Engineering

Report

All Stations

A Procedure for Probabilistic Creep-Fatigue Crack Initiation Assessment Consistent With R5 Volume 2/3

By: R.A.W. Bradford Date: 13/1/14

EDF Energy Nuclear Generation Limited

EDF ENERGY GENERATION LIMITED
CENTRAL ENGINEERING SUPPORT

SUMMARY SHEET

Power Station:	All Stations
Document Title:	A Procedure for Probabilistic Creep-Fatigue Crack Initiation Assessment Consistent With R5 Volume 2/3
Document Reference:	E/REP/BBAB/0028/GEN/13
Document Author:	R.A.W. Bradford
Verified by:	P.J.Holt / D.Tipping
Approved by:	D.W.Dean
Date:	13/1/14
Task File Number:	E/TSK/HPH/11239

Summary

A procedure is presented for the application of Monte Carlo probabilistic simulation to creep-fatigue crack initiation assessments to R5 Volume 2/3. This report does not address the 'precursor' assessments required by R5V2/3 but only the crack initiation assessment. The report addresses only probabilistic aspects, not the deterministic methodology to be used as the core of the probabilistic model, which is assumed to be that of R5 Volume 2/3. One exception relates to the adoption of primary creep reset versus continuous hardening in estimating creep relaxation (specific for 316H parent stainless steel).

Conclusion

Probabilistic creep-fatigue crack initiation assessments may be carried out to the procedure defined herein and with the limitations specified above.

Recommendation

To: Group Head, Assessment Technology

The probabilistic procedure defined herein should be considered for inclusion within R5 Volume 2/3.

Verification Certificate

Document Title:	A Procedure for Probabilistic Creep-Fatigue Crack Initiation Assessment Consistent With R5 Volume 2/3
Station:	All Stations
Document Reference:	E/REP/BBAB/0028/GEN/13
Revision:	000
Task File Number:	E/TSK/HPH/11239
QA Grade of document:	2
Author:	R.A.W. Bradford
Section/Group:	Structural Integrity Branch
Verifier(s):	P.J.Holt, Veritan Ltd
Principal Verifier:	D.J.Tipping
Section/Group:	Assessment Technology Group
Verification Statement:	This report has been verified in conformance with procedures BEG/SPEC/DAO/010 and E/PROC/ENG/BI/048 at QA Grade 2.
Technical Reviewer(s):	Not required
Section/Group:	-
Technical Review Details:	-
Approver:	D.W.Dean
Section/Group:	Assessment Technology Group

DOCUMENT ISSUE/AMENDMENT CERTIFICATE

Date	Author	Revision	Amendment
13/1/2014	R.A.W. Bradford	Rev.000	First issue

GLOSSARY (NOMENCLATURE)

Word, phrase or acronym	Definition
Aleatory	Uncertain quantities denoted “aleatory” are distributed variables whose randomness is deemed to be characterised by chance (see §9)
Correlation coefficient	A numerical measure of the degree of linear relationship between two distributed variables, see §4.9
CoV	coefficient of variation, the standard deviation normalised by the mean, see §4.8
cpf	Cumulative probability function, see §4.2
FSRF	Fatigue strength reduction factor
eol	End of life (end of the operating life of the plant in question)
Epistemic	Epistemic variables are distributed variables whose uncertainty arises simply from lack of knowledge (see §9)
Latin hypercube	A particular random sampling algorithm used in some Monte-Carlo simulations, described in §8
LB/UB	Lower Bound / Upper Bound (often identified with the 95%CL or the 98%CL)
Lognormal distribution	Another common form of pdf, appropriate for variables which cannot be negative, see §5.4
mean	See §4.3
MECT	Mean effective creep temperature
median	See §4.4
mode	See §4.5
Monte Carlo	A generic name given to probabilistic techniques based on random sampling
Non-σ distribution	A “non-σ distribution” is a pdf for which the standard deviation is not defined, i.e., the integral defining the standard deviation, equ.(4.5), is divergent. See also §6.
Normal distribution	A common form of pdf, appropriate for variables which can be negative, see §5.3
pdf	Probability density function, see §4.1
RCC-MR	French nuclear code: “Règles de Conception et de Construction des Matériels Mécaniques des Îlots Nucléaires RNR” (Design and Construction Rules for Mechanical Components of Nuclear Islands).
rms	Root-mean-square: the square root of the average of the squares of the quantity in question
R5V2/3	Volume 2/3 of the R5 procedure
sol	Start of life
standard deviation	See §4.6
trial	A deterministic calculation for a single random choice of the distributed variables (addressing the whole of life for all assessment locations of just a single specified component or a single randomly chosen component)
variance	See §4.7
With-σ distribution	A “with-σ distribution” is a pdf for which the standard deviation is well defined, i.e., the integral defining the standard deviation, equ.(4.5), is convergent. See also §6.
WSEF	Weld strain enhancement factor

95%CL	95% confidence level: a bound which there is a 95% probability that the variable in question will exceed (or be less than, according to context)

CONTENTS

1 INTRODUCTION

The traditional approach to structural assessments is deterministic. Adequate conservatism can be ensured in deterministic assessments by using suitable input assumptions in which at least some variables take "bounding" values (often key materials data). Other variables may be implicitly assumed to be known accurately, e.g., loads, stresses and geometry. However, if there is known to be significant uncertainty in these other variables, then some "bounding" values may be assumed for these also. This sort of approach is enshrined in custom and practice and forms the basis of most structural design codes as well as deterministic fitness-for-purpose assessments. This methodology has served the engineering community well over the years and will continue to be the basis of most structural integrity assessments.

However there are some circumstances in which an alternative approach based on probabilistic assessment may be desirable. One such circumstance occurs when failures (or leaks) have already occurred within a given population of components. In this case the apparent absolute guarantee provided by deterministic assessment is clearly undermined and attention will focus on the anticipated frequency of future failures. This is not addressed by deterministic assessments. Even if no failures (or leaks) have occurred, it may nevertheless be required to quantify the probability of failures (or leaks) in future. The existence of a deterministic margin does not mean that the failure probability is zero. Variables implicitly assumed to take their best estimate values may actually take more onerous values. Moreover, even variables assumed to take "bounding" values may take more onerous values since the usual bounds employed in deterministic assessments are typically 95% or 98% confidence levels. Consequently, there is a fairly high probability of more onerous data actually arising. Another motivation for probabilistic assessment may be that deterministic margins have not been demonstrated. This can happen if the deterministic assessment is obliged to assume rather conservative values for a number of parameters, the combined effect of which is unacceptable. The question then naturally arises as to how improbable such a combination of assumptions might be.

One of the particular benefits of probabilistic assessment is that it provides an objective view of the relative importance of uncertain input quantities in the form of the degree of correlation between these inputs and predicted failure events. For example, if a small number of boiler tubes have partially restricted flows, and hence elevated temperatures, this is likely to increase their individual failure probability. However, these partially restricted tubes may still make only a small contribution to the total probability of tube failure across the whole reactor, simply due to the far larger number of unrestricted tubes. Such a probabilistic perspective may influence the perceived requirement for plant modifications (e.g., tube plugging).

