6: Relaxation of the common failure rate assumption in modelling software reliability

T Giammo

Information Management and Technology Division
US General Accounting Office
Washington DC
US

A theorem recently published by Lindsay furnishes the conditions for solving maximum likelihood formulations which seek to find a maximising distribution of some parameter without placing any restriction on the distribution other than that it be a legitimate probability distribution. Reliance on this theorem in the context of the Jelinski/Moranda approach to the estimation of software reliability permits the formulation to be revised, replacing the assumption that each fault is equally likely to cause a failure with the less limiting assumption that each fault possesses its own potentially unique failure rate parameter chosen from some probability distribution common to all the faults. This paper introduces Lindsay's theorem and illustrates its use by applying it to Jelinski's and Moranda's original software reliability model.

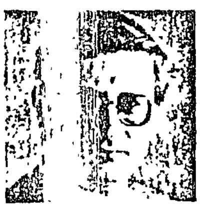

T Giammo

Thomas Giammo is currently an Associate Director in the General Accounting Office of the US Government where he is responsible for reviewing federal government policies and practices involving information technology. He has previously held several management positions at the Department of Health and Human Services and served as a member of the President's Reorganisation Project. He has also gained considerable technical experience as Technical Director of the Federal Computer Performance Evaluation and Simulation Center, Vice President of Colton Systems Research Group and as an Operations Research Analyst for TRW Inc. Mr Giammo has a BS from Rensselaer Polytechnic Institute and an MA from UCLA, both in Mathematics. He has published several papers on various aspects of queuing theory and computer performance modelling, including a contribution to a previous Pergamon Infotech State of the Art Report. He devotes a considerable amount of his personal time to independent research.

Relaxation of the common failure rate assumption in modelling software reliability

Introduction

From the time the first software reliability measurement models were introduced over a decade ago, the frequently made assumption that each fault in the system contributes equally at every point in time to the probability of system failure has been viewed as a serious limitation. Various approaches have been proposed to overcome this limitation by modifying this common failure rate assumption to permit changes, as faults are removed, in the estimated failure rate of the remaining faults. These 'reliability growth' models, however, have been based upon specific assumptions as to the 'growth process' that have little, if any, practical justification.

Based on recent theoretical work of Lindsay and of the author, this paper demonstrates how the common failure rate assumption may be directly replaced by the non-limiting and intuitively acceptable assumption that each fault possesses a potential unique failure rate of its own, independently drawn from some unknown distribution. No restrictions whatsoever need to be placed on the unknown distribution — any legitimate probability distribution is an admissable candidate. The times of failure which occur in some period of observation can then be used to find an estimator of the distribution of failure rates of all the faults. Having an estimator of this underlying distribution permits the calculation of various useful statistics regarding the reliability of the system at the termination of the observation period.

Until recently, no method was known of finding a maximum likelihood estimator of such an underlying distribution without the introduction of additional limiting assumptions. Lindsay's theorem, introduced in his 1983 article, provides such a method. This theorem is quite general and is applicable to a broad range of maximum likelihood formulations. To avoid adding complications to the explanation of the technique of utilising Lindsay's theorem, in this paper it is applied only to the straightforward, but somewhat simplistic, reliability model used by Jelinski and Moranda in their original work. Application of the theorem to the more sophisticated extensions of this formulation that have evolved since their original work can be accomplished following the general approach demonstrated here.

This paper introduces the theoretical basis in the form of Lindsay's and related theorems. It then runs through the Jelinski/Moranda (J/M) formulation in a manner compatible with the application of Lindsay's theorem. Various deficiencies in the J/M model are discussed so as to make clear the motivation for the relaxation of J/M's common failure rate assumption. The relaxation of this assumption and the application of Lindsay's theorem are then discussed in detail. There follows a brief description of the numerical algorithm used in the solution and the paper concludes with the presentation of the results of some simulations that illustrate the improvement brought about by the relaxation of the common failure rate assumption.

This paper is a summary of a considerable body of work, some of which has not been previously published. The intention is to introduce the basic concepts and techniques in a form that can be readily extended and applied. Although some effort has been directed to exhibiting the overall logic of the derivations of the new mate. l, some of the individual steps are not immediately obvious. In the interest of keeping this paper withi. easonable bounds, the lengthy analysis needed to support these steps has not been displayed.

Lin ty's theorem

In thi ection, the main theorem developed by Lindsay *(GIA1)* is stated without proof, along with several relat results useful in applying this theorem to software reliability models. Both the theorem and the relat results will be given in the form independently developed by the author *(GIA2)*. We shall begin with some efinitions. C is some compact subset of the 'reals', $P\{C\}$ is the space of all probability distributions on C. 'is the set of all functions on P of the following form:

$$\psi(\pi) = k\prod_{i=1}^{m} y_i(\pi) \tag{1}$$

wher ı is some finite integer, π belongs to $P\{C\}$, k is some constant independent of π and each $y_i(\pi)$ can be writt n in the form

$$y_i(\pi) = \int_C f_i(x)d\pi(x) \tag{2}$$

and, further, each f_i is a function on C into the reals such that:

- $f_i(x) \geqslant 0$ everywhere
- $f_i(x)$ is everywhere continuous and bounded
- $f_i(x)$ is not everywhere zero.

