Cost-Benefit Modelling for Reliability Growth

John Quigley*, Lesley Walls
Department of Management Science
University of Strathclyde
40 George Street
Glasgow UK
G1 1QE

(tel): +44 141 548 3152
(fax): +44 141 552 6686
e-mail: john@mansci.strath.ac.uk

Abstract

Decisions during the reliability growth development process of engineering equipment involve trade-offs between cost and risk. However slight, there exists a chance an item of equipment will not function as planned during its specified life. Consequently the producer can incur a financial penalty. To date, reliability growth research has focussed on the development of models to estimate the rate of failure from test data. Such models are used to support decisions about the effectiveness of options to improve reliability. The extension of reliability growth models to incorporate financial costs associated with 'unreliability' is much neglected.

In this paper, we extend a Bayesian reliability growth model to include cost analysis. The rationale of the stochastic process underpinning the growth model and the cost structures are described. The ways in which this model can be used to support cost-benefit analysis during product development are discussed and illustrated through a simple case.

Keywords: reliability; cost benefit analysis; stochastics

Introduction

The time-to-market of complex engineering product can be several years as a result of the need for long development processes to mature the design. Testing is an important part of the development process and usually involves both growth tests and demonstration tests. The former aims to expose the design to a range of stresses to expose weaknesses quickly so that corrective action can be taken to remove or mitigate the effect of faults thereby growing the reliability. While the latter aims to demonstrate that the mature reliability conforms to the specified customer or regulatory requirements. Both tests are important as they provide evidence of assurance that a design has improved and so unscheduled failure events are unlikely in operation.

Testing is costly in terms of time and resources required. This is particularly true for growth tests where the duration and conditions will be subject to considerable uncertainty given the reactive nature of the test decision process and the complexity of relationships between the failure mechanisms and the operating environment. As a result there is much scepticism surrounding the marginal benefits of growth testing in relation to its cost. Despite this, growth testing continues to be widely used because industry adopts a risk averse approach to reliability and believes that however expensive it is to remove a fault during design, it is likely to be much more expensive to incur the costs of its realisation as a failure in operation (see Hobbs[1]).

Many models exist to support reliability growth decision-making, see Jewell[2], Xie[3], Ansell et al[4] for a review and critical appraisal. However these models aim to provide information about the effectiveness of testing for improving reliability by measuring increases in the observed rate of failure. Both classical and Bayesian inference procedures have been proposed. However those models using classical inference demand statistical test data that can be, at best, late and, at worst, sparse in practice. Bayesian approaches intend to reduce the reliance on observed failures by combining what failure data that become available with prior data based on engineering belief about some aspect of expected reliability growth. While attractive in principle, because this approach allows us to acknowledge designs that are inherently good and so free from major defects, many models using Bayesian approaches tend to make strong assumptions about the form of the prior distributions for reasons of mathematical tractability that are not meaningful to practitioners (Walls and Quigley[5]). Therefore the data structures assumed lack credibility. To overcome these criticisms Quigley and Walls[6] proposed a reliability growth model which aims to combine engineering judgement about the inherent concerns with the design to be tested with observed failures on test that has nice theoretical properties and has been accepted by practitioners. Nevertheless this model again only aims to assist in assessing the effectiveness of reliability growth test. There is a need to extend such a model to support decisions concerning the financial worth of tests and allowing cost-benefit analysis to be conducted.

While there is considerable interest in the trade-offs between reliability and cost within the product lifecycle (Kumar et al[7]), there is no reported evidence of cost-benefit analysis of reliability growth tests and our experience indicates that practical analysis is rather informal and haphazard. In this paper we propose to extend earlier work (Quigley and Walls[6,8]) and develop a unified modelling framework to support effective and efficient decision making during a reliability growth test programme that addresses many of the shortcomings of existing models. To specify our framework we begin by describing the decision-making process inherent in reliability growth development programmes. Next the formulation of the mathematical model is introduced and an example application of the model within a Bayesian context is presented. We conclude by reflecting upon the strengths and weaknesses of the proposed model and directions for future research.

Product Development Decision-Making Process

The reliability plan provides the blue print for developing a product design that will meet customer reliability requirements in terms of functionality. In turn, the reliability plan will provide an input into the overall project plan that will take into account other customer requirements such as delivery schedules and cost of the product. Here we focus on the reliability plan and decisions about reliability related activities only. The typical activities included in a reliability plan and their purpose are described in O'Connor[9].

In the design phase one output of the reliability plan may be a risk register of potential concerns about weaknesses or faults in the design. This will evolve from a list of concerns about the product design constructed initially during concept phase and further articulated during detailed design. These may cover issues related to the anticipated strength of the design in relation to the stresses it is likely to experience. For example, use of new technology, changes in manufacturing processes in relation to more extreme environmental conditions such as higher/lower temperatures or increased vibration.

