Chapter 7: Software Reliability Measurement Experience

Allen P. Nikora (Jet Propulsion Laboratory)

Michael R. Lyu (Bellcore)

7.1 Introduction

The key components in the SRE process, as described in Chapter 6, include reliability objective specification, operational profile determination, reliability modeling and measurement, and reliability validation. These techniques are currently applied to several internal projects developed within Jet Propulsion Laboratory (JPL) and Bellcore. The project background, reliability engineering procedures, data collection efforts, modeling results, data analyses, and reliability measurements for these projects are presented in this chapter. Model comparisons for the software reliability applications, lessons learned with regard to the engineering effort, and directions for current and future software reliability investigations are also provided.

One major thing we observed is that for the failure data we analyzed, no one model was consistently the best. It was frequently the case that a model that had performed well for one set of failure data would perform badly for a different set. We therefore recommend that for any development effort, several models, each making different assumptions about the testing and debugging process, be simultaneously applied to the failure data. We also recommend that each model's applicability to the failure data be continuously monitored. Traditional goodness-of-fit tests, such as the Chi-Square or Kolmogorov-Smirnov tests, can be used. In addition, the model evaluation criteria described in Chapter 4 are also strongly recommended.

Another discovery is that of the software development efforts we studied, few had quantitative reliability requirements that were measurable. Strictly speaking, it is not necessary to have a reliability requirement for system in order to apply software reliability measurement techniques. It

is quite possible to measure a software system's reliability during test and make predictions of future behavior. However, the existence of a requirement is very helpful in that:

1. Specifying a reliability requirement helps the users and developers focus on the components of the system that will have the most effect on the system's overall reliability. Potentially unreliable components can be respecified or redesigned to increase their reliability.

2. A reliability requirement will serve as a goal to be achieved during the development effort. During the testing phases, software developers and managers can estimate software reliability and determine how close it is to the required value. The difference between current and required reliability can be converted into estimates of the time and resources that will be required to achieve the goal.

We also discovered that one of the most important aspects in an SRE program is identifying the data to be collected and setting up mechanisms to ensure that the data collected is complete and accurate. We found that development organizations generally have the capability to collect the type of data that is required to use software SRE techniques. Every software development effort that we studied has a mechanism for recording and tracking failures that are observed during the testing phases and during operations. Most projects also have requirements for the test staff to keep an activity log during the testing phases. Properly used, these data collection mechanisms would provide accurate failure data in a form that could easily be used by many currently-available software reliability models. However, since many software managers and developers are not aware of the types of analysis that can be done with this data, they do not devote the effort required to ensure that the collected data is complete and accurate.

Finally, we discovered that a properly-defined linear combination of model results produced more accurate predictions over the set of failure data that we analyzed than any one individual model[Lyu92a]. This linear combination modeling scheme is discussed in detail.

7.2 Measurement Framework

To enhance a company's ability to deliver timely, high-quality products through an application of SRE practices, as well as to help ensure that software vendors to deliver high-quality component products, several elements are included in our investigation. Figure 7.1 shows an SRE framework in our current practice. You can see that this framework is similar to that displayed in Figure 6.1; however, it is more focused on the product life-cycle phases during system test and post delivery.

First, customer usage is quantified by developing an operational profile. Second, quality is defined quantitatively from the customer's viewpoint by defining failures and failure severities, by determining a reliability objective, and by specifying balance among key quality objectives (e.g., reliability, delivery date, cost, etc.) to maximize customer satisfaction. We then advocate the employment of operational profile and quality objectives to manage resources and to guide design, implementation, and testing of software. Moreover, we track reliability during testing to determine product release. This activity may be repeated until a certain reliability level has been achieved. We also analyze reliability in the field to validate the reliability engineering effort and to introduce product and process improvements.

It can be seen from Figure 7.1 that there are four major components in this SRE process, namely, (1) *reliability objective*, (2) *operational profile*, (3) *reliability modeling and measurement*, and (4) *reliability validation*. A reliability objective is the specification of the reliability goal of a product from the viewpoint of the customer. If a reliability objective has been specified by the customer, that reliability objective should be used. Otherwise, you can select a reliability measure which is most intuitive and easily understood, and then determine the customer's "tolerance threshold" for system failures in terms of this reliability measure. For example, customer A might be mostly concerned with the total number of field failures product X may produce. Therefore, the reliability objective could be specified as, say, "product X should not produce more than 10 failures in its first 50 months of operation by customer A."

Operational profile concepts and technique are described in Chapter 5. It is a set of disjoint alternatives of system operation and their associated probabilities of occurrence. The construction of an operational profile encourages testers to select test cases according to the system's operational usage, which contributes to more accurate estimation of software reliability in the field.

Reliability modeling is an essential element of the reliability estimation process. It determines if a product meets its reliability objective and is ready for release. You are required to use a reliability model to calculate, from failure data collected during system testing (such as failure report data and test time), various estimates of a product's reliability as a function of test time. Several interdependent estimates make equivalent statements about a product's reliability. They typically include, as stated in Chapter 1 (Section 1.2.2), the product's failure intensity (failure rate, i.e, the number of failures per unit time) as a function of test time t, the number of failures expected up to test time t, and the mean time to failure (MTTF) at test time t. These reliability estimates can provide the following information useful for product quality management:

(1) The reliability of the product at the end of system testing.

(2) The amount of (additional) test time required to reach the product's reliability objective.

(3) The reliability growth as a result of testing (e.g., the ratio of the value of the failure intensity at the start of testing to the value at the end of testing).

(4) The predicted reliability beyond the system testing already performed. This can be, for example, the product's reliability in the field, if the system testing has already been completed, or the predicted reliability at the end of testing, if the system testing has not yet been completed.

Chapter 3 gives a comprehensive survey on existing reliability models. Despite the existence of more than 40 models, the problem of model selection and application is manageable. Guidelines and statistical methods for selecting an appropriate model for each application are developed in Chapter 4. Furthermore, experience has shown that it is sufficient to consider only a dozen models from among the 40 models, particularly when they are already implemented in software

tools (see Appendix A).

Using these statistical methods, "best" estimates of reliability are obtained during testing. These estimates are then used to project the reliability during field operation in order to determine if the reliability objective has been met. This procedure is an iterative process since more testing will be needed if the objective is not met. When the operational profile is not fully developed, application of a test compression factor can assist in estimating field reliability. A *test compression factor* is defined as the ratio of execution time required in the operational phase to execution time required in the test phase to cover the input space of the program. Since testers during testing are trying to "break" the software by searching through the input space for difficult execution conditions, while users during operation only execute the software at a normal pace, this factor represents the reduction of failure rate (or increase in reliability) during operation with respect to that observed during testing.

Finally, the projected field reliability has to be validated by comparing it with the observed field reliability. This validation not only establishes benchmarks and confidence levels of the reliability estimates, but also provides feedback to the SRE process for process improvement and better parameter tuning. For example, the model validity could be established, the growth of reliability could be determined, and the test compression factor could be refined, etc.

Various components in this SRE framework are discussed in detail below.

7.2.1 Establishing Software Reliability Requirements

Software reliability requirements are specified during earlier development phases, and SRE techniques are used to estimate the resources that will be required to achieve those requirements during test and operations. The resource requirements are translated into testing schedules and budgets. Resource estimates are compared to the resources actually available to make quantitative, rather than qualitative, statements concerning achievement of the reliability requirements.

7.2.1.1 Expressing Software Reliability

Reliability and reliability-related requirements can be expressed in one of the three following ways:

1. Probability of failure-free operation over a specified time interval.
2. MTTF.
3. Expected number of failures per unit time interval (failure intensity).

The first form, the basic definition of software reliability, is a probabilistic statement concerning the software's failure behavior. The other two forms can be considered relating to reliability. Reliability and reliability-related requirements must be stated in quantitative terms. Otherwise, it will not be possible to determine whether the requirements have been met. To help in understanding how to develop these requirements, examples of testable and untestable reliability requirements are given in the following paragraphs.

The following statements, paraphrased from a JPL software development effort, represent a requirement for which software SRE can be used to determine the degree to which that requirement has been met. "Reliability quantifies the ability of the system to perform a required function under the stated conditions for a period of time. Reliability is measured by the MTTF of a critical component. Under the expected operational conditions, documented elsewhere in this requirements document, the probability of the MTTF for the software being greater than or equal to 720 hours shall be 90%."

The above requirement is stated in a testable manner. If the expected operational conditions are stated in terms of the operational hardware configuration and the fraction of time each major functional area is expected to be used (the operational profile), the test staff can then design tests to simulate expected usage patterns and use reliability estimates made during these tests to predict operational reliability.

Confidence bounds should be associated with reliability or reliability-related requirements. If the above MTTF requirement had been stated as being simply 720 hours, it would have been possible to meet that requirement with a very wide confidence interval (e.g. 90% probability of the MTTF lying between 200 and 1240 hours). This could have resulted in the delivery of operational software whose MTTF was considerably less than the intended 720 hours. Yet, the end users of the delivered software would be told that the reliability requirement had been met. Not until the software was actually operated would the users realize the discrepancy. To avoid this problem, express the reliability requirement as the minimum value of the confidence interval. This will allow the end users to know the probability of the software meeting its reliability requirement, and permit them to plan accordingly. It is often needed to specify a tighter confidence interval. The price to pay for this improvement, though, is the need for extra validation effort to establish the tighter confidence interval.

An example of an untestable reliability-related requirement is now given. Again, the text is paraphrased from that found in a JPL development effort's system requirements document. "The system is designed to degrade gracefully in case of failures. As a first priority, system fault protection shall ensure that no system failures or component failures will compromise system integrity. As a second priority, minimum mission science objectives previously described in this document shall not be compromised. Accordingly, each instrument shall be designed so that if one fails (either through hardware or software failures), it will not jeopardize the safety of the system or damage adjacent instruments. This includes provision for isolation from the system via the instrument power supply. If a system fault occurs, the system will automatically stop any science data gathering and go to a safe state. After a safe state is achieved and subsystems are re-initialized, science can be resumed."

The foregoing type of requirement, frequently seen in industry for critical applications, does not provide a basis for measuring the reliability of the system under development, as it contains no quantitative statements concerning the system's failure behavior. Rather, it is a statement of

design constraints that are intended to localize damage resulting from a component failure to the immediate area (e.g., assembly, subsystem) in which the failure occurred. During subsequent phases of system development, it may indeed be possible to determine whether such constraints have been reflected in the system design and implementation. However, this information alone is not sufficient to make quantitative statements concerning the system's reliability. Although specifying constraints such as these is an important aspect of system specification, specific reliability requirements, similar in form to the first reliability requirement discussed in this section, would have to be provided if it were intended to use SRE techniques to determine compliance to a reliability requirement.

7.2.1.2 Specifying Reliability Requirements

To specify reliability requirements, use one or more of the three methods described below. The methods are[Musa87a]:

1. System balance.
2. Release date.
3. Life cycle cost optimization.

It is possible to use one of these methods for developing the requirements for one component of the system, and another for a separate component.

The system balance method is primarily used to allocate reliabilities among components of a system based on the overall reliability requirement for that system. The basic principle of this method is to balance the difficulty of development work on different components of the system. The components having the most severe functional requirements or being the most technologically advanced are assigned less stringent reliability requirements. In this way, the overall reliability requirement for the system is met while minimizing the effort required to implement the most complex components. For software, this might translate to assigning less stringent reliability requirements to functions never before implemented or functions based on untried algorithms.

This approach generally leads to the least costly development effort in the minimum time. The system balance method is frequently used in developing military systems.

The second approach is used when the release date is particularly critical. This is appropriate for flight systems facing a fixed launch time, or commercial systems aiming at delivery within a profit window. The release date is kept fixed in this approach. The reliability requirement is either established by available resources and funds, or is traded off against these items. With this approach, it is desirable to know how failure intensity trades off with release date. First, the way in which the failure intensity trades off with software execution time is determined. This execution time is then converted to calendar time. The following example uses the Goel-Okumoto Exponential Poisson Model (GO) model and Musa-Okumoto Logarithmic Poisson (MO) model for illustrations.

Example 7.1

For the GO model, the relationship between the ratio of failure intensity change during test and the execution time is given by (see Section 3.3.2):

$$\tau = \frac{1}{b} \ln \frac{\lambda_0}{\lambda_F} \tag{7.1}$$

τ = elapsed execution time
λ_0 = initial failure intensity
λ_F = required failure intensity
b = failure intensity decay parameter

This model also has a parameter "a", which specifies the number of failures that would be observed if testing were to continue for an infinite amount of time. For the MO model, the relationship between the ratio of failure intensity change during test and the execution time is given by (see Section 3.5.3):

$$\tau = \frac{1}{\theta \lambda_0} \left(\frac{\lambda_0}{\lambda_F} - 1 \right) \tag{7.2}$$

τ, λ_0, and λ_F as above
θ = failure intensity decay parameter

For this example, the failure history data from one of the testing phases of a JPL flight program (see J3 in Appendix B) is used. Applying the GO and MO models to this data, the following model parameter and failure intensity estimates are obtained:

Goel-Okumoto	Musa-Okumoto
λ_0 = .4403 failures / CPU hour a = 202.52 failures b = 6.044×10^{-7} per failure	λ_0 = .5064 failures / CPU hour $\theta = 1.693\times10^{-6}$ /failure

The above equations can be used to determine the amount of test time that will be needed for various failure intensity improvement factors:

Failure Intensity Improvement Factor λ_0 / λ_F	Execution Time (CPU Hours)	
	GO Model	MO Model
10	1058	2916
100	2118	32076
1000	3177	323,677
10000	4236	3,239,690

The relationship of execution time to calendar time will vary with the test phase, development methods, and the type of software under development. Knowledge of how the failure intensity improvement factor varies with execution time can be used to determine the general shape of the relationship between calendar time and failure intensity improvement factor. This relationship can be used to determine an attainable improvement factor, given the release date and available resources. The failure intensity requirement can then be obtained from this factor.