Additionally:

1 The probability distribution ϱ, belonging to P, is said to be a *maximising probability distribution with respect to some function* ψ belonging to Y if, and only if, $\psi(\varrho) = \max\{\psi(\pi)|\pi \in P\}$.

2 $R(\psi)$ is the subset of P consisting of all maximising probability distributions with respect to ψ.

3 $B(\pi)$ is the *support base of a probability distribution* π if, and only if, $B(\pi) = \{x|x \in C$ and every neighbourhood of x has a non-zero probability measure in $\pi\}$.

We can now state the basic theorems in terms of the above notation.

1 *Theorem 1*: for each ψ belonging to Y, there is at least one maximising distribution, that is $R(\psi)$ is non empty.

2 *Theorem 2* (Lindsay's theorem): for each ψ belonging to Y, there is associated:

- a unique set of positive, non-zero, real values, $\{z_i|i = 1,2,\ldots,m\}$
- a subset of C, $\beta(\psi)$, defined by

$$\beta(\psi) = \{x \mid \sum_{i=1}^{m} \frac{f_i(x)}{z_i} = m\} \tag{3}$$

such that

the following condition holds:

$$x \notin \beta(\psi) \Leftrightarrow \sum_{i=1}^{m} \frac{f_i(x)}{z_i} < m \tag{4}$$

a necessary and sufficient condition that a probability distribution π' be a maximising probability distribution with respect to ψ (that is $\pi' \in R(\psi)$ is that $B(\pi') \subseteq \beta(\psi)$ and $y_i(\pi') = z_i$ for each $i = 1,2,\ldots,n$.

3 *Theorem 3*: if for each open set in C there is no set of positive, non-zero real coefficients $\{\alpha_i|i = 1,2,\ldots,n\}$ such that the equation

$$\sum_{i=1}^{n} \alpha_i f_i(x) = n \tag{5}$$

is satisfied over the entire open set, then every member of $R(\psi)$ is a discrete distribution without a density component.

Jelinski/Moranda revisited

As mentioned previously, Lindsay's theorem is a general result that permits the reformulation of many problems originally posed in terms of finding specific values of parameters that maximise a likelihood of a set of observations. Using this theorem, one can often replace a limiting assumption that all of the observations of the set are governed by a common value of a parameter by the less restrictive assumption that each observation of the set is associated with a potentially unique random value of the parameter chosen from some common distribution. Lindsay's theorem offers the necessary and sufficient conditions to recognise the probability distribution that maximises the likelihood of a set of observations without the necessity of limiting the space of admissible underlying probability distributions in any way.

In order that the effect of applying Lindsay's theorem can be clearly illustrated in the context of the maximum likelihood techniques used in estimating software reliability measures, we shall first develop the 'classical' J/M formulation *(GIA3)* to serve as a base.

The 'classical' J/M formulation arises from the following assumptions:

1 There are initially m faults present in the system.

2 The time to failure of each fault is an independent, identically distributed random variable. The failure time distribution is negative exponential, with a rate of failure parameter μ.

3 When a failure occurs, the associated fault is immediately removed — or, equivalently, the clock is stopped until the fault is removed.

At some time t we wish to estimate m and μ based on the number of failures that have occurred and the times of their occurrence. Following J/M, we choose as estimators those values which maximise the likelihood of the observed number of failures and the failure times. As we shall see, this can be readily accomplished in two phases:

1 By developing the expression for the estimator of μ which maximises the likelihood of the observations, given a fixed value of m.

2 By substituting this expression, then finding the value of m which maximises the resulting likelihood expression.

Without loss of generality, we can simplify the expressions involved in the solutions if we assume that the time variable is scaled such that the period of observation ends at $t = 1$. For each of n observed failures, we would therefore have $0 \leq t_i \leq 1$, for $i = 1,2,\ldots,n$.

The individual likelihood factor y_i associated with an observed failure at time t_i is

$$y_i = \mu \exp(-\mu t_i) \tag{6a}$$

The likelihood factor z associated with each fault which did not fail during the observation period is

$$z = \exp(-\mu) \tag{6b}$$

Thus, given that the system initially has m faults, the likelihood of n failures occurring at $t_1, t_2, \ldots, t_n$ is

$$\psi = \frac{m!}{(m-n)!} \mu^n \exp[-\mu(m-n + \sum_{i=1}^{n} t_i)] \tag{7}$$

It is useful to introduce the following changes of variables so that cases in which $m \rightarrow \infty$ can be more readily handled:

$$x = \exp(-\mu) \tag{8a}$$

$$\sigma = (m-n)/m \tag{8b}$$

As a further simplification, we shall also admit values of σ corresponding to non-integer values of m. We shall, however, retain for the moment the restriction corresponding to keeping m finite, that is $\sigma < 1$.