As power plant is operated significantly beyond its original design life, but is subject to ongoing degradation by creep and/or fatigue, the desirability of probabilistic assessments to quantify the remaining plant lifetime is likely to increase.

2 SCOPE AND PURPOSE

This report addresses the application of probabilistic techniques to creep-fatigue crack initiation assessments to R5 Volume 2/3 (Ref.[23], referred to hereafter as R5V2/3). The scope of application is the same as that of R5V2/3. This report addresses only probabilistic aspects, not the deterministic methodology to be used as the core of the probabilistic model. Figure 1 provides an example of a possible flow diagram of the overall probabilistic approach.

Throughout, the term "initiation" is used to mean crack initiation by creep-fatigue. In common with a deterministic R5V2/3 assessment, this implies the initiation of a crack of some depth a_0. This depth is specified by the assessor but will always be a small fraction of the section thickness.

The objective of a probabilistic initiation assessment is to calculate, (i) the probability that initiation occurs by a given time in an individual component, or, (ii) the probability that initiation occurs in an

individual component in each year of some span of years, or, (iii) the number of cracks initiating, either annually or by a given time, across a population of components.

2.1 Limitations on Scope

- This report does not address the 'precursor' assessments required by R5V2/3 (i.e., the primary stress limits, the stress range limits, the shakedown assessment and the cyclically enhanced creep assessment).
- This report does not address the deterministic methodology for creep-fatigue crack initiation, which is assumed to be that of R5V2/3. One exception relates to the adoption of primary creep reset versus continuous hardening in estimating creep relaxation for 316H parent stainless steel (§13.1).
- This report does not recommend specific distributions for plant data or stress/loading data, which will depend upon the application and must be provided by the User on a case-by-case basis.
- This report does not recommend specific distributions of material properties, although some general remarks regarding how such distributions may be derived are included. One exception is for 316H parent stainless steel below ~580^{o}C for which provisional advice is indicated (§13).
- Some numerical data are provided within R5V2/3 itself, i.e., K_S and FSRF or WSEF. Possible approaches to the distribution of these quantities are discussed.
- No general purpose computer code for carrying out probabilistic R5V2/3 creep-fatigue crack initiation assessments exists and hence none can be recommended. However, some remarks pertinent to coding are included.

3 OVERALL APPROACH

The methodology advised in this report is Monte Carlo simulation. This is an umbrella term used to denote probabilistic techniques based on random sampling, Ref.[1]. The great strength of Monte Carlo simulation is that it can address large numbers of distributed variables. There is little detriment to the User, and little impact on computer run times, if additional distributed variables are included. A further advantage of the Monte Carlo method is that it can address correlations between variables.

The basic idea of Monte Carlo simulation is very simple. It consists of running a deterministic assessment a very large number of times with different combinations of randomly sampled input data. The key to obtaining a truly representative simulation is that the sampling process be representative. The best means of ensuring that the sampling process is representative is to use random trials which are equally likely. The way in which this is achieved is defined in §8.

Monte Carlo simulation can be used with many different sampling algorithms, Ref.[1]. In this report it is recommended that Latin hypercube sampling is employed. Recent applications to plant problems have proved this algorithm to be robust and simple to code, Refs.[2-10].

A probabilistic initiation assessment must be implemented via a computer code since the very large number of deterministic R5V2/3 assessments required make this a practical impossibility to do 'by hand'. This has implications for Verification (see §15). Probabilistic code will usually consist of three parts,

(i) The "deterministic core", which carries out the R5V2/3 assessment (for all required assessment locations on the chosen component) for a given set of input parameters and for one loading cycle (and dwell);

(ii) A loop over all load cycles defining the history and the future of the plant, applying (i) to each cycle in turn for the same (time independent) set of sampled parameters;

(iii) A loop over many sets of randomly sampled sets of parameters (trials), repeating (ii) for each trial, and implementing the Latin hypercube sampling algorithm.

A slightly more detailed break-down of a typical code is shown as the flow diagram of Figure 1.

A general purpose code for all applications may be used for (i), but see §15 for Verification requirements. In contrast, (ii) is application-specific. All applications will share a similar approach to (iii), the random sampling algorithm, but will differ as regards the number of distributed variables and the distributions chosen to represent them.

4 DEFINITIONS OF BASIC TERMS

4.1 Probability Density Function

The probability that a continuous variable, x, lie within the small range x to $x+dx$ is $P(x)dx$. This defines the probability density function (pdf), $P(x)$, for the variable x. A pdf necessarily integrates to unity, $\int_{-\infty}^{+\infty} P(x)dx = 1$, i.e., it is certain that the variable has some value. (The lower limit of $-\infty$ is replaced by 0 for variables which cannot be negative).

4.2 Cumulative Probability

The probability that a continuous variable, x, takes a value less than or equal to X defines the cumulative probability function (cpf), $P_{cum}(X)$. The cpf is found in terms of the pdf from,

$$P_{cum}(X) = \int_{-\infty}^{X} P(x)dx \tag{4.1}$$

(The lower limit of $-\infty$ is replaced by 0 for variables which cannot be negative). The reverse-cumulative distribution, $\widetilde{P}_{cum}(X)$, is the probability that x takes a value greater than or equal to X,

$$\widetilde{P}_{cum}(X) = \int_{X}^{+\infty} P(x)dx \tag{4.2}$$

so that $P_{cum}(X) + \widetilde{P}_{cum}(X) = 1$.

4.3 Mean

In this report the word “mean” is shorthand for the arithmetic mean, and is synonymous with “average” in common parlance. It is sometimes called the “expected value” or the “expectation value”. The algebraic definition of the mean of the variable x, written $\bar{x}$ or $\langle x \rangle$, whose pdf is $P(x)$, is,

$$\langle x \rangle = \int_{-\infty}^{+\infty} xP(x) \cdot dx \tag{4.3}$$

From here on it is to be understood that the lower integration limit of $-\infty$ is replaced by 0 for variables which cannot be negative.

4.4 Median

The median of a variable x with pdf $P(x)$ is the value $\hat{x}$ which splits the pdf into two regions whose probability is 0.5, i.e., is such that,

$$P_{cum}(\hat{x}) = \widetilde{P}_{cum}(\hat{x}) = 0.5 \tag{4.4}$$

4.5 Mode

The mode of a pdf $P(x)$ is the value of x at the maximum of $P(x)$. Hence the mode is the most probable value of x.

4.6 Standard Deviation

The standard deviation, σ_x, of a continuous random variable, x, is a measure of the spread of its pdf about its mean. It is defined as the root-mean-square (rms) of the deviation from the mean, thus,

$$\sigma_x = \sqrt{\int_{-\infty}^{+\infty} \left(x - \langle x \rangle\right)^2 P(x) \cdot dx} \tag{4.5}$$

The standard deviation may be divergent for some pdfs. For these pdfs, referred to here as "non-σ" distributions (see §6), the standard deviation is not a suitable measure of their spread.

