent failure problem. We use two more complicated approaches to solve it. Each of which is based on the effectiveness of the breakdown maintenance (minimal or opportunistic). Also we have studied the effects of multiple preventive maintenance activities on the expected total cost. To summarize our contributions, we have:

- formalized the power plant maintenance problem by using a decision-analytic approach and provided a basis for making informed and rational maintenance decisions;
- validated the above approach with three test cases;
- shown the advantages and disadvantages of the simple engineering approach in solving recurrent failure problems;
- used two more formal approaches to solve recurrent failure problems based on different assumptions about the effectiveness of the breakdown maintenance;
- identified problem types based on the assumption about the repair effectiveness and the distinctions regarding values, choices, and chances (see the following discussions).

The three PG&E test cases provide rich insights into the maintenance problems frequently encountered in conventional power plants. There are choices regarding operations, maintenance; chances about component's life, caused by maintenance (either perfect or imperfect repairs); value judgment about the consequences, etc.

In PG&E's test case 1, we have operational mode (turbine valve mode) and maintenance (new rotor installation) timing choices, both of which affect the component's life. Notice that there is still only one maintenance alternative: new rotor installation. All scenarios result from the chances of component's life (when the failing component will fail). The chances can be caused by either operational mode or a component's own deterioration process. We also have two modes of failures: minor and major. Test case 1 concerns both maintenance and production costs. Again the perfect repair assumption is employed. Once the maintenance alternative is implemented, the component will never fail again within the study period.

In test case 2, we have to make a decision from a set of maintenance alternatives. However, once a maintenance alternative is chosen, it is implemented immediately. That is, the timing is not of primary concern. In addition, no mode of operation affects the failing tubes' life and performance. The chances result from the failing component's own deterioration process. In calculating the economic value, only the planned and unplanned maintenance costs are considered. No production cost is involved. The repair is not perfect, though. Therefore, we have a recurrent failure problem. The system in test case 2, the boiler, is composed of 280 identical tubes, which makes it different from other test cases.

In test case 3, the choice is only about the maintenance alternatives. In addition to chances about the failing bearing's life, we are not certain about the success of a maintenance activity. For example, on-line oil change may cause an unexpected oil spill. Fortunately, it is assumed after the oil is changed (either on-line or off-line) the same kind of problem will not happen again within the study period, i.e., perfect repair. The cost in this case only involves the maintenance. No production cost is considered.

The types of problems coming from the above test cases are summarized in Table 7.1.

Table 7.1 Classification of Test Cases.

Problems		Test Case 1	Test Case 2	Test Case 3
Choices	Operational Modes	Yes	No	No
	Maintenance Alternatives	No	Yes	Yes
	Maintenance Timing	Yes	No	No
Chances	Component's Life	Yes	Yes	Yes
	Failure Modes	Yes	No	No
	Success of Maintenance	No	No	Yes
Values	Production	Yes	No	Yes
	Maintenance	Yes	Yes	Yes
Repair Effectiveness		Perfect	Imperfect	Perfect
System Configuration		Single Component	Collection of Identical Components	Single Component

7.2. Importance of Distinction

There has been some confusion in using the VA user interface. Most of the time it is caused by failure in recognizing the distinctions among choices, chances, and values.

For example, in test case 1, the timing of turbine valve modes and timing of rotor availability are different choices. However, one alternative mixes these two choices: 4-valve operation until a new rotor is available, then 2-valve operation. The timing of turbine valve mode is contingent on the availability of a rotor. Therefore, in building the VA user interface, the data input and interpretations become ambiguous.

Another example is test case 3, where we have two different chances: component's life and success of maintenance. These chances arise from different physical mechanisms. The mechanism of the former involves bearing's deterioration process, while the latter involves the maintenance process of on-line oil change. The interpretations of these two probabilities of failure are significantly different. One concerns the deterioration process

over time (probability of bearing failure), another concerns the success of maintenance implementation (probability of no oil spill).