Making these substitutions into (7) above and then taking the natural logarithm of each side yields

$$\ln(\psi) = n\ln[-\ln(x)] + \sum_{i=1}^{n} t_i \ln(x) + \frac{n\sigma}{(1-\sigma)}\ln(x) - n\ln(1-\sigma) + \sum_{k=1}^{n}\ln[k(1-\sigma) + n\sigma] \quad (9)$$

We can find the value of x associated with the maximum likelihood for any fixed σ by differentiating (9) and setting the result to zero. The maximising value x' is given as

$$x' = \exp[-\frac{n(1-\sigma)}{(1-\sigma)(\Sigma t_i) + n\sigma}] \quad (10)$$

Substituting this back into (9) produces the logarithm of the likelihood in terms of σ alone, which after some simplification of terms takes the following form:

$$\ln(\psi) = \sum_{k=1}^{n}\ln[k(1-\sigma) + n\sigma] + n[\ln(n) - 1 - \ln((1-\sigma)(\Sigma t_i) + n\sigma)] \quad (11)$$

Furthermore

$$\frac{\delta \ln(\psi)}{\delta\sigma} = \sum_{k=1}^{n}\frac{(n-k)}{k(1-\sigma) + n\sigma} - \frac{n(n - \Sigma t_i)}{(1-\sigma)(\Sigma t_i) + n\sigma} \quad (12)$$

It can be shown that (12) has no root in [0,1] if $2(\Sigma t_i) \geqslant n+1$. This would imply that there is no finite value of m which maximises the likelihood. This deficiency can be formally 'corrected' by admitting the point $\sigma = 1$, that is the limit point as $m \rightarrow \infty$. The other calculations related to the maximisation of likelihood are not perturbed by this addition to the domain of σ, since (10), (11) and (12) remain well defined at $\sigma = 1$ in all non-trivial cases (that is, at least one observed failure). If we make this addition to the domain of σ, we can show that the value of σ associated with the maximum likelihood is then given by

$$\sigma' = 1, \text{ if } 2(\Sigma t_i) \geqslant n + 1 \quad (13a)$$

$$\sigma' = 0, \text{ if } \sum_{k=1}^{n}\frac{1}{k} \leqslant \frac{n}{\Sigma t_i} \quad (13b)$$

If neither of the above two conditions is met, the maximising value of σ' is the solution obtained by setting (12) equal to zero and solving (implicitly) for σ.

We can then express the solutions for x' and σ' in terms of the maximising values of μ' and m' by employing the inverses of (8a) and (8b):

$$\mu' = -\ln(x') \quad (14a)$$

$$m' = n/(1-\sigma') \quad (14b)$$

A useful statistic, which we shall use in later comparisons, is the *estimated failure rate at the termination of observations* ζ, which is given by

$$\zeta = (m-n)\mu \quad (15)$$

For values of x' and σ' which maximise the likelihood, this becomes

$$\zeta' = \frac{n^2\sigma'}{(1-\sigma)(\Sigma t_i) + n\sigma'} \quad (16)$$

Of special interest are the cases in which (13a) holds, which we shall refer to as the 'infinite solution' cases. As can be seen from (14b), $\sigma' = 1$ implies that m' is undefined, that is $m' = \infty$. We should note, however, that ζ' remains bounded; specifically, $\zeta' = n$.

Criticisms of the Jeninski/Moranda model

The J/M formulation has been subject to much criticism. Before discussing the extensions to this model made possible by Lindsay's theorem, it is worthwhile to discuss the basis of this criticism for two reasons: first, to understand which criticisms are addressed by the extensions discussed in this paper and, second, to put the remaining criticisms into proper perspective.

A fundamental criticism made of the J/M and related formulations is that the use of maximum likelihood as the criterion for choosing the estimators for the model's parameters is intrinsically inferior to a Bayesian approach *(GIA4, GIA5)*. It is not clear that this line of criticism, on the philosophical level at least, is entirely valid. It can be shown that most maximum likelihood rules for choosing estimators can be recast as 'extended' Bayesian decision rules with suitable choices of a prior distribution and a risk function *(GIA6)*. This has been shown specifically in the case of the J/M model by several authors *(GIA7, GIA8)*. This type of criticism is not addressed by the extensions discussed in this paper.

More to the point is the criticism that, in certain cases, the use of the maximum likelihood criterion produces results that are undesirable, or even unacceptable, from a practical point of view and this can be avoided by the use of a suitable Bayesian decision rule *(GIA7-GIA9)*. The primary problems noted in this regard seem to be that the estimators, especially *m*, can be extremely sensitive to small changes in the observations and can often be 'meaningless', for example the 'infinite solution' cases in which $m = \infty$ and $\mu = 0$.

Some of these criticisms may be overstated. Firstly, many of the difficulties attributed to the use of the maximum likelihood criterion are intrinsic to the state of the problem and are not related to the choice of criteria for the estimators. In view of the fundamental assumption of the J/M model, that is that the individual failure times are random variables with a negative exponential distribution, one cannot with any confidence associate a given 'pattern' of observed failure times with a specific failure rate, *unless the number of observed failures is relatively large*. The apparent ability of some Bayesian formulations to sharply distinguish between failure rates when the number of observations is relatively small is strictly an attribute of the particular prior distribution used — specifically, a small variance in the prior distribution. In effect the Bayesian approach smoothly interpolates between some well-behaved prior distribution (when the number of observations is relatively small) and the J/M solutions (when the number of observations is relatively large). Failing a strong justification for a particular well-behaved prior distribution, however, it is not clear that this is anything more than a cosmetic improvement to cover over the fact that good estimators cannot be reliably found, based on the observations alone, when the number of observed failures is relatively small.