The risk register, or concerns list, will document the beliefs of engineers based on their expertise and experience about potential ways in which the design may fail to function, the reasons for such failure and the consequences of their occurrence. Processes used to capture such beliefs include FMEA (Gilchrist[10]), HAZOP (Kletz[11]) and elicitation of engineering judgement (Walls and Quigley[5]). Consequently the register provides an insight into expected failure scenarios and can help inform re-designs and development testing regimes used to prove the product design.

During development data collected from testing activities provide information about the observed failure events that can be compared with the expectations articulated during design. The reliability plan will stipulate that this assessment of reliability will be conducted at formal design reviews where the information will be used to predict whether the target product reliability required by the customer is likely to be achieved and whether the project plan for timely delivery of a mature product is sufficient.

The key decision at such a review will be informed by an assessment of reliability growth to establish if the planned development activities are maturing the product design. If reliability assessment concludes that expected growth is being achieved and faults are being exposed and corrected successfully then the decision-maker will progress the design to the next phase of development and, ultimately, release from test to manufacture. If reliability assessment concludes that reliability growth is not satisfactory then additional testing may be required to prove the design and uncover faults that may still lurk within the design or else provide evidence that the strength of the design is likely to exceed stresses beyond those to be experienced in use.

However the cost of changes to the product, and so the plan, should also be taken into account so that decisions made are both effective and efficient. In principle the decision will be influenced by the likelihood of faults existing in the design and their cost of occurrence if not corrected relative to the cost of corrective actions.

Figure 1 summarises our perspective that focuses on decisions to be taken at design reviews during the development phase, in particular we are interested in supporting the decision maker choose whether or not to continue test. To assess whether reliability targets have been met or whether additional testing is required, and if so how much, the decision maker must be able to evaluate the costs and benefits associated with further testing compared with progressing the design through the development process. The model described in the following section aims to provide a framework to support this decision.

FIGURE 1

Model Formation to Support Cost-Benefit Analysis

The cost model developed in this section aims to support informed decisions about the continuance or stoppage of tests by providing a framework in which the consequences of different test times can be analysed and the optimal length of test time to minimise expected total costs can be determined.

A diagrammatic representation of the model formulation is shown in Figure 2. After stating the assumptions underpinning this formulation we shall show the derivation of the key results.

FIGURE 2

Assumptions

It is assumed that a period of testing, τ, has been conducted. This may equal 0 if no testing has been conducted and events on test are assumed to occur at times $x_1, x_2, x_3 ...$

The testing costs comprise the fixed cost of running the test (e.g. test facilities and staff) and the cost of correcting faults detected on test which will vary depending on the fault and when it is found.

The cost of running the test is assumed deterministic charged at £*C* per unit of time. Costs incurred prior to release from test are accumulated at an assumed constant rate δ to assess their present value at τ test time units. Therefore the accumulated present value of running the test at test time τ is

$$\text{Expected cost of running test} = C\int_0^{\tau} e^{\delta y} dy = C\left[\frac{e^{\delta\tau} - 1}{\delta}\right]$$

The variable cost of fixing faults detected on test is assumed to be £*V* and using the same constant interest rate δ, the present value at τ of the cost associated with correcting the fault detected at test time x_1 is £$Ve^{\delta(\tau - x_1)}$. Clearly, this cost is a random variable dependent upon the time in which a fault is realised. Moreover, the total costs associated with testing will be dependent not only on the detection time but the number of faults detect. As such, we require information about the inherent model of reliability growth during development testing.

Given our earlier comments about the availability of some register of likely faults we advocate a Bayesian reliability growth model since the data captured in the risk register may be combined to form a prior distribution for the number of faults present in the design and the chance they might result in failure in operation (Quigley and Walls[6]). Therefore we assume that the design contains a fixed number of faults and that the time to realise these faults is independently and identically distributed.

Therefore the variable costs accumulated at the end of testing are given by

$$\text{Total variable cost of testing} = V\sum_{i=1}^{N(\tau)} e^{\delta(\tau - X_{(i)})}$$

where

$X_{(i)}$ is the time of the i^{th} fault detected on test

$N(\tau)$ is the number of faults detected by τ hours of test

The number of faults detected on test and the times of these realisations are the only random variables regarding costs associated with testing. The expected variable cost of testing is derived in the following.

$$\begin{aligned}\text{Expected variable cost of testing} &= VE\left[\sum_{i=1}^{N(\tau)} e^{\delta\left(\tau - X_{(i)}\right)}\right] \\ &= VE_{N(\tau)}\sum_{i=1}^{N(\tau)} E_{X_{(i)}}\left[e^{\delta\left(\tau - X_{(i)}\right)} \middle| N(\tau)\right]\end{aligned}$$

Conditioning on the number of faults that occur during τ hours of testing we can treat the failure times as independent and identically distributed random variables with the distribution truncated at time τ, allowing us to ignore the complications of the order statistics.