Note the differences between the predictions made by the two models. In the MO model, the relationship between additional execution time needed and the improvement factor is linear, while in the GO model it is logarithmic. At this point, a choice between the two models must be made. Since it is not possible to know a priori which model is best suited to the data, the applicability of models to a set of failure data must be evaluated while the models are being applied. Once the model most applicable to the failure data has been identified, that model's relationship between failure intensity improvement factor and execution time can be used in conjunction with the relationship between execution and calendar time to determine the failure intensity requirement. □

The basis of the third approach, life cycle cost optimization, is the assumption that reliability improvement is obtained by more extensive testing. Costs and schedules for non-testing phases are assumed to be constant. The part of development cost due to testing decreases with higher failure intensity requirements, while the operational cost increases. The total cost therefore has a minimum. This is shown below in Figure 7.2.

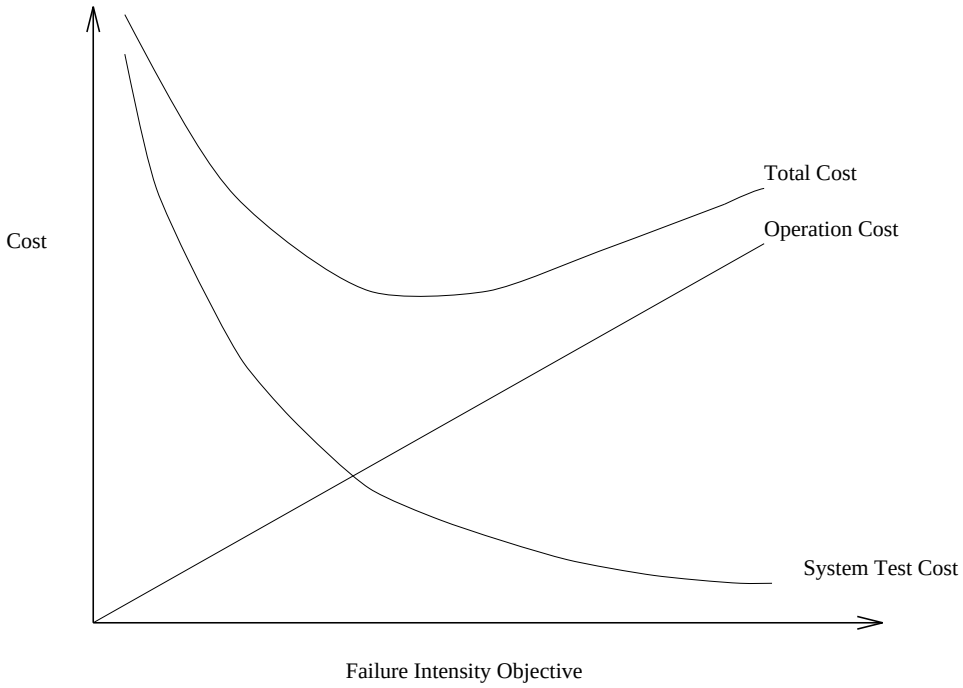

Figure 7.2: Reliability Objective by Life Cycle Cost Optimization

To find this minimum, testing cost as a function of failure intensity must be computed. If testing cost can be related to calendar time, and if the relationship between calendar and execution time is known, this calculation can be done for a specific model. Similarly, the operational cost as a function of failure intensity could be computed[Ehrl93a]. The following costs could be considered:

1. Terminating an improperly functioning program in an orderly manner.
2. Reconstructing affected databases.
3. Restarting the program.
4. Determining the cause(s) of the failure.
5. Developing procedures to prevent further failures of that type.
6. Repairing the fault(s) causing the failure if the severity and criticality of the failure warrants corrective action.

7. Testing the software to validate any repairs.
8. Effect of similar failures in the future on mission or program success.
9. Loss of customers' good will.
10. Other costs, such as settling lawsuits in the event of failures in life-critical systems (e.g., commercial avionics, medical systems).

As a result, the cost of operational failures depends on many complicated factors and is hard to determine accurately. In particular, the last two types of costs may be non-linear with increasing failure intensity. However, we may start by assuming that failures are equally costly in each severity class. Another simpler alternative is to assume all operational failures cost the same, in which case an average cost figure, normally available in each corporation, could be used. This number typically ranges from $10,000 to $100,000 depending on the size and the criticality of each project.

Note that the cost of testing could be expressed as an increasing function of time or number of test cases. Three conditions are normally considered: (1) The cost of testing is linear, i.e., it is constant per unit time; (2) It is linear up to time $t_{\max}$, and infinite beyond that, i.e., that project has to be finished by the time $t_{\max}$; and (3) initially it is linear, while later it increases exponentially due to the loss of credibility, lost of market window, penalties, etc.

Determination of the failure intensity requirement then becomes a constrained-minimum problem that can be solved analytically or numerically. A typical economic model[Dala88a, Dala90a] is illustrated in Example 7.2.

Example 7.2

This model includes costs and benefits derived from tradeoff between testing and operation cost due to failure. For simplicity, it is assumed that fault and failure have a one-to-one relationship, and could thus be referred simultaneously. Let

N = expected total number of faults in the program
$K(t)$ = number of faults observed up to time t

x = cost of fixing a fault when found while testing
y = cost of fixing a fault when found in the field
c = (y - x) = net cost of fixing a fault after rather than before release

Further, we assume that there is a known nonnegative monotone increasing function $g(t)$ that gives the sum of the cost of testing up to time t plus the opportunity cost of not releasing the software up to time t. Here we also assume that all the failures are equally costly.

Now the following total cost of testing up to time t could be formulated:

$$\begin{aligned} L(t, K(t), N) &= g(t) + xK(t) + y(N - K(t)) \\ &= g(t) - cK(t) + yN. \end{aligned} \tag{7.3}$$

It can be shown that if the amount of time it takes to find a fault X during testing is distributed with a known distribution function $F_X(t)$, and the failure times are independent, then the stop-testing rule turns out to be

$$\frac{g'(t)F_X(t)}{cf_X(t)} \geq K(t) \tag{7.4}$$

where $f_X(t)$ is the density function of $F_X(t)$.

Now if we take $g(t)$ to be linear, i.e., $g(t) = g \cdot t$, and if we apply the GO model (see Section 3.3.2) for $F_X(t)$, namely, $F_X(t)$ is $1 - e^{-bt}$ and $f_X(t)$ is be^{-bt}, then the stopping rule in Equation (7.4) reduces to

$$(g / bc)(e^{bt} - 1) \geq K(t). \tag{7.5}$$

Note that this stopping rule depends on $g(t)$ and c only through the ratio g / c. Also note that $K(t)$ can be estimated and predicted by

$$K(t) = N(1 - e^{-bt}) \tag{7.6}$$

□

7.2.2 Setting up a Data Collection Process

When you set up an SRE program, you should avoid the ambition to keep every bit of information about the project and its evolvement over the life cycle. Often people do not have a clearly defined objective for the data collection process. As a result, much effort is expended with little gain. There have been many instances in which large data collection efforts have been

implemented without any capability to analyze the data. Clearly defined objectives are necessary to help define the SRE requirements. In addition, when a large amount of data is required, the development staff is usually affected. Cost and schedule suffer because of the additional effort of collecting the data. Project management complains about the large amount of overhead involved in the data collection without any constructive feedback that could help the development process.

Therefore, we recommend you use the following sequence of steps to set up a data collection process.

1. Establish the objectives. Establishing the objectives is often the distinguishing point between successful and unsuccessful data collection efforts.
2. Develop a plan for the data collection process. Involve all of the parties that will be involved in the data collection and analysis. This includes designers, coders, testers, quality assurance staff, and line and project software managers. This insures that all parties understand what is being done and the impact it will have on their respective organizations. The planning should include the objectives for the data collection and a data collection plan. Address the following questions:
    a. How often will the data be gathered?
    b. By whom will the data be gathered?
    c. In what form will the data be gathered?
    d. How will the data be processed, and how will it be stored?
    e. How will the data collection process be monitored to insure the integrity of the data and that the objectives are being met?
    f. Can existing mechanisms be used to collect the data and meet the objectives?
    g. How much effort will be required to collect the data over the life of the project?
3. If any tools have been identified in the collection process, their availability, maturity and usability must be assessed. Commercially available tools must not be assumed to be

superior to internally developed tools. Reliability, ease-of-use, robustness, and support are factors to be evaluated together with the application requirements. If tools are to be developed internally, plan adequate resources – cost and schedule – for the development and acceptance testing of the tool.

4. Train all parties in use of the tools. The data collectors must understand the purpose of the measurements and know explicitly what data is to be collected. Data analysts must understand a tool's analysis capabilities and limitations.

5. Perform a trial run of the data plan to iron out any problems and misconceptions. This can save a significant amount of time and effort during software development. If prototyping is being done to help specify requirements or to try out a new development method, the "trial run" data collection could be done during the prototyping effort.

6. Implement the plan. Make certain that sufficient resources have been allocated to cover the required staffing and tool needs, and that the required personnel are available.

7. Monitor the process on a regular basis to provide assurance that objectives are met and that the software is meeting the established reliability goals.

8. Evaluate the data on a regular basis. Don't make the reliability assessment after software delivery. Waiting until after delivery defeats the usefulness of software reliability modeling because you have not used the information for managing the development process. Based on the experiences reported in[Lyu91a], weekly evaluation seems appropriate for many development efforts.

9. Provide feedback to all parties. This should be done as early as possible during data collection and analysis. It is especially important to do so at the end of the development effort. It is very important to provide feedback to those involved in data collection and analysis so they will be aware of the impacts of their efforts. Parties who are given feedback will be more inclined to support future efforts, as they will have a sense of efficacy and personal pride in their accomplishments.

7.2.3 Defining Data to be Collected

Many projects already have in place some data collection mechanisms for failure data. For example, JPL has Problem/Failure Report (P/FR), Bellcore has Modification Request (MR) database, and IBM has Authorized Program Analysis Report (APAR). These mechanisms track the date and time at which the failure was observed, a description of the failure, and some information about the system configuration at the time the failure was observed. Specific information needs to be collected is listed in the following subsections.

7.2.3.1 Time Between Successive Failures

Collect the execution time between successive failures. If execution time is unavailable, testing time between successive failures, measured by calendar time, can be used as a basis of approximation. Collect the start and completion time of each test session. Collect the times between failures (interfailure times).

If time-between-failures data cannot be collected, then collect test interval lengths and the number of failures encountered during each test interval. In many cases, this failure-count (or failure frequency) information seems to be more easily collected than the time-between-failure information. Test interval lengths should also be accurately recorded. If possible, collect the CPU utilization during the test periods to determine the relationship between CPU and calendar time.

For many development efforts, failure-count information is the only available type. However, some software reliability tools can use only time-between-failures as input. In this instance, the failure-count data can be transformed to time-between-failures data in one of two ways described in Section 1.2.4. Since the uncertainty in reported failure times affects the accuracy of modeling results, problem reporting mechanisms should be structured such that the mechanism’s resolution is greater than the average interfailure time throughout the test cycle.

7.2.3.2 Functional Area Tested

This can be done with reference to a software requirements document or a software build plan. Reliability predictions may be dramatically different when this information is or is not available. The importance of tracking this information is illustrated in the following example.

Example 7.3

Failure data set J3 in Appendix B was used for this example. The software reliability estimates were made using software reliability modeling tool SMERFS and CASRE (see Appendix A for descriptions). The Goel-Okumoto NHPP model was applied to the data. The software was assumed to be composed of two largely independent functional areas, and that each functional area would be executed 50% of the time during operations. In producing the estimates seen in Figure 7.3, the model was applied to the entire set of failure data. This yields an estimated failure rate of three failures per week at week 41 of the testing phase.

The actual failure rate curve, however, is bimodal. There is clearly a change in the test procedure after week 14 of the testing phase. If it is known that the software is composed of two distinct functional areas, and that after week 14, a different portion of the software is being tested than during the first 14 weeks, the reliabilities of the two functional areas can be separately modeled to yield a more accurate reliability estimate.

Figure 7.4 shows the reliability estimates for the two individual functional areas. By the end of week 14, the expected number of failures per week is 8 for the first functional area. During the interval between weeks 15 and 41, only the second functional area is tested. By the end of week 41, the expected number of failures per week is 1. If the software is delivered to operations at the end of week 41, and assuming that the functional areas are executed with equal frequency during each week of testing, it is seen that during operations, 4 failures per week can be expected while executing the first functional area, and .5 failures per week can be attributed to the second functional area. The resulting estimate of 4.5 failures per week is significantly different from the 3 failures per week that were estimated without taking the change in test focus into account.