Secondly, the 'infinite solution' ($m = \infty$ and $\mu = 0$), which has disturbed several critics, can in fact be given a 'meaningful' interpretation. As we have shown above, the estimated failure rate at the termination of the observation period remains finite and is equal to n/t, the number of observed failures divided by the duration of the observation period. This can be interpreted as meaning that the time between *system* failures is itself negatively exponentially distributed (with failure rate parameter n/t) and that the estimated number of initial faults is so large that the removal of the observed faults has had no measurable effect. One might plausibly argue that this interpretation still leaves the infinite solution with the 'unacceptable' characteristic that it is not improbable for actual cases with a *small* number of initial faults to give rise to observed failure times which result in the *infinite* J/M solution. However, this is another artefact of the same problem of finding good estimators based on a *small* number of observed failures. Stating from (13a), one can readily show that the probability of this 'contradiction' (that is of an actual case meeting the J/M assumptions and giving rise to the infinite solution) goes to zero as the period of observation and/or number of observed failures increase.

Thirdly, and perhaps most importantly, much of the criticism seems to be motivated by the fact that the approach has not given consistently good results in practice, even when the number of observed failures is relatively high. It is undeniable that this criticism is well founded. The real difficulty here, however, may well be 'unrealistic' assumptions of the underlying failure model rather than its use of the maximum likelihood criterion. As has been noted, previously, extensions of the model in the direction of imperfect removal of faults, introduction of new faults etc address some aspects of this problem *(GIA1[illegible]GIA13)*. These extensions, however, retain the limiting assumption that the time to failure of each fault is governed by some common failure rate parameter.

This common failure rate assumption has been singled out for criticism by several authors *(GIA4, GIA5, GIA14, GIA15)* on the basis that it runs strongly counter to the common experience that there is wide variation among software faults in their likelihood of being encountered and/or of causing a failure. One intuitively believes that a gross error lying in the main line of code will have a considerably shorter mean time to failure than a fault which resides in a seldom invoked subroutine and which would require some particular juxtaposition of data to cause a failure. Furthermore, one can expect that any method which ignores this variation in the failure rates of faults will produce poor results in practice. As can be seen from (10) and (12), the two quantities n and Σt_i completely determine the J/M choice of estimators, m' and σ'. Thus, any combination of failure rates among the initially present faults that produces the same n and Σt_i for the observation period will result in identical parameter estimates using the J/M model. Other approaches have been proposed *(GIA15-GIA19)* which have eliminated the common failure rate assumption, but only at the expense of introducing alternative assumptions which in turn seem to lack convincing justification.

As we shall show in the next section, Lindsay's theorem allows us to replace the common failure rate assumption with one that permits each fault to have its own, potentially unique, failure rate chosen from some unknown but unrestricted general distribution. This improvement alone should remove a frequent and significant cause of the observed poor performance of the J/M class of models. It should be noted that several other extensions appear to be compatible with the use of Lindsay's theorem, for example modelling the imperfect removal of faults, the introduction of new faults, changes in the definition of the 'time' metric etc.

The Jelinski/Moranda model extended by use of Lindsay's theorem

Using Lindsay's theorem, we can reformulate the assumptions underlying the J/M model by eliminating the common failure rate assumption:

1 There are initially m faults present in the system.

2 The time to failure of the ith fault is an independent random variable chosen from a negative exponential distribution with a rate of failure parameter μ_i.

3 Each rate of failure parameter μ_i is an independent, identically distributed random variable with an unknown distribution (which can be any one of the many probability distributions on the non-negative reals).

4 When a failure occurs, the associated fault is immediately removed or, equivalently, the clock is stopped until the fault is removed.

We can think of assumptions 2 and 3 in the following terms. At the time that each fault is embedded in the software system, an associated failure rate parameter μ_i (which characterises the likelihood that the particular fault will cause a failure at any given instant) is randomly chosen from some unknown general probability distribution π. The failure rate associated with each of the faults is chosen using the identical (unknown) probability distribution. The time of failure of each fault t_i is then independently chosen using a negative exponential distribution with that fault's associated failure rate parameter μ_i.

At some time t, we wish to estimate both π, the underlying general distribution which gives rise to the failure rate parameters, and m, the number of faults initially in the system. In the spirit of J/M, we choose as estimators that probability distribution π' and that number m' which maximise the likelihood of the observed number of failures and the failure times.

It is easy to see that this is not a traditional parameter estimation problem. Note that we are not seeking individual estimators of the μ_i parameters associated with the faults which have been observed to fail, but rather the underlying distribution from which the failure rates of *all* of the faults have been chosen.

Also note carefully, however, that this is not a traditional 'non-parametric' formulation in which we attempt to find, free of parametric constraints, the best (in some sense) probability distribution to directly fit the observed system failure times. We have already indirectly built in (by assumptions 2 and 3) a traditional parametric structure for the system failure time distribution; that is, the system failure time distribution is derived from the negative exponential failure time distributions (with parameters μ_i) of the individual faults. We are attempting to find a best 'fit' for the common distribution underlying the failure rate parameters of *all* the individual faults (both those with observed failures and those yet to fail) and not just the *observed* failure rate of the *system*, which is characteristic of traditional non-parametric formulations. This 'fit' to the common distribution underlying all the failure rate parameters is the unifying bridge that allows us, based on the observed failure times, to make estimates of the failure times of the remaining faults in the system.