$$\begin{aligned}\text{Expected variable cost of testing} &= Ve^{\delta\tau} E_{N(\tau)} \sum_{i=1}^{N(\tau)} E_{X_{(i)}} \left[e^{-\delta X_i} \middle| N(\tau) \right] \\ &= Ve^{\delta\tau} E_{N(\tau)} \sum_{i=1}^{N(\tau)} L_{X|\tau}(\delta) \\ &= Ve^{\delta\tau} L_{X|\tau}(\delta) E\left[N(\tau) \right]\end{aligned}$$

where

$$L_{X|\tau}(\delta) = \frac{\int_0^\tau e^{-\delta x} dF(x)}{F(\tau)}$$

is the Laplace Transform of the truncated distribution evaluated at the interest rate and

$$F(\tau) = \int_0^\tau dF(x)$$

is the probability that a fault is realised within the interval (0,τ) of testing.

In order to derive the expectation of $N(\tau)$ we first define an indicator variable Y_i as the following

$$Y_i = \begin{cases} 1, \text{ if } X_i \le \tau \\ 0, \text{ if } X_i > \tau \end{cases}$$

We assume that the times of fault realisation are i.i.d. with the number of expected faults detected after an infinite number of hours test, i.e. $N(\infty)$, being λ. Therefore,

$$\begin{aligned} E\left[N(\tau)\right] &= E_{N(\infty)} E_{Y_1} .. E_{Y_{N(\infty)}} \left[\sum_{i=1}^{N(\infty)} Y_i \middle| N(\infty) \right] \\ &= E_{N(\infty)} \sum_{i=1}^{N(\infty)} E\left[Y_i \middle| N(\infty)\right] \\ &= E_{N(\infty)} \left[N(\infty) F(\tau) \right] \\ &= \lambda \mathrm{F}(\tau) \end{aligned}$$

Combining the expected costs of running the test with those of detecting faults we obtain the expected total costs of testing

$$\text{Expected Total Costs of Testing} = C\left[\frac{e^{\delta\tau} - 1}{\delta}\right] + Ve^{\delta\tau}\ \lambda \mathrm{F}(\tau) \mathrm{L}_{X|\tau}(\delta)$$

We assume that manufacture does not introduce faults and so once operational it is assumed the product will fail J times at times $t_1, t_2, \ldots, t_J$, respectively. The cost to the manufacturer at the time of each failure is £P. However, the present value of this penalty, evaluated at the time of releasing the item to market is $Pe^{-\delta t}$ for each failure, where δ is a measure of the interest rate of the value of money.

Therefore the total cost associated with the in-service failures of the item, evaluated at the time of releasing the item to market is

$$\text{Total In-service Cost} = P\sum_{i=1}^{J} e^{-\delta t_i}$$

Taking the expectation we obtain

$$\begin{aligned}\text{Expected Total In-service Cost Given J} &= P\sum_{i=1}^{J} E\left[e^{-\delta t_i}\right] \\ &= PJL_T\left(\delta\right)\end{aligned}$$

where $L_T(\delta)$ is the Laplace Transformation of the assumed distribution of the time to realise a fault in operation, evaluated at the continuous rate of interest δ.

Again we assume the prior distribution on the number of faults that may be realised as failures is a Poisson distribution with mean λ and if we release the item from test then the expected number of faults remaining in the design that is released to operation will be $\lambda(1-F(\tau))$.

Consequently the expectation of the total in-service costs with respect to *J* is given by:

$$\text{Expected Total In-service Cost} = PL_T\left(\delta\right)\lambda\left(1-F\left(\tau\right)\right)$$

Cost model

Following the above assumptions, the total costs, *TC*, resulting from test and operation is given by

$$TC = C\left[\frac{e^{\delta\tau}-1}{\delta}\right]+Ve^{\delta\tau}\ \lambda F(\tau)L_{X|\tau}(\delta)+PL_T(\delta)\lambda\big(1\text{-}F(\tau)\big) \tag{1}$$

By differentiating (1) with respect to test time τ we obtain

$$\frac{\partial\ TC}{\partial\tau} = \ Ce^{\delta\tau}+\delta Ve^{\delta\tau}\ \lambda F(\tau)L_{X|\tau}(\delta)+V\lambda f(\tau)\text{-}PL_T(\delta)\lambda f(\tau) \tag{2}$$

Setting (2) to 0 and re-arranging we obtain

$$e^{\delta\tau}\left[C+V\delta\lambda F(\tau)L_{X|\tau}(\delta)\right]=\lambda f(\tau)\left[PL_T(\delta)-V\right]$$

This can be interpreted as equating the marginal costs of testing one more unit of time with the expected marginal benefit associated with fault detection at time τ.