As a numerical comparison, the mean square error for the predictions in Figure 7.3 is 42.9, while that for the predictions in Figure 7.4 is 28.3, a significant improvement. Note that the mean square error for the predictions in weeks 15-41 in Figure 7.4 further drops to 14.2. The close fit in this period can been seen in the figure. □

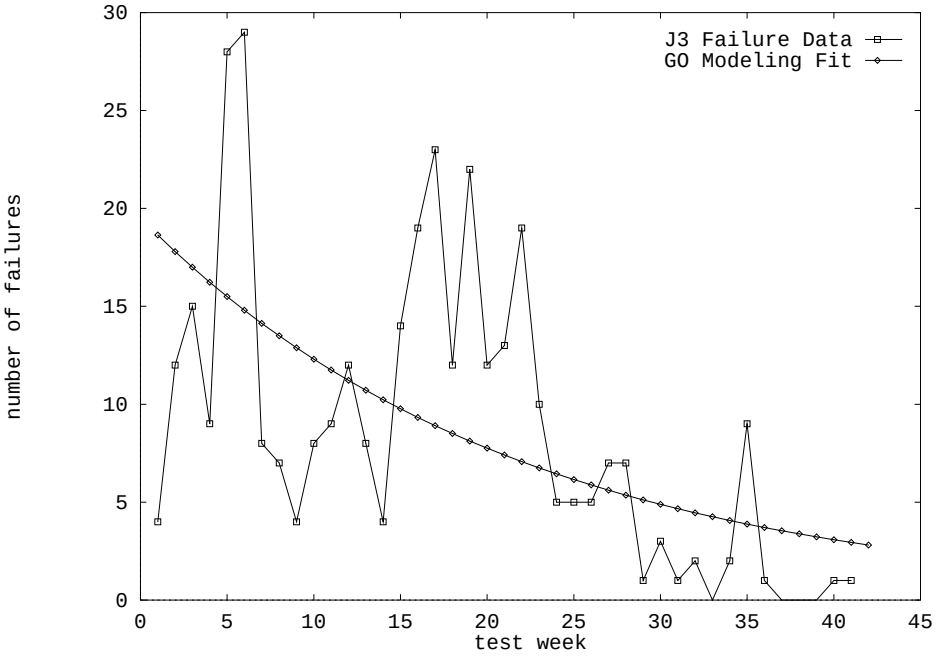

Figure 7.3: Application of GO Model to Entire Data Set

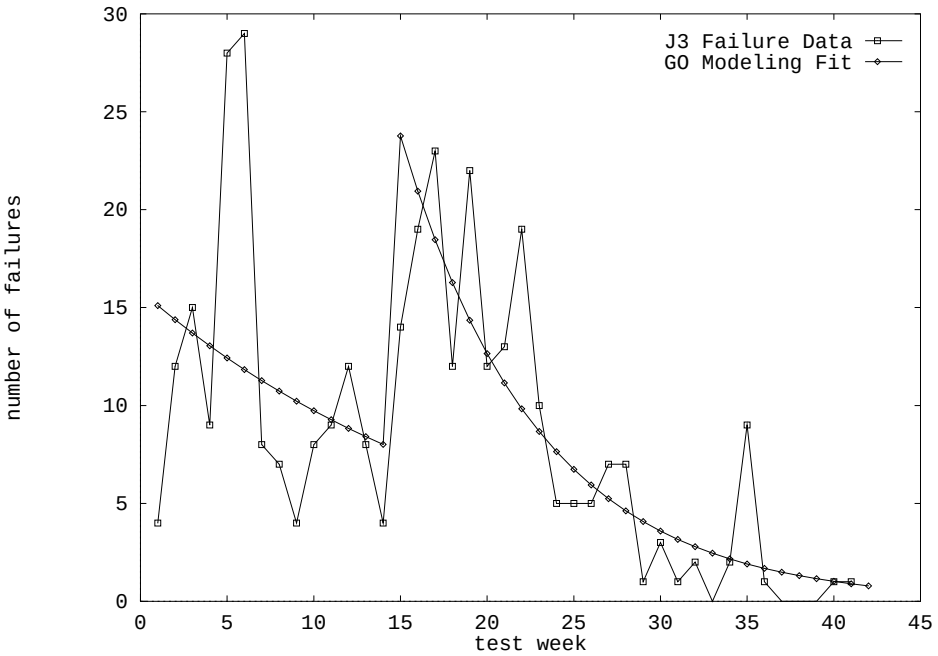

Figure 7.4: Separate GO Model Application of Distinct Functional Areas

7.2.3.3 Changes During Testing

Significant events that may affect the failure behavior during test include:

a. Addition of functionality to the software under test or significant modification of existing functionality. If the software under test is still evolving, the failure intensity may be underestimated during the early stages of the program's development, yielding overly optimistic estimates of its reliability.

b. Increases or decreases in the number of testers. This will increase or reduce the failure-count data (expressed in calendar time) as testers are added or taken away from the development effort. The time spent by each tester in exercising the software must be recorded so that the failure count or time-between-failures inputs to the models are accurate.

c. Changes in the test environment (addition/removal of test equipment, modification of test equipment). If the test equipment is modified during a test phase to provide greater throughput, the time-between-failures and failure-count data recorded subsequently to the modification will have to be adjusted to be consistent with the failure data recorded prior to the modification. For instance, if the clock speed in the test computer is increased by a factor of two, the test intervals subsequent to the clock speed increase will need to be half as long as they were prior to the speedup if failure-count information is being recorded. If time-between-failures information is being recorded, the interfailure times recorded subsequent to the speedup will have to multiplied by 2 to be consistent with the times-between-failures recorded before the speedup occurred.

d. Changes in the test method (e.g. switching from "white box" to "black box" testing, changing the stress to which the software is subjected during test). If the test method changes during a testing effort, or if the software is exercised in a different manner, new estimates of the software's reliability will have to be made, starting at the time when the testing method or testing stress changed.

7.2.3.4 Other Considerations

Interfailure times expressed in terms of CPU time are the preferred data. However, failure-count data is also recommended since existing problem reporting mechanisms can often be used. The relative ease of collecting this information will encourage the use of SRE techniques. Currently, most problem reporting systems collect the number of failures per unit test time interval. If your projects have existing mechanisms for collecting software failure data during developmental testing, use this data to obtain time-between-failures or failure-count data.

If failure-count data is used, a useful length for the test interval must be determined. This is influenced by such considerations as the number of testers, the number of available test sites, and the relative throughputs of test sites. Many development efforts summarize their findings on a weekly basis. For many development efforts, a week during subsystem or system-level testing is a short enough period of time that the testing method will not change appreciably. In general, enough failures can be found in a week's time during the early stages of test to warrant recomputing the reliability on a weekly basis.

Many development projects require that test logs be kept during developmental and system-level testing, although the information recorded in these logs is generally not as accurate as that tracked by the problem reporting system. Used as intended, these logs can be utilized to increase the accuracy of the failure-count or time-between-failure data available through the problem tracking system. Without much effort beyond that is required to record failures, the following items can be collected:

a. Functionality being tested. The functionality can be related to items in a software build plan or requirements in a software requirements document. The reliability for each functional area should be modeled separately.

b. Test session start date and time.

c. Test session stop date and time.

In addition, it may be possible to gather CPU utilization data from the test bench's accounting facilities for each test period recorded.

If only one functional area is to be tested during a session, record only one start and stop time. If more than one functional area is to be tested, however, start and stop times should be recorded for each functional area. If testing is being done at more than one test site, keep a log at each test site. To determine test interval lengths, use the test logs from all test sites to determine the amount of testing time spent in a fixed amount of calendar time. Count the number of unique failure reports from all test sites written against that functional area in the chosen calendar interval to determine the failure-count data. These failure counts and test interval lengths can then be used as inputs to software reliability models. Note that the reliability of each functional area is separately determined.

7.2.4 Choosing a Preliminary Set of Software Reliability Models

After specifying the software reliability requirements, you need to make a preliminary selection of software reliability models. Examine the assumptions that the models make about the development method and environment to determine how well they apply to your project. For instance, many models assume that the number of faults in the software has an upper bound. If software testing at the subsystem level does not occur until the software is relatively mature, and if there is a low probability of making changes to the software actually being tested, models making this assumption can be included in the preliminary selection (e.g. Goel-Okumoto model, Musa Basic model). If, on the other hand, significant changes are being made to the software at the same time it is being tested, it would be more appropriate to choose from those models that do not assume an upper bound to the number of faults (e.g. Musa-Okumoto and Littlewood-Verrall models). Many models also assume "perfect debugging." If previous experience on similar projects indicates that most repairs do not result in new faults being inserted into the software, choose from those models making this assumption (e.g. Goel-Okumoto model, Musa-Okumoto model). However, if a significant number of repairs result in new faults being inserted into the

software, it is more appropriate to choose from those models that do not assume perfect debugging (e.g. Littlewood-Verrall model).

It is important to note that there is currently no known method of evaluating these assumptions to determine a priori which model will prove optimal for a particular development effort[Abde86a]. You are advised that this preliminary selection of models will be a qualitative, subjective evaluation. After a model has been selected, its performance during use can be quantitatively assessed (see Chapter 4). However, these assessment techniques cannot be applied to the preliminary selection.

If a model has been shown valid for a similar project or the early release of the same project, use that model continuously and consistently. If you have no practical experience with software reliability models, you are advised to use the following models, as recommended by AIAA[AIAA93a]. The order of this list is arbitrary:

(1) Generalized Exponential model (see Section 3.7.1), which include Jelinski-Moranda model (JM)[Jeli72a] and Shooman model[Shoo73a] (also in Section 3.3.1), Musa Basic model (MB)[Musa79a] (also in Section 3.3.4), and Goel-Okumoto model (GO)[Goel79a] (also in Section 3.3.2);

(2) Schneidewind model (SM)[Schn75a] (also in Section 3.3.3);

(3) Musa-Okumoto Logarithmic Poisson model (MO)[Musa84a] (also in Section 3.5.3);

(4) Littlewood-Verrall model (LV)[Litt73a] (also in Section 3.6.1).

7.2.5 Choosing Reliability Modeling Tools

Many software reliability tools are available to model and measure software reliability automatically. See Appendix A for a comprehensive survey.

In our study we use two tools: SRMP[Litt86a] and CASRE[Lyu92b]. The CASRE tool calculates a product's reliability (the present reliability as well as future predictions of reliability) as a

function of test time, and represents it in terms of several interrelated reliability measures, such as cumulative number of failures, failures per time interval, the product's reliability function. This enables us to analyze a product's reliability from several points of view. CASRE is capable of providing product reliability estimates not only during system testing but during the product's field operation as well. In the latter case product reliability is expressed as a function of field operation time. CASRE allows users to select and apply existing software reliability models to the data displayed in the work space. These models come from the model library of the SMERFS tool[Farr88a], and consist of two categories based on their input data: time-between-failures models take the sequence of times between failures as the input, while failure-count models take number of failures per interval as the input.

7.2.6 Model Application and Application Issues

After setting up a data collection mechanism and selecting the model(s) and tool(s), measurement of software reliability can be started. Do not attempt to measure software reliability during unit test. Although failures may be recorded during this testing phase, the individual units of code are too small to make valid software reliability estimates. Our experience indicates that the earliest point in the life cycle at which meaningful software reliability measurements can be made is at the subsystem software integration and test level. Experience gained in our study, as well as empirical evidence reported in[Musa87a], suggests that software reliability measurement should not be attempted for a software system containing fewer than 2000 lines of uncommented source code. Instead, other measures (e.g., statement coverage, data and control flow coverage, data definitions and uses) could be investigated. Chapter 13 provides some emerging techniques in this area.

We now turn to the assumptions made by some of the more widely-used software reliability models. Chapter 3 discusses the model assumptions in detail. These assumptions are made to cast the models into a mathematically tractable form. However, there may be situations in which the assumptions for a particular model or set of models do not apply to a development effort. In

the following paragraphs, specific model assumptions are listed and the effects they may have on the accuracy of reliability estimates are described.

a. *During testing, the software is operated in a manner similar to the anticipated operational usage*. This assumption is often made to establish a relationship between the reliability behavior during testing and the operational reliability of the software. In practice, the usage pattern during testing can vary significantly from the operational usage. For instance, functionality that is not expected to be frequently used during operations (e.g. system fault protection) will be extensively tested to ensure that it functions as required when it is invoked.

   When the operational usage distribution is not obtainable, one way of dealing with this issue is to model the reliability of each functional area separately, and then use the reliability of the least reliable functional area to represent the reliability of the software system as a whole. Predictions of operational reliability that are made this way will tend to be more pessimistic than the reliability that is actually observed during operations, provided that the same inputs are used during test as are used during operations. If the inputs to the software during test are different from those during operations, there will be no easily identifiable relationship between the reliability observed during test and operational reliability.

b. *There are a fixed number of faults contained in the software*. Because the mechanisms by which faults are introduced into a program during its development are poorly understood at present, this assumption is often made to make the reliability calculations more tractable. Models making this assumption should not be applied to development efforts during which the software version being tested is simultaneously undergoing significant changes (e.g. 20% or more of the existing code is being changed, or the amount of code is increasing by 20% or more). Among the models making this assumption are the Jelinski-Moranda, the Goel-Okumoto, and the Musa Basic Models. However, if the major source of change to the software during test is the correction process, and if the corrections made do not significantly change the software, it is generally safe to make this assumption. In practice,

this would tend to limit application of models making this assumption to subsystem-level integration or later testing phases.

c. *No new faults are introduced into the code during the correction process*. Although there is always the possibility of introducing new faults during debugging, many models make this assumption to simplify the reliability calculations. In many development efforts, the introduction of new faults during the correction process tends to be a minor effect. If the volume of software, measured in source lines of code, being changed during correction is not a significant fraction of the volume of the entire program, and if the effects of repairs tend to be limited to the areas in which the corrections are made, this assumption is deemed acceptable. In the event that code is changing quickly or new faults are introduced when trying to fix old faults[Leve90a], reliability models are still adaptable by examining the code "churns"[Dala94a].

d. *Detections of failures are independent of one another*. This assumption is not necessarily valid. Indeed, there is evidence that detections of failures occur in groups, and that there are some dependencies in detecting failures. The reason for this assumption is that it enormously simplifies the estimation of model parameters. Determining the maximum likelihood estimator of a model parameter, for instance, requires the computation of a joint probability density function (pdf) involving all of the observed events. The assumption of independence allows this joint pdf to be computed as the product of the individual pdfs for each observation, keeping the computational requirements for parameter estimation within practical limits.

