1. INTRODUCTION

1.1 Objective

The data analysis portion of a nuclear power plant probabilistic risk assessment (PRA) provides estimates of the parameters used to determine the frequencies and probabilities of the various events modeled in a PRA. The objective of this handbook is to provide methods for estimating the parameters used in PRA models and for quantifying the uncertainties in the estimates.

1.2 Background

Probabilistic risk assessment is a mature technology that can provide a quantitative assessment of the risk from accidents in nuclear power plants. It involves the development of models that delineate the response of systems and operators to accident initiating events. Additional models are generated to identify the component failure modes required to cause the accident mitigating systems to fail. Each component failure mode is represented as an individual "basic event" in the systems models. Estimates of risk are obtained by propagating the uncertainty distributions for each of the parameters through the PRA models.

During the last several years, both the U.S. Nuclear Regulatory Commission (NRC) and the nuclear industry have recognized that PRA has evolved to the point where it can be used in a variety of applications, including as a tool in the regulatory decision-making process. The increased use of PRA has led to the conclusion that the PRA scope and model must be commensurate with the applications. Several procedural guides and standards have been and are being developed that identify requirements for the PRA models. This handbook was generated to supplement these documents. It provides a compendium of good practices that a PRA analyst can use to generate the parameter distributions required for quantifying PRA models.

The increased use of risk assessment has also helped promote the idea that the collection and analysis of event data is an important activity in and of itself. In particular, the monitoring of equipment performance and evaluation of equipment trends can be used to enhance plant performance and reliability. The guidance provided in this handbook can support those efforts.

1.3 Scope

This handbook provides guidance on sources of information and methods for estimating parameter distributions. This includes determination of both plant-specific and generic estimates for initiating event frequencies, component failure rates and unavailabilities, and equipment non-recovery probabilities.

This handbook provides the basic information needed to generate estimates of the parameters listed above. It begins by describing the probability models and plant data used to evaluate each of the parameters. Possible sources for the plant data are identified and guidance on the collection, screening, and interpretation is provided. The statistical techniques (both Bayesian and classical methods) required to analyze the collected data and test the validity of statistical models are described. Examples are provided to help the PRA analyst utilize the different techniques.

This handbook also provides advanced techniques that address modeling of time trends. Methods for combining data from a number of similar, but not identical, sources are also provided. These are the empirical and hierarchical Bayesian approaches. Again, examples are provided to guide the analyst.

This handbook does not provide guidance on parameter estimation for all of the events included in a PRA. Specifically, common cause failure and human error probabilities are not addressed. In addition, guidance is not provided with regard to the use of expert elicitation. For these topics, the PRA analyst should consult other sources, such as the following references:

Common cause failures

- NUREG/CR-5497 (Marshall et al. 1998),
- NUREG/CR-6268 (Kvarfordt et al. 1998),
- NUREG/CR-5485 (Mosleh et al. 1998),
- NUREG/CR-4780 (Mosleh et al. 1988), and
- EPRI NP-3967 (Fleming, 1985).

Human errors

- NUREG/CR-1278 (Swain and Guttman, 1983),
- NUREG/CR-4772 (Swain, 1987),

- NUREG-1624 (NRC, 2000b), and
- EPRI TR-TR-100259 (Parry et al. 1992).

Expert Judgement

- NUREG/CR-6372 (Budnitz et al. 1997) and
- NUREG/CR-1563 (Kotra et al. 1996).

This list is not meant to be a comprehensive list of all of the methodologies available for performing these types of analyses.

1.4 Contents of the Handbook

This section provides a road map of the contents of the handbook and an overview discussion on how to use the handbook to perform the elements of a data analysis. The basics of probability and statistics described in Appendices A and B, respectively, are provided as reference material for the analyst. Appendix C provides statistical tables for selected distribution types that can be used in the data analysis.

1.4.1 Identification of Probability Models

The handbook provides guidance on the evaluation of five types of parameters that are included in a PRA:

- initiating events,
- failures to start or change state,
- failures to run or maintain state,
- durations, and
- unavailability from being out of service.

A description of each of these parameters along with examples, is provided in Chapter 2. Chapter 2 is fundamental reading for all users of this handbook.

The first step in a data analysis is to determine the appropriate probability models to represent the parameter. Chapter 2 provides a detailed description of the standard probability models for each event. This includes a discussion of the assumptions on the physical process inherent in the models and a description of the kind of data that can be observed. The type of data required to estimate the model parameter(s) are described and example data sets are examined in the light of the model assumptions. These examinations illustrate the kind of thinking necessary for the data analyst. Finally, a short discussion of related issues is presented for the analyst to consider.

1.4.2 Collection of Plant Specific Data

Once probability models have been defined for the basic events, plant-specific data should be evaluated for the purpose of quantifying estimates of the probability model parameters. Plant-specific data, if available in sufficient quantity and quality, is the most desirable basis for estimating parameter values. Chapter 5 discusses the process by which plant-specific data should be identified, collected, screened, and interpreted for applicability to the basic events defined in the systems analysis and to their probability models. To ensure that the collection and evaluation of plant-specific data is thorough, consistent, and accurate, the steps laid out in Chapter 5 should be followed for events defined in a PRA. The identification and evaluation of appropriate sources of plant-specific data for the basic events are discussed in Section 4.1.

The process for collecting and evaluating data for initiating events is discussed in Section 5.1. Guidance is provided for screening the data, for grouping the data into appropriate categories of initiating events, and for evaluating the denominator associated with the data.

The process for collecting and evaluating data for component failures is discussed in Section 5.2. It is critical that data be collected and processed accurately according to the definition of the component boundary. For example, it should be clearly noted whether or not a pump's control circuit is within or without the physical boundaries of the component for purposes of systems modeling. If failure of the control circuit has been modeled separately from hardware failures of the pump, then data involving failure of the pump should be carefully evaluated to ensure that actuation failures and other pump faults are not erroneously combined. This process could result in some iteration between the systems analysis task and the data collection task. It is possible that system models may be simplified or expanded based on insights derived during the data collection. Chapter 3 describes the difference between faults and failures, and discusses component boundary definitions and failure severity as it relates to data collection and analysis.

Other aspects of data collection for component failures discussed in Section 5.2 include classification and screening of the data, allocation of the data to appropriate component failure modes, and exposure evaluation (determining the denominator for parameter estimates).

The collection of data for recovery events is described in Section 5.3. Guidance is provided on where to find recovery-related data and on how to interpret such data.

1.4.3 Quantification of Probability Model Parameters

Once appropriate probability models have been selected for each basic event, estimates for the model parameters must be quantified. There are two basic approaches: 1) statistical estimation based on available data; and 2) utilization of generic parameter estimates based on previous studies. Both approaches can incorporate generic data. Several generic data sources currently available and used throughout the nuclear PRA industry are identified in Section 4.2.

1.4.3.1 Parameter Estimation from Plant-Specific Data

If the plant-specific data collection process yields data of sufficient quantity and quality for the development of parameter estimates, the statistical methods in Chapter 6 can be applied to the data to derive and validate parameter estimates for the basic events.

Chapter 6 discusses the statistical methods for estimating the parameters of the probability models defined in Chapter 2. Note that Appendix B discusses basic concepts of statistics that will help the user to understand the methods presented in Chapter 6.

For each type of event, two fundamental approaches are presented for parameter estimation: classical (frequentist) and Bayesian. An overview and comparison of these two approaches are presented in Section 6.1. The Bayesian approach is more commonly used in PRA applications, but classical methods have some use in PRA, as discussed in Section 6.1.

The probability models discussed in Chapter 2 for each type of event are applicable for most applications. However, erroneous results can occur in some cases if the assumptions of the model are not checked against the data. In some applications (e.g., if the impact of casual factors on component reliability is being examined) it is imperative that the probability model chosen for each basic event be validated given the available data. It may seem sensible to first confirm the appropriateness of the model and then estimate the parameters of the model. However, validation of a model is usually possible only after the model has been assumed and the corresponding parameters have been estimated. Thus, estimation methods are presented first in Chapter 6 for each type of probability model; then methods for validating the models against the available data are presented.

1.4.3.2 Parameter Estimation from Existing Data Bases

If actual data are unavailable or of insufficient quality or quantity then a generic data base will have to be used. Several generic data sources currently available and used throughout the nuclear PRA industry are identified in Section 4.2. Section 4.2.6 provides guidance on the selection of parameter estimates from existing generic data bases.

1.4.4 Advanced Methods

The last two chapters of the handbook describes some methods for analyzing trends in data and Bayesian approaches for combining data from a number of similar sources.

1.4.4.1 Analyzing Data for Trends and Aging

Data can be analyzed to assess the presence of time trends in probability model failure rates and probabilities (i.e., λ and p). Such trends might be in terms of calendar time or in terms of system age. Ordinarily, the analysis of data to model time trends involves complex mathematical techniques. However, the discussion of Chapter 7 presents various approaches that have been implemented in computer software. The discussion in Chapter 7 focuses on the interpretation of the computer output for application in PRA.

1.4.4.2 Parameter Estimation Using Data from Different Sources

Two Bayesian approaches for combining data from a number of similar, but not identical, sources are discussed in Chapter 8.

1.5 How to Use This Handbook

This handbook is intended for workers in probabilistic risk assessment (PRA), especially those who are concerned with estimating parameters used in PRA modeling. Broadly speaking, three groups of readers are anticipated: **data collectors**, who will be finding, interpreting, and recording the data used for the estimates; **parameter estimators**, who will be constructing the parameter estimates from the data and quantifying the uncertainties in the estimates; and (to a lesser extent) **PRA analysts**, who will be using the estimated parameters. These three groups will find their primary interests in different portions of the handbook, as discussed below.

The major sections of the handbook can be grouped into several areas:

- Foundation: Chapters 1 and 2;
- Data Collection: Chapters 3, 4, and 5;
- Parameter Estimation: Chapters 6, 7, and 8; and
- Supporting Material: Appendices, References, Index.

These sections are shown in Figure 1.1, a schematic representation of the contents of the handbook.

PRA analysts will be most interested in the foundational material. Data collectors will need to read much of the foundational material, and then read the chapters on data collection. Parameter estimators will need to read the foundational chapters, but may then wish to skip directly to the relevant sections on parameter estimation. The supporting material can be read by anyone at any time.