4.7 Variance

The variance is the square of the standard deviation, σ_x^2.

4.8 Coefficient of Variation

The coefficient of variation (CoV) is the standard deviation normalised by the mean,

$$\text{CoV} = \frac{\sigma_x}{\langle x \rangle} \tag{4.6}$$

4.9 Correlation Coefficient

The correlation coefficient between two random variables x and y, denoted C_{xy}, expresses the extent to which they are linearly related. Variables with non-zero correlation cannot be described by independent pdfs like $P_x(x)$ and $P_y(y)$. Instead they have a joint probability density function, $P(x, y)$, in which the probability of the variables being within the ranges x to $x + dx$ and y to $y + dy$ is $P(x, y)dxdy$. The correlation coefficient is defined by,

$$C_{xy} = \frac{1}{\sigma_x \sigma_y} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left(x - \langle x \rangle\right)\left(y - \langle y \rangle\right) P(x, y) \cdot dxdy \tag{4.7}$$

In practice the correlation coefficient might be estimated from a random set of pairs of data (x_i, y_i) in which case (4.7) becomes,

$$C_{xy} \approx \frac{\sum_i \left(x_i - \langle x \rangle\right)\left(y_i - \langle y \rangle\right)}{\sqrt{\sum_i \left(x_i - \langle x \rangle\right)^2 \sum_i \left(y_i - \langle y \rangle\right)^2}} \tag{4.8}$$

In applying (4.8) the true means would in practice be replaced by the means over the available dataset. (4.8) is then an estimate of (4.7) which will become accurate in the limit of a very large randomly sampled dataset.

If there is a strict deterministic linear relationship between x and y then C_{xy} will be ± 1 (the negative sign applying if the x, y graph has negative slope). If the variables have no underlying relationship then $C_{xy} = 0$. Intermediate values, $0 < \left|C_{xy}\right| < 1$, indicate an imperfect linear relationship. It is important to recognise that the correlation coefficient will only identify ***linear*** relationships. For example, if the variables are deterministically related by $x - \langle x \rangle \propto \left(y - \langle y \rangle\right)^2$ they will nevertheless have $C_{xy} = 0$ (assuming a y distribution symmetrical about its mean).

4.10 Trial

A "trial" is a single assessment for a given set of randomly sampled inputs. This will include assessing all cycles over plant life to obtain the final damage (or the time of initiation). A single trial covers the assessment of all the required assessment locations of just one component (which may either be specified by the User or sampled from the population of similar components, according to whether the failure statistics of a single particular component or of the whole population is being addressed (see §10.4)).

5 INPUT DISTRIBUTIONS

5.1 Conservatism and the Probabilistic Philosophy

It is essential that a structural integrity assessment be suitably conservative, containing implicit or explicit reserve margins commensurate with the application. In traditional deterministic assessments this conservatism is usually inherent in the use of "bounding" data in some parts of the assessment, usually the material "strength" type properties, such as the proof stress, the creep rupture stress, the creep ductility, the fatigue endurance, etc. This approach may be characterised as "input conservatism" in the sense that the conservatism lies in (some of) the assumed input data to the assessment. Provided that these input assumptions are regarded as adequately bounding, an end-of-life damage less than unity might be considered to provide sufficient assurance that cracking will not initiate. In particular this is the case for crack initiation assessments carried out using the R5V2/3 procedure for the purpose of component life assessment to the procedure of Ref.[22] which explicitly defines a damage of unity as being the acceptable limit. In the case of crack initiation assessments this is partly in recognition of the implicit margin which will exist between crack initiation and structural failure. No additional R5V2/3 reserve margin is expected in the assessed damage, i.e., no "output conservatism" is required.

The situation for probabilistic assessment is the reverse. The ideal objective in probabilistic assessment is for all inputs to be as accurate as possible – including, of course, an accurate description of their distributions. If this ideal were achieved, and assuming that the assessment methodology is sound, then the calculated initiation probability would also be accurate. So the ideal is that there be no input conservatism. Instead the required margin resides in the probability which is regarded as acceptable. An "output conservatism" can be imposed via the acceptable probability of crack initiation adopted. What initiation probability is regarded as acceptable will depend upon many issues, such as the nuclear safety duty of the component and upon whether the safety case obtains support

from other factors. The functional difference between the "input conservatism" and "output conservatism" approaches is that in the latter case the analyst need not be concerned with the required degree of margin. Instead the analyst can concentrate upon supplying as accurate an estimate of the initiation probability as practicable. The required conservatism is applied at the point of application, e.g., in the safety case, in the form of some maximum acceptable initiation probability or frequency.

The ideal that all inputs and assumptions to a probabilistic assessment be "best estimate" (including the distributions) will rarely be achievable. Some uncertain factors are likely to defy quantification. In such cases it will continue to be necessary to make conservative assumptions. In practice, therefore, the calculated initiation probability is likely to be conservative. Nevertheless, the philosophical ideal described above should be the guiding principle.

5.2 Where Are The PDFs?

At the present time, materials data, plant operating data, etc., are not generally supplied as pdfs. Instead what will readily be available will be only the best estimate (mean) and some "bounding" value or standard deviation. The analyst will have two options, as follows,

- The first option is to go back to the dataset from which the mean and standard deviation were derived. The analyst can then fit his own pdf to this dataset. Clearly this option is only possible if the dataset itself is available. Even then the analyst may be reluctant to take this route, either because of the labour involved, or, especially in the case of materials data, due to a valid concern regarding his competence to understand correctly the raw data.
- The second option is to assume a distribution with the given mean and standard deviation. If this option is chosen, there may be little information on which to base the choice between different pdfs with the same mean and variance (other than as noted in §5.5). This need not be unduly concerning. It should be appreciated that any reasonable pdf with the correct mean and variance will be a far closer representation of reality than a single deterministic value. Guidance on whether the assessor should be concerned that the assumed form of pdf might have introduced errors is given in §12.1.

5.3 The Normal Distribution

A normal distribution of x is specified by the pdf,

$$P(x) = \frac{1}{\sigma_x \sqrt{2\pi}} \exp\left\{ -\frac{\left(x - \langle x \rangle\right)^2}{2\sigma_x^2} \right\} \tag{5.1}$$

The corresponding cumulative distribution, $P_{cum}(x)$, has no closed form expression in terms of standard functions other than as the integral of (5.1) given by (4.1). Although $P_{cum}(x)$ is closely related to the error function, $erf(x)$, defined in Ref.[19], this function is also defined only through a similar integral. The probability that x is smaller than z standard deviations below the mean, or equivalently that x is larger than z standard deviations above the mean, is $P_{cum}\left(\langle x \rangle - z\sigma_x\right) = \widetilde{P}_{cum}\left(\langle x \rangle + z\sigma_x\right)$. Sample values of these cumulative probabilities, or confidence levels, are given in Table 1.

