Reliability Engineering and System Safety 128 (2014) 66–83

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Bounds on survival probability given mean probability of failure per demand; and the paradoxical advantages of uncertainty

Lorenzo Strigini *, David Wright

Centre for Software Reliability, City University London, Northampton Square, London EC1V 0HB, United Kingdom

ARTICLE INFO

Article history:
Received 29 August 2012
Received in revised form
19 February 2014
Accepted 22 February 2014
Available online 1 March 2014

Keywords:
Safety critical systems
Software reliability
Parameter uncertainty
Epistemic uncertainty
System acceptance
Regulatory decision making

ABSTRACT

When deciding whether to accept into service a new safety-critical system, or choosing between alternative systems, uncertainty about the parameters that affect future failure probability may be a major problem. This uncertainty can be extreme if there is the possibility of unknown design errors (e.g. in software), or wide variation between nominally equivalent components.

We study the effect of parameter uncertainty on future reliability (survival probability), for systems required to have low risk of even only one failure or accident over the long term (e.g. their whole operational lifetime) and characterised by a single reliability parameter (e.g. probability of failure per demand – *pfd*). A complete mathematical treatment requires stating a probability distribution for any parameter with uncertain value. This is hard, so calculations are often performed using point estimates, like the expected value.

We investigate conditions under which such simplified descriptions yield reliability values that are sure to be pessimistic (or optimistic) bounds for a prediction based on the true distribution. Two important observations are (i) using the expected value of the reliability parameter as its true value guarantees a *pessimistic* estimate of reliability, a useful property in most safety-related decisions; (ii) with a given *expected pfd*, broader distributions (in a formally defined meaning of "broader"), that is, systems that are *a priori* "less predictable", lower the risk of failures or accidents.

Result (i) justifies the simplification of using a mean in reliability modelling; we discuss within which scope this justification applies, and explore related scenarios, e.g. how things improve if we can test the system before operation. Result (ii) not only offers more flexible ways of bounding reliability predictions, but also has important, often counter-intuitive implications for decision making in various areas, like selection of components, project management, and product acceptance or licensing. For instance, in regulatory decision making dilemmas may arise in which the goal of minimising risk runs counter to other commonly held priorities, like predictability of risk; in safety assessment using expert opinion, the commonly recognised risk of experts being "overconfident" may be less dangerous than their being *under*confident.

1. Introduction

Predictions of reliability and safety through probabilistic modelling depend on the values of model parameters, e.g. component failure rates, which are often uncertain.

The main application scenario that motivates our research involves decisions on accepting a software product for use in a safety critical application requiring low accident probability over the operational life of the system in which it is embedded. For instance, an explicit requirement in civil aviation is that "catastrophic failure conditions" be "so unlikely that they are not anticipated to occur during the entire operational life of all airplanes of one type" [16]. Nuclear power protection systems may have required or claimed *pfd* bounds like 10^{-7} or 10^{-9} [23,22], to assure low probability of even one failure during operational life. Formally, the system's predicted reliability function, concerning those failures that may cause accidents, must be close to 1 at the end of the intended operational life.

With software, decisions are made especially difficult by uncertainties about whether design faults are present, and about their effect on probability of failure. The same difficulty arises with respect to the probability of any system failures due to design faults. Similar decision problems may also arise regarding physical failures of components, if there is a concern about broad variation in reliability parameters, as for instance with the current alarm about electronic component supplies "contaminated" with unreliable "counterfeit" components.

Our reference scenario is a system S with high required confidence of operating until the end of its service life without

* Corresponding author.
E-mail address: strigini@soi.city.ac.uk (L. Strigini).