The arrows in Figure 1.1 help the reader find the quickest way to the sections of interest. For example, the figure shows that Chapters 3-5 and Chapters 6-8 do not refer to each other or assume material from the other section, so it is possible to read from one section and not the other. The only strong dependencies are shown by the arrows: read Chapter 2 before starting Chapter 3 or 6, read Chapter 3 before starting Chapter 4 or 5, and so forth. In practice, data collectors, data analysts, and PRA analysts must work together, giving feedback to each other. The handbook, on the other hand, is formed of distinct segments, each of which can be read in isolation from the others.

The material for PRA analysts and data collectors is intended to be accessible by anyone with an engineering background and some experience in PRA. The material for data analysts, on the other hand, begins with elementary techniques but eventually covers advanced models and methods. These advanced topics will not be needed in most cases, but are included as reference material.

To aid the reader, Appendices A and B summarize the basics of probability and statistics, and Appendix C provides useful statistical tables. A glossary of terms is provided in Appendix D. Persons who have no previous experience with probability or statistics will need a more thorough introduction than is provided in these sections of the handbook.

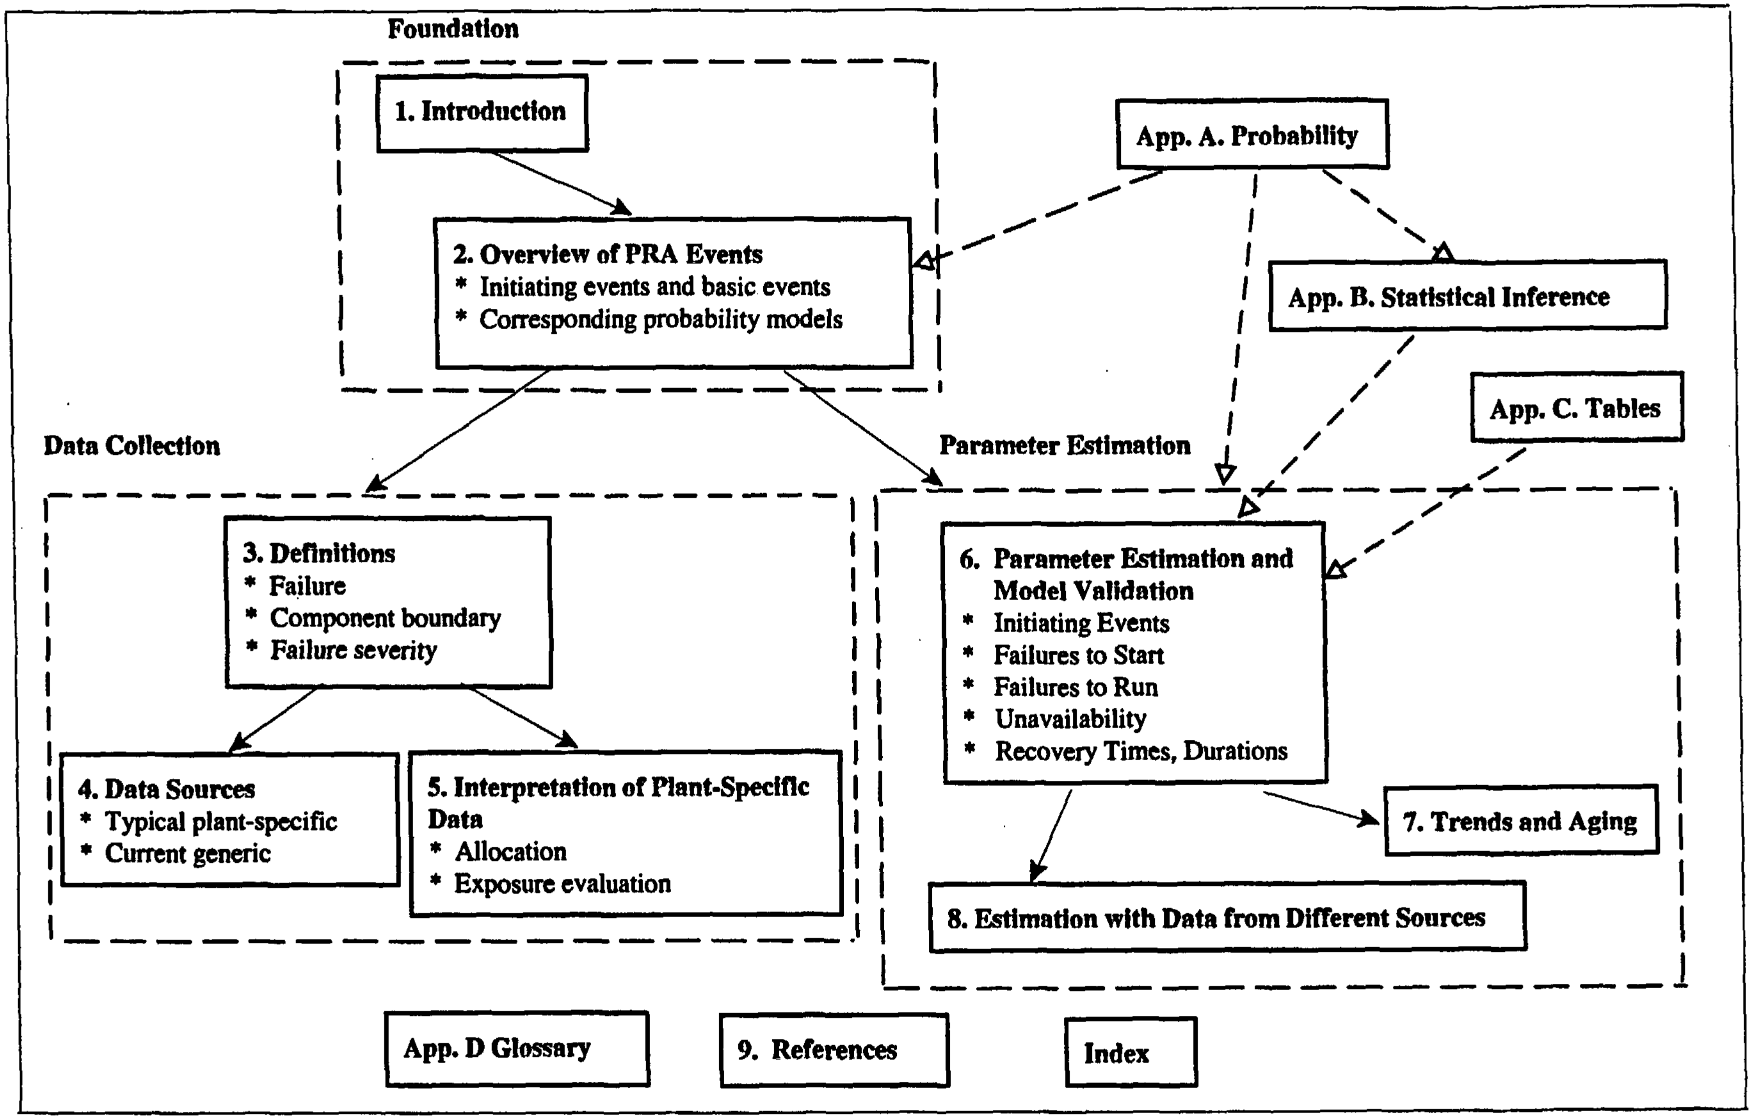

Figure 1.1 Relation of sections of this handbook. An arrow with a solid head indicates a strong dependence, with required information flowing in the direction of the arrow. An dashed arrow with a hollow head indicates that background or supplementary material flows in the direction of the arrow. The glossary, references, and index supply assistance for the entire handbook, so the arrows are not shown.

2. BASIC EVENT PROBABILITY MODELS

2.1 Overview

This chapter introduces the models used for basic events and for initiating events. This first section is an overview, and the remaining sections of the chapter give further detail.

Probabilistic risk assessment (PRA) considers various possible accident sequences. An accident sequence begins with an **initiating event** which challenges the safety of the plant. Typically, one or more standby safety systems are then demanded, and other, normally operating, systems must continue operating to ensure that no serious undesirable consequences occur. For the systems to fail to bring the situation under control, several components must either fail or be unavailable. The logic events in the PRA model that represent these failures or modes of unavailability are called **basic events**.

It is not possible to predict precisely when an initiating event or a component failure will occur, because the processes that lead to their occurrences are complex. Therefore, the initiating events and basic events are modeled as resulting from random processes.

The first step in the data analysis task is, therefore, to determine the appropriate probability model to represent the initiating event or basic event. (Probability is reviewed in Appendix A, and the probability models introduced here are presented more fully there.) These probability models typically have one or more parameters. The next major step is to estimate the values of these parameters. This estimation is based on the most applicable and available data. The process of choosing data sources, extracting the data in an appropriate form, and using it to estimate the parameters is the main subject of this handbook.

Basic events are customarily divided into unavailability (because the equipment is undergoing testing or maintenance), failure to start or change state, and failure to run (after successfully starting) or maintain state to the end of the required mission time. Unavailability and failure to run are each modeled in a single way. On the other hand, two different probability models have been used to represent a failure to start or to change state. The first method is to model the failures as having a constant probability of **failure on a demand**. The second method is to model the failures as occurring, in an unrevealed way, randomly in time. The failed condition is then discovered at the time of the demand. This is usually called the **standby failure-rate model**. Both models are discussed here.

The above events are the typical ones considered in a PRA. In addition, one must occasionally analyze durations, such as the time to restore offsite power or time to recover a failed component. Although such an analysis is not needed for a typical accident sequence, it is discussed in this handbook. Also, methods for analyzing durations can be used when estimating unavailability.

In summary, five topics are considered in the rest of this chapter:

- initiating events,
- failures to start or change state (modeled in two possible ways),
- failures to run or maintain state,
- durations, and
- unavailability from being out of service.

These topics are the subjects of Sections 2.2 through 2.6. Each section begins with examples of the data that might be analyzed. This is followed by a brief subsection presenting the assumptions of the usual model for the random process (the result of underlying physical mechanisms) and describing the kind of data that can be observed. The next subsection summarizes the data required to estimate the model parameter(s). The example data sets are then examined in the light of the model assumptions. These examinations illustrate the kind of thinking necessary for the data analyst. Finally, the section may conclude with a short discussion of related issues.

As a preview, Table 2.1 indicates the models, the parameters, and the data needed for each of the topics in the above five bullets. The top line of the table also indicates which section of Chapter 2 treats the topic.