The Excel spreadsheet function $NORMSDIST(z)$ returns the value of the confidence level (as a fraction), i.e., $1 - \widetilde{P}_{cum}\left(\langle x \rangle + z\sigma_x\right)$. The Excel spreadsheet function $NORMDIST(x, \langle x \rangle, \sigma_x, TRUE)$ returns the same result for $x = \langle x \rangle + z\sigma_x$. The Excel spreadsheet function $NORMDIST(x, \langle x \rangle, \sigma_x, FALSE)$ returns the normal pdf, Equ.(5.1).

5.4 The Lognormal Distribution

A variable is lognormally distributed if its logarithm is normally distributed. Consequently a lognormal distribution of a variable y is achieved by using (5.1) with $x = \log y$ and interpreting the mean as $\langle x \rangle = \langle \log y \rangle$ and similarly σ_x is the standard deviation of $\log y$. The logarithm may be natural (ln) or $\log_{10}$ provided that the same choice is used consistently. The confidence levels follow from Table 1. For example, the 95%CL upper bound of y is,

$$10\text{^}\left[\langle \log_{10} y \rangle + 1.6449\sigma_{\log_{10} y}\right] \equiv \exp\left[\langle \ln y \rangle + 1.6449\sigma_{\ln y}\right]$$

It is often possible to confine attention to the mean and standard deviation of the logarithm of the variable. Adopting natural logarithms and denoting these by $\mu = \langle \ln(y) \rangle$, $\sigma = \sigma_{\ln(y)}$. The statistical measures of the lognormally distributed variable y are,

- The median of y is $\exp(\mu)$
- The mean of y is $\exp\left(\mu + \frac{1}{2}\sigma^2\right)$
- The mode of y is $\exp\left(\mu - \sigma^2\right)$
- The standard deviation of y is $\sqrt{\exp\left(\sigma^2\right) - 1} \cdot \exp\left(\mu + \frac{1}{2}\sigma^2\right)$
- The coefficient of variation of y is $\sqrt{\exp\left(\sigma^2\right) - 1}$

5.5 Is a Normal or a Lognormal Distribution More Appropriate?

A variable which can be negative cannot be lognormally distributed.

A variable which cannot be negative cannot strictly be normally distributed.

Hence, a bending moment, for example, which could be of either sign, might be represented by a normal distribution but a lognormal distribution would not be appropriate.

On the other hand, creep ductility, which is necessarily positive, might be chosen to be lognormally distributed but a normal distribution would not be appropriate. The same remark applies in principle to all material properties, which cannot be negative. However it is most important to use a lognormal rather than a normal distribution when the CoV is large, typically for creep ductility, creep rupture life, creep strain rates, fatigue endurance and cyclic stress-strain parameters (e.g., the Ramberg-Osgood A parameter).

In some cases there will be little difference. For example, Young's modulus, and possibly some tensile strengths, might have standard deviations which are small compared with their means (i.e., small CoVs). In such cases, random sampling will, in practice, never produce negative (or even especially small) values if a normal distribution is assumed. Lognormal distributions would strictly be more appropriate but it is unlikely that assuming normal distributions in such cases will make very much difference (i.e., normal and lognormal distributions are very similar for small CoV).

The difference between normal and lognormal distributions for the same CoV is illustrated by Figures 2a-d. For CoVs of 0.1 or 0.2, the difference is minor, but for a CoV of 0.4 or greater the difference is large with the normal distribution producing a significant probability of negative values. Hence normal distributions should be avoided for variables which cannot be negative if the CoV is larger than ~0.2.

5.6 Other Distributions

There are a great many possible pdfs which could be employed. In principle any can be used if they provide a reasonable fit to the data. In terms of their implementation within the Monte Carlo program the main requirement is to ensure equal probability trials. This means that equal probability intervals ("bins") must be devised for whatever distribution is used. In the case of normal and lognormal distributions, the manner of accomplishing this is specified in §8. If the analyst wishes to use another form of distribution and is writing his own code then he will need to devise equal probability intervals in a bespoke manner for each type of distribution. If other distributions are required it is advantageous to use proprietary software to implement the Latin hypercube sampling (see §11.3) since this software will carry out this task automatically.

The motivation to use other types of distribution will be because they are believed to better represent the data or the expected physical behaviour of the variable. The range of pdfs available is too large to attempt a review here. However examples of reasons to use other than a normal or lognormal distribution include,

- The variable in question may have a strict minimum and maximum, so that the pdf is required to be exactly zero outside this finite range. Refs.[5,7] have made use of the PERT distribution in order to impose such a requirement. (The PERT distribution is a special case of the beta distribution, see Ref.[1]).
- Two other, but simpler, distributions which impose a strict maximum and minimum are the rectangular and triangular distributions. The former has uniform probability density between the minimum and maximum of the distributed variable. The latter has a probability density which increases linearly with the distributed variable between the minimum and some maximum of the pdf, and similarly reduces linearly from the maximum of the pdf to the maximum of the distributed variable – the graph of the pdf therefore forming a triangle.
- The Student-t distribution (Ref.[1]) can have advantages in representing small datasets. It may also be desirable if a distribution with a "fatter tail" is deemed appropriate. Ref.[2] has made use of the Student-t distribution in analysing inspection data.

5.7 Histograms

If you write your own code, it is not essential to represent a given dataset by a continuous pdf. Instead a histogram of the data can be used directly. For example, Figure 3 shows a histogram of minimum temperatures of the item being assessed during hot standby, derived by obtaining plant data for a large number of hot standby conditions. Clearly Figure 3 would not be well represented by any continuous pdf. Instead Figure 3 can be input to the Monte Carlo code in tabular form and used directly.

Within the code, correct sampling is achieved simply by ensuring that the histogram of samples reproduces Figure 3. Hence, for this example, if 6720 trials were used then a standby temperature of 270^{o}C should be sampled 1620 times, etc.

Using a histogram of data directly can be the simplest, as well as the most accurate, way of coding the distribution.

If proprietary software is used to implement the Latin hypercube sampling (see §11.3) it is unlikely to support directly the use of arbitrary histograms. However, arbitrary histograms can be used with such software by doing some simple coding.

5.8 K_s and WSEF

The shakedown factor, K_s, and the weld strain enhancement factor (WSEF) are unusual in that advice on their numerical magnitude is given within R5V2/3 itself. Consequently some remarks about their uncertainty are appropriate here.

Data from shakedown tests, e.g., Ref.[20], are usually interpreted in a conservative manner, i.e., a low estimate of K_s is recommended. It is unlikely, at the present time, that sufficient data exist on any material at a single temperature to provide a meaningful measure of the scatter in K_s. Differences between test results at different temperatures have been interpreted as genuine temperature variations rather than scatter. Consequently it is recommended that K_s is treated as deterministic (but temperature dependent) using the R5V2/3 recommended values (or such updates as may be approved). This does not preclude the possibility of treating K_s as distributed should sufficient test data become available to provide an indication of its scatter.

The recommended WSEF values are given in Ref.[21] against three weldment types. At the time of writing only the WSEF values for austenitic weldments have been deconvoluted from the original fatigue strength reduction factor (FSRF) values. The three different WSEF values provide only the coarsest discrimination between different weldments, many of which will be awkward to allocate uniquely to one of the three standard types. On the other hand, the results of the assessment are quite likely to be sensitive to the assumed WSEF. Consequently it is important to treat the WSEF as distributed.