As in the J/M model, we can proceed in two phases:

i Find the estimator for π which maximises the likelihood of the observations, given a fixed value of m.

ii Treat this estimator as a function of m, then find the value of m which maximises the resulting likelihood.

To simplify the notation somewhat, we shall again adopt the convention that time is scaled such that the period of observation ends at $t = 1$. We therefore must also have $0 \leq t_i \leq 1$, for $i = 1,2,\ldots,n$.

Under the 'extended' assumptions, the individual likelihood factor y_i, associated with an observed failure at time t_i, is the weighted average of the likelihood over the range of admissible μ:

$$y_i = \int_0^{\infty} \mu \exp(-\mu t_i)\, d\pi(\mu) \tag{17a}$$

Similarly, the likelihood factor z, associated with a fault which did not fail during the observation period, is

$$z = \int_0^{\infty} \exp(-\mu)\, d\pi(\mu) \tag{17b}$$

Thus, given that the system has initially m faults, the likelihood of n failures occurring at $t_1,t_2,\ldots,t_n$ is

$$\psi = \frac{m!}{(m-n)!} \quad z^{(m-n)} \quad \prod_{i=1}^{n} \; y_i \tag{18}$$

Before applying Lindsay's theorem, it is again useful to substitute the change of variable for μ previously used in our J/M analysis:

$$x = \exp(-\mu) \tag{19}$$

In applying this change of variable to (17a) and (17b) we shall take the limits of integration to be the closed interval [0,1], thus in effect including the limit point, $\mu = \infty$, as a potential support point for π. We can now confirm that the premises of theorems 1, 2 and 3 are satisfied. (18) is of the form of (1), the z and y_i factors are each of the form (2) with the integration (after the change of variables) being over a compact subset of the reals and the integrands of (17a) and (17b) satisfy the conditions following (2).

Theorem 1 can now be applied to show that some probability distribution, $\pi' \in P\{[0,1]\}$, exists, which maximises the likelihood defined by (18). We can also apply theorem 3 immediately. We first note that the f_i values of theorem 3 correspond to the integrands of (17a) and (17b) (after the change of variables):

$$f_i = -\, x^{t_i} \ln(x) \qquad \text{for } i = 1, 2, \ldots, n \tag{20a}$$

$$f_i = x \qquad \text{for } i = n+1, n+2, \ldots, m \tag{20b}$$

It can be easily shown that no set of non-zero, positive, real coefficients can exist which satisfy (5) everywhere over any open set of $C = [0,1]$. Thus, π' must be a discrete distribution without a density component. Accordingly, we can write

$$\pi' = \{(w_r,x_r) | r = 1,2,\ldots, v\} \tag{21}$$

Using theorem 2 (Lindsay's theorem), we can then identify B, the support base of π', as a subset of β':

$$\beta' = \{x_r \,|\, \sum_{i=1}^{n} \frac{f_i(x_r)}{y_i(\pi')} + (m-n)\, \frac{x}{z(\pi')} = m\,\} \tag{22}$$

where:

$$y_i = \sum_{r=1}^{\nu} w_r f_i(x_r) \quad \text{for } i = 1, 2, \ldots, n \tag{23a}$$

$$z = \sum_{r=1}^{\nu} w_r x_r \tag{23b}$$

Since the defining equation of (22) can only have a finite number of roots in [0,1], the number of points, ν, in β' must also be finite.

Without loss of generality, we can treat β' itself as the support base of π' by adopting the convention that $w_r = 0$ for each $x_r \in \beta'$ where $x_r \notin B$. As an additional convenience, we shall further assume that the points of β' are ordered: $x_1 < x_2 < \ldots < x_\nu$.

In addition to (22) and (23) above, the necessary and sufficient condition for π' to be a maximising distribution is then, by Lindsay's theorem

$$\sum_{i=1}^{n} \frac{f_i(x)}{y_i(\pi')} + (m-n)\frac{x}{z(\pi')} < m \qquad \text{for } x \notin \beta' \tag{24}$$

If the number of initial faults m is fixed, (22), (23) and (24) can be used to find π', the failure rate probability distribution which maximises the likelihood as defined in (18). The estimated failure rate at the termination of observations, ζ, is then

$$\zeta = -(m-n)\frac{\Sigma\, w_r x_r \ln(x_r)}{\Sigma\, w_r x_r} \tag{25}$$

Before discussing the choice of a maximising value of m, let us first consider some of the characteristics of the solution. As shown by (21), each initial fault can fall into any one of ν categories, each category distinguished by a failure rate μ_r and a probability w_r that a specified fault belongs to that category. Note that the solution furnishes both the number of categories ν and the values of the x,w pairs. Also note that, since the failure rate associated with each fault is an independent random choice (by assumption 3), the solution offers no assurance that the actual proportion of faults in the rth category equals w_r.