Obviously, the distinctions are important to framing and computation of the maintenance problems. To facilitate the decision process, we must articulate those distinctions. Otherwise, each time when a new maintenance problem emerges, we have to overhaul the VA again, which is not suitable for a general purpose maintenance decision support system.

7.3. Statement of the Maintenance Problem

For a recurrent failure (or an imperfect repair) problem, a maintenance policy should address the following questions.

1. When should I inspect the component's condition next and in the future?

The answer depends on the difficulty and cost of observations. If they are fairly easy and cheap, we don't have the inspection problem. In the present analysis, it is assumed that all the critical components are monitored continuously by electronic devices and at low cost. As a result, inspection time is not an issue under current study.

2. When should I perform the next and future maintenance activities?

Note that the time between consecutive maintenance activities needs not be a constant.

3. What are the costs and benefits of the next and future maintenance activities?

Once again, maintenance activities are not necessarily identical. Besides, the maintenance level will affect the component's life distribution.

4. Given a sequence of planned maintenance activities, how should the sequence of plans be updated when new information flows in? That is, what is a sequential maintenance strategy?

Ideally, we want to analyze the maintenance problem to optimize a pre-specified objective, e.g., expected benefit. The optimality is achieved by adjusting the maintenance frequencies, maintenance levels, and updating strategies.

The analysis of the maintenance policy for the Heater Tube Leaks Problem (Test Case 2) shows great potential for improving the quality of maintenance by studying the recurrent failure more rigorously and considering wider set of maintenance policies. Even though the two-step maintenance policy in the example is not sequential, it is still superior to any of the one-shot maintenance policies. Since the sequential maintenance policy takes into account more information, e.g., continued updating of the component's failure rate when we have incomplete information, we expect it to be more adaptive, which allows us to seek an even better outcome.

To sum up, a systematic way to analyze the recurrent failure problem can add value to the plant's operation by addressing

- Increased candidate decision option set.
- Effects of repair (both planned and unplanned) to lower the probability of failure.
- Bias caused by the simple calculation.
- Effects of information update on maintenance policies.

Another challenge is to automate the decision process. In analyzing the above three test cases, most of the efforts were spent on problem formulation: what are the decision maker's values? what are the available alternatives? what are the consequences given each alternative? Some of them need creative thinking, which is hard to automate. However, others are routine or concern functional relationships that are well understood. For example, it is feasible to predict the behavior of a failing component through automatic computation. If we can automate the latter through a decision support system, the decision process can be accelerated significantly.

A decision support system that handles the problem domain of maintenance management should be able to

- simplify the process of problem specification, e.g., values, choices, and chances.
- communicate the solution process with the user easily, e.g., using the decision tree or some graphs.
- compute the solution efficiently.

The first property involves capturing the knowledge of what chances are possible and what choices are available. The former is represented by a knowledge base of modeled system operation. While the latter is represented by a maintenance knowledge base. In addition to these two knowledge bases, we also need a *virtual* decision maker who is capable of identifying decision problems.

7.4. Future Work

To solve this frequently encountered maintenance problem, we need to find a way that is both efficient and practical. We propose two approaches to tackle this problem.

First, we want to build a decision support system, based on reliability and control theories, using established engineering knowledge about both the system behavior and maintenance activities. To start with, we analyze two simple systems: single component and two-component. Then we will analyze the three PG&E test cases through this approach. As mentioned in section 7.2, the distinctions of Table 7.1 must be clarified unambiguously in a decision support system.

The second approach is to study different maintenance policies by Monte Carlo simulation. The problems are the same as those in the first approach. The simulation should be much faster than exhaustive search of a big set of possibilities. The simulation results are used to check the plausibility of a given maintenance policy.

We have developed a simulator. We need to extend it to handle the recurrent failure problem. This is much easier than the effort to modify the VA, as mentioned before.

To justify the optimal maintenance policies derived from the above two approaches, we need to address the following questions.