Practitioners using any currently-available models have no choice but to make this assumption. Almost all the models analyzed in Chapter 3 make this assumption. Nevertheless, practitioners from AT&T, Hewlett Packard, and Cray Research report that the models produce fairly accurate estimates of current reliability in many situations [Ehrl90a, Rapp90a, Zinn90a]. If the input data to the software are independent of each

other, failure detection dependencies may be reduced.

When the above assumptions are deemed necessary for analytical approaches to reliability modeling, other techniques have been developed to relieve some of these assumptions. Simulation approaches (see Chapter 16) and Neural Networks (see Chapter 17) are two of the promising attempts.

7.2.7 Dealing with Evolving Software

Most models described in Chapter 3 assume that the software being tested will not be undergoing significant changes during the testing cycle. This is not always the case. A software system undergoing test may be simultaneously undergoing development, with changes being made to the existing software or new functionality being added periodically. To accurately model software reliability in this situation, changes made to the software have to be taken into account. There are three ways of handling changes to a program under test. These approaches are[Musa87a]:

1. Ignore the change.
2. Apply the component configuration change method.
3. Apply the failure time adjustment technique.

Ignoring changes is the simplest method, and is appropriate when the total volume of changes is small compared to the overall size of the program. In this case, the continual re-estimation of parameters will reflect the fact that some change is in fact occurring.

The component configuration change approach is appropriate for the situation in which a small number of large changes are made to the software, each change resulting from the addition of independent components (e.g. addition of the telemetry gathering and down-linking capability to a spacecraft command and data subsystem). The reliability of each software component is modeled separately. The resulting estimates are then combined into a reliability figure for the overall system.

The failure time adjustment approach is most appropriately used when a program cannot be conveniently divided into separate independent subsystems and the program is changing rapidly enough to produce unacceptable failures in estimating the software's reliability. The three principal assumptions that are made in failure time adjustment are:

1. The program evolves sequentially. At any one time, there is only one path of evolution of the program for which reliability estimates are being made.
2. Changes in the program are due solely to growth. Differences between version k and version k+1 are due entirely to new code being added to version k. In practice, there may be reductions in one area of the code between versions k and k+1, while growth in other areas occurs. If the reductions are small in comparison to the growth, as often occurs in repairing faults, they can usually be ignored.
3. The number of faults introduced by changes to the program are proportional to the volume of new code.

We illustrate this approach by the following example.

Example 7.4

Figure 7.5 provides an example of a software system to which failure time adjustment techniques could be applied. This figure represents the cumulative number of failures for the system whose failure frequencies are shown in Figures 7.3 and 7.4. Recall that Figures 7.3 and 7.4 show abrupt changes in the failure-count data at week 14. For this example, we assume this change is attributed to the addition of new functionality to the software under test. The testing method remains the same during the two stages.

When testing proceeds in two stages, the expected number of failures as a function of time will follow a known curve during the first stage (weeks 0-14 in Figure 7.5). The parameters of this curve will depend on the fault content and the total amount of code being executed in this stage. After the 14th week of testing (week 14 in Figure 7.5, denoted by T in the text), additional code that implements the remainder of the system is added. At this point, the curve representing the expected number of failures will switch to the one that would have occurred for a system in its final configuration. The curve, however, is temporally translated, the amount of translation depending on the number of failures that were experienced during the first test stage. The translation can be determined by modeling the first and second stages independently.

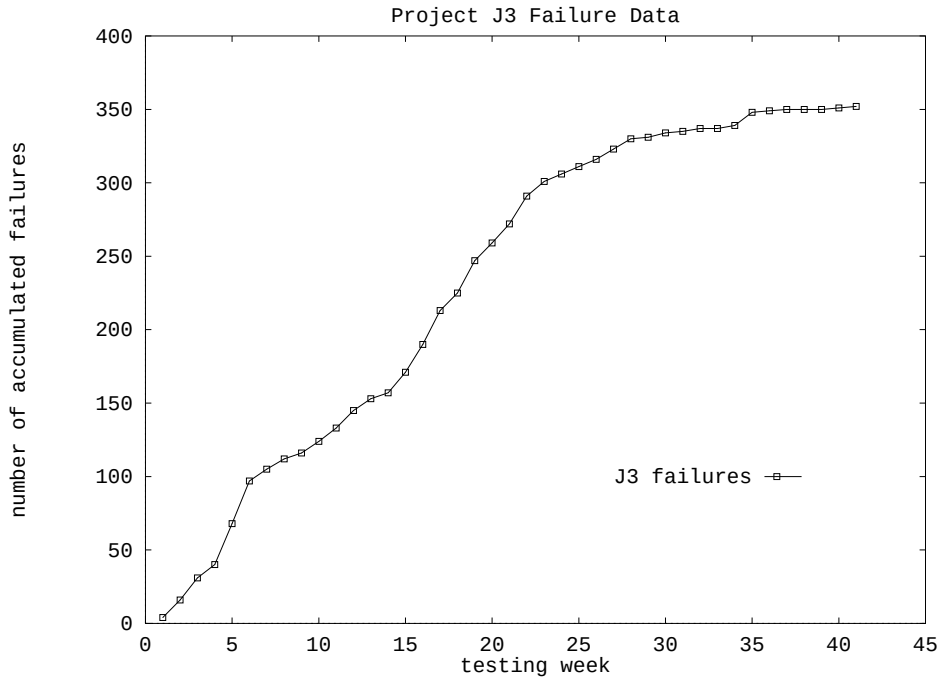

Figure 7.5: Cumulative Number of Failure for J1 Project Data

If the Goel-Okumoto model is used, the parameters for the first and second stages are as follows:

measurement period	a	b
First Stage (weeks 1-14)	a_1 = 317	b_1 = .0487787
Second Stage (weeks 15-40)	a_2 = 413	b_2 = .0461496

In the first stage, then, the expected number of failures is given by $\mu_1(t) = a_1(1 - e^{-b_1 t})$. Substituting the values of a and b in the table above, and using a value of 14 for t, the expected number of failures is 156. Now assume that testing had started using the final configuration, represented by the second stage. The expected number of failures would be given by $\mu_2(t) = a_2(1 - e^{-b_2 t})$, using the values of a and b given for the second stage. For the second stage, the number of failures expected to be observed during the first stage (156) would be observed in 10 time units, denoted by $\tilde{T}$, rather than in 14. Therefore, in going from the first to the second stage of testing, the expected number of failures, $\mu(t)$, will be a translated version of the expression for $\mu_2(t)$. The amount of translation is given by $T - \tilde{T}$, which in this case is 4. Namely,

$$\mu(t) = a_2(1 - e^{-b_2(t - (T - \tilde{T}))}), \text{ or} \tag{7.7a}$$

$$\mu(t) = a_2(1 - e^{-b_2(t - 4)}) \tag{7.7b}$$

□

7.2.8 Practical Limits in Modeling Ultra-Reliability

It is important to note a limitation of applying software reliability modeling techniques to verify systems for ultra-high reliability (e.g., one failure per 10^7 hours of operation). It could be shown[Butl91a] that quantification of software reliability in the ultra-reliable regime is infeasible, since the required amount of testing time exceeds practical limits. For example, a system having a required probability of failure of 10^{-7} for a 10-hour mission implies that MTTF of the system (assuming exponentially distributed) T_F is approximately 10^8. There are two basic approaches: testing with replacement, and testing without relacement. In either case, testing continues until r failures have been observed. In the first case, when a device fails, a new device is put on test in its place. In the second case, the failed device is not replaced. For the first case, the expected time on test, D_t, is given by:

$$D_t = T_F \frac{r}{n} \tag{7.8}$$

where n is the number of items place on test. For the second case, the expected time on test is:

$$D_t = T_F \sum_{j=1}^{r} \frac{1}{n - j + 1} \tag{7.9}$$

If r is set to 1, this gives the shortest test time possible. Table 7.1 shows the expected test duration as a function of the number of test replicates, n. The expected test time with or without replacement is almost the same in this case.

No. of Replicates (n)	Expected Test Duration D_t
1	10^8 hours = 11415 years
10	10^7 hours = 1141 years
100	10^6 hours = 114 years
1000	10^5 hours = 11.4 years
10000	10^4 hours = 1.14 years

Table 7.1: The Expected Test Duration as a Function of n

To get satisfactory statistical significance, larger values of r are required, which translates to

even more testing. Given that economic considerations rarely allow the number of test replicates to be greater than 10, life-testing of ultra-reliable systems looks quite hopeless.

We should note, however, that the critical software function which requires such an ultra reliability seldom execute the complete period of the 10-hour mission. If fact, it could only requires a small portion of the CPU time (e.g., the last landing approach of a long airplane flight). If we assume that only 0.1 CPU hours of the 10-hour mission requires this 10^{-7} failure probability, then T_F becomes approximately 10^6 CPU hours for the critical function. Consequently, the numbers in Table 7.1 become those in Table 7.2.

No. of Replicates (n)	Expected CPU Test Duration D_t
1	10^6 hours = 114 years
10	10^5 hours = 11.4 years
100	10^4 hours = 1.14 years
1000	10^3 hours = 42 days
10000	10^2 hours = 4.2 days

Table 7.2: The Expected CPU Test Duration as a Function of n

It is noted that the critical function could be tested in a more powerful CPU which is equivalent to an increase of the number of replicates. For example, if we use 10 replicates, each running a CUP which is 100 times faster than the original one, then the goal could practically be achieved in 42 CPU days. There are, of course, obstacles to overcome before this kind of testing and validation can happen. But its achievement would not be complete infeasible.

7.3 Investigation at JPL

For project applications in SRE practice, we first introduce a feasibility study on SRE techniques to JPL historical projects. The objectives of this study are:

1. Examine the applicability of software reliability models to real world projects.
2. Apply model selection criteria and compare models.

3. Determine if there is the best model suitable for all applications.

4. Evaluate the cost-effectiveness of SRE techniques.

7.3.1 Project Selection and Characterization

Data Set J1

The project J1 was one of the first spacecraft in which a significant fraction of the functionality was provided by software. This software system, totaling approximately 14,000 lines of uncommented assembly language, was divided among three real-time embedded subsystems - the Attitude and Articulation Control Subsystem (AACS), the Command and Control Subsystem (CCS), and the Flight Data Subsystem (FDS). The failure data we analyzed comes from spacecraft system testing, at which point the AACS, CCS, and FDS had been integrated into the spacecraft. Among the items recorded on the Problem/Failure Reports during system test are a) Time of failure, b) Failure type, and c) Subsystem in which the failure occurred. During J1 system test, approximately 9.5 faults per KLOC were discovered.

Data Set J2

Launched in 1989, project J2 was developed as a planet orbiter carrying an atmospheric probe. As with the project J1, a large fraction of project J2's functionality was provided by software. Approximately 7,000 uncommented source lines of HAL/S were implemented for the AACS, while about 15,000 source lines of assembly language were developed for the Command and Data Subsystem (CDS). Project J2 failure data comes from spacecraft system testing. During J2 system test, approximately 10.2 faults per KLOC were discovered.

Data Set J3

Failure data for the project J2 CDS during one phase of subsystem-level integration testing was available for analysis. We were able to reconstruct some elements of the testing profile. For example, it was known to us that the number of hours per week during which testing occurred

was nearly constant throughout this phase, which was composed of two testing stages. In addition, the main functional areas of the software received roughly the same amount of testing every calendar week. This information resulted in the failure data being more accurate than that for other projects. During J3 subsystem test, approximately 10.1 faults per KLOC were discovered.

Data Set J4

Like project J2, project J4 has an AACS and a CDS, and the number of uncommented source lines of code for each is roughly the same as that for project J2. As with projects J1 and J2, the failure data comes from the spacecraft system test period. During J4 system test, approximately 8.0 faults per KLOC were discovered.

Data Set J5

Project J5 is a facility for tracking and acquiring data from Earth resources satellites in high-inclination orbits. Totaling about 103,000 uncommented source lines of code, the software is written in a mixture of C, Fortran, EQUEL, and OSL. About 14,000 lines were reused from previous efforts. The failure data reported here was obtained from the development organization's anomaly reporting system during software integration and test. During J5 system test, approximately 3.6 faults per KLOC were discovered.

This variety of project data would give us a chance to see whether the reliability measurement techniques developed for one type of development effort would work well for another.

7.3.2 Characterization of Available Data

For all of the JPL efforts, the following data was available:

1. Date on which a failure occurred.
2. Failure description.
3. Recommended corrective action.
4. Corrective action taken.

5. Date on which failure report was closed.

For each of the flight projects, the severity of each failure was also available.

Note that the following items were not systematically recorded, and were generally unavailable for use in the modeling effort:

1. Execution times between successive failures, or comparable information (e.g., total time spent testing during a calendar interval).
2. Test interval lengths. It was therefore necessary to assume that they were constant.
3. Operational profile information (e.g., functional area being tested, referenced to requirements or design documentation; subsystem being tested; points at which the testing method may have changed.)

The data collected from these development environments tend to be very noisy, and the assumptions of most software reliability models do not necessarily hold under the described circumstances. Nevertheless, failure data collected based on calendar time are typically under similar circumstances in many historical projects.

7.3.3 Experimental Results

In the reliability analysis of the JPL project data J1 through J5, we have to assume that the test time per unit interval of calendar time was relatively constant, and the testing method remained constant, since this information was not systematically recorded. Largely because of this lack of information, we decided to model the reliability of the facility as a whole, rather than attempted to model the component reliabilities. Subsequently, we applied the SRMP reliability tool to the five JPL projects and obtained the following model comparison results. Note that all the JPL project data were collected in failure-count format, and we had to convert the data into time-between-failures format for proper execution by SRMP. Random distribution of the grouped failure data was assumed for the conversion.