The ideal would be to obtain fatigue test data for the type of weldment being assessed and to obtain a bespoke WSEF, including its scatter, from these tests. However this is unlikely to be a practical possibility in most cases. The second option is to examine the raw fatigue test data which underlies the WSEF recommendations of Ref.[21] to determine what sub-set of the data is actually most relevant. These data will also provide an indication of scatter. However, assessors may be reluctant to attempt a re-analysis of fatigue test data, or may simply have insufficient time to do so. The simplest option is along the following lines. Suppose the weldment to be assessed is judged to be best represented as Type 2, but there is an outside chance that Type 3 might be relevant. This provides some indication of the credible scatter. For example, the Type 2 WSEF might be adopted as the mean, whilst the Type 3 WSEF is interpreted as applying at some higher confidence level, e.g. 95%CL or 98%CL.

The interim WSEF values given in Ref.[21] for ferritic weldments are simply the previous FSRFs. They will therefore be very conservative and should not be adopted as mean values. Expert advice should be sought for ferritic weldment WSEFs. One option might involve using the austenitic WSEFs from Ref.[21] as best estimates for ferritic weldments, with the ferritic WSEFs from Ref.[21] applying at some higher confidence level.

6 NON-σ DISTRIBUTIONS

A "non-σ distribution" is a pdf for which the standard deviation is not defined, i.e., the integral in (4.5) is divergent. It should not be imagined that non-σ distributions are exotic or rare. A simple example is the standard Cauchy distribution,

$$P(x)=\frac{1}{\pi\left(1+x^2\right)} \qquad P_{cum}(x)=\frac{1}{2}+\frac{1}{\pi}\tan^{-1}(x) \qquad (6.1)$$

This is a perfectly reasonable pdf, with well defined median and mode (both being 0), despite not possessing a standard deviation or a mean. (A special case of the Student-t distribution is the Cauchy distribution, hence the familiar Student-t distribution can be non-σ in some cases). To be a "with-σ" distribution (i.e., to have finite variance) the tail of the pdf (the value of $P(x)$ at large x) must reduce faster than $1/x^3$. Conversely, if $LIM_{x\to\infty} P(x) \geq O\left(\frac{1}{x^3}\right)$ then the distribution is non-σ. What this means is that distributions with "fat tails", i.e., which decrease less fast than $1/x^3$, are non-σ. Consequently, the possibility that a distribution might be non-σ becomes important if the tail of the distribution determines the initiation probability, which is most likely for extremely low probabilities, see §12.

Just as non-σ distributions are not at all exotic mathematically, they are not uncommon in real world phenomena either. There are a great many non-σ pdfs, for example the Levy distributions, which have been shown to arise naturally in applications, for example see Ref.[11].

A non-σ pdf can trivially be made into a "with-σ" pdf by truncating its value to zero outside some finite range of x (in mathematics jargon, "compactifying the support" of $P(x)$). However, whether this is a valid representation of the physical variable is another matter since it removes the tails completely.

In reality, non-σ distributions are common but their use in simulations is uncommon. Their use tends to be artificially vetoed by the description of a finite set of data in terms of a mean and a standard deviation. The temptation is then to employ a "with-σ" distribution with these "measured" values of mean and standard deviation. The error is essentially one of extrapolation beyond the available data. Imagine a dataset consisting of 100 random values of x from a distribution which is actually a Cauchy pdf, Equ.(6.1). The mean and standard deviation of these 100 data are well defined, and one is tempted simply to assume (say) a normal distribution with the same mean and standard deviation. Within the range of the 100 samples the assumption of a normal distribution may be reasonable, i.e., it may deviate only slightly from the underlying Cauchy distribution, (6.1). However, for values of x larger than the greatest sampled, the exponentially thin tail of the normal pdf will differ by an increasingly large factor from the fat tail of the Cauchy distribution, (6.1). In other words, the nature of the distribution is uncertain when extrapolating beyond the data. The assumption that "with-σ" pdfs apply is common in engineering assessments. However it should be borne in mind that, if wrong, the assumption is non-conservative in the region of the tails. Fortunately, in cases where the calculated probabilities are not too small, the tails of the individual distributions may not be important. However, for the higher nuclear safety duty components where extremely small failure (initiation) probabilities are expected, the result may be crucially dependent upon the assumed form of the tails. In this case any assumption that "with-σ" pdfs apply beyond the region of available data must be supported by physical arguments.

The Monte Carlo / Latin hypercube methodology recommended here can be used with both "with-σ" and "non-σ" distributions.

7 ILLUSTRATIVE DISTRIBUTED VARIABLES

An example set of distributed variables is listed below. It is intended only for illustration. In a given application there may be other variables, and many of the variables listed below may be unnecessary or be adequately treated as deterministic.

7.1 Illustrative Component Feature

To motivate the example list of variables, a particular structure will be assumed for illustration. This is shown in Figure 4. It consists of a butt weld connecting a pipe to a cast feature. They are of different materials (or, at least, different casts). The assessment needs to be carried out at three locations: (1) the weldment, on the pipe side; (2) the weldment on the casting side, and, (3) a stress raising radius feature in the casting parent material. The computer code will need to track the accumulation of damage at all three locations separately, and in parallel. Crack initiation is conceded when any one of the three locations reaches a damage of unity. Assessment of all three locations in parallel (i.e., within the inner-most loop) is required for this reason: that cracking is conceded when any one location initiates for the current sample of random parameters.

7.2 Time Independent Distributed Variables

Time independent distributed variables are those variables which are randomly sampled at the start of a trial and are constant over all cycles / plant life. The bulk of variables will be time independent. The

specific set of variables which are chosen to be distributed will vary according to the application. An illustrative list is as follows.

1) Pipe thickness at sol (e.g., spread based on drawing tolerance)
2) Cast stub thickness at sol (e.g., spread based on drawing tolerance)
3) Parameter controlling metal loss on steam side
4) Parameter controlling metal loss on reactor coolant side
5) Deadweight system moment, in-plane
6) Deadweight system moment, out-of-plane
7) Steady operating thermal system moment, in-plane
8) Steady operating thermal system moment, out-of-plane
9) Steady operating pressure differential
10) Steady operating reactor coolant temperature at this location
11) Steady operating steam temperature at this location
12) Interpolation factor for deriving metal temperature (based on distribution of flow rates)
13) Young's modulus (parent)
14) 0.2% proof stress (pipe)
15) 0.2% proof stress (casting)
16) The cyclic stress-strain Ramberg-Osgood A parameter (assuming β is determinate) - pipe
17) The cyclic stress-strain Ramberg-Osgood A parameter (assuming β is determinate) – casting
18) The cyclic stress-strain Ramberg-Osgood A parameter (assuming β is determinate) – weld/HAZ
19) The scatter parameter for the creep strain rate – pipe
20) The scatter parameter for the creep strain rate – casting
21) The creep ductility – pipe side HAZ
22) The creep ductility – casting and cast-side HAZ
23) The creep ductility – weld
24) Fatigue endurance – pipe parent
25) Fatigue endurance – casting parent
26) The elastic follow-up factor, Z
27) A parameter used to refine the estimate of creep relaxation, defined in Ref.[15] where it is called "zeta" (ζ), see §13.1 for further discussion
28) WSEF (weld strain enhancement factor)
29) A parameter defining the future cycling frequency.