1. Are the results from the above two approaches consistent?

The answer should be yes. We expect the simulation approach to implement both consistency and efficiency checks of any maintenance policies generated from the decision support system or any other sources. Therefore, this does not void the benefit from the first approach, which provides insight into how each maintenance policy affects the system behavior and guideline to construct more sophisticated one.

2. When and how much is the sequential maintenance better than the breakdown maintenance?

The preventive maintenance policy should be better when we have frequent failure (as compared with the length of the study period), increasing failure rate, or the cost of preventive maintenance is cheaper.

3. When and how much are the sequential maintenance better than "back-of-envelope calculation"?

With the example given above, we know the back-of-envelope calculation is correct only when the failure rate is constant over time.

4. How much is sequential maintenance better than time-invariant periodic maintenance policies?

Since a time-invariant periodic maintenance policy is constructed a priori, it does not take into account the information unfolding with time. The sequential maintenance policy takes advantage of new information to choose subsequent optimal maintenance activity. Intuitively, the results from the latter policy should be better.

Criteria for Selecting Test Cases

To validate our approaches, we need more test cases. Some criteria below are suggested to facilitate selecting good test cases.

1. For each recurrent failure problem, there should be large variation in costs and time among available maintenance alternatives.

2. There should be large variation in the effects of maintenance on the component's life distribution.

3. Each study period should have at least two scheduled outages, whose durations are not necessarily longer than those of maintenance activities.

4. Each test case should document how maintenance people solve the problem and the main concern and rationale in solving the problem.

5. The objectives of all test cases need not be the same. For example, in one case we may want to maximize the expected benefit while in another we may want to minimize expected system down time.

6. As evidenced in Table 7.1, the system under study is still very simple. Maintenance problems about multiple-component system are thus desirable.

APPENDIX A

PG&E Three Test Case Statements IRTMM Meeting Notes 9/21/93

Decision Method and Description

Case 1: CC Unit 7 Tri-Pin Blade and Two Valve Operation Decisions

October 10, 1993

Problem Statement

Opportunity for Savings

Based on historical and forecast information, the plant staff expects that CC Unit 7 will be operated between 30 and 60 MW for about 3000 hours per year for the next several years. Reducing the number of main steam governor valves which are used at this load allows operating the valves which are used at a more open position. The resultant reduction in throttling losses will provide a fuel cost savings of \$76,000 per year if 4 valves are used (instead of all 5) and a fuel cost savings of \$250,000 per year if 2 valves are used (instead of all 5). The "4 valve mode" is presently employed.

Associated Risk

The savings is not achieved without risk. The use of a reduced number of valves results in non uniform distribution of steam around the annulus of the HP turbine, which in turn results in cyclic loading of the turbine blades, and consequent fatigue cracking. This is considerably more severe for 2 valve mode than for 4 valve mode. The possibility of this mode of failure is completely eliminated by operating in 5 valve mode. Westinghouse says the probability of blade failure is significant over a several year period. The probability of a failure in 2 valve mode is thought to be about 2.2 times the probability of a failure in 4 valve mode.

Failures can occur at two levels of severity; Minor, where damage is confined to one or a few blades, and Major, where a blade, upon breaking free, causes secondary damage. If the new rotor (see next section) were available when a failure occurred, it would be installed. If not, several options are possible to respond to major and minor failures. Each has an associated duration and cost. If a failure occurs, the unit would revert to 5 valve operation after the repair, until a tri-pin rotor is installed.

Option to Eliminate Risk

A spare rotor is available, which can be fitted with blades attached with pins ("tri-pin blades") by Westinghouse at a cost of about \$1.2 million, and with 14 months to do the work. If this is done, the rotor at CC 7 will be interchangeable with the rotors at Morro Bay units 3 and 4. The rotor is not usable by unit 7 in its' present form. This rotor would not be susceptible to fatigue cracking. Also, its performance would degrade with age less than with the present design. The new rotor could be installed when it is ready, or during an outage scheduled in 1997, or during the next major outage in 2007.