We evaluated a number of models surveyed in Chapter 3 with respect to the seven selection criteria listed in Section 1.3.2, Six models, JM, GO, MO, Duane model (DU) [Duan64a] (also in Section 3.5.1), Littlewood model (LM)[Litt81a], and LV, were judged to perform sufficiently well[Lyu91b] and were appled to these five projects. The following tables, Tables 7.3-7.7, summarize the analysis of model applicability for the JPL efforts. For each development effort, the models applied were evaluated with respect to prequential likelihood, model bias, bias trend, and model noise. The value for each criterion is given in the tables, while the corresponding ranking is given in parenthesis. Each of these criteria was given equal weighting in the overall ranking.

Measure	JM	GO	MO	DU	LM	LV
Accuracy	894.7(6)	573.7(3)	571.5(2)	586.6(4)	829.9(5)	549.1(1)
Bias	.2994(5)	.2849(3)	.2849(3)	.2703(2)	.2994(5)	-.0793(1)
Trend	.0995(5)	.0965(3)	.0957(2)	.2551(6)	.0994(4)	.0876(1)
Noise	∞(5)	13.81(3)	9.225(2)	8.402(1)	∞(5)	24.51(4)
Overall Rank	(6)	(3)	(2)	(4)	(5)	(1)

Table 7.3: Model Rankings for J1 Data

Measure	JM	GO	MO	DU	LM	LV
Accuracy	1074(2)	1075(4)	1078(5)	1098(6)	1074(2)	1051(1)
Bias	.3378(3)	.3378(3)	.3379(5)	.1944(1)	.3382(6)	-.2592(2)
Trend	.4952(3)	.4954(4)	.5041(6)	.4618(2)	.4954(4)	.1082(1)
Noise	2.607(3)	2.593(2)	2.395(1)	4.541(5)	2.624(4)	23.33(6)
Overall Rank	(2)	(3)	(5)	(4)	(6)	(1)

Table 7.4: Model Rankings for J2 Data

Measure	JM	GO	MO	DU	LM	LV
Accuracy	643.0(3)	639.3(2)	681.1(5)	728.5(6)	643.0(3)	612.3(1)
Bias	.1783(3)	.1783(3)	.1700(1)	.1748(2)	.1784(5)	-.2581(6)
Trend	.3450(3)	.3408(2)	.4262(5)	.4282(6)	.3450(3)	.2426(1)
Noise	4.042(5)	3.908(4)	2.673(3)	2.287(1)	4.042(5)	2.564(2)
Overall Rank	(3)	(2)	(3)	(5)	(6)	(1)

Table 7.5: Model Rankings for J3 Data

Measure	JM	GO	MO	DU	LM	LV
Accuracy	627.1(3)	627.1(3)	627.1(3)	616.0(1)	627.1(3)	622.9(2)
Bias	.2968(4)	.2968(4)	.2969(2)	.1858(1)	.2969(2)	-.3483(6)
Trend	.2399(3)	.2399(3)	.2399(3)	.2180(2)	.2399(3)	.1429(1)
Noise	1.007(1)	1.007(1)	1.007(1)	2.003(5)	1.009(4)	5.563(6)
Overall Rank	(3)	(3)	(1)	(1)	(5)	(6)

Table 7.6: Model Rankings for J4 Data

Measure	JM	GO	MO	DU	LM	LV
Accuracy	915.7(1)	915.8(4)	915.7(1)	925.5(6)	915.7(1)	920.5(5)
Bias	.3023(1)	.3023(1)	.3023(1)	.4249(6)	.3023(1)	-.3672(5)
Trend	.0606(1)	.0615(3)	.0620(4)	.0918(5)	.0606(1)	.1009(6)
Noise	1.587(3)	1.540(2)	1.395(1)	1.650(5)	1.589(4)	3.189(5)
Overall Rank	(1)	(4)	(2)	(5)	(2)	(5)

Table 7.7: Model Rankings for J5 Data

From Tables 7.3-7.7 we found that there was no one "best" model for the development efforts that were studied. This is consistent with the findings reported in Chapter 4. It is easy to see that a model that performs well for one development effort may do poorly in another. For instance, the Littlewood-Verrall model performs very well for the first three data sets – in fact, it outperforms all of the other models. However, it comes in last for the remaining two projects. This inconsistency is repeated for the other five models as well. There were no clear differences between the development processes for the five JPL applications, certainly none that would favor the selection of one model over another prior to the start of test. These findings suggest that multiple models be applied to the failure data during the test phases of a development effort, preferably models making different assumptions about the failure detection and fault removal processes. In addition, the models should be continually evaluated for applicability to the failure data. The model(s) ranking highest with respect to the evaluation criteria should then be chosen for use in predicting future reliability.

7.4 Investigation at Bellcore

After the study of JPL historical project data, we learned the lesson that data collection plays a crucial role in SRE applications. We also learned that multiple models should be applied to project data and the selection of best model(s) should be exercised in a continuous fashion. To this need, we implemented the CASRE tool[Niko92a, Lyu93a] for a systematic and automatic approach to SRE. Consequently, another study was performed to investigate Bellcore projects for SRE applications. The objectives of this study are:

1. Apply a better data collection effort and observe the effect.
2. Search better model(s) for particular projects as *a posteriori*.
3. Observe the growth of reliability during testing.
4. Classify the characteristics of reliability models.

7.4.1 Project Characteristics

Data Set B1

Project B1 is a key telecommunication software system for daily telephony operations. This system has been in existence for over 10 years and it is deployed by all the Regional Bell Operating Companies. The whole system includes around 1 million lines of C source code, but the main application is composed of 700K lines of code. During testing of various project B1 releases, complex interactions with other large telecommunication systems are involved. Due to test metrics collected through an automatic testing environment, some approximations to software execution time are available, which include number and volume of messages received per day, as well as total and unique test cases executed per day. The failure reports, collected from 6/17/93 to 9/9/93, represent testing activities conducted by a group of 15 staff members that were involved with testing this release, which includes about 250K lines of new and changed lines of code.

Data Set B2

Project B2 is a large data base application in the telecommunication market. This project was initiated 20 years ago. The current system is composed of over 1 million lines of code, written in more than ten different languages including Assembly, PL1, COBOL, C/C++ and others. During testing the project B2 releases, CPU time has been hard to collect, and staff time has to be collected manually. Project B2 involves mostly stand-alone software testing, and the complete software is available to the testers from the first day of testing. The project B2 test failure data represents testing efforts of seven testers from 9/8/93 to 11/16/93 for a software release, which includes 15.2K new and changed source lines of code. Failure data were collected with respect to calendar testing time and total staff time. There is also a classification of failure reports to distinguish those whose software faults were found located in the new and changed code instead of the old code.

Data Set B3

Project B3 is composed of two major parts: B3a is a traditional large software system for telephony operation, while B3b is a new system combining a major feature with system B3a. The testing of B3a and B3b involves different testing schemes. B3a system testing invokes mostly batch jobs which are CPU-bounded, while B3b system testing uses screen navigations for non-CPU-bounded executions. It is hard to collect software execution time for either system. Moreover, it is hard to equate test cases of various sizes, since they vary tremendously. As a result, various staff time components are collected. It is also noted that code churn situation exists for both projects, since not all the software code is available to the tester from testing day one. The current data under analysis include 8/2/93 – 11/16/93 for B3a system, and 10/4/93 – 11/16/93 for B3b system.

7.4.2 Data Collection

Failure data for all three projects are collected from Bellcore internal problem tracking systems.

This database stores information about all the failures found during testing and operation. The discovery date, description, originator, severity and other tracking information are associated with each failure. A query is available to return the number of failures found on each testing day. Staff time related information is collected manually for each project. The data collection form for this process is shown in Figure 7.6. In addition, project B1 is able to automatically extract some time-related records to represent the intensity of testing.

From several months' effort of software failure data collection during system testing, the following scenarios were made available for the first Bellcore project (B1):

(1) B1.calendar: this data set collects software failure data reported during B1 system testing, based on calendar time.

(2) B1.staff: this data set records B1 system testing failures based on staff time reported by each tester during testing.

(3) B1.all_test: this data set records B1 failures based on the total test cases executed per testing day. The test case information is automatically collected.

(4) B1.uniq_test: this data set records B1 failures based on the number of unique test cases executed each day. Repeated test cases run on the same day are not counted. This information is automatically collected.

(5) B1.message: this data set records B1 failures based on the messages the system sends out. The number of messages is considered as an indication of the intensity of software execution. This information is automatically collected.

For the second Bellcore project (B2), an effort was devoted to filter out reported failures which are not against the current software release under testing:

(6) B2.calendar: similar to B1.calendar.

(7) B2.staff: similar to B1.staff.

(8) B2.calendar_rf: This data set records "real faults" of software and uses calendar time. Real

faults are counted only when the reported failures were caused by software faults residing in the new or changed code of the current release.

(9) B2.staff_rf: This data set records "real faults" of software and uses staff time.

For B3 project, only calendar time and staff time information are collected for the reported failures of B3, which includes B3a system and B3b system. The failure data collected by calendar time and staff time are called B3a.calendar, B3b.calendar and B3a.staff, B3b.staff, respectively.

(10) B3a.calendar: similar to B1.calendar for B3a system.

(11) B3b.calendar: similar to B1.calendar for B3b system.

(12) B3a.staff: similar to B1.staff for B3a system.

(13) B3b.staff: similar to B1.staff for B3b system.

7.4.3 Application Results

We use the CASRE tool to apply software reliability models to the Bellcore data sets. There are two types of models in CASRE: time-between-failures (TBF) models and failure-count (FC) models. Table 7.8 presents a comparison of TBF models, including LV, Geometric model (GEO)[Farr83a] (also in Section 3.5.2), MB, JM and MO, for the project B1 data using messages as a time measure (B1.message). Table 7.9 shows comparable results for FC models, which include Generalized Poisson model (GP)[Farr83a], Brooks and Motley Poisson model (BMP) and Binomial model (BMB)[Farr83a], and Yamada Delayed S-Shaped model (YSS)[Yama83a] (also in Section 3.4.2).

The first column lists the models being compared. These are followed by the columns which record results from several model evaluation criteria, including the prequential likelihood measure (normalized by its negative natural log value which result in a value between 0 and 1, denoted by -ln PL), model bias ("u-plot"), bias trend ("y-plot"), model noise ("deviance"), and

goodness-of-fit or Chi-Square. Rank ordering of the measure for each criterion is listed in parentheses. Overall ranks are provided in the last column. We can see for B1.message data, LV performs the best among the TBF models, while among the FC models, the GP model performs the best.

Model Name	-ln PL	Model Bias	Bias Trend	Model Noise	Goodness-of-Fit	Overall Rank
LV	526.8 (1)	0.127 (5)	0.0940 (1)	1.85 (2)	0.045 (1)	1
GEO	529.5 (2)	0.113 (3)	0.1490 (2)	1.78 (1)	0.107 (4)	2
MB	529.8 (4)	0.080 (2)	0.1644 (5)	2.76 (4)	0.104 (2)	3
JM	529.5 (3)	0.067 (1)	0.1643 (4)	2.80 (5)	0.111 (5)	4
MO	529.9 (5)	0.120 (4)	0.1511 (3)	1.87 (3)	0.105 (3)	5

Table 7.8: TBF model comparisons for B1.message data

Model Name	-ln PL (rank)	Chi-Square (rank)	Degree of Freedom	Rank
GP	99.87 (1)	42.91 (1)	16	1
BMP	104.12 (3)	50.09 (2)	19	2
NHPP	103.63 (2)	55.42 (4)	19	3
BMB	104.80 (4)	50.29 (3)	19	4
YSS	118.03 (5)	60.38 (5)	12	5

Table 7.9: FC model comparisons for B1.message data

Tables 7.10 and 7.11 list, for B1.message data, the estimated times between failure, failure rates, and the factor of reliability growth for the TBF and FC models, respectively. To capture the growth of reliability from each model's viewpoint, we define a Reliability Growth Factor (RGF) to be

$$RGF = \frac{final\ time\ between\ failures}{initial\ time\ between\ failures} \quad (for\ TBF\ models) \tag{7.10a}$$

$$= \frac{initial\ failure\ rate}{final\ failure\ rate} \quad (for\ FC\ models\ except\ the\ YSS\ model) \tag{7.10b}$$

In other words, RGF is the same as the failure intensity improvement factor discussed in Example 7.1. However, for the YSS model, we define

$$RGF = \frac{highest\ failure\ rate}{final\ failure\ rate} \tag{7.10c}$$

Model Name	Time Between Failures		RGF
	Initial	Final	
GEO	798	3277	4.11
JM	896	4213	4.70
LV	817	3790	4.64
MB	909	4047	4.45
MO	823	3269	3.97
Average	849	3719	4.37

Table 7.10: Time between failures and reliability growth estimated by TBF Models

Model Name	Failure Rates		RGF
	Initial	Final	
BMB	3.34	0.65	5.16
BMP	3.38	0.64	5.31
GP	2.92	0.46	6.29
NHPP	3.30	0.66	5.02
YSS	0.47	0.34	4.32
Average	3.24	0.60	5.45

Table 7.11: Failure Rates and Reliability Growth Estimated by FC Models

The overall comparison of the 13 Bellcore data sets is summarized in Table 7.12 for the TBF models, and in Table 7.13 for the FC models. From these two tables we can see that for the Bellcore project data, LV is the best TBF model. For the FC type models, the YSS model is the best one. Noted, however, that there are some data dependencies among data sets from the same project, since they all track the same set of failures. For example, the LV model performs very well among the TBF models for B1 data sets, while in the FC models, the YSS model is always the best in data sets from B2 project, and the GP model is very consistent for the data sets of B1 project. Nevertheless, this overall rank still tracks the three project results very well. The ranking would have remained the same order if only one data set (say, the average) had been used for each project instead of all data sets.