The term system is used to denote the set of hardware for which data are collected; it may be an entire nuclear power plant (NPP), or a system in the traditional sense, such as the auxiliary feedwater (AFW) system, or a train, component, or even piece part. This reduces the need for phrases such as "system or component."

The lengthiest part of each section below consists of the examination of examples to see whether the

Table 2.1 Kinds of models considered.

2.2 Initiating Events	2.3 Failures to Start or Change State (2 models)		2.4 Failures to Run or Maintain State	2.5 Durations	2.6 Unavailability
Typical Event					
Event occurs initiating an accident sequence	Standby system fails on demand		System in operation fails to run, or component changes state during mission	A condition persists for a random time period	System is unavailable, intentionally out of service, when demanded
Parameter(s) to Estimate					
λ, event frequency	For failure on demand: p, probability of failure on demand	For standby failure: λ, rate of occurrence of standby failures	λ, rate of failure to run	Parameters of assumed probability distribution of duration time	q, fraction of time when system will be out of service
Data Required to Estimate Parameters[a]					
Number of events, x, in total time, t	Number of failures, x, in total number of demands, n	Number of failures, x, in total standby time, t	Number of failures, x, in total running time, t	Depends on model, but typically the lengths of the observed durations	Onset times and durations of observed out-of-service events; OR observed fractions of time when system was out of service

[a] The data here are the minimal requirements to estimate the parameter. More detailed data are needed to check the model assumptions.

assumptions of the probability model appear to be satisfied. Verifying model assumptions is an important part of good data analysis. Ways to investigate the appropriateness of assumptions are considered in Chapter 6, along with parameter estimation. The present chapter, however, only introduces the assumptions and illustrates their meanings through examples. If the assumptions are clearly not satisfied, some mention is given of ways to generalize the model, although such generalizations are not presented until Chapters 7 and 8 in this handbook.

Also, examples and extended discussion of examples are printed in Arial font, to distinguish them from the more general material.

2.2 Initiating Events

2.2.1 Examples

In the context of a nuclear-power-plant PRA, an initiating event is any event that perturbs the steady state operation of the plant, thereby initiating an abnormal event such as a transient or a loss-of-coolant accident within a plant. Initiating events begin sequences of events that challenge plant control and safety systems. Failure of these systems can lead to core damage and a release of radioactivity to the environment. However, the consideration of the potential plant response to initiating events is irrelevant when estimating their frequencies.

Here are several examples of data sets counting such initiating events.

Example 2.1 Unplanned reactor trips

A U.S. commercial nuclear power plant had 34 unplanned reactor trips in 1987 through 1995. It had its initial criticality on Jan. 3, 1987, and experienced a total of 64651 critical hours, or 7.38 critical years (Poloski et al. 1999a).

Example 2.2 Shutdown loss of offsite power

In U.S. commercial nuclear power plants in 1980-1996, there were 80 plant-centered loss-of-offsite-power (LOSP) events during shutdown. In that period, the plants experienced 455.5 reactor-shutdown years (Atwood et al. 1998).

Example 2.3 Through-wall pipe leaks

In world-wide experience of western-style pressurized water reactors (PWR)s (3362 calendar years of operation), a single through-wall leak event has been reported in large-diameter piping (Poloski et al. 1999a, Appendix J).

The final example of this section does not have initiating events in the usual sense. However, the model assumptions and the form of the data are exactly the same as for initiating events. Therefore, such data can be analyzed just like initiating-event data.

Example 2.4 Temperature sensor/transmitters

Eide et al. (1999a) report that temperature sensor/transmitters in the reactor protection system (RPS) of Westinghouse NPPs had 32 failures in 2264.1 component-years. These sensor/transmitters operate continuously, and when they fail they are repaired or replaced in a relatively short time. The number of failures is conservatively estimated from sometimes incomplete Nuclear Plant Reliability Data System (NPRDS) data, and the number of component years is based on an estimated number of components per loop.

These examples have several elements in common. First, they involve a **number of events** that occurred, and an **exposure time, or time at risk**, when the events could have occurred. The next subsection will present a simple probability model that gives rise to random events in time. In addition, in each of the above examples corrective action is taken after any event, so that the system then resumes operation (the system is **repairable**.) This means that the recorded operating history consists of a sequence of random event occurrences, which is summarized as a count of events in some fixed time. This type of data will direct us to a particular type of analysis, presented in Chapter 6.

The events may be the initiating events of an ordinary PRA (Example 2.1), initiating events of a shutdown PRA (Example 2.2), failures in a passive system (Example 2.3), which incidentally happen to be initiating events in a PRA. As mentioned above, Example 2.4 does not describe initiating events in the traditional PRA sense. However, the example may be analyzed in the same way as the first three examples, because the sensor/transmitter failures occur in a continuously running system and they initiate quick repair action. A PRA analyst would distinguish among the examples based on their safety consequences. The present discussion, however, adopts the viewpoint of probability modeling, in which the important fact is not the consequence of the events, but the way that they occur randomly in time. Reactor trip initiators are the prototypical examples of such events, but are not the only examples.

The exposure time is the length of time during which the events could possibly occur. In Example 2.1, the exposure time is reactor-critical-years, because a reactor trip can only occur when the reactor is at power. Because only one plant is considered, "critical years" can be used as shorthand for "reactor-critical-years." In Example 2.2, the event of interest is LOSP during shutdown, so the exposure time must be the number of reactor-shutdown-years in the study period. In Example 2.3, reactor-calendar-years are used, primarily because more detailed worldwide data could not be easily obtained. The model therefore assumes that a crack in large-diameter piping could occur with equal probability during operation and during shutdown. The model also does not consider differences between plants, such as differences in the total length of large-diameter piping at a plant. In Example 2.4, the exposure time is the number of component-years, because the components operate constantly.

The possible examples are endless. The events could be unplanned demands for a safety system, forced outage events, or many other kinds of events that resemble initiating events.

The data given in the above examples are expressed in the crudest summary terms: a count of events in a total exposure time. This is sufficient for the simple model of this section. Section 2.5 will consider more sophisticated models using the exact event times.

The data could also be broken down into smaller pieces. For example, the initiating event data could be summarized for each calendar year, with an event count and an exposure time reported separately for each year from 1987 through 1995. This additional information allows one to look for trends or other patterns, as discussed in later chapters.

2.2.2 Probability Model

The assumptions concerning the physical process are given here, along with a description of the kind of data that can be observed.

It is standard to assume that the event count has a Poisson distribution. As listed in Section A.6.2, the usual assumptions (following Thompson 1981) for a Poisson process are:

1. The probability that an event will occur in any specified short exposure time period is approximately proportional to the length of the time period. In other words, there is a rate $\lambda > 0$, such that for any interval with short exposure time Δt the probability of an occurrence in the interval is approximately $\lambda \times \Delta t$.
2. Exactly simultaneous events do not occur.
3. Occurrences of events in disjoint exposure time periods are statistically independent.

In addition, it is worthwhile to spell out the kind of data that can be observed.

- A random number of events occur in some prespecified, fixed time period. As a minimum, the total number of events and the corresponding time period are observed.

Under the above assumptions, the number of occurrences X in some fixed exposure time t is a Poisson distributed random variable with mean $\mu = \lambda t$,

$$\Pr(X = x) = e^{-\mu} \mu^{x} / x! \ . \qquad (2.1)$$

The **probability distribution function** (p.d.f.) is sometimes used to abbreviate this: $f(x) = \Pr(X = x)$. (Throughout this handbook, upper case letters are used for random variables and lower case letters are used for particular numbers.)

The parameter λ is a **rate** or **frequency**. To make things more clear, the kind of event is often stated, that is, "initiating event rate" in Example 2.1, "through-wall-crack occurrence frequency" in Example 2.3, and so forth. Because the count of events during a fixed period is a unitless quantity, the mean number of occurrences μ is also unitless. However, the rate λ depends on the units for measuring time. In other words, the units of λ are per unit of time, such as 1/year or 1/reactor-critical-hour.

This model is called a **Poisson process**. It is extremely simple, because it is completely specified by the exposure time, t, and the one unknown parameter, λ. Assumption 1 implies that the rate λ does not change over time, neither with a monotonic trend, nor cyclically, nor in any other way. Assumption 2 says that exactly simultaneous events do not occur. The only way that they could occur (other than by incredible coincidence) is if some synchronizing mechanism exists – a common cause. Therefore, the operational interpretation of Assumption 2 is that common-cause events do not occur. Assumption 3 says that the past history does not affect the present. In particular, occurrence of an event yesterday does not make the probability of another event tomorrow either more or less likely. This says that the events do not tend to occur in clusters, but nor do they tend to be systematically spaced and evenly separated.

As stated above, a common cause that synchronizes events violates Assumption 2. However, some common-cause mechanisms do not exactly synchronize the events. Instead, the second event may occur very soon after the first, as a slightly delayed result of the common cause. In this case, Assumption 3 is violated, because the occurrence of one event increases the probability of a second event soon after. One way or the other, however, common-cause events violate the assumptions of a Poisson process, by violating either Assumption 2 or Assumption 3.

2.2.3 Data Needed to Validate the Model and Estimate λ

Suppose that the Poisson model holds. Then any reasonable estimator of λ needs only two pieces of information: the total exposure time, t, in the data period, and the number of events, x, that occurred then.

However, more information is needed to investigate whether the Poisson model is valid. For example, the data might cover a number of years or a number of plants, and λ might not be constant over time or the same at all plants. These possibilities are not allowed by the listed model assumptions. To study whether they occur, the times and locations of the initiating events should be recorded, or at least the data should be partitioned into subsets, for example corresponding to plants or years. Then the event count and exposure time, x_i and t_i, should be given for each subset.

2.2.4 Case Studies: Validity of Model Assumptions in Examples

Let us examine the reasonableness of the Poisson model assumptions for Examples 2.1 through 2.4. Chapter 6 will address this issue by performing data analysis. Here we will merely cite the results of published studies and use critical thinking.

Example 2.1 Initiating Events

An initiating event is an event with the reactor critical, causing an unplanned reactor trip. Assume that any time interval starts on some date at some time and ends on some date at some time, and that the length of the interval, Δt, is the number of critical years contained between the start and stop of the time interval. For example, if the time period is two 24-hour days and the reactor was critical for half of that time, then Δt = 1/365 critical years.