Alternatively, we can write this as a decision rule whereby testing is stopped if the intensity function (or rate of occurrence of faults detected) falls below a critical level.

$$\lambda f(\tau) \leq \frac{e^{\delta\tau}\left[C + V\delta\lambda F(\tau) L_{X|\tau}(\delta)\right]}{PL_T(\delta) - V}$$

There are two possible local extrema obtained through this approach. To examine the optimatisation properties we consider the second partial derivative with respect to τ.

$$\frac{\partial^2 TC}{\partial \tau^2} = C\delta e^{\delta\tau} + Ve^{\delta\tau}\ \ \delta^2\lambda F(\tau) L_{X|\tau}(\delta) + V\delta\lambda f(\tau) + \left(V - PL_T(\delta)\right)\lambda \frac{df(\tau)}{d\tau}$$

Therefore, we will have obtained a local minimum if the following condition is met

$$C\delta e^{\delta\tau} + Ve^{\delta\tau}\ \ \delta^2\lambda F(\tau) L_{X|\tau}(\delta) + V\delta\lambda f(\tau) + \left(V - PL_T(\delta)\right)\lambda \frac{df(\tau)}{d\tau} > 0 \tag{3}$$

Re-arranging (3) we obtain the following.

$$\frac{df(\tau)}{d\tau} < \frac{C\delta e^{\delta\tau} + Ve^{\delta\tau}\ \ \delta^2\lambda F(\tau) L_{X|\tau}(\delta) + V\delta\lambda f(\tau)}{\left(V - PL_T(\delta)\right)\lambda} \tag{4}$$

If the variable cost of testing, *V*, are greater than the present value of the penalty associate with detecting the fault during operation then clearly no testing would be done. As such, we consider the right hand side of (4) must be negative. This implies that the local extrema will be a cost minimisation if the rate of change of the intensity function is decreasing at a suitably fast rate. In short, from the two possible roots of the first order conditions, the test time that minimises cost will typically be the greater test time.

Illustrative Example

Statement of problem

As an example consider the use of the cost model based on a reconstruction of a real reliability growth development test programme. The test regime is essentially test, analyse and fix, whereby the item is tested until it fails. The failure is investigated, appropriate corrective action and the resultant reliability growth modelled and assessed. The decision-maker is the project manager who wants to determine test duration for a prototype item manufactured to build standard at the specified stress level given the expected costs of test and future failures in relation to the potential benefits gained in terms of reliability growth. The original project plan allocated 1000 hours to test, however now that the detailed design is complete and information is available about the engineers concerns about potential weaknesses, the project manager wishes to re-assess the test plan duration.

Data selection and validity of assumptions

A Poisson distribution with mean, λ, equal to 7 was found to be an adequate description of for the prior distribution of the concerns of engineers about potential faults that was elicited prior to the start of test. Reflections on the process used to elicit this prior are reported in Hodge et al[12].

In the absence of test data, we assume that the distribution function of the time to realise failures on the proposed test can be approximated by an estimated distribution function for data from a nominally identical test of an earlier generation of the design. This data set comprises 12 failures and a total test time of 41183 hours. The best fit model is an exponential distribution with mean time to failure of μ = 3432 hours. Therefore,

$$F(\tau) = \frac{1}{3432}\exp\left(-\frac{\tau}{3424}\right)$$

We use an empirical Bayes approach to re-scale this model in the light of the prior distribution elicited for the proposed test. The posterior distribution is given by

$$\pi(\mu / \underline{x}) \propto L(\mu, \underline{x})\pi(\mu)$$

where a prior of ignorance is assumed setting $\pi(\mu) = 1$ and giving a Gamma posterior distribution

$$\pi(\mu / \underline{x}) = \frac{41183^{13}\mu^{12}e^{-41183\mu}}{\Gamma(13)}$$

which is shown in Figure 3 for a selection of values of the hazard rate, μ. Hence the predictive distribution is

$$\begin{aligned} F(\tau) &= \int_0^\infty F(\underline{x} / \mu)\pi(\mu / \underline{x})d\mu \\ &= 1 - \frac{41183^{13}}{(41183 + \tau)^{13}} \end{aligned}$$

The corresponding intensity function is given by

$$z(\tau)=\frac{13\lambda(41183)^{13}}{(41183+\tau)^{14}}$$

As might be expected this function exhibits an exponential decay in the rate of fault realisation during the reliability growth test.

FIGURE 3

The general form of the Laplace transform of this predictive distribution is messy and as we simply illustrate it as a function of test time as shown in Figure 4. This was calculated using Maple 7 for an assumed interest rate $\delta = 0.05\text{p.a.}$.

FIGURE 4

We know the cost of running a test is C = £500 per hour and the cost associated with fixing faults on test are estimated to be roughly V = £50000 on average.

The penalty associated with realising faults in operation are much harder to cost and so we may want to think about the magnitude of these costs in relation to the cost of fixing the fault during development. This leads us to define $P = kV$ where k is a constant multiplier. It is of interest to examine what values of k lead to a switch in decisions between continuing and stopping test.