Present agreements with Westinghouse provide that, unless a decision about fitting the spare rotor with tri-pin blades is made by the end of October 1993, \$160,000 will be forfeit to Westinghouse. An effort is underway to extend this deadline to April 1994. There is some (unclear) level of commitment to some sort of blade replacement anyway. This commitment could be credited to the replacement with tri-pin blades, so that the real cost is the increment to achieve the tri-pin design.

Structure of the Decisions

In summary, the "decision space" consists of 12 possibilities, arranged in a 4 X 3 matrix. This allows each combination of the following strategies to be considered.

Governor valve strategies (between now and whenever a new rotor is installed):

1. 2 valve operation.
2. 4 valve operation until a new rotor is available, then 2 valve operation.
3. 4 valve operation.
4. 5 valve operation.

Possible turbine rotor strategies include:

1. Refit rotor and install as soon as possible.
2. Refit rotor and install in 1997.
3. Refit rotor and install in 2007.

Since the probability of a failure is the primary uncertain element, it will be helpful to characterize each combination of valve operation and rotor refit strategies as follows:

Where $p(f)$ is the probability that the rotor will fail in two valve mode,

- Identify combinations having any interval for $p(f)$ for which each is the best combination. Call these "candidates", and identify the interval for which each is the best.
- For each Candidate, at what value of $p(f)$ is the candidate better (has lower cost or higher savings) than the next best by the widest margin (the optimal point), and how what is this margin? What is the benefit of the Candidate at this point, compared with 5 valve operation and 2007 rotor refit?

Present the above for the base case RPV schedule, and for the High Thermal RPV schedule.

Background

This represents a large expenditure for the plant, compared with most maintenance decisions. Most decisions fall in the range of \$25k - \$200k. They find that EASOP is satisfactory for determining the economics of small value, simply structured decisions. Tools are desired for large issues like this, which involve uncertainty.

Presentation of break even probabilities and probability ranges, is a useful way to compare alternatives.

One general goal is that all expenditures have a payback period of 5 years or less.

They typically evaluate an alternative based on the "baseline" RPV, and then compare this with the "High Thermal" RPV case, to evaluate sensitivity to future rainfall. For short term decisions, they use Power Controls' 7 day values for RPV, from Mr. Raymond.

Input Available

- Costs represented as total financial (fully burdened).
- Forecast hours spent in each load point range, for 1994. This is available from the plant, and also Mike Palatroni may have forecasts.
- Load vs. Operating Cost per Hour, for 2 valve, 4 valve and 5 valve operation.
- The Cost of repair and time to repair, for a Minor Failure and for a Major Failure. This data will be provided for each of the options for Minor Failure (Cut blades, or order reblading), and for a major failure (Cut Blades, initiate expedited rotor refit and install rotor as soon as it is received, or else drop steeple, which is not likely for Unit 7).

Reference Material

1. White Paper, "HP Turbine First Stage Blade Replacement Cursory Evaluation for Tri-Pin Modification", 5 Pages.
2. List, "Questions Related to Unit 6 & 7 TriPin Modification, M. L. Jones, August 11 / August 27, 1993, 2 Pages.

3. Memo, Bob Deutschman (Westinghouse), "Experience List for TriPin Modification, and Answers to Other Questions", FAXed to Mike Jones September 13, 1993, 7 Pages including cover.

Decision Method and Description
Case 2: PPP Unit 6 Secondary Superheater Outlet Header Stub Tube Welds
Level Of Repair

October 10, 1993

Problem Statement

The Secondary Superheater Outlet Header Stub Tube welds, where the stub tubes (total 180) connect to the header, have a history of cracking, which currently results in about \$100,000 per year outage and repair costs, due to an average of two leaks per year. If nothing is done, it is anticipated that the failure rate could increase. Repair upon failure is also undesirable because the vestibule where the header is located is quite hot immediately after a shutdown and cool down, so there are safety and productivity issues.