Assumption 1 is violated in two ways. First, in the industry as a whole, and presumably in individual plants, the probability of an initiating event in an interval of length Δt (such as one critical day) has not been constant. Instead, the probability dropped substantially from 1987 to 1995. Equivalently, the event rate, λ, dropped from 1987 to 1995. This violation can be eliminated by considering only a short time period for the study, such as one calendar year instead of nine years. If, however, the whole nine-year period is of interest, a more complicated model must be used, such as one of the trend models described in Chapter 7.

A second violation of Assumption 1 arises because this particular plant was new at the start of the study period, with initial criticality on January 3, 1987, and commercial start on May 2, 1987. Many new plants seem to experience a learning period for initiating events, and this plant had 15 of its 34 initiating events during the first six months of 1987. After that initial period with a high event rate, the event rate dropped sharply. This violation of Assumption 1 can be resolved by eliminating data before the plant reached a certain age. That is, not counting either the operating time or the initiating events from the plant until it has reached a certain age — excluding that portion of the plant's history from the universe being studied.

Assumption 2 says that exactly simultaneous initiating events do not occur. This is reasonable for events at a single plant.

Assumption 3 says that the probability of an initiating event in one time period does not depend on the presence or absence of an initiating event in any earlier time period. This assumption may be challenged if the plant personnel learn from the first event, thus reducing the probability of a second event. This kind of dependence of one event on another is not allowed by Assumption 3. Suppose, however, that the learning is modeled as a general kind of learning, so that the event rate decreases over time but not as a clear result of any particular event(s). This may justify using a Poisson model with a trend in the event rate, as considered in detail in Chapter 7.

There is a length of time when the reactor is down after a reactor trip when an initiating event cannot possibly occur. This does not violate Assumption 3 because during that time the plant has dropped out of the study. Its shutdown hours are not counted in the exposure time. Only when the reactor comes up again does it begin contributing hours of exposure time and possible initiating events.

Example 2.2 Shutdown LOSP

Just as with the previous example, consider the three assumptions of the Poisson model. In this case, because data come from the entire industry, λ is interpreted as the average rate for the entire industry.

First consider Assumption 1. The report that analyzed this data (Atwood et al. 1998) found no evidence of a trend in the time period 1980 through 1996. It did find evidence of differences between plants, however. These differences can affect the industry average, because plants enter the study when they start up and leave the study when they are decommissioned. When a plant with an especially high or low event rate enters or leaves the study, this will affect the industry average. However, the event rate at the worst plant differed from the industry average by only a factor of about 3.4, and the best plant differed from the average by less than that. Many plants (116) were considered. Therefore, the effect of a single plant's startup or decommissioning should be small. Therefore, it appears that the overall industry event rate was approximately constant, as required by Assumption 1.

Assumption 2 rules out exactly simultaneous events. In this example, however, events at sister units at a single site are somewhat dependent, because a common cause can result in LOSP events that are simultaneous or nearly simultaneous at both units.

Of the 80 events in the data, two pairs of events occurred together at sister units, each pair from a common cause. Thus, simultaneous events do occur, but they are not frequent. This departure from Assumption 2 is probably not large enough to be serious.

Assumption 3 requires statistical independence of the number of events in disjoint time intervals. As with Example 2.1, there may be some learning, although the lack of trend indicates that any learning is minimal.

In summary, the assumptions for the Poisson model seem to be approximately satisfied.

Example 2.3 Through-Wall Leaks

This differs from the other examples in that the number of events is very small. Any departures from the Poisson assumptions cannot be seen in the data, because so few events have occurred. With no theoretical reason to postulate a trend or other nonconstancy, or a high rate of multiple events, or dependence between events, we accept the Poisson assumptions. The assumptions may not be perfectly true, and a different model may be more accurate, but the Poisson model is simple, and good enough for analyzing such a sparse data set.

Example 2.4 Temperature Sensor/Transmitters

A report by Eide et al. (1999a) divides the total study time for instrumentation failures into two halves, and finds a difference between λ in 1984-1989 and λ in 1990-1995. The example here is for 1990-1995 only. Within this time period the report does not see strong evidence of a trend. That is, a small trend may be present, but the time period is too short, and the failures too few, for any trend to be clear. Further, because the components are regularly maintained, it is reasonable to assume that the failure rate, λ, is roughly constant, as required by Assumption 1.

Assumption 2 requires that common-cause failures be negligible. However, the report states that 14 of the 32 component failures occurred during four common-cause events. Thus, Assumption 2 is seriously violated.

Finally, Assumption 3 requires independence of the number of events in disjoint time intervals. The report does not address this issue, but independence appears plausible.

In summary, the example violates Assumption 2, but probably satisfies the other two assumptions. One way to deal with the violation of Assumption 2 would be to model the independent failures and the common-cause failures separately, although Eide et al. do not do this.

2.2.5 Discussion

2.2.5.1 More General Models

The model considered thus far is a **homogeneous Poisson process (HPP)**, which has a constant event occurrence rate, λ. The number of events in time t is a Poisson random variable with parameter $\mu = \lambda t$. A generalization is a **nonhomogeneous Poisson process (NHPP)**, in which λ is a function of t. Such a model is useful for analyzing trends. Chapter 6 includes ways to test the assumptions of a homogeneous Poisson process, and Chapter 7 includes ways to analyze data where a trend is present.

When data come from the industry, one may consider the differences between plants. Ways to model such differences are discussed in Chapter 8 of this handbook. The present chapter's interest is restricted to λ when no such variation is present. Of course, if the data come from only one plant, λ refers to that plant and the issue of differences typically does not arise.

Any mathematical model, such as the model for a homogeneous Poisson process given here, is an imperfect approximation of the true process that generated the data. Data are used to validate or refute the adequacy of the model. The data set may be sparse — in the present context, this means that the data set contains few events. In this case, two consequences typically result: (a) it is difficult or impossible to see evidence of departures from the model, and (b) the data set contains too little information to allow realistic estimation of the parameters of a more complicated model. If, instead, the data set has many events, departures from the model become visible, and typically a more complicated model is appropriate. These statements have been illustrated by the small and large data sets given as examples.

2.2.5.2 Non-randomness of t

In the model considered here, the exposure time is treated as fixed, and the number of events is treated as random. This is a common type of data found in PRA work. Sometimes, however, a fixed number of events is specified by the data collector, and the corresponding total time is random, as in the following two examples.

One example occurs when equipment is tested until it fails. That is, a predetermined number of items are tested, say x items. Each item is run until it fails, and the total running time of the items is random. The second example occurs in a PRA context if the analyst thinks that the event frequency has changed over time and that only the recent history fully represents current conditions. The analyst may then decide to consider only the most recent events. If there are four recent events, x is fixed at 4, and the corresponding time, measured backwards from the present to the 4th event in the past, is random.

These are examples of **duration** data with exponentially distributed durations, discussed in Section 2.5. The probability model is the Poisson process presented above, but the data collection, and resulting data analysis, are different. Because the time t until the xth event can be called a waiting time, these models are also sometimes called **waiting time models**.

2.3 Failure to Change State

This section considers two probability models, in Subsections 2.3.2 and 2.3.3. First, however, example data sets are given.

2.3.1 Examples

Here are four examples of failure to change state, three with failure to start and one with failure to close.

Example 2.5 HPCI failures to start

At 23 BWRs in the 1987-1993 time period, the high pressure coolant injection (HPCI) system had 59 unplanned attempts to start. The system failed to start on 5 of these demands (Grant et al. 1999a). The failures were typically erratic starts, which the operator stabilized manually. These demands occurred during 113.94 reactor-critical-years.

Example 2.6 EDG failures to start

Emergency diesel generators (EDGs) are sometimes demanded because of unplanned loss of power to a safety bus, and they are also tested periodically, with one set of tests during each operating cycle and another set of tests monthly. In addition, a return-to-service test is normally performed after maintenance of an EDG. At one plant over an 18-month time period, the number of such demands is counted, and the number of failures to start is counted.

Example 2.7 Steam binding in AFW

Between demands, steam binding can develop in the AFW system, so that one or more pumps cannot function when demanded. This is mentioned by Wheeler et al. (1989), and by Nickolaus et al. (1992).

Example 2.8 Failures of isolation valves

Nickolaus et al. (1992) review the causes of about 45 failures of air-operated and motor-operated isolation valves. Some of the principal causes are corrosion, instrument drift, and moisture in instrument and control circuits. Other causes include contamination and corrosion products in the instrument air system, and debris in the system. These are all conditions that can develop while the valves are not being used.

2.3.2 Failure on Demand

All these examples involve a number of demands and a number of failures, where the terms "demand" and "failure" can be defined according to the purposes of the study. Non-PRA contexts provide many other examples of failures on demand. A simple example in elementary probability or statistics courses is tossing a coin n times, and counting the number of heads. Count either a head or a tail as a "failure." Just as in the PRA examples, this example has a number of demands, with a random number of the demands resulting in failures.

2.3.2.1 Probability Model

The standard model for such data assumes that the number of failures has a binomial distribution. The assumptions are listed in Appendix A.6.1. These assumptions can be restated as two assumptions about the physical process and one about the observable data:

1. On each demand, the outcome is a failure with some probability p, and a success with probability $1 - p$. This probability p is the same for all demands.

2. Occurrences of failures for different demands are statistically independent; that is, the probability of a failure on one demand is not affected by what happens on other demands.

The following kind of data can be observed:

- A random number of failures occur during some fixed, prespecified number of demands. As a minimum, the total number of failures and number of demands are observed.

Under these assumptions, the random number of failures, X, in some fixed number of demands, n, has a binomial (n, p) distribution.

$$\Pr(X = x) = \binom{n}{x} p^x (1 - p)^{n-x}, \quad x = 0, \ldots, n \tag{2.2}$$

where

$$\binom{n}{x} = \frac{n!}{x!(n - x)!} .$$

This distribution has two parameters, n and p, of which only the second is unknown. (Although n may not always be known exactly, it is treated as known in this handbook. Lack of perfect knowledge of n, and other uncertainties in the data, are discussed briefly in Section 6.1.3.2)

2.3.2.2 Data Needed to Validate the Model and Estimate p

Suppose that the binomial model is appropriate. Then any reasonable estimator of p needs only two pieces of information: the number of demands, n, in the data period, and the number of failures, x, that then occurred.