It is emphasised that this list is merely an illustration. What variables are required will vary. As regards whether the required variables are treated as distributed the golden rule is to do so if any meaningful estimate of their spread is available. There is little detriment to including more distributed variables in the Monte Carlo simulation, either in terms of coding complexity or computer run times.

7.3 Time Dependent Distributed Variables

Examples of possible time dependent distributed variables are given below.

- In general there will be different types of load cycle, e.g. reactor cycles to cold shutdown, reactor cycles to hot standby, boiler cycles, etc. The plant life (both past and future) may be modelled as some deterministic sequence of such cycles, if plant records exist to provide this information. Alternatively, the assessment of each cycle may start by choosing the cycle type randomly. If there are (say) 300 cycles, this implies 300 random samples to determine the cycle sequence for each trial.
- If there are transient thermal stresses during start-up or shut-down which cause the total stressing to exceed that in steady operation it will be necessary to include these transient conditions in the hysteresis cycle construction (within the R5V2/3 procedure). But start-up and shut-down transients are likely to be variable, so these transient stresses will differ from cycle to cycle and hence should be sampled independently on each cycle – hence 300 times per trial if there are 300 cycles over life.
- There may be other parameters which differ between cycles even of the same nominal type, for example the minimum temperature and/or the minimum pressure differential during a hot standby condition (the former may control the system stress range).

There are three options for coding the time dependent variables. The first is to treat them in exactly the same manner as the time independent variables, sampling them at the start of each trial and including them in the Latin hypercube. This means that each cycle is treated as having its own separate set of variables. For example, if there are four types of time dependent variable and 300 cycles over life, this requires 1200 variables. Hence the dimension of the Latin hypercube would be 1200 plus the number of time independent variables.

The second method is a means of accomplishing the same algorithm as option 1 but without cluttering the Latin hypercube with a large number of extra variables. This is done by adding just one extra variable, I, to the Latin hypercube, this taking integer values from 1 to some large integer N_I. Prior to entering the loop over trials, the set of 4 x 300 variables required to define a specific cycle sequence is randomly sampled N_I times and the results stored in an array (of dimensions 4 x 300 x N_I). The single variable I is then treated as a time-independent variable with integer values between 1 and N_I being equally probable. The sampled I value is then used to select the pre-determined cycle sequence from the stored array. The array will be quite large (~10 MB for N_I = 1000 for this example) but this is unlikely to be a problem.

The third option is to define the Latin hypercube to contain only the time independent variables. There would then be (in the above example) four additional time dependent variables which would not be included in the Latin hypercube. These four time dependent variables would be re-sampled and over-written at the start of each of the 300 cycles. This avoids cluttering the Latin hypercube with a large number of additional variables and dramatically reduces the amount of sampling data being stored. In principle this method bypasses the correct sampling algorithm. However, so long as the time dependent variables are sampled on the basis of equal probabilities the three methods should produce indistinguishable results.

7.4 Same Distribution versus Same Value

It is important to recognise the distinction between the following two things,

(i) Two variables having the same distribution;

(ii) Two variables being perfectly correlated.

For an example of (i) consider that the pipe and the cast materials in Figure 4 are of the same generic type, say 316H stainless steel, and that their creep deformation is described by the same RCC-MR law. Nevertheless, they are independently sourced materials (different casts, different product forms).

Consequently their position within the scatter of creep deformation behaviour will be described by two independent, uncorrelated, normalised error variables, say z_x and z_y. Random creep rates are sampled for the two materials from the same distribution but at these differing points in the scatter band, z_x and z_y. Despite having the same distribution of creep deformation behaviour, the fact that (say) the pipe happens to be sampled in a given trial as having a particularly fast deformation rate implies nothing about where the cast item may lie in the scatter band. They are uncorrelated.

As an example of (ii) consider the system stress at the weld and the system stress at the radius feature of the casting in Figure 4. The list of variables in §7.2 contained only one set of bending moments, because the analyst has decided that no distinction need be made between the moments at the weld and at the cast radius (since they are very close). The stresses at the two locations will, of course, be different because of the different stress concentration factors applicable at the two locations. However, the same set of randomly sampled systems moments is to be applied to both. Consequently the ratio of the stresses at the two locations will be the same in every trial. They are perfectly correlated by virtue of being controlled by the same random variables.

8 LATIN HYPERCUBE SAMPLING

8.1 Overview of the Latin Hypercube Sampling Algorithm

The essence of a Monte Carlo probabilistic assessment is to carry out a deterministic assessment many times with different combinations of the distributed parameters. Each deterministic assessment (of the whole of plant life for one component) constitutes a "trial". Provided that all trials are equally probable, the fraction of trials which reach a damage of unity provides an estimate of the probability of crack initiation.

For sake of exposition, suppose N_b different values of each variable are considered. When there are a large number of distributed variables, N_v, the number of possible combinations, $N_b{}^{N_v}$, rapidly becomes prohibitively large so that an exhaustive consideration of all combinations is not possible. Hence it is necessary to sample a small subset of all possible combinations. The generic name for probabilistic simulation by sampling a subset of possible parameter combinations is "Monte Carlo simulation". Care is required to ensure that this sampled subset is representative. A systematic sampling methodology which has proved robust is the Latin hypercube. This is described in sufficient detail to permit coding in §8.2 (and in even greater detail in Ref.[1]).

The dimension of the hypercube is the number of distributed variables, N_v, which would be 29 for the illustrative list in §7.2 (excluding time dependent variables). Each side of the hypercube is divided into N_b "bins", i.e., ranges of values from which a discrete set of possible sampled values for the variable is devised. Each of the $N_b{}^{N_v}$ cells into which the hypercube is thus divided defines one of the $N_b{}^{N_v}$ possible combinations of parameters (i.e., a trial).

A Latin hypercube is a hypercube in which exactly N_b cells are "occupied" and such that no two occupied cells share a bin in any variable. To use a chess analogy, a conventional chess board provides a 2D Latin hypercube if eight rooks are arranged on the board such that none are *en prise*. An example for a 5 x 5 chessboard is shown in Figure 5.

Hence a Latin hypercube represents N_b different combinations of the N_v parameters which together sample every bin in every variable (though not, of course, in every combination). This ensures equitability of sampling between the variables and the bins.

The robustness of the Latin hypercube sampling algorithm relies upon all the resulting trials (combinations of N_v parameters) having equal probability. This is ensured by choosing the bins to have equal probability. Consider a normally distributed variable, z. Equal probability bins implies

choosing wider bins when z lies a long way from its mean, and narrower bins near to its mean, so that the area under the probability density curve (i.e., the bin probability) is the same for each bin. Hence careful definition of the bins is crucial to the outcome of a Monte Carlo simulation.