The final input data required is the net present value of costs incurred fixing faults in operation. Assuming the same interest rate of $\delta = 0.05\text{p.a.}$ and using field failure time data for a variant product already in service which has a exponential distribution with mean time between failure of 30000 hours, we find $L_T\left(0.05\right) = 0.84$. The project manager estimates that the cost of realising a fault in the field will be at least ten times the cost of detecting a fault in test, therefore we use k = 10. That is, the ratio of penalty associated with realising a fault in operation (P) to the cost of fixing a fault detected during test (V).

Model implementation and interpretation

Figure 5 shows that prior to about 1500 hours it is better to continue testing as the net benefits of testing outweigh the costs of finding faults in the field. However after about 1500 hours then testing is no longer efficient and should be stopped. Therefore it would appear that the plan to terminate test at 1000 hours may be premature.

FIGURE 5

For this design, the planned duration of the test is worthwhile if the cost of testing is less than 7 times that of the costs of detecting a fault within operation. If this is not acceptable to the project manager then again it reinforces the need for longer test.

Most uncertainty in specifying parameter values for such a cost model will be in estimating the penalty of faults being realised in use, even when it is expressed as a multiple (k) of the cost of detecting a fault on test. Table 1 also shows the relationship between total costs, the optimal test time and the multiple k for increments of k when the expected number of faults in the system (λ) is 7, the costs associated with an hour of test (C) is £500, the cost of correcting a fault during test (V) is £50000, and the mean time to realise a fault during operation is 30000 hours. Constructing such a table of information for given test parameters should assist decision-making because it permits the project manager to assess the sensitivity of the total cost to relative costs of realising failures in the field.

TABLE 1

The relationships between test time and the multiple k on total costs can be explored more generally through Figure 6. This shows that a flatter cost profile for small values of k compared with larger values where clear minima are exhibited. This is intuitive as small differentials in the costs of faults being realised on test or field would imply that testing has little value. Whereas larger values of k imply longer test durations are economic.

FIGURE 6

The sensitivity associated with the inputs can be assessed through Table 2. If the cost of testing, i.e. C, were halved to £250 per hour the optimal test time would be about 2.5 times greater than at the assumed input values for other parameters, while an increase of 50% to £750 per hour would result in a 83% reduction in planned testing. These are illustrated in Figure 2a.

Figures 2b and 2c illustrate the sensitivity associated with test time and costs for variations in the cost associated with the detection of a fault during test, i.e. V, and the expected number of faults within the design, i.e. λ, respectively. It is worth noting that the cost of detecting a fault in operation is expressed as being proportional to V through k. The optimal length of testing is much more sensitive to underestimates in V, in so far as reducing V by 50% resulted in no testing, while an increase by 50% resulted in 80% more planned testing. Not surprisingly, decreasing the expected number of faults reduces the optimal test time while increasing reverses this effect. Changes in the mean time to detection of a fault in operation had little effect on the optimal test time, where a mean time of 15000 hours resulted in a 20% increase in planned test, while an increase to 45000 hours resulted in a 12% decrease in planned testing.

TABLE 2

Discussion on Use of Model and Conclusions

We have presented a simple expected present value model to evaluate the financial benefits of reliability growth testing. We accept that the decision to terminate testing will be based upon criteria in addition to costs. In this section we highlight and discuss some shortcomings of the proposed model and discuss how these might be overcome.

We do not consider severity of failures within the model. The faults that we are addressing are ones, which can exist within an item that we release into operation. Catastrophic faults can be identified in two ways during development. Firstly, relevant experts can anticipate them during an elicitation exercise. Secondly, they can be exposed during testing. For those faults that are identified during the first situation we propose that this framework would be applied to support decisions concerning which test is likely to expose these faults with the minimal costs. The decision would not be when to stop testing as this would be decided once the probability of safety critical faults existing within the item being below a threshold, however, the sequencing and durations of various reliability tests must be decided and costs will be a consideration.

We have assumed that the cost associated with detecting a fault within operation is fixed, known and applies to all faults. We could extend this model to make the costs a random variable, but a more effective manner to address this issue would be through categorising types of faults following the elicitation process. Following from this, we could group the possible faults, which are similar in penalties if exposed and similar in lifetime characteristics.

We have assumed that the cost of testing is constant per unit of time. The model can be extended to support a more complex expression for the costs associated with testing. For example, the costs associated with diagnostics once a fault has been exposed increases the total costs as personnel are moved from one job to devote time to the item under test.

We have assumed that the manufacturer is risk neutral and hence their utility towards risk can be measured through total costs. Future research in this area would investigate and develop a loss function, capturing vital characteristics of the risk aversion inherent in a company for the release of particular items with the possible existence of identified faults.

Finally, although the model has been proposed for use within the development phase when conducting test, the same principles underpin cost-benefit modelling during design. Currently we are working with industrial partners to operationalise the modelling framework in this broader context.