Model	GEO	JM	LV	MB	MO
B1.calendar	(5)	(3)	(1)	(2)	(4)
B1.staff	(5)	(3)	(1)	(2)	(4)
B1.all_test	(5)	(1)	(1)	(3)	(4)
B1.uniq_test	(5)	(1)	(2)	(3)	(4)
B1.message	(2)	(4)	(1)	(3)	(5)
B2.calendar	(2)	(2)	(1)	(4)	(4)
B2.staff	(4)	(5)	(2)	(2)	(1)
B2.calendar_rf	(4)	(1)	(1)	(3)	(5)
B2.staff_rf	(5)	(2)	(2)	(1)	(4)
B3a.calendar	(2)	(1)	(5)	(3)	(4)
B3a.staff	(1)	(2)	(5)	(3)	(4)
B3b.calendar	(2)	(3)	(1)	(5)	(3)
B3b.staff	(2)	(5)	(2)	(2)	(1)
Sum of Rank	44	33	25	36	47
Total Rank	(4)	(2)	(1)	(3)	(5)

Table 7.12: Overall Model Comparisons of TBF models

Model	BMB	BMP	GP	NHPP	YSS
B1.calendar	(4)	(2)	(3)	(1)	(4)
B1.staff	(3)	(3)	(1)	(5)	(1)
B1.all_test	(4)	(3)	(1)	(5)	(1)
B1.uniq_test	(4)	(3)	(1)	(5)	(2)
B1.message	(4)	(2)	(1)	(3)	(5)
B2.calendar	(5)	(2)	(3)	(4)	(1)
B2.staff	(3)	(3)	(2)	(3)	(1)
B2.calendar_rf	(2)	(4)	(2)	(5)	(1)
B2.staff_rf	(5)	(2)	(3)	(3)	(1)
B3a.calendar	(5)	(4)	(2)	(2)	(1)
B3a.staff	(4)	(1)	(4)	(3)	(1)
B3b.calendar	(1)	(1)	(1)	(5)	(4)
B3b.staff	(2)	(1)	(4)	(5)	(3)
Sum of Rank	46	31	28	49	26
Total Rank	(4)	(3)	(2)	(5)	(1)

Table 7.13: Overall model comparisons of FC models

Table 7.14 summarizes the overall indication of level of data convergence and reliability growth for each data set, across all TBF models. Table 7.15 presents an analogous summary for the FC models. The second column in these tables, first convergence point, represents the number of failure observations upon which at least one model converges, divided by the total number of failures. This percentage ratio is given in parentheses. The third column, common convergence

point, is given by the number of points at which all models converge divided by the total number of points. The percentage ratio is also in parentheses. A high ratio indicates the parameter estimation did not converge until very late. This would happen when the data set is very noisy.

Project Data	First Convergence Point	Common Convergence Point	RGF
B1.calendar	31/150 (20.6%)	102/150 (68.0%)	5.10
B1.staff	34/150 (22.7%)	130/150 (86.7%)	3.57
B1.all_test	34/150 (22.7%)	130/150 (86.7%)	3.46
B1.uniq_test	77/150 (51.3%)	130/150 (86.7%)	4.20
B1.message	32/150 (21.3%)	90/150 (60.0%)	4.37
B2.calendar	45/64 (70.3%)	60/64 (93.8%)	1.10
B2.staff	34/64 (53.1%)	59/64 (92.2%)	1.23
B2.calendar_rf	24/38 (63.2%)	34/38 (89.5%)	2.80
B2.staff_rf	14/38 (17.5%)	34/38 (89.5%)	3.03
B3a.calendar	70/75 (93.3%)	71/75 (94.7%)	2.25
B3a.staff	70/75 (93.3%)	70/75 (93.3%)	3.15
B3b.calendar	21/80 (26.3%)	67/80 (83.8%)	2.03
B3b.staff	7/80 (8.8%)	74/80 (92.5%)	3.49
Average	(43.4%)	(86.0%)	3.05

Table 7.14: Overall RGF of TBF models for each data set

Project Data	First Convergence Point	Common Convergence Point	RGF
B1.calendar	5/85 (5.9%)	33/85 (38.8%)	6.52
B1.staff	5/85 (5.9%)	55/85 (64.7%)	4.67
B1.all_test	4/68 (5.9%)	44/68 (64.7%)	4.35
B1.uniq_test	4/68 (5.9%)	44/68 (64.7%)	5.81
B1.message	5/71 (7.0%)	14/71 (19.7%)	5.45
B2.calendar	3/91 (3.3%)	78/91 (85.7%)	2.04
B2.staff	2/75 (2.7%)	55/75 (73.3%)	2.14
B2.calendar_rf	3/91 (3.3%)	65/91 (71.4%)	5.31
B2.staff_rf	2/75 (2.7%)	52/75 (69.3%)	5.75
B3a.calendar	4/87 (4.6%)	68/87 (78.2%)	2.63
B3a.staff	4/81 (4.9%)	58/81 (71.6%)	3.40
B3b.calendar	1/42 (2.4%)	27/42 (64.3%)	2.96
B3b.staff	1/25 (4.0%)	14/25 (56.0%)	3.53
Average	(4.5%)	(63.3%)	4.20

Table 7.15: Overall RGF of FC models for each data set

From Table 7.14 and Table 7.15, we can make following observations:

(1) When a reliability model is applied to a data set, it usually requires some failure observations in the estimation process for the model to get convergence. On the average, TBF

models take at least 40% of the data to get initial convergence, while FC models only requires 5% of the data. However, even if a model can converge very early, we still recommand that you use at least 30 observations (failures for TBF models or intervals for FC models) or 30% of the total data points for parameter estimation in the model application.

(2) It usually takes quite a few data points before all the models converge, particularly for TBF models. In some cases a TBF model would not converge until a large amount (say, 90%) of the failure data is observed.

(3) Not only do FC models converge earlier than TBF models, but they normally have a higher RGF than TBF models. In general, the data sets with an earlier converging point would have a larger RGF.

(4) The failure data based on real software bugs in the B2 project experience better reliability growth both on calendar time basis and on staff time basis.

(5) The B1 project tracks reliability growth well during testing, using either calendar time, staff time, or test-related time (test cases executed, messages sent by the system). This consistency indicates that the particular B1 release has a well-controlled testing procedure and a smoothly-conducted testing activity, which is confirmed when evaluating the software engineering process of the project.

(6) For the B2 and B3 projects, there is better tracking of reliability growth when time is measured as "staff time" rather than "calendar time." For the B1 project, however, this trend is reversed. In general, staff time information is more expensive to collect since it is normally done manually.

Table 7.16 summarizes the overall RGF, average through all the 13 data sets in this section, for all the models. From Table 7.16 we can see that JM and MB are more optimistic than the other TBF models. Moreover, the TBF models are in general more pessimistic than the FC models.

Model	RGF
GEO	2.89
JM	3.62
LV	2.84
MB	3.22
MO	2.73
Average for TBF models	3.05
BMB	4.19
BMP	4.63
GP	4.18
NHPP	3.87
YSS	4.18
Average for FC models	4.20

Table 7.16: Overall RGF from each model for all data set

7.5 Linear Combination of Model Results

Our other finding is that linear combinations of model results, even in its simplest format, appear to provide more accurate predictions than the individual models themselves[Lyu91c]. Basically, we adopt the following strategy in forming combination models:

1. Identify a basic set of models (the component models). If you can characterize the testing environment for the development effort, select models whose assumptions are closest to the actual testing practices.
2. Select models whose predictive biases tend to cancel each other. As previously described, models can have optimistic or pessimistic biases.
3. Separately apply each component model to the data.
4. Apply criteria you have selected to weight the selected component models (e.g. changes in the prequential likelihood) and form the combination model for the final predictions. Weights can be either static or dynamically determined.

In general, this approach is expressed as a mixed distribution,

$$\tilde{f}_i(t) = \sum_{j=1}^{n} \omega_j(t)\, \tilde{f}_i^{j}(t) \qquad (7.11)$$

where n represents the number of models, $\tilde{f}_i^j(t)$ is the predictive probability density function of the jth component model, given that $i-1$ observations of failure data have been made. Note that $\sum_j \omega_j(t) = 1$ for all t's.

The linear combination model tends to preserve the features inherited from its component models. Also, because each component model performs reliability calculations independently, the combination model remains fairly simple. The component models are plugged into the combination model only at the last stage for final predictions.

Selecting appropriate component models is, of course, important to the success of the combination model. The parameter-estimation method you select to implement the component models may, to a certain extent, affect the combination model's prediction validity. We felt that the Goel-Okumoto (GO), Musa-Okumoto (MO), and Littlewood-Verrall (LV) models are best candidates for our linear combination models. We select them because in our investigations, we found that their predictions were valid[Lyu92a]. Other practitioners have also found that they perform well, and they are widely used[AIAA93a]. Another reason is that they represent different model categories. GO, which is similar to JM, SM and MB, represents the exponential-shape NHPP model, MO represents the logarithmic-shape NHPP model, and LV represents the inverse-polynomial-shape Bayesian model. Finally, at least with the data set we analyzed, the biases of these models tend to cancel out. GO tends to be optimistic, LV tends to be pessimistic, and MO might go either way.

We experiment with three types of combinations. The goal of each is to reduce the risk of relying on a specific model, which may produce grossly inaccurate predictions, while retaining much of the simplicity of using the component models.

1. Statically-weighted combinations.
2. Dynamically-weighted combinations, in which weights are determined by comparing and ranking model results.

3. Dynamically-weighted combinations, in which weights are determined by changes in model evaluation criteria.

These types of combinations are further described in the following sections.

7.5.1 Statically-Weighted Linear Combinations

This type of model is the simplest combination to form. Each component model has a constant weighting which remains the same throughout the modeling process. The main statically-weighted combination is the Equally-Weighted Linear Combination (ELC) model, which is formed by the arithmetic average of all the component models' predictions. Namely, ELC = $\frac{1}{3}GO + \frac{1}{3}MO + \frac{1}{3}LV$. This model follows a strategy similar to that of a Delphi survey, in which authorities working independently are asked for an opinion on a subject, and an average of the results is taken.

7.5.2 Weight Determination Based on Ranking Model Results

Combination models may produce more accurate results if the weights are dynamically assigned rather than remaining static throughout the modeling process. One way of dynamically assigning weights is based on simply ranking component model results. If a combination model contains "n" components, choose a set of "n" values that can be assigned to the components based on a ranking of model results. One of the combinations is the Median-Weighted Linear Combination (MLC), formed by the following: For each failure, the component models would be run, and the results of the models would then be compared. The models predicting the highest and lowest times to the next failure would then be given weights of 0 in the combination, while the prediction in the middle would be given a weight of 1.

The other combination of this type is the Unequally-Weighted Linear Combination (ULC) model. This model is similar to MLC except that optimistic and pessimistic predictions contribute to the final prediction. The prediction is not determined solely by the median value. Here

we use weightings similar to those in the *Program Evaluation and Review Technique*: $\frac{1}{6} O + \frac{4}{6} M + \frac{1}{6} P$, where O represents an optimistic prediction, P, a pessimistic prediction, and M, the median prediction.

7.5.3 Weight Determination Based on Changes in Prequential Likelihood

The last type of combination model is the one in which weights were both dynamically determined and assigned. The basis for determining and assigning weights was changes in the prequential likelihood (see Chapter 4) over a small number of observations.

In the DLC model, we assume that the applicability of any individual model to the project data may change as testing progresses. Therefore, the component models' weights will change according to changes in a model's applicability. Here, we use changes in prequential likelihood – a measure that denotes a model's accumulated accuracy – to assign weights to the component models, which could be taken over a few or many time frames. As a baseline, we formed the simplest DLC model by choosing an observation window of one time frame before each prediction as the reference in assigning weights.

7.5.4 Sample Modeling Results

In order to determine the validity of the linear combination modeling scheme, we use six models selected in Section 7.3.3 as a reference group to compare with the experimental group of linear combination models. We use data set SYS3 in Appendix B as a sample to compare modeling results. Table 7.17 show the result of this application.

Data Set 3 (207 data points/starting data – 60)										
Measure	JM	GO	MO	DU	LM	LV	ELC	ULC	MLC	DLC
Accuracy	-811.1 (4)	-811.2 (7)	-811.1 (4)	-814.3 (10)	-811.3 (8)	-812.7 (9)	-810.8 (2)	-810.8 (2)	-811.1 (4)	-809.1 (1)
Bias	.0835 (8)	.0761 (6)	.0586 (1)	.0994 (10)	.0829 (7)	-.0845 (9)	.0640 (4)	.0594 (3)	.0586 (1)	.0649 (5)
Trend	.0623 (7)	.0663 (9)	.0487 (5)	.0740 (10)	.0602 (6)	.0630 (8)	.0467 (2)	.0474 (3)	.0480 (4)	.0462 (1)
Noise	5.384 (9)	5.209 (8)	4.088 (5)	2.426 (1)	6.002 (10)	3.714 (2)	4.224 (7)	4.196 (6)	4.073 (4)	3.901 (3)
Rank	(6)	(8)	(4)	(9)	(9)	(6)	(4)	(3)	(2)	(1)

RECOMMENDED MODELS: 1. DLC 2. MLC 3. ULC 4. ELC 4. MO

Table 7.17: Model Comparisons for Data Set SYS3

In Table 7.17, numbers in each row represent the computed measure under each criterion, with ranks in parentheses corresponding to the models in columns. The last row, "Rank", was determined by equally treating all the four criteria. Note that the "starting data" indicate when the model predictions began; previous data points were used for parameter estimations. This starting point was chosen such that a small but reasonable set of data points could be used for the parameter estimations. It is observed from this table that the proposed linear combination models performed relatively well compared with the other six models.