However, more information is needed to investigate whether the binomial model is valid. For example, Assumption 1 assumes that p is the same on all demands. If the data cover a number of years or a number of systems or plants, p might not be constant over time or the same at all systems or plants. To study whether this is true, the times and locations of the demands and failures should be recorded, or at least the data should be partitioned into subsets, for example corresponding to systems, plants, or years. Then the failure and demand counts, x_i and n_i, should be given for each subset.

2.3.2.3 Case Studies: Validity of Model Assumptions in Examples

Let us examine Examples 2.5 through 2.8 to see if the assumptions appear to be valid.

Example 2.5 HPCI Failures to Start

Assumption 1 says that the probability of failure on demand is the same for every demand. If data are collected over a long time period, this assumption requires that the failure probability does not change. Likewise, if the data are collected from various plants, the assumption is that p is the same at all plants.

In the HPCI example, the five failures do not reveal any clear trend in time. However, one Licensee Event Report (LER) mentions that a better-designed switch had already been ordered before the HPCI failure. This gives some evidence of a gradual improvement in the HPCI system, which might be visible with more data.

As for differences between plants, it happens that three of the five failures occurred at a single plant. Therefore, it might be wise to analyze that one plant (three failures in nine demands) separately from the rest of the industry (two failures in 50 demands). In fact, Grant et al. (1995) did not analyze the data that way, because they considered two types of failure to start, and they also considered additional data from full system tests performed once per operating cycle. However, the high failure probability for the one plant was recognized in the published analysis.

Assumption 2 says that the outcome of one demand does not influence the outcomes of later demands. Presumably, events at one plant have little effect on events at a different plant. However, the experience of one failure might cause a change in procedures or design that reduces the failure probability on later demands at the same plant. One of the five LERs mentions a permanent corrective action as a result of the HPCI failure, a change of piping to allow faster throttling. This shows some evidence of dependence of later outcomes on an earlier outcome at that plant.

Example 2.6 EDG Failures to Start

Assumption 1 says that every demand has the same probability, p, of failure. This is certainly not true for return-to-service tests, because such tests are guaranteed to result in success. If the EDG does not start on the test, maintenance is resumed and the test is regarded as a part of the maintenance, not as a return-to-service test. Therefore, any return-to-service tests should not be used with the rest of the data.

As for the other demands, one must decide whether the unplanned demands, operating-cycle tests, and monthly tests are similar enough to have the same value of p. Can plant personnel warm up or otherwise prime the diesel before the test? Can an operator stop the test if the EDG is clearly having trouble, and then not consider the event as a test? If so, the different types of demands do not have the same *p*, and they should not be analyzed as one data set. For PRA purposes, one is normally most interested in the failure probability on an actual unplanned demand. To estimate this, one should use only data from unplanned demands and from tests that closely mimic unplanned demands.

If the EDGs in the data set differ in some way, such as having different manufacturers, this may also lead to different values of *p* on different demands. Analyzing the data while ignoring differences between the individual EDGs will allow us to estimate the average *p*, corresponding to failure to start for a random EDG. However, this average *p* is not the same as the *p* for a particular EDG.

Assumption 2 says that the outcome on one demand does not affect the probability of failure on a different demand. When the plant is very new there may be some learning from individual failures, but when the plant is mature, failure or success on one demand should not change the chances of failure or success on later demands. The only way for such dependence to arise is if the first failure results from a common cause. If the plant is mature and common-cause failures are rare, then Assumption 2 is approximately satisfied.

Example 2.7 Steam binding in AFW

Assumption 1 says that every demand corresponds to the same probability of failure. If the steam comes from backflow through a check valve, it will build up, and become more of a problem when the AFW system has been unattended longer. Technically, this is a violation of Assumption 1. However, ignoring the differences between demands results in estimating *p* for an average demand, and this may be adequate for many purposes.

Assumption 2 says that the AFW pumps fail independently of each other. However, steam-binding of the AFW system was a recognized common-cause mechanism in the 1970s and 1980s. This means that Assumption 2 may be plausible if interest is in the performance of a single AFW pump, but not if interest is in an interconnected set of pumps.

Section D-1 of Poloski et al. (1998) says that steam binding has not been seen in 1987-1995 AFW experience. Therefore, Example 2.7 is probably no longer relevant, although it received great attention at one time.

Example 2.8 Failures of isolation valves

The causes of valve failures postulated in this example are degradations, so the probability of failure increases over time, violating Assumption 1. If failures from such causes are rare, then the increase in failure probability may not be a problem. In general, ignoring the differences results in estimating an average *p*, averaged over components that have been allowed to degrade for different amounts of time. This may be acceptable.

As in Example 2.7, some of the mechanisms for valve failure are common causes, violating the independence required by Assumption 2. The seriousness of the violation depends on how many multiple failures occur.

2.3.2.4 Discussion

2.3.2.4.1 More General Models

The model considered above has a constant failure probability, *p*. A generalization would let *p* be a function of time. Such a model is useful for analyzing trends. Chapter 6 includes ways to test the assumptions of the model assumed above, and Chapter 7 includes ways to analyze data where a trend is present.

When data come from the industry, one might consider the differences between plants, just as for initiating events. Ways to model such differences are discussed in Chapter 8. The present section's interest is restricted to *p* for the industry as a whole, the average of all the plants. Of course, if the data come from only one plant, *p* refers to that plant and the issue of differences typically does not arise.

Any mathematical model is an imperfect approximation of the true process that generated the data. When the data set is sparse (few demands, or few or no failures, or few or no successes), (a) it is difficult or impossible to see evidence of departures from the model, and (b) the data set is too small to allow realistic estimation of the parameters of a more complicated model. When the data set has many events, departures from the model become visible, and a more complicated model may be appropriate.

2.3.2.4.2 Non-randomness of *n*

One could argue that the numbers of demands in the examples are not really fixed in advance. That is, no one decided in advance to look at the outcomes of 59 unplanned HPCI demands. Instead, Grant et al. decided to look at seven years of data from 23 plants, and they observed that 59 demands had taken place. The response to this argument is that we are actually conditioning on the number of demands, that is, dealing with conditional probabilities assuming that 59 demands take place. Conditioning on the number of demands enables us to focus on the quantity of interest, p. Treating both the number of failures and the number of demands as random is needlessly complicated, and yields essentially the same conclusions about p as do the simpler methods in this handbook.

In the model considered here, the number of demands is treated as fixed, and the number of failures is treated as random. Sometimes, however, the number of failures is specified in advance and the corresponding number of demands is random. For example, the analyst may believe that p has been changing, and that only the most recent history is relevant. In this case, the analyst might decide to consider only the most recent failures and to treat the corresponding number of demands as random. For example, if only the four most recent failures are included, one would count backwards from the present until $x = 4$ failures were seen in the plant records, and record the corresponding number of demands, n, regarded as an observation of a random variable. This is a **waiting time model**, with n equal to the waiting time until the 4th failure. Bayesian analysis of such data is discussed briefly in Section 6.3.2.6.

2.3.3 Standby Failure

As stated in the introduction to this chapter, failure to change state can be modeled in two ways. One way was given in Section 2.3.2. The second way is given here, in which the system (typically a component) is assumed to transform to the failed state while the system is in standby. This transition occurs at a random time with a constant transition rate. The latent failed condition ensures that the system will fail when it is next demanded, but the condition is not discovered until the next inspection, test, or actual demand.

2.3.3.1 Probability Model

The underlying assumption is that the transition to the failed condition occurs randomly in time. Two settings must be distinguished:

1. the *data*, the operational experiences in the past that allow us to estimate λ, and
2. the *application* to PRA, in which the estimate of λ is used to estimate the probability that a component will fail when demanded.

These two settings are discussed in the next two subsections.

2.3.3.1.1 Probability Model for the Data

It is customary to consider only the simplest model.

1. Assuming that the system is operable at time t, the probability that the system will fail during a short time period from t to $t + \Delta t$ is approximately proportional to the length of the exposure period, Δt. The probability does not depend on the starting time of the period, t, or on anything else.

2. Failures of distinct systems, or of one system during distinct standby periods, are independent of each other.

The kind of observable data is spelled out here. It is obvious, but is written down here for later comparison with the data for similar models.

- At times unrelated to the state of the system, the condition of each system (failed or not) can be observed. As a minimum, the total number of failures and the corresponding total standby time are observed.

The times mentioned here can be scheduled tests or unplanned demands.

Assumption 1 is essentially the same as for a Poisson process in Section 2.2.2. It implies that there is a proportionality constant, λ, satisfying

$$\lambda \Delta t \approx \Pr(t < T \le t + \Delta t \mid T > t),$$

where T is the random time when the system becomes failed. Then the probability that the system is failed when observed at time t is

$$\Pr(\text{system is in failed state at time } t) = 1 - e^{-\lambda t}. \quad (2.3)$$

This follows from Equation 2.6, given in Section 2.5 for the exponential distribution. The parameter λ is called the **standby failure rate**. It is so named because the failed condition develops while the system is in standby, waiting for the next demand.

2.3.3.1.2 Application of the Model to PRA

The model is used to evaluate the probability of failure on an unplanned demand. For this, one assumes that there are periodic tests and the unplanned demand occurs at a random time within the testing cycle. Then the probability of failure on demand is approximated by

$$p \approx \lambda t_{test}/2 \ , \qquad (2.4)$$

where λ is the standby failure rate and t_{test} is the time interval between tests.

A more accurate expression is the average of terms from Equation 2.3, averaging over all the possible demand times in the test interval:

$$p = \frac{1}{t_{test}} \int_0^{t_{test}} \left(1 - e^{-\lambda s}\right) ds$$
$$= 1 - \left(1 - e^{-\lambda t_{test}}\right) / \left(\lambda t_{test}\right)$$

This equation is approximated by Equation 2.4, as can be verified by use of the second-order Taylor expansion:

$$\exp(-\lambda t) \approx 1 + (-\lambda t) + (-\lambda t)^2/2! \, .$$

When more than one system is considered, the formulas become more complicated. For example, suppose that two systems (such as two pumps) are tested periodically and at essentially the same time. Suppose that we are interested in the event that both fail on an unplanned demand. This is:

$$\begin{aligned} &\Pr(\text{both fail}) \\ &= \frac{1}{t_{test}} \int_0^{t_{test}} (1 - e^{-\lambda s})^2 ds \qquad (2.5) \\ &\approx (\lambda t_{test})^2 / 3 \end{aligned}$$

When more systems are involved, or when testing is staggered, the same ideas can be applied.