Acknowledgements

We would like to thank the UK Department of Trade and Industry and the REMM consortium for financial and industrial support during the development of this work.

References

1. Hobbs GK (2000) *Accelerated Reliability Engineering – HALT and HASS*, Wiley, Chichester.

2. Jewell, W (1984) "A General Framework for Learning Curve Reliability Growth Models", *Operations Research*, **32**, 547-558.

3. Xie M. (1993) "Software Reliability Models – a Selected Biography", *Software Testing, Verification and Reliability*, **3**, 3-28.

4. Ansell J., Walls, L.A., and Quigley J.L. (1999) "Achieving Growth in Reliability", *Annals of Operational Research*, **91**, 11-24.

5. Walls L.A. and Quigley J.L. (2001) "Eliciting Prior Distributions to Support Bayesian Reliability Growth Modelling – Theory and Practice", *Reliability Engineering and System Safety*,**74**, 117-128

6. Quigley, J.L. and Walls L.A. (1999) "Measuring the effectiveness of reliability growth testing", *Quality Reliability Engineering International*,**15,** 87-93

7. Kumar U.D., Crocker J., Knezevic J. and El-Haram M. (2000) *Reliability, Maintenance and Logistic Support: A Lifecycle Approach*, Kluwer Academic Publishers, Boston.

8. Quigley J.L. and Walls L.A. (2003) "Confidence Intervals for Reliability Growth Models with Small Sample Sizes*", IEEE Transactions in Reliability,* **52**, 257-262.

9. O'Connor P.D.T. (1996) *Practical Reliability Engineering*, Wiley, New York.

10. Gilchrist, W (1993) Modelling Failure Mode and Effects Analysis, *International Journal of Quality and Reliability Management*, **10**, 16-23.

11. Kletz T.A. (1997) "Hazop – Past and Future", *Reliability Engineering and System Safety*, **55**, 263-6.

12. Hodge, R, Evan E, Marshall J, Quigley J and Walls L (2001) "Eliciting Engineering Knowledge about Reliability during Design – Lessons Learnt From Implementation", *Quality and Reliability Engineering International*, **17**, 169-179.