7.5.5 Overall Project Results

Tables 7.18 and 7.19 list the performance comparisons for the three data sets from[Musa80a] and the five data sets from JPL (J1 through J5). The overall comparison is done by using all four measures in Table 7.18, or by using the prequential likelihood measure (the "Accuracy" criterion) alone in Table 7.19, since it was judged to be the most important one. In general, we consider a model as being satisfactory if and only if it is ranked 4 or better out of the 10 models for a particular project. To extend this idea, we define a "handicap" value, which is calculated by subtracting 4 (the "par" value) from the rank of a model for each data set before its ranks being summed up in the overall evaluation. (Or subtract 32 from the "Sum of Rank" row in Tables 7.18 and 7.19.) A negative handicap value represents satisfactory overall performance for the eight data

sets.

Summary of Model Ranking for Each Data by All Four Criteria										
Model	JM	GO	MO	DU	LM	LV	ELC	ULC	MLC	DLC
SYS1	(10)	(9)	(1)	(6)	(8)	(6)	(4)	(2)	(3)	(5)
SYS2	(9)	(10)	(6)	(7)	(8)	(1)	(4)	(5)	(2)	(2)
SYS3	(6)	(8)	(4)	(9)	(9)	(6)	(4)	(3)	(2)	(1)
J1	(10)	(7)	(6)	(7)	(9)	(2)	(2)	(4)	(5)	(1)
J2	(5)	(7)	(10)	(6)	(9)	(4)	(1)	(3)	(8)	(2)
J3	(8)	(6)	(6)	(8)	(10)	(1)	(1)	(1)	(4)	(5)
J4	(5)	(5)	(8)	(1)	(9)	(10)	(1)	(5)	(4)	(3)
J5	(1)	(5)	(1)	(9)	(3)	(10)	(8)	(7)	(3)	(6)
Sum of Rank	54	57	42	53	65	40	25	30	31	25
"Handicap"	+22	+25	+10	+21	+33	+8	-7	-2	-1	-7
Total Rank	(8)	(9)	(6)	(7)	(10)	(5)	(1)	(3)	(4)	(1)

Table 7.18: Overall Model Comparisons Using All Four Criteria

Summary of Model Ranking for Each Data Using the Accuracy Measure										
Model	JM	GO	MO	DU	LM	LV	ELC	ULC	MLC	DLC
SYS1	(10)	(9)	(2)	(8)	(6)	(7)	(5)	(4)	(3)	(1)
SYS2	(7)	(9)	(4)	(10)	(7)	(1)	(4)	(4)	(3)	(2)
SYS3	(4)	(7)	(4)	(10)	(8)	(9)	(2)	(2)	(4)	(1)
J1	(10)	(7)	(6)	(8)	(9)	(2)	(3)	(4)	(5)	(1)
J2	(5)	(7)	(9)	(10)	(5)	(4)	(2)	(3)	(8)	(1)
J3	(6)	(5)	(8)	(10)	(6)	(2)	(3)	(4)	(8)	(1)
J4	(6)	(6)	(6)	(2)	(6)	(5)	(3)	(4)	(6)	(1)
J5	(2)	(6)	(2)	(10)	(2)	(9)	(8)	(7)	(2)	(1)
Sum of Rank	50	56	41	68	49	39	30	32	39	9
"Handicap"	+18	+24	+9	+36	+17	+7	-2	0	+7	-23
Total Rank	(8)	(9)	(6)	(10)	(7)	(4)	(2)	(3)	(4)	(1)

Table 7.19: Overall Model Comparisons by the Accuracy Measure

There are several important points we can observe from these summary tables:

1. There are two sets of models under investigation here: the set of single models, and the set of combination models. In general, the set of combination models perform better than the

set of single models. The acceptable models (those with a negative "handicap"), when considering all four measuring criteria (Table 7.18), are exactly the four linear combination models. When considering the Accuracy criterion alone (Table 7.19), the three acceptable models, DLC, ELC, ULC, also belong to the combination model set. By evaluating the handicap value, we also note that the combination models usually beat the other single models by a significant number of "strokes."

2. By weighting or averaging the predictions from the three well-known component models, GO, MO, and LV, the combinational models appear to be less sensitive to potential data noise than their component models and other single models. This is reflected in the investigated data sets which include both execution-time based data and calendar-time based data. Moreover, when we examine all project data for the evaluation criteria, we can see that the combination models could *sometimes* outperform all their component models, but they *never* perform worse than the worst component model.
3. The DLC and ELC Models perform rather consistently. Most other models seem to perform well for a few data sets but poorly for other data sets, and the fluctuation in performance is significant. By preserving good properties from the three well-known models with equal weightings, the ELC model achieves a good overall performance as expected. On the other hand, since the DLC model is allowed to dynamically change its weightings according to the outcome of the accuracy measure, it is not surprising to see it consistently produce the best accuracy measure for almost every data set. This consistency suggests that, if you use whatever accuracy measures you deem the most important as the weighting criterion in forming the DLC model, you will get the best results.

7.5.6 Extensions and Alternatives

You can extend or alter our basic approach in the following ways:

1. Extend the DLC model by increasing the size of the observation window from one time frame to N time frames. The DLC model consistently produces the best accuracy measure,

but with only one observation window, it might fail to note a global measurement trend. Thus, a natural extension is to enlarge the window.

2. Try to apply models other than GO, MO, and LV as component models. If some models perform well in a particular data set, they should be the candidate component models to form a combination model.

3. Use more than three models as component models. We believe that the more component models you apply, the better the prediction. However, more computations are required, and the returns may diminish as more models are added.

4. Apply alternative weighting schemes that are based on project criteria and engineering judgments. Our approach is flexible enough that you can decide how you want to form a combination model.

5. Use the combination models themselves as component models to form another combination model.

6. As the original assumptions behind each model become lost through the layers of linear combinations, a distribution free (nonparametric) modeling technique may emerge.

In our investigation, the most promising approach was to extend the DLC model. We considered a DLC model with a fixed N window, DLC/F, and a DLC model with a sliding N window, DLC/S. Figure 7.7 shows how the two models differ.

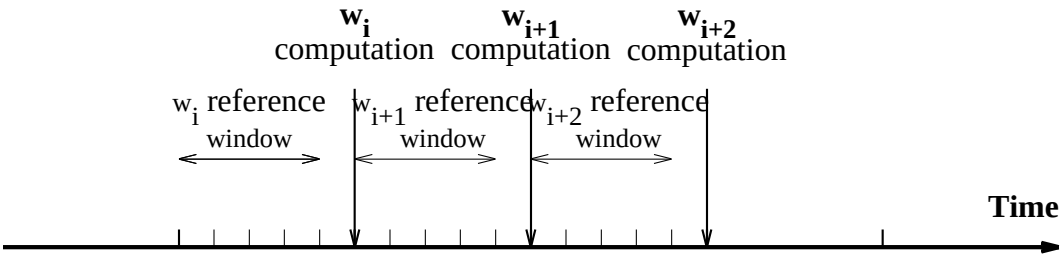

Figure 7.7(a): The DLC Model with a Fixed Observation Window

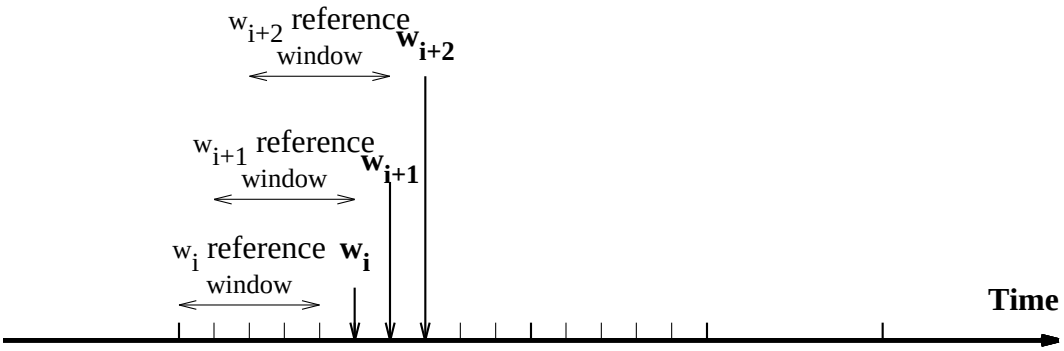

Figure 7.7(b): The DLC Model with a Sliding Observation Window

In the DLC/F window, the weight assignments for each model are based on changes in the accuracy measure over the last N observations. The weight assignment for each model remains fixed for the next N predictions. At the end of that time, the weights are recomputed according to the changes in accuracy over the last N observations. To compute the weight of a component model, you first determine the amount of change in component model A's accuracy measure over the last N observations. You then identify component model B, the component model whose accuracy measure changed the most. The unnormalized weight for A is simply the ratio of the change in its accuracy measure to the change in B's accuracy measure.

In the DLC/S model, you recompute the weight assignments for each model at each data point, using changes in the accuracy measure over the last N observations as the basis for determining each model's weight. To compute weights for component models, the procedure is the same as that in the DLC/F model.

Figure 7.8 summarizes the accuracy measure of the DLC/F and DLC/S type models, normalized with respence to the number of measured points in each data set before being summed up for the eight data sets.

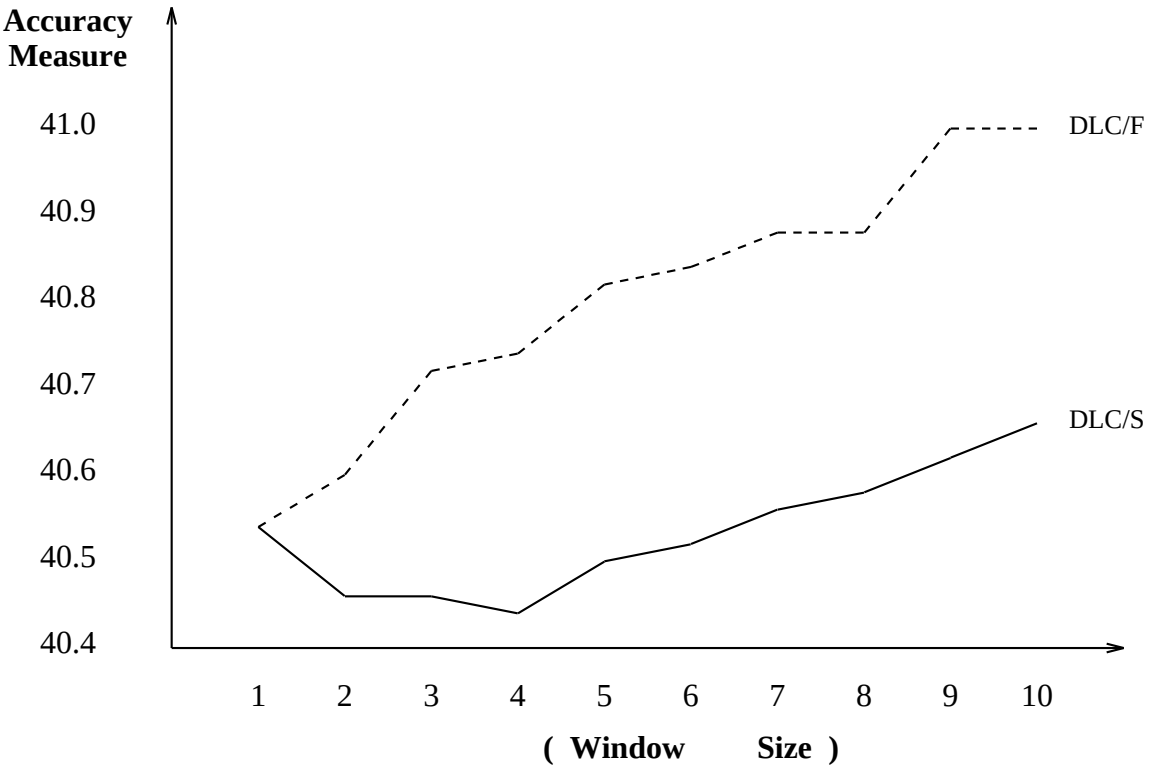

Figure 7.8: Summary of the DLC/F and DLC/S Models for Windows up to 10 Time Frames

As Figure 7.8 shows, the DLC/S model is generally superior to the DLC/F model. This result is not surprising, since DLC/S allows the observing window to advance dynamically as step-by-step prediction moves ahead. In general, the accuracy of the DLC/F model deteriorates when the window becomes larger. The DLC/S model's performance, on the other hand, improves when the window becomes larger, but only slightly larger. We found that a window size of three to four time frames is optimal.

Of course, the best window size depends on your development environment, testing scheme, and operational profile, but, in general, the window size should be fewer than five time frames, since the model is then able to catch fast shifts in model applicability among the component models.

The accuracy measure in Figure 7.8 is the prequential likelihood, but other accuracy measures, such as the Akaike information criterion – another criterion to denote how close a prediction is to the actual data [Khos89a] – or mean square error, are also feasible. The main strength of the DLC models is that they combine component models in a way that lets the output be fed back for model adjustment.

The fundamental approach of the linear combination models is simple. However, by applying more complicated procedures, we risk losing the individual model's assumptions about the physical process. It then becomes harder to get insight into the process of reliability engineering. Most reliability models view software as a black box, form which to observe failure data and make predictions. In that context, our combination models do not degrade any properties assumed in current reliability-modeling practices.

7.5.7 Long-Term Prediction Capability

Our results showed that the combination models performed well in making step-by-step predictions – in which you can adjust the model's parameters for each prediction – but we also wanted to determine how they performed in making long-term predictions, say 20 failures ahead. For this evaluation, we selected the ELC and DLC models and compared them with the GO, MO, and LV component models. Figure 7.9 shows the prediction curve for each model for the data set J3.