The next opportunity for preventative repair is the outage scheduled March 1, 1994, and projected to take 6 weeks. Eight preventative repair alternatives have been identified, in addition to doing nothing. There are 5 main repair approaches, 3 with two alternative levels. The entire list of choices is:

1. Do Nothing
2. Grind Out and Replace All Stub Welds
3. Grind Out and Replace 60 Stub Welds
4. Replace All Stub Tubes with Standard Stub Tubes
5. Replace 60 Stub Tubes with Standard Stub Tubes
6. Replace All Stub Tubes with New B&W Sleeves and Insert
7. Replace 60 Stub Tubes with New B&W Sleeves and Insert
8. Ultrasonic Test Welds and Replace Only Cracked Welds
9. Replace the Entire Header

There is not sufficient lead time to obtain a new header by the time of the next outage (March 1994). The header is thought to have a limited remaining life, possibly as short as 5 years, due to header cracking (a different problem than stub tube weld cracks). An inspection of Unit 5 header may support inference of the condition and remaining life of the Unit 6 header. A possible measure to extend the life of the header is to derate the unit temperature and pressure, balancing this cost against the cost of replacement of the header.

Opportunity for Savings, and Associated Costs

- By repairing the welds during the outage (which is expected to be 6 weeks duration), the presently experienced two outages per year due to stub tube leaks may be reduced

or eliminated, along with the attendant \$100,000 in associated outage and repair costs. Otherwise, the failure rate and costs can be expected to increase.

- There are costs associated with each repair strategy. Also, some strategies reduce but do not eliminate the probability of leaks.
- It has been said that Alternative 6 will extend the life 40 years, but this method has not been used before.
- Alternative 8 involves ultrasonic testing of unproved effectiveness.
- Some alternatives require a startup strainer ahead of the turbine to catch possible contaminants. After a brief run, the unit must be shut down and the strainer pulled. This adds about \$30,000 to the cost of the outage.
- The next scheduled outage is in 1996.

Structure of the Decisions

In summary, the "decision space" consists of choosing one of the 9 alternatives from the list above, for action in the March 1994 outage. The estimated construction cost of each is shown in the following table, along with other parameters. The alternative titles have been abbreviated in the table.

R(f) is the failure rate of the Secondary Superheater Outlet Header Stub Tube welds.

Table 1: Alternatives vs. Costs and Consequences

Alternative	Construction Cost (\$ x 1000)	R(f)	Repair Time (weeks)	Require Startup Strainer
1. Do Nothing	0	Will Increase	0	No
2. All Welds	164	0 for 4 years	6	No
3. 60 Welds	50	Reduce	<6	No
4. All Standard Tubes	181	0 for 7 years	8	Yes
5. 60 Standard Tubes	72	Reduce	<6	Yes
6. All B&W	214	0 for 7 years	8	Yes
7. 60 B&W	87	Reduce	<6	Yes
8. UT / Reweld Cracks	60	Unclear	?6	No
9. Replace Header	1,900	0 for ? Years	10	Yes

The primary uncertain elements are the period for which the leaks will be eliminated and the amount by which $R(f)$ will be reduced. The required time is also somewhat uncertain.

The approach will be to estimate the expected total cost $C(n)$, $n = 1$ to 9 , of each of the alternatives, including construction cost, additional outage time to complete repair, and subsequent outages due to tube failures. This should be calculated for the base case RPV schedule, and for the High Thermal RPV schedule, to determine sensitivity. Since 7 years is the longest interval of influence, the study period will be 7 years, starting in April 1994.

$C(4)$ and $C(6)$ can be calculated from the data given. To allow for uncertainty in $R(f)$ in some alternatives, the "break even" $R(f)$ will be calculated for each of the alternatives 1,2,3,5,7 and 8. This is the value of $R(f)$ (if such a value exists) for which the total evaluated cost of the alternative is equal to the lesser of $C(4)$ and $C(6)$. These can be reviewed against the criteria "is it reasonable to assume that this repair will reduce $R(f)$ to a level no greater than this?" Those for which the answer is "no" can be removed from consideration.