8.2 Details of the Latin Hypercube Sampling Algorithm

This section gives sufficient details of the Latin hypercube algorithm to permit the analyst to code the method for the case of normal distributions, lognormal distributions and arbitrary histograms.

8.2.1 Normal Distributions

The algorithm is formulated here in terms of dimensionless, normalised error parameters, z_i. These are the errors in the physical parameters (e.g., temperature, yield stress, etc.) divided by their standard deviation. Hence z_i is the number of standard deviations by which the quantity deviates from its mean (greater than the mean when z_i is positive). Hence the relevant normal probability density function (pdf) becomes the standard normal distribution (i.e., with zero mean and unit variance),

$$P(z) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{z^2}{2}\right\} \tag{8.1}$$

The cumulative probability is,

$$P_{cum}(z) = \int_{-\infty}^{z} P(x)dx \tag{8.2}$$

The algorithm addresses N_v distributed variables, x_i, where $i \in [1, N_v]$. Each parameter takes one of N_b possible values, each of which is defined by the mean of the parameter and the value taken by its error variable, z_i, for the particular random sample in question. Thus,

$$x_i = \langle x_i \rangle + \langle z \rangle_i * \sigma_i \tag{8.3}$$

where $\langle x_i \rangle$ is the mean of x_i, σ_i is the standard deviation of x_i, and $\langle z \rangle_i$ is one of the N_b possible values of the dimensionless error parameter, z_i. This notation refers to the fact that $\langle z \rangle_i$ is some mean, or central, value for one of the N_b 'bins' into which the parameter space is divided. The bin ranges are as follows,

$$\text{I}^{\text{th}} \text{ Bin: } \quad z \in [\xi_{I-1}, \xi_I] \tag{8.4}$$

Capital subscripts such as $_{\text{I}}$ will be used to denote bin numbers, to distinguish them from indices representing the variables, x_i.

Note that the Latin Hypercube methodology constrains the number of bins, N_b, to be the same for all of the variables x_i. Moreover, the bins, (8.4), are also the same for all (normally distributed) variables, x_i.

In Latin hypercube sampling, the bins $z \in [\xi_{I-1}, \xi_I]$ are defined so as to represent equal probabilities. Since there are N_b bins this probability must be $1/N_b$. This means that,

$$P_{cum}(\xi_I) - P_{cum}(\xi_{I-1}) = 1/N_b \tag{8.5}$$

The left-most boundary is chosen to be $\xi_0 = -\infty$ so that $P_{cum}(\xi_0) = 0$, and hence (8.5) allows all the bins to be found from,

$$1 \le I \le N_b \qquad \xi_I = P_{cum}^{-1}\left(P_{cum}\left(\xi_{I-1}\right) + \frac{1}{N_b} \right) \qquad (8.6)$$

From this it follows that $\xi_{N_B} = +\infty$, so the entire parameter space from $-\infty$ to $+\infty$ is spanned by unequal sized bins with equal probabilities. In practice the bin boundaries, (8.6), must be found numerically since there is no closed-form expression for P_{cum} other than its definition as the integral (8.2). Consequently a finite value for ξ_N will necessarily result. The User should check that this is a sufficiently large number (of standard deviations) to not prejudice the accuracy of the simulation. This may depend upon the precision of the numerical estimate of P_{cum} and its inverse.

For each bin, a representative value of z must be determined. This is taken to be the mean value of z within the bin, i.e.,

Bin I:
$$\left\langle z_I \right\rangle = \frac{\int_{\xi_{I-1}}^{\xi_I} zP(z)dz}{\int_{\xi_{I-1}}^{\xi_I} P(z)dz} \qquad (8.7)$$

The denominator is just equal to (8.5), whereas the numerator can be evaluated explicitly for a normal distribution, (8.1), to give,

$$\left\langle z_I \right\rangle = \frac{N_b}{\sqrt{2\pi}}\left(\exp\left\{ -\frac{\xi_{I-1}^2}{2} \right\} - \exp\left\{ -\frac{\xi_I^2}{2} \right\} \right) \qquad (8.8)$$

Note that the use of (8.8) is particularly important for the first and last bins since it assigns a finite mean $\left\langle z_I \right\rangle$ to a bin of theoretically infinite width. The values of $\left\langle z_I \right\rangle$ for the first and last bins define the extremes of the sampling, i.e., the minimum and maximum values.

A Latin hypercube is defined as a choice of N_b cells none of which share any row/column/rank/…. etc., covering all N_v directions. The Latin hypercube algorithm consists of randomly selecting a Latin hypercube and then using all N_b trials which the Latin hypercube represents. The advantage of this approach is that it ensures that every bin of every parameter is used in just N_b trials (albeit in only a very small sub-set of possible combinations). Note that this means that the number of trials equals the number of bins, N_b.

The value chosen for N_b determines the greatest number of standard deviations away from the mean which is sampled. It is not possible with the Latin hypercube algorithm to sample a large number of standard deviations using only a small number of trials – because the number of trials equals the number of bins, and this would conflict with the requirement for bins of equal probability. The relationship between the number of standard deviations sampled (and hence the confidence interval for each variable) and the number of trials (= the number of bins) is given in Table 2.

8.2.2 Lognormal Distributions

The algorithm for the normal distribution can also be used for variables distributed lognormally since a lognormally distributed variable, y, can be defined in terms of a normally distributed normalised error variable, z, as follows. If logarithms to base 10 are used, and if a and b are the mean and standard deviation of $\log_{10} y$ respectively, then sampling a normally distributed variable, z, in accord with equ.(8.1), will provide a lognormally distributed y by setting,

$$y = 10\text{^}\{a + bz\} \tag{8.9a}$$

Alternatively, if natural logarithms are used, and a and b are the mean and standard deviation of $\ln y$ respectively, then sampling a normally distributed variable, z, in accord with equ.(8.1), will provide a lognormally distributed y by setting,

$$y = \exp\{a + bz\} \tag{8.9b}$$

The probabilistic coding may therefore be carried out using the normally distributed variable, z, and the corresponding 'bins', as specified in §8.2.1, and (8.9a) or (8.9b) used to find the actual (lognormally distributed) variable, y.

8.2.3 Other Continuous Distributions

There is no difficulty in principle in accommodating other pdfs. If code is to be written to do so then the bin ranges $z \in [\xi_{I-1}, \xi_I]$ and mean bin values $\langle z_I \rangle$ must be derived for the pdf in question. It is greatly advantageous to use proprietary software to do this automatically (see §11.3). For example, Holt has made use of the PERT distribution in Ref.[7] and the Student-t distribution in Ref.[2] via RISKAMP, Ref.[12].

8.2.4 Histograms

There is no difficulty in including probabilities represented by arbitrary histograms of data in purpose written code. The guiding rule is simply that,

- Over the N_b trials of one Latin hypercube, the histogram of sampled values shall exactly reproduce the desired histogram.

This rule ensures two things: firstly that all trials are equally probable, and secondly that the sampling covers all (implicit) "bins" exactly once. Hence the two key requirements of the sampling process are respected.