Figure 7.9: Long-Term Predictions for Data Set J3 from Several Models

We used the first 152 data points J3, or up to 777 cumulative test hours as indicated by the dashed line, to estimate each model's parameters. Immediately following this estimation stage is the prediction stage. For the project J3, these two stages follow the project's natural breakdown into two testing stages.

For the DLC model, we computed model preferences and weights in the estimation phase, and fixed the weight assignments in the prediction phase.

LV's prediction curve is too pessimistic, and GO's and MO's are too optimistic. In fact, all three curves for the component models are out of the actual project data curve. ELC and DLC, on the other hand, compensate these extremes and make rather reasonable long-term predictions.

To show quantitative comparisons of long-term predictions, we use mean square error instead of

prequential likelihood. Prequential likelihood is more appropriate for comparing step-by-step predictions, while the mean square error provides a more widely understood measure of the distance between actual and predicted values.

Table 7.20 shows the summary of long-term predictions. The value in each project raw represents the mean square errors for a model applied to the project data in that raw, and the model ranking is shown in parentheses. The values under "Sum of Mean Square Errors" and "Sum of Ranks" indicate that the ELC and DLC models generally perform better than the component models. Even though the component models make a better prediction than the ELC and DLC models on several occasions, they also perform significantly worse on others. The ELC and DLC models, on the other hand, never make the worst long-term predictions.

Summary of Model Ranking for Long-Term Predictions Using Mean Square Errors					
Data Model	GO	MO	LV	ELC	DLC
SYS1	2117(5)	687.4(4)	567.7(3)	266.7(2)	169.7(1)
SYS2	1455(5)	1421(4)	246.1(1)	930.5(2)	955.7(3)
SYS3	480.0(2)	253.2(1)	2067(5)	745.5(3)	779.8(4)
J1	1089(4)	782.9(2)	5283(5)	130.1(1)	876.7(3)
J2	4368(4)	4370(5)	539.3(1)	2171(3)	1791(2)
J3	4712(5)	3073(3)	4318(4)	1322(2)	1141(1)
J4	3247(4)	3248(5)	219.5(1)	1684(3)	1354(2)
J5	60.22(3)	60.12(1)	104.45(5)	68.44(4)	60.15(2)
Sum of Mean Square Errors	17528.5	13896.6	13345.0	7317.3	7128.3
Sum of Ranks	(32)	(25)	(25)	(20)	(18)
Overall Rank	(5)	(4)	(3)	(2)	(1)

Table 7.20: Summary of Long-Term Predictions

7.6 Summary

We have initiated SRE programs at JPL and Bellcore for better specification and tracking of software reliability for different projects. We lay out the framework of this SRE process, which includes the most current state of practice in industry. In applying the reliability-objective-

setting method and the software reliability modeling and measurement techniques, we obtain some modeling results which look promising.

These results indicate that model performances are dramatically different depending on the context of different projects, and the use of multiple models is deemed necessary. Moreover, the reliability growth phenomenon is better demonstrated when the SRE mechanism is put in place. It is noted that the data collection process is the key to successful measurement of software reliability. In particular, failure data should be scrutinized for better classification of real software failures against the current release, and more data should be collected automatically for accuracy and reasons of economy. The application of software reliability tools also greatly simplifies the tedious job of reliability measurement and model comparison. Finally, the linear combination modeling scheme is introduced as a simple technique to increase accuracy in software reliability measurement.

In summary, we believe that when the SRE application receives more attention and wider implementation in industry, more insights leading to improvement in product quality and software development process will gradually emerge.

Problems

1. Analyze the failure reporting and tracking mechanisms used by a software development organization with which you're familiar. Compare the data collected by these mechanisms with the minimum set of failure data described in this chapter. Report on the suitability of your failure reporting and tracking mechanisms for software reliability estimation, and describe any changes that would result in a more suitable mechanism.

2. How has your development organization, or an organization you know, set reliability requirements? How do the requirements compare to the two examples given in this chapter (Section 7.2.1.1)? Which of the three methods of setting requirements (Section 7.2.1.2) has

the organization used, or was another method used? Why was the method chosen?

3. How does your development organization, or the organization you know, estimate and forecast software reliability? How does the organization choose the model(s) used? How does the method of selection compare to that presented in this chapter?

4. Using the guidelines given in this chapter, write a plan that your organization could use to guide them in collecting data for software reliability measurement.

5. What will the function $\mu(t)$ be if the following parameters are given in Example 7.4?

measurement period	a	b
First Stage (weeks 1-14)	310	.050
Second Stage (weeks 15-40)	400	.045
Third Stage (weeks 41-60)	450	0.035

6. Given a project data set, the estimate of b is 1/57.6 in Equation (7.5) for the applied GO model. The loaded salary for a single tester is 1000 dollars per day, and there are 6 full-time testers in the project. Further suppose that the cost of fixing a fault during testing is 200 dollars, and that for fixing a fault in the field is 2200 dollars. It is estimated that there are a total of 1300 faults in the software.

   (a) Draw the curves similar to Figure 7.2. Write down the equations for each curve.

   (b) Draw the two curves representing the two sides of Equation (7.5).

   (c) When should the testing be stopped?

   (d) Repeat (b) and (c) for 10 testers and 2 testers, respectively.

7. Comparing with prequential likelihood, what is the advantages and disadvantages in using Reliability Growth Factor (RGF) to determine the model validity during testing?

8. Why do combination models provide better results, on average, than individual models? Can you think of other methods to increase the prediction accuracy of models? What limitations do combination models have? Under what circumstances are these limitations important?

9. The linear combination models are applied mainly to one type of data. Is it TBF data type or FC data type? Why is it difficult for the other type? For this data type, discuss which linear combination models could be formed easily, and which could not.

References

Abde86a.

Abdel-Ghaly, A.A., Chan, P.Y., and Littlewood, B., ‘‘Evaluation of Competing Software Reliability Predictions,’’ *IEEE Transactions on Software Engineering*, vol. SE-12, pp. 950-967, September 1986.

AIAA93a.

AIAA,, in *Recommended Practice for Software Reliability*, AIAA, February 1993.

Butl91a.

Butler, R.W. and Finelli, G.B., ‘‘The Infeasibility of Experimental Quantification of Life-Critical Software Reliability,’’ in *Proceedings of ACM SIGSOFT’91 Conference on Software for Critical Systems*, pp. 66-76, New Orleans, Louisana, December 1991.

Dala88a.

Dalal, S.R. and Mallows, C.L., ‘‘When Should One Stip Testing Software,’’ *Journal of the American Statistical Association*, vol. 83, no. 403, pp. 872-879, September 1988.

Dala90a.

Dalal, S.R. and Mallows, C.L., ‘‘Some Graphical Aids for Deciding When to Stop Testing Software,’’ *IEEE Journal on Selected Areas in Communications*, vol. 8, no. 2, pp. 169-175, February 1990.

Dala94a.

Dalal, S.R. and McIntosh, A.A., ‘‘When to Stop Testing for Large Software Systems with Changing Code,’’ *IEEE Transactions on Software Engineering*, vol. 20, no. 4, pp. 318-323, April 1994.

Duan64a.

Duane, J.T., ‘‘Learning Curve Approach to Reliability Monitoring,’’ *IEEE Transactions on Aerospace*, vol. AS-2, pp. 563-566, 1964.

Ehrl90a.

Ehrlich, W., Stampfel, J., and Wu, J., ‘‘Application of Software Reliability Modeling to Product Quality and Test Process,’’ in *Proceesing 12th International Conference on Software Engineering*, Nice, France, March 1990.

Ehrl93a.

Ehrlich, W., Prasanna, B., Stampfel, J., and Wu, J., ‘‘Determining the Cost of a Stop-Test Decision,’’ *IEEE Software*, pp. 33-42, March 1993.

Farr83a.

Farr, W.H., ‘‘A Survey of Software Reliability Modeling and Estimation,’’ Technical Report 82-171, NSWC, 1983.

Farr88a.

Farr, W.H. and Smith, O.D., ‘‘Statistical Modeling and Estimation of Reliability Functions for Software (SMERFS) User’s Guide,’’ TR 84-373, Revision 1, NSWC, December 1988.

Goel79a.

Goel, A.L. and Okumoto, K., ‘‘Time-Dependent Error-Detection Rate Model for Software Reliability and Other Performance Measures,’’ *IEEE Transactions on Reliability*, vol. R-28, pp. 206-211, 1979.

Jeli72a.

Jelinski, Z. and Moranda, P.B., ‘‘Software Reliability Research,’’ in *Statistical Computer*

Performance Evaluation, ed. W. Freiberber, pp. 465-484, Academic, New York, 1972.

Khos89a.

Khoshgoftaar, T.M. and Woodcock, T.G., ‘‘A Simulation Study of the Performance of the Akaike Information Criterion for the Selection of Software Reliability Growth Models,’’ in *Proceedings the 27th Annual South East Region ACM Conference*, pp. 419-423, April 1989.

Leve90a.

Levendel, Y., ‘‘Reliability Analysis of Large Software Systems: Defect Data Modeling,’’ *IEEE Transactions on Software Engineering*, vol. SE-16, no. 2, pp. 141-152, February 1990.

Litt73a.

Littlewood, B. and Verrall, J.L., ‘‘A Bayesian Reliability Growth Model for Computer Software,’’ *Journal Royal Statistics Society C*, vol. 22, pp. 332-346, 1973.

Litt81a.

Littlewood, B., ‘‘Stochastic Reliability Growth: A Model for Fault-Removal in Computer Programs and Hardware Designs,’’ *IEEE Transactions on Reliability*, vol. R-30, pp. 313-320, October 1981.

Litt86a.

Littlewood, B., Abdel-Ghaly, A.A., and Chan, P.Y., ‘‘Tools for the Analysis of the Accuracy of Software Reliability Predictions,’’ in *Software System Design Methods*, pp. 299-335, Springer-Verlag, Heidelberg, 1986.

Lyu91b.

Lyu, M.R., ‘‘Measuring Reliability of Embedded Software: An Empirical Study with JPL Project Data,’’ in *Proceedings International Conference on Probabilistic Safety Assessment and Management*, pp. 493-500, Beverly Hills, California, February 1991.

Lyu91c.

Lyu, M.R. and Nikora, A., ‘‘A Heuristic Approach for Software Reliability Prediction: The Equally-Weighted Linear Combination Model,’’ in *Proceedings 1991 International Sympo-*

sium on Software Reliability Engineering, pp. 172-181, Austin, Texas, May 1991.

Lyu91a.

Lyu, M.R. and Nikora, A., ‘‘Software Reliability Measurements Through Combination Models: Approaches, Results, and A Case Tool,’’ in *Proceedings the 15th Annual International Computer Software and Applications Conference (COMPSAC’91)*, pp. 577-584, Tokyo, Japan, September 1991.

Lyu92a.

Lyu, M.R. and Nikora, A., ‘‘Using Software Reliability Models More Effectively,’’ *IEEE Software*, pp. 43-52, July 1992.

Lyu92b.

Lyu, M.R. and Nikora, A.P., ‘‘CASRE – A Computer-Aided Software Reliability Estimation Tool,’’ in *CASE 92 Proceedings*, pp. 264-275, Montreal, Canada, July 1992.

Lyu93a.

Lyu, M.R., Nikora, A., and Farr, W., ‘‘A Systematic and Comprehensive Tool for Software Reliability Modeling and Measurement,’’ in *Proceedings FTCS-23*, pp. 648-653, Toulouse, France, June 1993.

Musa79a.

Musa, J.D., ‘‘Validity of Execution-Time Theory of Software Reliability,’’ *IEEE Transactions on Reliability*, vol. R-28, no. 3, pp. 181-191, August 1979.

Musa80a.

Musa, J.D., ‘‘Software Reliability Data,’’ RADC Technical Report, 173 pp., DACS, Rome Air Development Center, 1980.

Musa84a.

Musa, J.D. and Okumoto, K., ‘‘A Logarithmic Poisson Execution Time Model for Software Reliability Measurement,’’ in *Proceedings Seventh International Conference on Software Engineering*, pp. 230-238, Orlando, Florida, 1984.

Musa87a.

Musa, J.D., Iannino, A., and Okumoto, K., *Software Reliability – Measurement, Prediction, Application*, McGraw-Hill Book Company, New York, New York, 1987.

Niko92a.

Nikora, A.P., Lyu, M.R., and Antczak, T.M., ‘‘A Linear Combination Software Reliability Modeling Tool with A Graphically-Oriented User Interface,’’ *Proceedings of Symposium on Assessment of Quality Software Development Tools*, pp. 21-31, New Orleans, Louisana, May 1992.

Rapp90a.

Rapp, B., ‘‘Application of Software Reliability Models in Medical Imaging Systems,’’ in *Proceedings 1990 International Symposium on Software Reliability Engineering*, Washington, D.C., April 1990.

Schn75a.

Schneidewind, N.F., ‘‘Analysis of Error Processesin Computer Software,’’ in *Proceedings International Conference on Reliable Software*, pp. 337-346, Los Angeles, 1975.

Shoo73a.

Shooman, M., ‘‘Operational Testing and Software Reliability During Program Development,’’ in *Proceedings 1973 IEEE Symposium on Computer Software Reliability*, pp. 51-57, New York, April 1973.

Yama83a.

Yamada, S., Ohba, M., and Osaki, S., ‘‘S-Shaped Reliability Growth Modeling for Software Error Detection,’’ *IEEE Transactions on Reliability*, vol. R-32, pp. 475-478, December 1983.

Zinn90a.

Zinnel, K.C., ‘‘Using Software Reliability Growth Models to Guide Release Decisions,’’ in *Proceedings 1990 International Symposium on Software Reliability Engineering*, Washington, D.C., April 1990.