Note: for alternative 2, assume $R(f) = 0$ for the first 4 years, and find break even $R(f)$ for the last 3 years of the 7 year period.

Inputs

- Material received describes some of the distinct costs and benefits of each alternative.
- Other required inputs TBD.

Decision Method and Description
Case 3: Pittsburg Unit 7 FD Fan Bearing Maintenance

October 10, 1993

Problem Statement

On Aug. 29, the operator's received a high temp alarm on 7-2 FD fan bearing. Upon investigation, they found the bearing cap very hot (over 200 F). They reduced load, put water on top of the cap and topped up the bearing reservoir with oil. This brought the bearing temperature within normal limits.

The following week, engineers discussing the problem, and suggested collecting an oil sample and looking for wear metals. About a week later, results were received that showed very high levels of contamination and wear metal particles. Continued operation with contaminated oil will almost certainly lead to a bearing failure eventually. The unit cannot be operated without the fan operating, so shutdown of the fan or failure of the bearing will each incur an outage. In order to clean up the oil, we considered the following alternatives:

1. Shut the unit down within a day or two and change the oil.
2. While the fan is running, drain a portion of the oil out, refill, and repeat the process until acceptable oil quality was achieved.
3. While the fan is running, install a small, kidney loop type filter system on the reservoir and filter the oil.
4. Keep running with no modifications until the next unit shutdown which provides enough time to change the oil.

Options 2 & 3 are difficult because the reservoir has no drain valve, only a drain plug. These options would have required installation of a drain valve while the reservoir was full of oil, resulting in the risk of major oil loss, and possibly requiring a unit shutdown as a consequence of the oil loss.

After discussing the problem, we decided to thoroughly check the bearing temperature and vibration levels. These were found to be normal. Given the likelihood of a fuel economy outage in the near future, we opted for periodic vibration analysis, along with Alternative 4. The unit was shut down on fuel economy on 9/16 and the oil was changed at that time.

Opportunity for Savings, and Associated Costs

The opportunities for savings are:

For Alternatives 1, 2 and 3: As soon as unit is shut down, or the oil is cleaned up "on line", the probability of a bearing failure becomes zero. This saves the possible cost of

bearing replacement. It is assumed that any of these will be done soon enough that the bearing will not fail before the oil is cleaned up.

For Alternatives 2, 3 and 4: The cost of a special outage is avoided. However, for alternative 4, the risk of an outage due to bearing failure is incurred.

For Alternative 4: The cost of the on line modifications, as well as the risk of spilled oil, is eliminated.

Structure of the Decisions

In summary, the "decision space" consists of 4 alternatives. The expected cost of each is the sum of the component costs.

In the following table, an X indicates that for the associated alternative, there is a possibility of incurring the indicated cost.

Table 1: Alternatives vs. Possible Cost Elements

Alternative	Outage Cost	Oil Spill Cost	On Line Mod Cost	Bearing Repair Cost
1	X			
2	X	X	X	
3	X	X	X	
4	X			X

The primary uncertain elements are the probability of failure of the bearing during a given interval $P(b)$, and the probability that the oil will be spilled if the on line modification is attempted $P(s)$. another uncertain element is the time until the next outage which will be of sufficient length to replace the oil (N weeks).

The approach will be to estimate the expected total cost of each of the alternatives, by multiplying each cost component by the probability of its occurrence. For each of several assumed time intervals until the next outage, a matrix will be calculated. Each element of the matrix contains A, the number of the least cost alternative, and c, the cost difference between A and the next lowest cost alternative, for the specified values of P(b) and P(s):

Table 2-N: Preferred Alternatives Over the Range of P(s) and P(b) for N Weeks Until Next Outage

	P(s)	
P(b)	0.1	0.2
0.1	A, c	
0.2		
..		

This set of tables should be calculated for the base case RPV schedule, and for the High Thermal RPV schedule, to determine sensitivity.