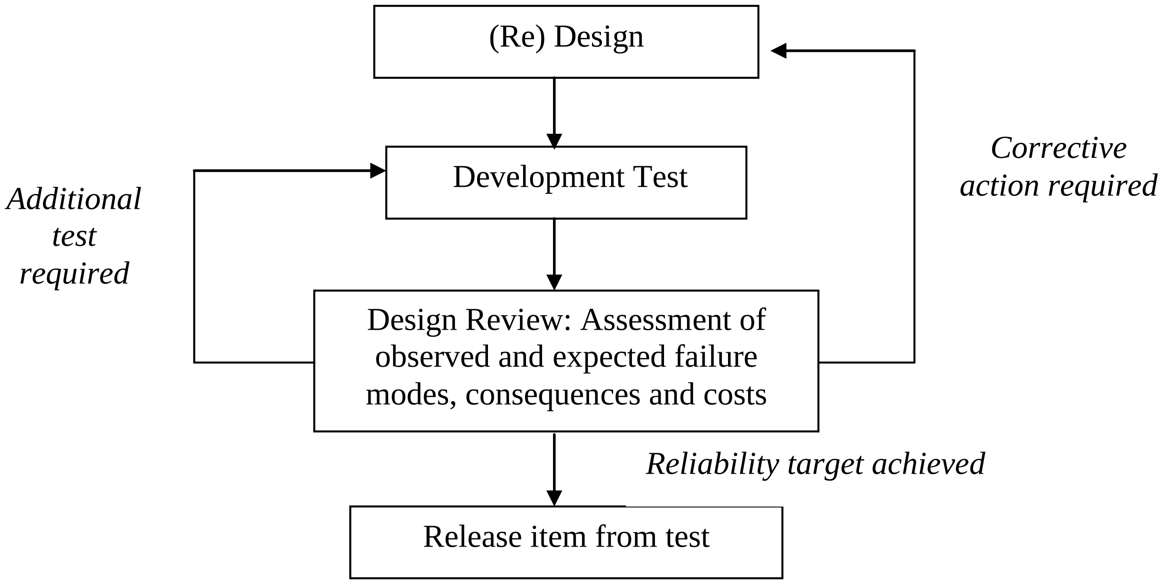

Figure 1 Decision points during development process

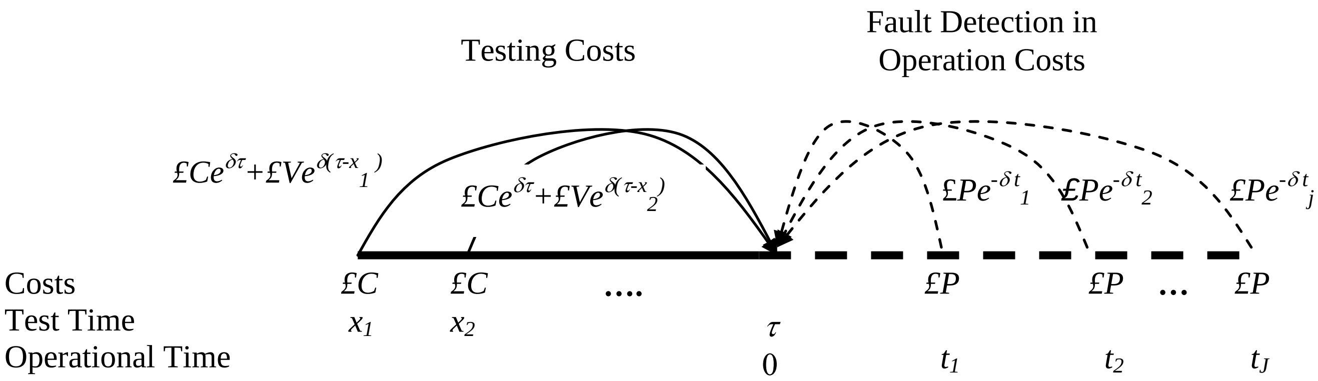

Figure 2 Cost model formulation

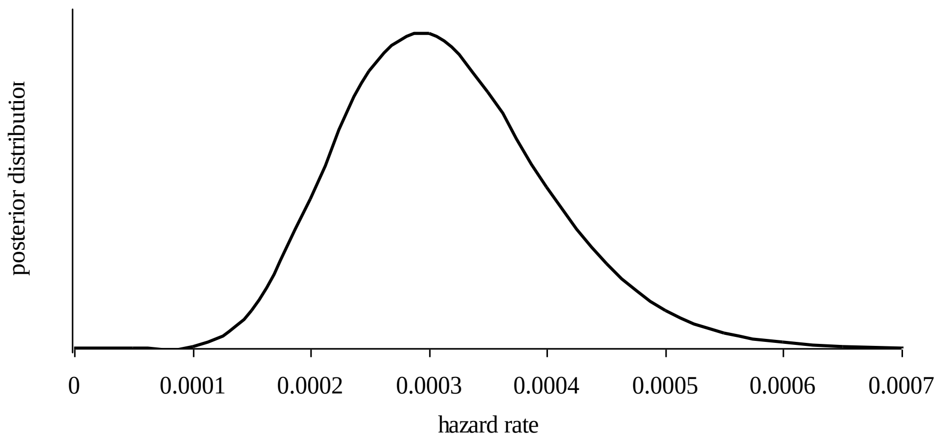

Figure 3 Posterior distribution of hazard rate

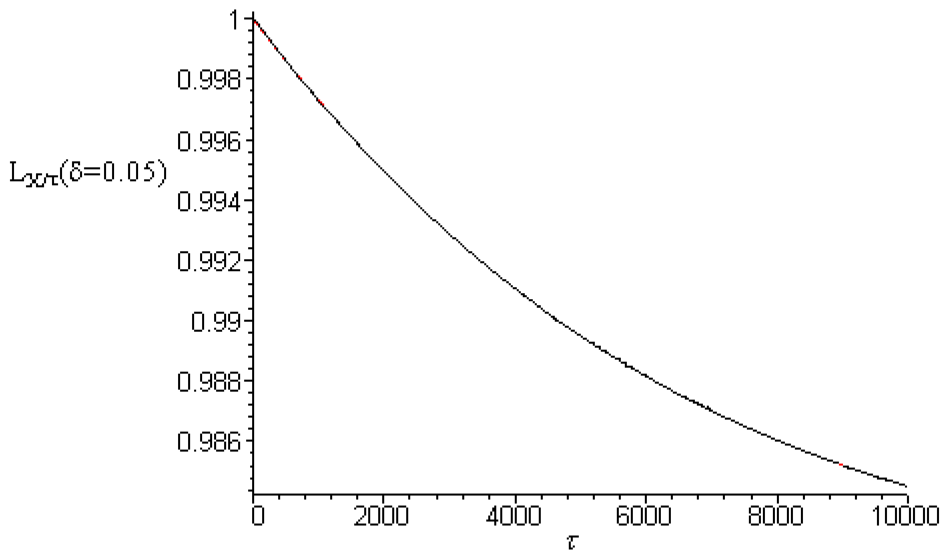

Figure 4 Laplace transform of truncated distribution at 5% p.a. interest rate as function of test time