It is also of some interest to note that $x_\nu = 1$ (that is $\mu_\nu = 0$) can be a member of the solution set for finite m. Since this implies that any associated fault has a zero failure rate, we have the somewhat paradoxical result that a fault of this category can neither have failed in the period of observation or fail at any time thereafter. In effect, the solution removes all such faults from the active set of initial faults.

Unfortunately, some of the coefficients of (22), (23) and (24) are ill-behaved as $m \to \infty$ and these equations are not suitable for investigating the behaviour of the solution for large values of m. Specifically one can show that, as $m \to \infty$

$$w_r(\pi') \to 0, \text{ for } r < \nu \tag{26a}$$

$$w_\nu(\pi') \to 1 \tag{26b}$$

$$x_\nu(\pi') \to 1 \tag{26c}$$

$$y_i(\pi') \to 0, \text{ for } i = 1,2,\ldots,n \tag{26d}$$

$$z(\pi') \to 1 \tag{26e}$$

However, one can also show that as $m \to \infty$ the products $mw_r(\pi')$, $m[1-x_r(\pi')]$ and $my_i(\pi')$ remain finite and can limit to some value greater than zero. We can take advantage of this fact to perform a change of variables that will result in well-behaved coefficients for all values of m:

$$\sigma = (m-n)/m \tag{27a}$$

$$\alpha_i = y_i/(1-\sigma), \text{ for } i = 1,2,\ldots,n \tag{27b}$$

$$\lambda_r = w_r(1-x_r)/(1-\sigma), \text{ for } r = 1,2,\ldots, \nu \tag{27c}$$

After some algebraic manipulation, we can replace (22), (23) and (24) with the following equivalent set of equations:

$$\beta' = \{x_r \,|\, \sum_{i=1}^{n} \frac{f_i(x_r)}{\alpha_i(\pi')} + \frac{n\sigma x_r}{\gamma(\pi')} = n\} \tag{28}$$

where

$$\alpha_i = \sum_{r=1}^{\nu} \frac{\lambda_r f_i(x_r)}{(1-x_r)} \qquad \text{for } i = 1, 2, \ldots, n \tag{29a}$$

$$\gamma = [1 - (1 - \sigma)(\sum_{r=1}^{\nu} \lambda_r)] \tag{29b}$$

and

$$\sum_{i=1}^{n} \frac{f_i(x_r)}{\alpha_i(\pi')} + \frac{n\sigma x_r}{\gamma(\pi')} < n \text{ for } x \notin \beta' \tag{30}$$

Using the same change of variables, we also obtain

$$\zeta = -\frac{n\sigma}{\gamma} \sum_{r=1}^{\nu} \frac{\lambda_r x_r \ln(x_r)}{(1-x_r)} \tag{31}$$

$$\ln(\psi) = \sum_{i=1}^{n} \ln(\alpha_i) + \frac{n\sigma}{(1-\sigma)} \ln(\gamma) + \sum_{k=1}^{n} \ln[k(1-\sigma) + n\sigma] \tag{32}$$

For $\sigma = 1$, (32) reduces to

$$\ln(\psi) = \sum_{i=1}^{n} \ln(\alpha_i) - n\,(\Sigma\lambda_r) + n \ln(n) \tag{33}$$

Taking the partial differential of (23) with respect to σ and evaluating at π' produces

$$\left.\frac{\delta \ln(\psi)}{\delta\sigma}\right|_{\pi'} = \frac{n}{(1-\sigma)^2}\left(\sum_{k=1}^{n} \frac{(n-k)}{k(1-\sigma)+n\sigma} + \ln[\gamma(\pi')]\right) \tag{34}$$

It is conjectured, by analogy with (12), that (34) has at most one zero in [0,1]. If so, we can show that the value of σ associated with the maximum likelihood is then given by

$$\sigma' = 1, \qquad \text{if } \sum_{i=1}^{n} \frac{(2t_i-1)}{\alpha_i(\pi')} \geqslant \frac{1}{\lambda_\nu^2(\pi')} \text{ at } \sigma = 1 \tag{35a}$$

$$\sigma' = 0, \qquad \text{if } \sum_{k=1}^{n} \frac{1}{k} \leqslant \ln[\gamma(\pi')] \qquad \text{at } \sigma = 0 \tag{35b}$$

If neither of these conditions is true, the maximising value of σ' is obtained by solving (34) for the value of σ that results in the equation being equal to zero. (If our conjecture regarding the single zero of (34) is not true, all the local maxima would have to be compared.)

A numerical algorithm

We can characterise the solution of (28) to (30) for some fixed value of σ as the following set:

$$G(\sigma) = \{(x_r, \lambda_r) | r = 1,2,\ldots,\nu\} \tag{36}$$

formed from the members of the support set of π' and transformations of the associated probability weights w in accordance with (27c). In representing the solution in this form, we should take note of the fact that (27c) implies that $x_\nu = 1 \Rightarrow \lambda_\nu = 0$, except in the case where $\sigma = 1$. If $\sigma = 1$, it can be shown that $x_\nu = 1$ is always a member of the support set of π' and that λ_ν can take on any value from zero to one.