Inspection of the matrices will determine if more than one alternative can be the lowest cost, within the reasonable range of the variables P(s), P(b) and N, and if the frequency with which more than one is the lowest is sufficient to warrant further pursuit of an estimate of any of the probabilities. Note that the "Reasonable Range" assessment of P(b) is probably affected by the observations of bearing temperature and vibration level. However, this assessment is done by the reviewer of the results, and need not be incorporated into the results structure, since the results are presented for all values of P(b).

Inputs

- Costs of bearing replacement, modifications and oil spill, represented as total financial (fully burdened).
- The outage time to repair the bearing, and the outage time to replace the oil.

APPENDIX B

Note on Minimal Repair

In maintenance management, engineers usually use a simple calculation to obtain the expected cost of breakdown maintenance on a system. If the planning horizon is T , and is divided into m equal-spaced periods. That is,

$$T = m\Delta t \tag{1}$$

Then the expected cost of breakdown maintenance on a system of age a over m periods is calculated from the following equation.

$$C(T, a) = \sum_{i=1}^m \beta^i C_f r(a + i\Delta t) \Delta t \tag{2}$$

where β : discount factor;

C_f : cost of breakdown maintenance;

$r(t)$: system's failure rate at time t .

Equation (2) (compare with the third term of equation (6.2)) appears intuitively correct. However, the underlying assumption about the effectiveness of breakdown maintenance on the failing system is not clear. We will show how the above engineering approach is related to one of the maintenance policies in Barlow and Proshcan's book (1965: 84-118) and the justification for it.

Consider a system composed of a large number of identical components. For example, in PG&E's test case 2, it is a boiler having 280 superheater tubes in a conventional power plant (see Section 6). Most of the time, an unexpected failure of any of the constituting components is fixed by breakdown maintenance, which is either a local repair or a replacement. However, the number of fixed components, each of which has a

reduced failure rate, is small compared with the number of overall components. Therefore, in modeling the life distribution of this system, it is often assumed that the breakdown maintenance only recovers the system's functionality, but does not disturb the system's age. That is, a minimal repair policy, as defined by Barlow and Proschan (1965: 96-98), is used.

Now consider a system of age a , with an underlying life distribution F , is maintained by a minimal repair policy over a period of T . Suppose F is continuous and differentiable, the expected cost $C(T, a)$ can be expressed as

$$C(T, a) = \int_0^T \beta^\tau [C_f + C(T - \tau, a + \tau)] \frac{f(a + \tau)}{\bar{F}(a)} d\tau \quad (3)$$

where f : probability density function of F .

If we discretize T into m periods with equal duration Δt , then equation (3) can be rewritten as

$$C(T, a) = \sum_{i=1}^m \beta^i [C_f + C(T - i\Delta t, a + i\Delta t)] \frac{f(a + i\Delta t)}{\bar{F}(a)} \Delta t \quad (4)$$

After one period has passed, the system's age has increased to $a + \Delta t$, whether an unexpected failure happens or not, and the remaining time under consideration is $T - \Delta t$. Therefore,

$$\begin{aligned} C(T - \Delta t, a + \Delta t) &= \sum_{i=1}^{m-1} \beta^i [C_f + C(T - \Delta t - i\Delta t, a + \Delta t + i\Delta t)] \frac{f(a + \Delta t + i\Delta t)}{\bar{F}(a + \Delta t)} \Delta t \\ &= \frac{1}{\bar{F}(a + \Delta t)} \left\{ \sum_{i=1}^{m-1} \beta^i [C_f + C(T - \Delta t - i\Delta t, a + \Delta t + i\Delta t)] f(a + \Delta t + i\Delta t) \Delta t \right\} \\ &= \frac{1}{\bar{F}(a + \Delta t)} \left\{ \sum_{i=2}^m \beta^{i-1} [C_f + C(T - i\Delta t, a + i\Delta t)] f(a + i\Delta t) \Delta t \right\} \\ &= \frac{1}{\beta \bar{F}(a + \Delta t)} \left\{ \sum_{i=2}^m \beta^i [C_f + C(T - i\Delta t, a + i\Delta t)] f(a + i\Delta t) \Delta t \right\} \end{aligned} \quad (5)$$