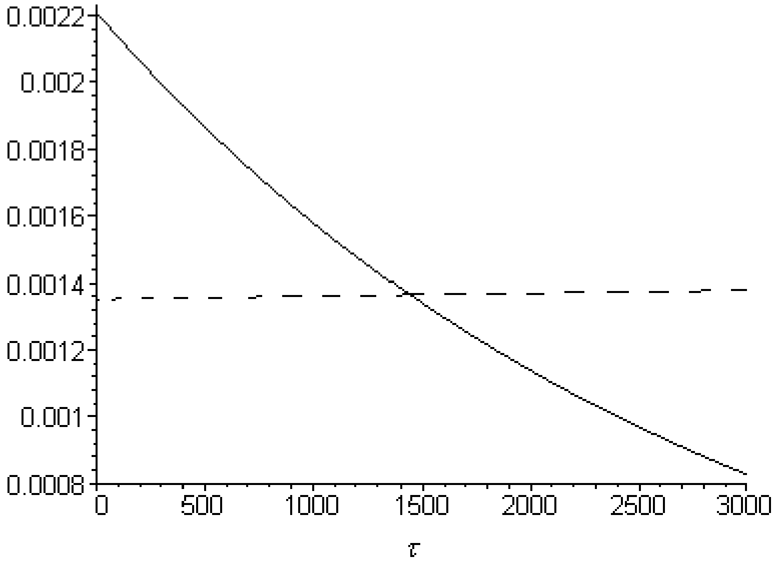

Figure 5 Comparison on intensity function (solid) and cost ratio (dashed) for k =10

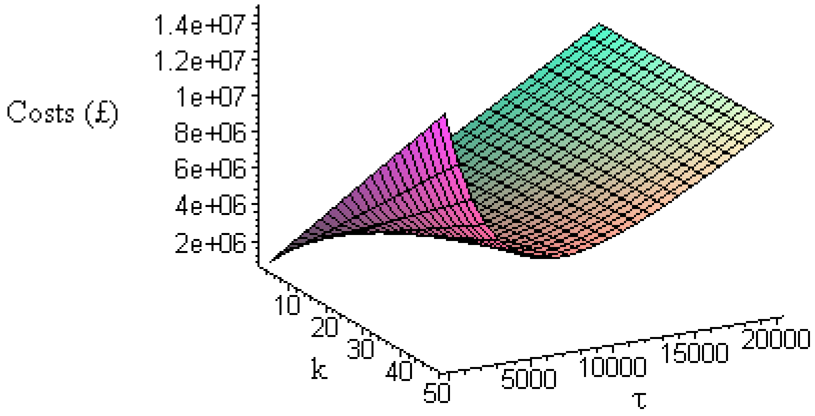

Figure 6 Total cost plane as a function of test time and *k*

Table 1 Minimum Total Costs and Optimal Test Time as a function of *k*

Ratio of Cost of Detecting a Fault in Operation to Cost of Detecting a Fault on Test (k)	Optimal Test Time (hrs)	Minimum Total Costs (£M)
5	0	1.5
7.5	458	2.2
10	1443	2.7
12.5	2195	3.1
15	2806	3.5
17.5	3322	3.8
20	3769	4.0
22.5	4164	4.3
25	4518	4.5

Table 2a Optimal costs for variations on cost of testing C

C	£250/hr		£750/hr	
Ratio of Cost of Detecting a Fault in Operation to Cost of Detecting a Fault on Test (k)	Optimal Test Time (hrs)	Minimum Total Costs (£M)	Optimal Test Time (hrs)	Minimum Total Costs (£M)
5	1008	1.4	0	1.5
7.5	2523	1.8	0	2.2
10	3556	2.1	250	2.9
12.5	4345	2.4	982	3.5
15	4985	2.5	1576	4.0
17.5	5526	2.7	2078	4.4
20	5995	2.8	2512	4.8
22.5	6409	2.9	2896	5.1
25	6780	3.1	3240	5.4

Table 2b Optimal costs for variations on cost of detecting a fault on test V

V	£25,000		£75,000	
Ratio of Cost of Detecting a Fault in Operation to Cost of Detecting a Fault on Test (k)	Optimal Test Time (hrs)	Minimum Total Costs (£M)	Optimal Test Time (hrs)	Minimum Total Costs (£M)
5	0	0.7	170	2.2
7.5	0	1.1	1655	3.0
10	0	1.5	2668	3.6
12.5	140	1.8	3441	4.0
15	723	2.2	4070	4.4
17.5	1215	2.4	4600	4.7
20	1641	2.7	5060	4.9
22.5	2017	2.9	5465	5.2
25	2355	3.1	5829	5.4

Table 2c Optimal costs for variations on expected number of faults λ

λ	5		9	
Ratio of Cost of Detecting a Fault in Operation to Cost of Detecting a Fault on Test (k)	Optimal Test Time (hrs)	Minimum Total Costs (£M)	Optimal Test Time (hrs)	Minimum Total Costs (£M)
5	0	1.1	0	1.9
7.5	0	1.6	1196	2.7
10	450	2.1	2198	3.3
12.5	1185	2.5	2964	3.7
15	1783	2.8	3586	4.0
17.5	2287	3.1	4110	4.3
20	2724	3.3	4565	4.6
22.5	3110	3.6	4967	4.8
25	3455	3.8	5327	5.0