http://dx.doi.org/10.1016/j.ress.2014.02.004
0951-8320/

Nomenclature and abbreviations

A, B	labels denoting the two components of a simple series or parallel composite system
cdf	cumulative distribution function of a random variable
dem.	demands
E	expectation operator as applied to random variable
f_Q	probability density function of *pfd*
i,I	lower case i is the density function of leftward-moved mass in the broadening operation of Section 5.1; upper case I is an interval containing all of this mass before it is moved left, under-barred to denote its infimum (left-hand endpoint), and over-barred to denote its supremum (right hand endpoint)
j,J	lower case j is the density function of leftward-moved mass in the broadening operation of Section 5.1; upper case J is an interval containing all of this mass before it is moved right, under-barred to denote its infimum (left-hand endpoint), and over-barred to denote its supremum (right hand endpoint)
k	ratio of the sizes of the two masses moved apart and also of their respective distances moved, so that mean is preserved in broadening operation of Section 5.1
$\Lambda, \lambda, \lambda^*$	Continuous-time failure rate parameter: as random variable, instantiated value, or mean – analogously to Q below
MTTF	mean time to failure
pfd	probability of failure per demand of component or system
Pr	probability
PRA	probabilistic risk assessment
Q, q, q^*	*pfd*, upper case Q when regarded as an uninstantiated random variable; or lower case q to denote a particular realised value; or starred q^* to denote its mean value $E(Q)$
$R(t)$	reliability function R evaluated at usually discrete time t demands. $R(t) = \Pr(\text{no failure occurs over first } t \text{ demands})$; sometimes with further semicolon-separated parameter arguments. Also used for analogous continuous time reliability
surv.	survives
$\mathcal{S}$	a system subject to discrete demands on each of which it may succeed or failure
t	system operating time, usually discrete (number of demands $t = 0, 1, 2, \ldots$) unless otherwise stated
U	$1 - Q$, where Q is the *pfd* as a random variable
w.r.t.	with respect to

failure causing accident. Subject to discrete demands, $\mathcal{S}$'s failure process is completely characterised by a constant *probability of failure per demand* (*pfd*). Examples are failure of software due to design faults (the original motivation of our work), or hardware without aging or maintenance.

Mathematically similar scenarios exist regarding reliability even without safety implications, e.g. when a component should last for the lifetime of the system of which it is part, because it cannot be replaced or repaired, by either design (as in many consumer products) or necessity (e.g. in spacecraft).

We must predict $\mathcal{S}$'s probability of surviving t future discrete, independent demands – its reliability $R(t)$ in discrete time – with t an upper bound on the lifetime number of demands, if accident-free.[1] This would be straightforward except for uncertainty about the *pfd* value [3], arising e.g. because *pfd* is

- inferred from reliability databases on components that are similar, but not identical, to the one for which a prediction is sought, and/or that operate in potentially different conditions, affecting their reliability differently. If the details of which systems failed and when are missing or not released;
- guessed using indirect evidence, as e.g. often done for *pfd*s due to software design faults.

This uncertainty can in theory be rigorously described by a subjective probability distribution for the value of each parameter. However, an assessor has seldom a clear idea of this distribution, and many calculations are *de facto* performed by treating their available $E(pfd)$ estimate as though it were the true *pfd*. Sensitivity analysis may be used to check that small variations in the estimate only cause acceptable prediction error, but in practice much reasoning among practitioners only deals with a point estimate $E(pfd)$, without acknowledging that, in fact, the shape of the probability distribution of the *pfd* may also have a substantial effect on the predicted value sought (e.g. reliability over a given period of operation), and this effect may be non-obvious.

Thus, using a point *pfd* estimate to calculate a system's lifetime survival probability may lead to errors of various kinds [1,9].

Uncertainty about parameter values is a typical case of *epistemic* uncertainty in predictions (i.e., uncertainty arising from lack of knowledge rather than from an "inherent randomness" of the process studied). Epistemic uncertainty is widely studied [19,20] and many formal mathematical methods have been proposed for dealing with it, but much normal practice does not use them. The practical approaches, e.g. in the nuclear industry [13,15], are essentially of two kinds: qualitative criteria for accepting evidence (e.g. requiring that parameter value be derived from evidence that is more clearly pertinent to the specific plant, the more critical the parameters in question are) and numerical methods for performing either sensitivity analysis or calculations taking into account the complete probability distributions that describe uncertainty on the parameters. To cite the NUREG guidance document [13].

> Because the impact of parameter uncertainty can be addressed in terms of a probability distribution on the numerical results of the PRA, it is straightforward to compare a point value, be it the mean, the 95th percentile, or some other representative value with an acceptance guideline or criterion … For most regulatory applications, that value is specified to be the mean […] The mean values referred to are the arithmetic means of the probability distributions that result from the propagation of the uncertainties on the input parameters.

Uncertainty propagation methods will in theory produce accurate results for any given distribution; but their application is hard: apart from computational complexity, their fundamental

[1] This number of demands is usually a random variable, but we will treat the problem with reference to a fixed t. Conclusions for a random number of demands can be derived if required.