Combine equations (4) and (5), we have

$$\begin{aligned}
 \bar{F}(a)C(T, a) &= \sum_{i=1}^m \beta^i [C_f + C(T - i\Delta t, a + i\Delta t)]f(a + i\Delta t)\Delta t \\
 &= \beta[C_f + C(T - \Delta t, a + \Delta t)]f(a + \Delta t)\Delta t + \sum_{i=2}^m \beta^i [C_f + C(T - i\Delta t, a + i\Delta t)]f(a + i\Delta t)\Delta t \\
 &= \beta[C_f + C(T - \Delta t, a + \Delta t)]f(a + \Delta t)\Delta t + \beta\bar{F}(a + \Delta t)C(T - \Delta t, a + \Delta t) \quad (6)
 \end{aligned}$$

If Δt is very small, then

$$\bar{F}(a) \simeq f(a + \Delta t)\Delta t + \bar{F}(a + \Delta t)$$

Equation (6) can be rewritten as

$$C(T, a) = \beta C_f \frac{f(a + \Delta t)}{\bar{F}(a)} \Delta t + \beta C(T - \Delta t, a + \Delta t)$$

By the assumption about Δt ,

$$\frac{f(a + \Delta t)}{\bar{F}(a)} \simeq r(a + \Delta t)$$

where $r(\cdot) = \frac{f(\cdot)}{\bar{F}(\cdot)}$: system's failure rate.

We have

$$C(T, a) \simeq \beta C_f r(a + \Delta t)\Delta t + \beta C(T - \Delta t, a + \Delta t) \quad (7)$$

Apply equation (7) repeatedly, we have

$$C(T, a) = \sum_{i=1}^m \beta^i C_{fr}(a + i\Delta t)\Delta t \quad (8)$$

Equation (8) is exactly the same as equation (2). Therefore, the underlying assumption of the engineering approach mentioned above is performing a minimal repair policy on a system of a large number of identical components. If the number of identical components in the system under study is small or the number is big but some components other than those failed are repaired opportunistically, then the calculation fails. Therefore, if we want to quickly calculate the expected cost of performing breakdown maintenance on a system of 2 light bulbs, the above approach does not fit. However, using it on a system of 280 superheater tubes is plausible when no opportunistic breakdown maintenance is employed.

REFERENCES

- Barlow, R.E. and Proschan, F. *Mathematical Theory of Reliability* John Wiley & Sons, 1965.
- Barlow, R.E. and Proschan, F. *Statistical Theory of Reliability and Life Testing* 2nd printing, To Begin With, 1981.
- Hogg, R.V. and Craig, A.T. *Introduction to Mathematical Statistics* 3rd ed., Macmillan, 1970.
- Jin, Y. and Levitt, R.E. Levitt, "Model-Based and Interactive Planning for Predictive Plant Maintenance Management," *5th International Conference on Computing in Civil Engineering* 1993.
- Jin, Y., Kunz, J.C., Levitt, R.E., and Winstanley, G., "Design of Project Plans from Fundamental Knowledge of Engineered Systems," *AAAI Fall Symposium - Design from Physical Principles* 1992.
- Law, A.V. and Kelton, W.D. *Simulation Modeling and Analysis* 2nd ed., McGraw-Hill, 1991.
- Lloyd, D.K. and Lipow, M. *Reliability: Management, Methods, and Mathematics* 2nd ed., published by the authors, 1977.
- Rice, J.A. *Mathematical Statistics and Data Analysis* Wadsworth & Brooks/Cole Advanced Books & Software, 1988.
- Simpa, H., "PROSPER - IRTMM Knowledge Sources," Unpublished Working Paper, 1993.
- Simpa, H. and Hayes-Roth, B., "IRTMM: The Situation Assessment Module," Unpublished Working Paper, 1992.