2.3.3.2 Data Needed to Validate the Model and Estimate λ

Suppose that the standby failure rate model holds. If the standby times are all similar, then an estimator of λ needs only two pieces of information: the number of failures, x, in the data period, and the corresponding total standby time, t. If, instead, the standby times vary substantially, then the total standby times should be recorded separately for the failures and the successes, as explained in Section 6.4.

To validate the model, the data could be partitioned. As with initiating events, if the data come from various years or plants, the data could be partitioned by year and/or by plant, and the above information should be given for each subset.

2.3.3.3 Case Studies: Validity of Model Assumptions in Examples

Let us now examine the applicability of the model assumptions in the examples given above. Much of the discussion in Section 2.3.2.3 applies here as well. In particular, when Section 2.3.2.3 sees a violation of an assumption and suggests a remedy, an analogous violation is probably present here, with an analogous remedy.

Example 2.5 HPCI Failures to Start

Assumption 1 says that the probability of becoming failed in a short time period is proportional to the length of the time period, and on nothing else. As discussed in Section 2.3.2.3, there is no clear evidence of a trend in time. It may be, however, that the probability of failure is higher at one plant than at the other plants. If true, this would violate Assumption 1, and suggests that the outlying plant be analyzed separately from the others.

Assumption 2 says that failures in distinct time periods and locations are independent of each other. As discussed in Section 2.3.2.3, there may be a very small amount of learning, causing fewer failures later in the history.

Example 2.6 EDG Failures to Start

Assumption 1 says that the probability of becoming failed in a short time period is proportional to the length of the time period, and on nothing else. Section 2.3.2.3 discusses different types of tests of EDGs. That discussion is applicable here as well. If an EDG fails on one type of test more readily than on another type of test, Assumption 1 is violated. Another interpretation of this situation is that the bulleted assumption on the data is false: it is not true that a failed condition is always discovered on a test. Some tests discover only major failed conditions while other, more demanding tests discover less obvious failed conditions. Just as mentioned in Section 2.3.2.3, if the primary interest is the probability of failure on an unplanned demand then one should use only data from unplanned demands and from tests that closely mimic unplanned demands.

Assumption 2 says that failures in distinct time periods and locations are independent of each other. As discussed in Section 2.3.2.3, this is probably true if the plant is mature and if common-cause failures are rare.

Example 2.7 Steam Binding in AFW

Assumption 1 says that the failed-condition event is as likely to hit the system in one time interval as in another of the same length. As discussed in Section 2.3.2.3, steam binding can result from a gradual buildup, and become more of a problem when the AFW system has been unattended longer. In this case, Assumption 1 is violated. Ignoring this fact is equivalent to treating the average of AFW conditions.

As discussed in Section 2.3.2.3, steam binding is a common-cause mechanism. Therefore Assumption 2, independence of distinct AFW pumps, is violated.

Example 2.8 Failures of Isolation Valves

Just as discussed in Section 2.3.2.3, the causes listed for Example 2.3 are degradations, violating Assumption 1. However, it may be acceptable to ignore the changes over time, and estimation of an average parameter λ. Also, as discussed in Section 2.3.2.3, some of the mechanisms for valve failure are common causes, violating the independence required by Assumption 2. The seriousness of the violation depends on how many multiple failures occur.

2.3.4 Comparison of the Two Models for Failure to Change State

Two models have been presented for failure to change state, the failure-on-demand model and the standby-failure model. Several aspects of the models are compared here.

2.3.4.1 Ease of Estimation

One great appeal of the standby-failure model is that the analyst does not need knowledge of the number of demands. Standby time is normally much easier to obtain than a count of demands.

2.3.4.2 Use in PRA Cut Sets

The two models differ in their application to cut sets in a PRA model. Consider failure of two redundant components, each having the same probability of failure. When the failure-on-demand model is used, we have

$$\Pr(\text{both fail}) = p^2 = [\Pr(\text{one fails})]^2.$$

On the other hand, when the standby-failure model is used and the two components are tested periodically at the same time, with time t between tests, Equations 2.4 and 2.5 show that

$$\Pr(\text{one fails}) \approx \lambda t_{test}/2$$

$$\Pr(\text{both fail}) \approx (\lambda t_{test})^2/3$$

so that

$$\Pr(\text{both fail}) \neq [\Pr(\text{one fails})]^2.$$

This fact is often ignored.

2.3.4.3 Estimates Obtained

The two models can produce different estimates of basic event probabilities. For example, suppose that an EDG is tested monthly by starting it. In 100 monthly tests, 2 failures have been seen. A simple estimate of p, the probability of failure on demand, is $2/100 = 0.02$. A simple estimate of λ, the standby failure rate, is 0.02/month. Now suppose that a basic event in a PRA is that the EDG fails to start, when demanded at a random time. Based on the estimate of p, the estimated probability of the basic event is

$$\Pr(\text{EDG fails to start}) = p \approx 0.02\ .$$

Based on the estimate of λ and Equation 2.4, the estimated probability of the basic event is

$$\begin{aligned}\Pr(\text{EDG fails to start}) &\approx \lambda t/2 \\ &\approx (0.02/\text{month})\times(1\ \text{month})/2 = 0.01\ .\end{aligned}$$

The two models give estimates that differ by a factor of two, with the failure-on-demand model being more pessimistic than the standby-failure model. The reason is simple: All, or virtually all, of the failures and demands in the data occur at the end of test intervals. However, unplanned demands might occur at any time between tests. The standby-failure model says that demands soon after a successful test have smaller probability of failure. The failure-on-demand model says that all demands have the same probability of failure.

The differences can be more extreme. For example, suppose that two EDGs are tested monthly, and tested at essentially the same time rather than in a staggered way. According to the failure-on-demand model, the probability that both EDGs fail to start is p^2, which is estimated by $(0.02)^2$. On the other hand, according to the standby-failure model, Equation (2.5) shows that the same probability is approximately $(\lambda t_{test})^2/3$, which is estimated by $(0.02)^2/3$. The two models give estimates that differ by a factor of three. More extreme examples can be constructed.

It might be mentioned that these numerical differences between estimates disappear if only unplanned demands are used in the data. However, unplanned demands are rare, and so most analysts prefer to use test data if possible.

2.3.4.4 A Model Involving Both Terms

The model described next postulates two reasons for the observed randomness of failures.

One reason for the randomness of failures is that demands are not all equally stressful. When a demand occurs that is unusually harsh, the system will fail. From the viewpoint of an outside observer, it appears that failures just occur randomly with some probability p, but the underlying cause is the variability in the severity of the demands.

The other reason for randomness of the failures is that the unattended system degrades, and becomes inoperable at unpredictable times. This is simplified in the standby-failure model by supposing that the system changes suddenly from perfectly operable to completely failed, with these transitions occurring at random times. This leads to the standby-failure model, with failure-transition rate λ, and with probability of failure λt at time t after the last system restoration.
If just one of the two mechanisms described above is considered, we are led to either the failure-on-demand model or the standby-failure model. It is possible, however, to construct a model that involves both terms, corresponding to the two kinds of variation. In this two-parameter model, the probability of failure is $p + \lambda t$ at time t after the last system restoration. (For example, see Section 5.2.10 of Samanta et al. 1994.)

Lofgren and Thaggard (1992) state "it is virtually impossible to directly determine from work maintenance record descriptions whether the component has failed from standby or demand stress causes." However, they look for patterns in data from EDGs and motor-operated valves (MOVs) at a small number of plants that use different test intervals. Their data suggest that the standby-failure-rate model is most appropriate for MOV failures, and the two-parameter model is best for EDGs.

In a similar spirit, the T-Book (TUD Office and Pörn Consulting, 2000) uses the two-parameter model for many components. The T-Book does not attempt to identify which mechanism applies to which failures, but instead estimates the two parameters from overall patterns in the data. Some of the resulting estimates have large uncertainties; for example, at a typical plant the estimate of p for EDG failure to start has an error factor of about 13. For components that cannot be analyzed in this way, the T-Book uses the standby-failure model. For details, see Pörn (1990).

2.3.4.5 Choosing a Model

No consensus exists among PRA workers as to which model is most advantageous. In particular, the typical mechanisms of failure are not understood well enough to justify a theoretical basis for a model. Most current work uses one of the two simple models given here: failure on demand or standby failure. Therefore, this handbook presents only these two models. The user may choose between them.

2.4 Failure to Run during Mission

Aspects of this type of failure closely resemble the initiating events of Section 2.2. One important difference is in the kind of data normally present. The difference is summarized here.

Example 2.4 of Section 2.2 is an example of continuously running components (temperature sensor/transmitters) that occasionally failed to run. When a component failed, it was repaired or replaced in a relatively short time, and resumed operation. That is, the component was **repairable**. The present section considers components or systems that do not run continuously. Instead, they are occasionally demanded to start, and then to run for some mission time. If they fail during the mission, they are **nonrepairable**, that is, they cannot be repaired or replaced quickly. Two points deserve clarification:

- Some failures may be recoverable. They would not be modeled as failures in the sense of causing mission failure. Unrecoverable failures cause mission failure, however.

- Given enough time, almost any system can be repaired. During a mission, however, time is not available. Because the component or system cannot be repaired *within the time constraints*, it is called "nonrepairable."

As stated earlier, the word **system** is used in this handbook for any piece of hardware for which data are taken. In particular, components and trains are kinds of systems.

2.4.1 Examples

Here are two examples of failures to run during missions.

Example 2.9 EDG failures to run

Grant et al. (1999b) report that in 844 demands of 30 minutes or more for EDGs to run, there were approximately 11 unrecovered failures to run in the first 30 minutes. The count is approximate because a few failure times were not given and had to be inferred.

Example 2.10 AFW turbine train failures to run

Poloski et al. (1998) report that in 583 unplanned demands of AFW system turbine trains, the train failed to run 2 times, and the total running time was 371 train-hours. The information is taken from LERs, only 17% of which report running times for the train. The total running time of 371 hours is an extrapolation from the LERs with reported run times.

These examples are typical, in that hardly any of the demands to run resulted in a failure. Therefore, for most demands the time when failure would eventually have occurred is unknown.