Numerical solutions determining this set present some difficulties. In contrast to (10) of the J/M model, (28) to (30) have no closed-form solution. Furthermore, any straightforward approach to an iterative solution is hampered by the fact that we initially have no information regarding the number of points ν in the support of π'.

Following the general approach suggested by Lindsay *(GIA1)*, the solution algorithm we have used to find $G(\sigma)$ proceeds by the following steps:

1 Initialise G to a single pair (x_1, λ_1) where x_1 is the J/M solution given by (10) and λ_1 equals $(1-x_1)/(1-\sigma)$. (If $\sigma = 1$, then $\lambda_1 = 1$.)

2 Add a pair to G for each local maximum point of the defining equation of (28) where the equation equals or exceeds n. The x value of the pair should be set to the location of the maximum and the λ value of the pair should be set to zero.

3 Holding the x_r values of the pairs of G constant, find the λ_r values which maximise (32) (or (33), if appropriate), subject to either

$$\sum_{r=1}^{\nu} \lambda_r = 1 \qquad \text{if } 1 \in G \tag{37a}$$

or

$$\sum_{r=1}^{\nu} \frac{\lambda_r(1-\sigma)}{(1-x_r)} = 1 \quad \text{otherwise} \tag{37b}$$

and

$$0 \leq \lambda_r \tag{37c}$$

	Scenario	J/M	Giammo
Initial faults, m			
Mean	165.0	134.1	161.8
Std dev	0.0	5.0	36.0
Failure rate at term, ζ			
Mean	33.1	8.3	26.9
Std dev	4.4	3.5	9.5
ζ(estimate) – ζ(scenario)			
Mean	–	−24.8	−6.2
Std dev	–	6.9	11.9
Root n of likelihood, $\psi^{1/n}$			
Mean	–	78.4	83.1
Std dev	–	5.0	4.9

Figure 1: Summary results

		Number of initial faults			Failure rate at termination			Root *n* of likelihood	
Case	Fails	Scenario	J/M	Giammo	Scenario	J/M	Giammo	J/M	Giammo
1	130	165.0	131.2	181.1	35.0	5.0	27.2	81.21	85.44
2	127	165.0	128.6	133.9	38.0	6.6	17.3	76.24	78.66
3	134	165.0	138.2	199.7	31.0	14.1	42.5	71.26	76.26
4	133	165.0	134.4	144.8	32.0	6.1	21.1	82.68	85.80
5	134	165.0	136.3	157.2	31.0	8.8	29.2	78.13	83.20
6	135	165.0	139.2	152.7	30.0	14.3	30.0	72.13	74.63
7	131	165.0	132.2	148.1	34.0	5.1	24.7	82.75	88.39
8	131	165.0	132.4	136.3	34.0	6.1	14.7	81.60	85.45
9	136	165.0	138.3	157.0	29.0	8.9	29.2	78.97	82.78
10	132	165.0	134.2	164.2	33.0	8.6	28.9	76.53	79.11
11	133	165.0	137.2	328.2	32.0	14.0	50.7	70.93	82.77
12	133	165.0	134.4	143.8	32.0	6.2	20.8	83.09	87.48
13	142	165.0	144.7	153.0	23.0	10.3	24.9	81.84	88.22
14	127	165.0	131.2	191.5	38.0	14.1	42.7	67.07	72.60
15	132	165.0	133.2	142.4	33.0	5.1	18.9	82.83	85.86
16	123	165.0	124.8	141.7	42.0	7.3	26.5	72.88	77.92
17	122	165.0	122.1	123.6	43.0	0.6	5.9	87.20	90.95
18	131	165.0	131.4	136.5	34.0	1.9	13.3	90.91	93.87
19	122	165.0	124.1	141.9	43.0	7.9	26.7	70.47	73.83
20	133	165.0	137.4	182.3	32.0	14.8	41.7	70.55	75.61
21	126	165.0	127.4	141.2	39.0	5.8	22.2	78.56	82.01
22	133	165.0	135.2	168.4	32.0	8.8	36.0	77.83	87.29
23	133	165.0	135.2	191.5	32.0	8.7	40.0	76.98	87.46
24	137	165.0	140.1	175.1	28.0	11.5	37.7	76.95	82.81
25	140	165.0	143.5	271.0	25.0	12.5	44.3	76.86	83.62
26	123	165.0	123.4	127.2	42.0	1.8	11.2	85.97	88.80
27	135	165.0	136.2	139.5	30.0	5.2	13.1	85.10	88.4[illegible]
28	140	165.0	142.6	153.5	25.0	10.1	26.5	80.57	84.62
29	133	165.0	136.1	159.7	32.0	11.1	32.9	74.22	78.24
30	130	165.0	131.4	165.8	35.0	5.9	28.2	79.56	83.71
31	129	165.0	131.2	147.2	36.0	8.5	29.1	75.24	84.02
32	134	165.0	136.3	152.0	31.0	8.8	27.6	78.13	81.89
33	138	165.0	140.1	158.5	27.0	8.3	28.5	82.73	87.75
34	128	165.0	129.9	243.2	37.0	7.6	32.9	75.82	81.23
35	133	165.0	135.5	153.9	32.0	9.6	32.0	76.73	88.74
36	127	165.0	128.6	132.0	38.0	6.7	14.0	77.13	78.26
37	133	165.0	134.4	141.7	32.0	6.1	19.3	82.12	86.07
38	131	165.0	133.5	151.0	34.0	9.5	29.4	75.70	80.74
39	138	165.0	142.0	150.2	27.0	13.9	26.9	74.58	79.00
40	130	165.0	130.6	134.5	35.0	3.1	12.1	88.85	91.23
41	130	165.0	131.2	181.1	35.0	5.0	27.2	81.21	85.44
42	127	165.0	128.6	133.9	38.0	6.6	17.3	76.24	78.66
43	134	165.0	138.2	199.7	31.0	14.1	42.5	71.26	76.26
44	133	165.0	134.4	144.8	32.0	6.1	21.1	82.68	85.80
45	134	165.0	136.3	157.2	31.0	8.8	29.2	78.13	83.20
46	135	165.0	139.2	152.7	30.0	14.3	30.0	72.13	74.63
47	131	165.0	132.2	148.1	34.0	5.1	24.7	82.75	88.39
48	131	165.0	132.4	136.3	34.0	6.1	14.7	81.60	85.45
49	136	165.0	138.3	157.0	29.0	8.9	29.2	78.97	82.78
50	132	165.0	134.2	164.2	33.0	8.6	28.9	76.53	79.11