2.4.2 Probability Model

In principle, the times to failure are **durations**. Section 2.5 deals with duration data, in the context of recovery times. That section mentions various possible distributions of time to failure, of which the simplest is the exponential distribution.

Data for this section differ from data of Section 2.5, however, because nearly all of the observed times in this section are truncated before failure. This is illustrated by the above examples. Therefore, the full distribution of the time to failure cannot be observed. In Example 2.9, no information is given about the distribution of failures times after the first 30 minutes. In Example 2.10, the average run time was only 38 minutes, and most AFW missions lasted for less than one hour. In such cases the exponential distribution, restricted to the observed time period, is a simple, reasonable approximation of the observable portion of the distribution.

Two assumptions are made concerning the physical process:

1. Assuming that no failure has occurred by time t, the probability that a failure will occur in a short time period t to $t + \Delta t$ is approximately proportional to the length of the exposure period, Δt. The probability does not depend on the starting time of the period, t, or on anything else.

2. Failures of distinct systems, or of one system during distinct missions, are independent of each other.

The kind of observable data is as follows:

- For each observed mission, the run time is observable. Also, it is known whether the run terminated in failure or in successful completion of the mission. As a minimum, the total run time and the number of failures to run are observed.

Assumption 1 implies that the time to failure is exponentially distributed with parameter λ. The interpretation of λ is that if the system is running, the probability of failure in the next short interval of length Δt is approximately $\lambda\Delta t$. That is

$$\lambda\Delta t \approx \Pr(t < T \le t + \Delta t \mid T > t),$$

where T is the random time until failure. When defined this way, λ is sometimes called the **failure rate**, or **rate of failure to run**. Many authors use the term **hazard rate**, denoted by h, and discussed in Appendix A.4.4. Note, the definition of λ is different for repairable systems (Section 2.2) and nonrepairable systems (the present section), even though it is represented by the same Greek letter and is called "failure rate" in both cases. See Thompson (1981) for a reasonably clear discussion of the subtle differences, and the glossary of this handbook for a summary of the definitions. The topic is discussed further in Appendix A.4.4.

It is instructive to compare the models for failure to run and standby failure. The physical process is essentially identical, but the observable data differs in the two models. That is, Assumptions 1 and 2 in the two sections agree except for small differences of wording. However, the time of failure to run is observable, whereas the time of transition to a standby failure is never known.

It may also be somewhat instructive to compare the Assumptions 1 and 2 here with the Assumptions 1-3 of the Poisson process in Section 2.2.2. For the standby-failure model and the failure-to-run model, Assumptions 1 and 2 do not explicitly include an assumption ruling out simultaneous failures. The reason is that simultaneous failures are ruled out by the other two assumptions: it is not meaningful for a system to fail twice simultaneously; and distinct systems are assumed to fail independently of each other, and therefore not exactly simultaneously.

2.4.3 Data Needed to Validate the Model and Estimate λ

Suppose that the time to failure has an exponential distribution. Then, any reasonable estimator of λ needs only two pieces of information: the total running time, t, in the data period, and the number of failures to run, x, that occurred then.

However, more information is needed to investigate whether the exponential distribution is valid. Assumption 1 says that λ is constant during the mission. To investigate this, the analyst should know the failure times, that is, how long the failed pumps ran before failing. The analyst should also know the mission times, that is, how long the system ran when it did not fail; often, however, this information is not recorded and can only be estimated or approximated.

Implicit in Assumption 1 is that λ is the same over all the years of data, at all the plants where the data were collected. To investigate this, the data should be divided into subsets, corresponding to the different plants and years. Then the failure count and running time, x_i and t_i, should be given for each subset. This is the exact analogue of what was said in Section 2.2.3 for initiating events.

2.4.4 Case Studies: Validity of Model Assumptions in Examples

Consider now whether the assumptions of the model are plausible for the two examples.

Example 2.9 EDG Failures to Run

Assumption 1 says that a running EDG is as likely to fail in one short time interval as in any other time interval of the same length. That is, the EDG does not experience burn-in or wear-out failures. The reference report (Grant et al. 1999b) says that this is not true over a 24-hr mission. Indeed, that report divides the EDG mission into three time periods (first half hour, from one-half hour to 14 hours, and from 14 to 24 hours) to account for different failure rates during different time periods. Within the first half hour, however, the data do not give reason for believing that any short time interval is more likely to have a failure than any other time interval. Therefore, Assumption 1 can be accepted.

Assumption 2 is violated by common-cause failures. It is also violated if a failure's root cause is incorrectly diagnosed, and persists on the next demand. If these two conditions are rare the assumption may be an adequate approximation. More subtle dependencies are difficult to detect from data.

Example 2.10 AFW Turbine Train Failures to Run

Assumption 1 says that a running turbine train is as likely to fail in one short time interval as in any other time interval of the same length. The data are too sparse – only 2 observed failures – to confirm or refute this assumption. The data are also too sparse to confirm or refute Assumption 2, although failures in separate plants are virtually certain to be independent. In such a situation, it is common to accept the simple model as adequate. A more complicated model is justified only when a larger data set is available.

2.4.5 Discussion

The exponential time to failure can also be derived as the time to *first* failure in a Poisson process of Section 2.2. This is possible because the time to first failure and the times between subsequent failures are all exponentially distributed when the failures follow a Poisson process. The present context is simpler, however, because the process ends after the first event, failure to run. The Poisson-process assumptions about hypothetical additional failures are irrelevant.

2.5 Recovery Times and Other Random Duration Times

This section is about modeling of time data. Often, a measurement of interest is a random duration time, such as the time required to return a failed system to service or the lifetime of a piece of hardware. The distinction between random duration times here and events in time in Sections 2.2 and 2.4 is that here the individual times are measured on a continuous scale with units such as minutes or hours, while the earlier data sets involve discrete counts of the number of events occurring in a total length of time.

2.5.1 Examples

Here are some examples involving random duration times. They are only summarized here. Actual examples, with lists of durations times, will be analyzed in Chapter 6.

Example 2.11 Recovery times from loss of offsite power

A plant occasionally loses offsite power. When this happens, the plant reports the time until power is restored. Atwood et al. (1998) present such durations for LOSP events in 1980-1996.

Example 2.12 Repair times for turbine-driven pumps

A turbine-driven pump must occasionally be taken out of service for unplanned maintenance. The duration of time out of service for maintenance may be extractable from maintenance records.

Example 2.13 Time to failure of a component

A typical power plant will have many individual components such as compressors. When a component is put into service, it operates intermittently until it fails to perform its required function for some reason. Høyland and Rausand (1994) give an example of such data.

Example 2.14 Times to suppress fires

When a fire occurs in a nuclear power plant, the time until the fire is suppressed is of interest. Nowlen et al. (2002) report on analysis of such suppression times. One difficulty is that the time of fire onset often is not exactly known.

Example 2.15 Gradual degradation until failure

Examples 2.7 (steam binding) and 2.8 (failure of isolation valves) involve gradual degradation, which builds up until the system is inoperable. The time until the system is inoperable can be modeled as a duration time.

The common element in these examples is a duration time that varies in an unpredictable way. In Examples 2.11 and 2.12, the recovery time is composed of several factors such as the time to diagnose, perform and test repairs, and the time to complete documentation required before returning the plant to normal operating conditions. Example 2.13 is a failure-to-run example, similar to those of Section 2.4. This example differs from that of Section 2.4, however, because here it is assumed that virtually all of the times to failure are recorded. In Section 2.4, on the other hand, most of the systems did not fail during the test period or operational mission. The severe truncation of the data in Section 2.4 meant that only a simple model could be considered. The more complete data here allows analysis of a more complex model. Example 2.14 is complicated by the lack of exact knowledge of the duration time. Finally, Example 2.15 gives a realistic conceptual way to model the gradual degradations encountered in Section 2.3.1, although good data are unobtainable.

All five examples involve a duration time that is uncertain due to random factors. Consequently, the duration times are modeled as continuous random variables.

2.5.2 Duration-Time Models

The duration, *T*, is random, following some probability distribution. Two assumptions are made about the process:

1. Each duration is statistically independent of the others, and

2. All the random durations come from the same probability distribution.

The data description is simple:

- The individual durations are observable. As a bare minimum, the number of durations and the total duration time are observed.

Assumptions 1 and 2 can be summarized by saying that the durations are **independent** and **identically distributed**. Independence means that one duration does not influence the probability of any other duration. The assumption of identical distributions means that each random duration is as likely as any other to be long or short. If the durations are from distinct systems, the systems are assumed to be identical and to act independently. If the durations are in sequence, as for a system that alternates being up and down, the assumption implies that no learning or long-term aging takes place, and that each repair restores the system to a condition as good as new. Such a process is called a **renewal process**.

The assumptions do not require a particular distribution for the time between events. The most important such distributions in PRA applications are:

- lognormal,
- exponential,
- Weibull, and
- gamma.

These distributions are summarized in Appendix A.7. An important part of the data analysis consists of deciding on the form (or several plausible forms) of the distribution. This will be discussed in Chapter 6. For now, we simply note that these and other distributions are possible.

There are different ways to specify a probability distribution, and the next material summarizes some of the concepts: their definitions, how to interpret them, and how they are related to each other. The data-analysis techniques of Chapter 6 will use these ways of characterizing distributions. The usual convention is to denote the random variables using capital letters, T, and observed times as lower case, t. The letter T is used, rather than some other letter such as X, because the random quantities are times. As seen from the examples, the durations may be times to repair, times to failure, or other times. However, the concepts and formulas are valid for any application.

The **cumulative distribution function** (c.d.f.) of a real-valued random variable T is defined as

$$F(t) = \Pr(T \le t)$$

for all real numbers t. The name is sometimes abbreviated to **distribution function**. The c.d.f. is the probability that the random variable T will assume a value that is less than or equal to t. The c.d.f. is a monotonically increasing function of t, with the limiting properties $F(0) = 0$ and $F(+\infty) = 1$. [For random variables that, unlike durations, can take negative values, the limiting properties are $F(-\infty) = 0$ and $F(+\infty) = 1$. That general case has few applications in this handbook.]

The distribution is commonly used to characterize the lifetimes, or recovery times, or some other kind of durations, of a whole population of systems. The population might be a large set of identical systems that are operating in similar applications and with durations that vary due to random influences. $F(t)$ is the fraction of items that have durations t or less, in a hypothetical infinite population.