Figure 2: Individual case results

(37b) follows from the requirement that $\Sigma w_r = 1$, which can be evaluated directly using the λ_rs if $1 \notin G$. (37a) can be derived by evaluating the defining equation in (28) at $x = 1$ with $\sigma = 1$.

4 Eliminate pairs from G if maximisation of the likelihood would tend to cause λ_r to violate (37c).

5 Repeat steps 2, 3 and 4 until the maximum value of the defining equation in (28) is within some ε of n.

Numerical results

A test scenario was constructed to explore the degree of improvement in various estimators afforded by the relaxation of the common failure rate assumption of the basic J/M formulation; 50 test cases were independently generated. Estimators for m, the total number of faults, ζ, the failure rate at the termination of observations and ψ, the likelihood of the observations, were developed using the J/M equations and the corresponding equations of the J/M model extended by Lindsay's theorem. Various comparisons were then made between these results and the 'true' results from the generating scenario.

The scenario used to generate the test cases assumed the following:

1 The software system initially contained 165 faults.

2 The time to failure of each fault was a random variable with a negative exponential distribution.

3 The failure rate parameter associated with 75 of the faults was 10 failures per unit time; the failure rate parameter associated with 90 of the faults was one failure per unit time.

4 Failures were observed for one unit of time. The faults associated with each failure occurring in this period were immediately removed.

The scenario was not constructed to be a 'typical' case. The parameters were chosen to represent a situation where it was felt that the common failure rate assumption of the J/M formulation would result in a significant misestimate of both the initial number of faults in the system and the failure rate at the termination of observations. It was hoped that an improvement in these estimators due to the relaxation of the assumption would be clearly shown by a comparison of the results.

On the other hand, the scenario should not be considered as necessarily atypical. The number of faults should not be considered as excessive for a large complex system. The pattern of failure times could well correspond to a situation in which 75 of the 165 faults are in commonly encountered routines and are removed in the early phases of system testing, while the remaining 90 are more rarely encountered. The period of observation should correspond to a point in time at which the rate of failures has fallen to a level (less than five per cent of the initial failure rate) at which a decision involving the current failure rate and/or the number of remaining faults might be called for.

Fifty cases were independently generated and evaluated. Figure 1 presents a statistical summary of the results for the 50 cases as a whole. Figure 2 gives individual results for each of the 50 cases.

In general, the results confirmed our expectations. The relaxation of the common failure rate assumption brought about a significant improvement in the estimates, both of the total number of initial faults and of the failure rate at the termination of observations. As shown in Figure 1, the mean difference between the estimated number of initial faults and the actual number improved from −30.9 to −3.2, indicating a significant reduction in the bias of the estimator. The variance of the estimate remained high (standard deviation equal to 36.0), however, the distribution of the difference appearing highly skewed with a large tail on the positive side. The mean difference between the estimated and actual failure rates at the termination of observations improved from −24.8 to −6.2, again showing a significant reduction in the bias of the estimator. The variance of this estimate was significantly lower (standard deviation equal to 11.9) than in the case of the estimate of the number of initial faults, with the distribution of the difference between estimated and actual values being much more symmetric.

The improvement was also reflected in the increase in the likelihood of the observations. The statistic shown in Figures 1 and 2 is the nth root (where n is the number of observed failures) of the likelihood ψ, as defined by means of (11) and (32). Since the J/M solution is also a candidate solution in the revised formulation (and not vice versa), the likelihood of the J/M solution must always be less than or equal to the likelihood of the solution under the revised assumptions. This is clearly shown in Figure 2.

In summary, the relaxation of the common failure rate assumption has shown the expected improvements, as compared to the standard J/M formulation, in removing the bias associated with this assumption in estimating two useful statistics, the number of initial faults and the failure rate at the termination of the observations. However, the resultant estimators still exhibit the characteristically high variance of estimators of this type.