A related function, denoted by $f(t)$, is called a **probability density function** (p.d.f.) for a continuously distributed positive-valued random variable T. It is related to the c.d.f. by

$$f(t) = \frac{d}{dt}F(t) \quad \text{and}$$

$$F(t) = \int_0^t f(u)du \quad .$$

The variable u is a dummy variable of integration, and t is the upper limit of the integral. An example of a p.d.f. and the associated c.d.f. are shown in Figure 2.1.

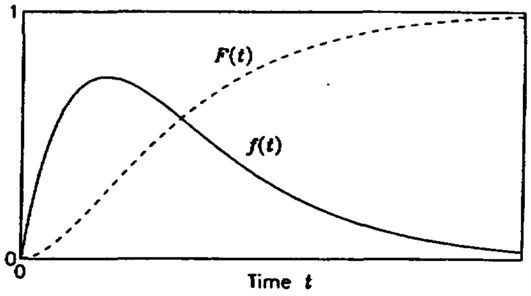

Figure 2.1 Probability density function (p.d.f.) and cumulative distribution function (c.d.f.).

It follows that probabilities corresponding to occurrences in a small interval of time are approximately proportional to the p.d.f.,

$$\Pr(t < T \le t + \Delta t) \approx f(t)\Delta t.$$

Therefore, the ordinate of a p.d.f. has units of "probability density" and not probability (as for a c.d.f.). Thus, a p.d.f. determines how to assign probability over small intervals of time. Now consider an arbitrary interval from a to b. In this case we have

$$\Pr(a < T \le b) = \int_a^b f(t)dt \quad .$$

The simplest distribution is the **exponential** distribution. It arises when Assumption 1 of Section 2.4.2 is satisfied. (That assumption is phrased as if T is a time until failure.) In that case, the probability distribution is exponential, and determined by a single parameter, λ. The p.d.f. and c.d.f. are given by

$$f(t) = \lambda e^{-\lambda t}$$

$$F(t) = 1 - e^{-\lambda t} . \qquad (2.6)$$

When deriving the distribution mathematically from Assumption 1, it is necessary to assume that $F(0) = 0$, that is, failures at time 0 have zero probability. Although not stated explicitly, this assumption is implicit in the context of failure to run, because any failures at time 0 would be counted as failures to start, not failures to run.

2.5.3 Data Needed to Estimate Distribution of Durations and Validate Model

In general, a sample of observed durations is needed to estimate the distribution of duration times. These durations must independent and identically distributed, that is, they must be generated by a process satisfying the two assumptions given at the beginning of Section 2.5.2.

The special case when the times are assumed to have an exponential (λ) distribution is simpler. Only the number of durations and the total duration time are needed to estimate λ,. However, the individual durations are still needed to investigate whether the distribution is exponential or of some other form. Incidentally, when the distribution is assumed to be exponential, the model given here differs from the standby-failure model (Section 2.3.3.1.1) and from the failure-to-run model (Section 2.4.2) *only* by the kind of data that can be observed.

To validate whether the distribution is the same for all the data, extra information should be recorded for each duration, the relevant circumstances of each duration. The circumstances of interest are those that might affect the durations, such as time of the event, system location, and system condition just before the event.

2.5.4 Case Studies: Validity of Model Assumptions in the Examples

Examples 2.11 through 2.13 all appear to satisfy the assumptions of Section 2.5.2. Example 2.14 also does, except that the durations are not observed exactly.

In each case, all the distributions come from some distribution. Discovering the form of that distribution is a task for the data analyst.

One might ask whether the durations are statistically independent. For example, does a long repair time for a turbine-driven pump add an extra benefit to the pump, so that the next few repair times will be short?

One might also ask, for each example, whether the durations all come from the same probability distribution. For example, if the data cover a period of years, has there been any long-term learning, so that recovery times or repair times tend to be shorter than at the start of the data period? Are different durations associated with different systems for the turbine-driven pumps, with different causes of loss of offsite power, or with different kinds of fires?

The above are questions that could be investigated during the data analysis, if enough durations have been observed.

Example 2.14 is complicated by lack of exact measurements of the durations. Bounds can be given, and the analysis must be based on these upper and lower bounds rather than on exact times.

Example 2.15 is different because the durations are not observable at all. It might be theoretically interesting to model the time until the system is in a failed condition as a duration, but there is no monitor on the pump or valve that says, "At this time the system just became inoperable." Therefore, the durations are not directly observable, not even in principle. Therefore, the methods of this handbook are not applicable to this example.

Fortunately, degradation mechanisms have become minor contributors to risk. When a degradation mechanism is recognized as important, the natural response is not to collect data to better estimate the rate of degradation. Instead, the natural response is (a) to shorten the interval between preventive maintenance activities, and so to identify and correct incipient degradation, or (b) to modify the plant to mitigate or eliminate the problem. Examples are the apparent elimination of steam-binding in AFW pumps, mentioned above, and of intergranular stress corrosion cracking (IGSCC) in BWR piping (Poloski et al. 1999a, Appendix J).

2.6 Unavailability

This section considers **test-and-maintenance unavailability**, corresponding to intentional removal of the equipment from service for testing and/or maintenance. This section does not consider unavailability resulting from the hardware being in an unrecognized failed condition; that topic was treated in Section 2.3.3.

The discussion here is presented in terms of trains, although other hardware configurations, such as individual components, could be considered equally well. A standby train, such as the single train of the HPCI system or a motor-driven train of the AFW system, is normally available if it should be demanded, but sometimes it is out of service for planned or unplanned maintenance. The event of a train being unavailable is called an **outage**, and the length of time when it is unavailable is called an **outage time** or **out-of-service time**. In a data set, the **exposure time** is the time (e.g. number of hours) when the train should have been available. The **unavailability** is the long-term ratio of outage time to exposure time – the fraction of time that the system is out of service when it should be available. More precisely, the **planned-maintenance unavailability** is the fraction of time that the system is out of service for planned testing and maintenance, and the **unplanned-maintenance unavailability** is defined similarly. In summary, outage times are random but the unavailability is a parameter, an unknown constant, denoted here by q. Subscripts such as "planned" and "unplanned" can be attached to q for clarity if needed.

2.6.1 Example

Example 2.16 CVC unavailability for test and maintenance

Train outages of various durations occurred during 15 calendar months at a plant with two trains in the chemical and volume control (CVC) system. For each month, the outage durations are given by Atwood and Engelhardt (2003).

A way to picture the status of a standby train or other repairable system uses a **state variable**, defined as $S(t) = 1$ if the system is up at time t, and $S(t) = 0$ if it is down at time t. A particular system history is illustrated in Figure 2.2, from Engelhardt (1996). This figure shows when a particular system was operating ($S = 1$) or shut down ($S = 0$). A nominally identical system would have a somewhat different history for the same period, or the same system would have a different history over a different time period of the same length.

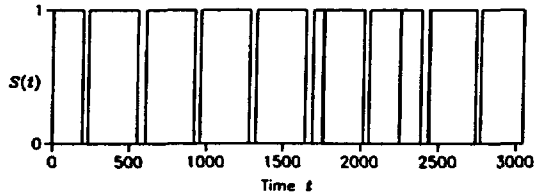

Figure 2.2 Uptime and downtime status for one system.

As stated above, the long-term fraction of time when the system is down is called the system **unavailability**.

2.6.2 Probability Model

The assumed underlying model is an **alternating renewal process**. At any point in time a system is in one of two states: "up" or "down," corresponding in our application to being available or out of service. Initially, the system is up, and it remains up for a random time Y_1; it then goes down, and stays down for a random time Z_1. Then it then goes up for a time Y_2, and then down for a time Z_2, and so forth. The assumptions needed for the data analysis methods in Chapter 6 are the following:

1. The random variables Y_i have one distribution that is continuous with a finite mean, and so do the random variables Z_i.
2. All the random variables are independent of each other

The sum of the down times, ΣZ_i, is the total outage time in the data. The sum of all the times, $\Sigma Y_i + \Sigma Z_i$, is the exposure time – the time when the system should be available. Time when the system is not required to be available is not counted in either the up time or the down time.

Two kinds of data can be considered:

- **Detailed data**: the onset time and duration of each individual outage are recorded, as well as the total time when the train should have been available; and
- **Summary data.** Data totals are given for "reporting periods," such as calendar months. For each reporting period, the total outage time and exposure time are recorded.

Section 6.7 describes how to analyze both types of data.

2.6.3 Data Needed to Validate the Model and Estimate q

The unavailability, q, can be estimated from either kind of data. Enough data should be collected so that any periodic, lengthy, planned outages are appropriately represented – neither over-represented nor under-represented.

In addition, if summary data are used, the methods given in Chapter 6 combine reporting periods into larger subsets of the data, at the very least so that the aggregated subsets do not contain outage times of zero. Therefore, a large enough set of summary data is needed so that it consists of at least two (as a bare minimum) subsets of approximately equal exposure time, with each subset containing nonzero outage time.

To validate the model, any information that might be related to unavailability should be recorded. For example, if a motor-driven pump has most of its scheduled maintenance during the plant's refueling outages, and the pump's availability during shutdown is of interest, then the data should indicate which outages and exposure times correspond to reactor shutdown. Separate analyses will probably need to be performed for the time when the reactor is up and when the reactor is down, to keep Assumption 1 from being violated.

2.6.4 Case Study: Validity of Model Assumptions in Example

The ideas here are applicable to virtually any system, with Example 2.16 being just one example.

The trains may undergo periodic, infrequent, lengthy testing and maintenance, and less lengthy testing and maintenance at more frequent intervals. This periodicity of planned maintenance means that Assumption 2 cannot be exactly true. The lengthiest outages tend to be evenly spaced, not random as assumed. However, more realistic assumptions would be very difficult to work with.

It seems plausible that this deterministic periodicity should lead to conservative estimates. That is, analysis methods that assume pure randomness will tend to overestimate the variance, so that the resulting uncertainty in q is overestimated. However, this conjecture has not been carefully investigated, and the 15 months of data in Example 2.16, analyzed in Section 6.7, do not support the conjecture.

Assumption 1, on the other hand, is surely correct. The distributions are continuous, and it is inconceivable that the durations for an operating power plant would have infinite means.