Proprietary software may not include the facility to address arbitrary histograms directly but this can be accomplished with additional coding by the User.

9 ALEATORY VERSUS EPISTEMIC VARIABLES

It is common in reliability theory to distinguish between two types of uncertain variables: aleatory and epistemic variables. Aleatory variables are distributed variables whose randomness is characterised by chance. Their randomness is deemed to be irreducible. In contrast, epistemic variables are distributed variables whose uncertainty arises simply from lack of knowledge. Note, however, that the classification as aleatory or epistemic is one of convenience and may vary with the context. The distinction is not entirely objective. An example will clarify how the categorisation might be used in practice.

Consider some material property, perhaps tensile strength, or creep ductility, it does not matter. Suppose some test data exists, all of which are regarded as equally applicable to all components within a population of similar components whose integrity is being assessed. This material property is an aleatory variable under these conditions. The scatter in the test data reflects the uncertainty in the material property applicable to each individual component. However, it may be that several different casts of material were used in the construction of the components, and the material property in question may be cast dependent. If cast-specific test data were available, and if the cast used for individual components were known, then cast-specific assessments could be carried out for individual components. More likely, it may be that complete cast-specific information is not available. The test data might be for unspecified casts and the plant knowledge may extend only to the percentages of various casts used, rather than identification of the casts used for each individual item. Nevertheless,

every individual item is of a definite cast, and hence its material property is specific to this cast. The overall uncertainty in the material property can now be considered as comprising two parts: an epistemic part due to cast-to-cast differences, and an aleatory part due to scatter within a given cast.

It might be objected that, in principle, all uncertain variables are really epistemic since there is ultimately a cause for every instance of scatter. In the above example, every individual item has a definite yield strength, rupture strength, etc., we just do not know what it is. Even the toss of coin, the usual analogy for aleatory uncertainty, is deterministic in principle. However, the distinction between aleatory and epistemic variables is not made for reasons of philosophical nicety but for practical purposes. Perhaps the chief benefit is in addressing plant survival/failure data. Even if a given population of components has suffered no failures in 30 years it is *not* true in general to conclude that no failures will ever occur. On the other hand it is perfectly valid to deploy the absence of past failures as a constraint on predicted future failure rates. This can be quantified by restricting the possible range of values for chosen epistemic variables to be consistent with the plant experience. This in turn has implications for predicted future failure rates. A specific example of the use of epistemic variables in this way is Refs.[2,3]. Here the unknown initial distributions of crack depths and initiation times were used to obtain acceptable agreement between the Monte Carlo model and the historic inspection data and history of steam leaks. This permitted predictions of future defect incidence to be made. The initial crack depths and initiation times were the epistemic variables in this example.

It is obviously desirable to deploy all the information available, which includes the plant experience. In this respect it should not be forgotten that survival, i.e., the absence of failures, is also plant experience. The utility of epistemic variables lies in providing the degrees of freedom required to implement the constraints arising from plant experience.

The distinction between aleatory and epistemic variables, and in particular the use of plant data to reduce epistemic uncertainty, is discussed in far greater depth by Chevalier, Ref.[18].

10 SIMULATION OPTIONS

10.1 Cycle Sequence

It is common in deterministic R5V2/3 assessments to idealise the cycling of the plant to just one or two bounding cycles. Hence, in the simplest case, the total plant life may be represented by (say) 300 cycles, with all cycles being assumed to be the same and equal to some notionally bounding cycle. This might be refined a little by considering (say) 100 cycles to cold shutdown plus 200 cycles to hot standby conditions, but again with all ‘cold’ cycles being assumed bounding and all ‘hot’ cycles also being bounding for that type. In probabilistic simulation, retaining this simplified approach to the cycling history (and future) remains an option. However, it is contrary to the spirit of the probabilistic approach (namely to be “best estimate” where ever possible, §5.1).

The ideal probabilistic approach is to obtain the complete plant cycling history data. Some cycles may be shown to be minor (e.g., on-load refuelling cycles). It will often be the case that only reactor and boiler cycles need be addressed in the R5V2/3 assessment, possibly with a small increase in the actual number of cycles to accommodate the effects of the minor cycles (see, for example, Ref.[13]). The cycling history for a given reactor may often be adequately described by,

- The chronological sequence of cycle type (e.g., reactor cycle to cold shutdown, reactor cycle to hot standby, or a boiler cycle);
- The operating hours (creep dwell) associated with each cycle;
- Data specific to variable hot standby conditions, e.g., the minimum temperature and minimum steam pressure.

If the simulation is to be carried out quad-by-quad, or boiler-unit-by-boiler-unit, then boiler cycles may be deployed on a quad or boiler specific basis. Otherwise it may be simplest to base the assessment on the largest number of boiler cycles for any single boiler, or the average thereof. Applying the largest number of cycles for any boiler to all boilers departs from the ideal of being 'best estimate' (§5.1) and is one of the many compromises of this ideal which might be adopted for convenience. For applications other than to nuclear power plant, similar remarks will apply whenever the whole plant consists of a number of broadly equivalent sub-units.

The future cycling is, of course, not known. Nevertheless, the analyst has the option to treat the future cycling as determinate or to attempt a probabilistic description. In either case the future cycling should be represented realistically based on the *recent* cycling history (unless there is reason to believe that this will change). It is commonly the case that the frequency of cycling has changed over plant life, with more frequent shutdowns often occurring in early life. There will be a degree of subjectivity in judging what period is the best guide to the future. The objective is to be "best estimate". Given that there is genuine and significant uncertainty about future cycling rates, it may be better to attempt to include a distributed variable representing the future cycling rate. On the other hand, coding is simplified if the number of cycles is fixed, and a safety case customer may have a specific requirement as regards assumed future cycling rates.

10.2 Cycle Interaction

The advice of §10.1 leads to the need to assess, for each trial, a sequence of unequal cycles. This raises some issues which are not fully addressed in R5V2/3. The first is that each individual stress-strain hysteresis cycle is not closed, even in principle – because the starting condition (e.g., cold shutdown) will in general be different from the end condition (e.g., hot standby). This is important but relatively easy to accommodate within the spirit of the R5V2/3 procedure (see, for example, the detailed algorithm of Ref.[6] Appendix C for hysteresis cycle construction in this circumstance).

The greater problem concerns the absolute positioning of a hysteresis cycle along the stress axis. Where a sequence of identical cycles is presumed, the R5V2/3 methodology is to assume "symmetrisation" of the cycle (i.e., the hysteresis cycle is assumed to protrude equally beyond $K_S S_y$ at both ends of the cycle). However, if there is a sequence of unequal cycles it will not be true that every cycle is individually "symmetrised". Instead, the absolute positioning of cycle 2 will depend to some degree on the position of cycle 1. Cycle 2 will start where cycle 1 leaves off (in stress-strain space). On the other hand, if cycles 2 to 10 are identical, but differ significantly from cycle 1, by cycle 10 the hysteresis cycle will probably be symmetrised, the influence of cycle 1 having been "forgotten". An algorithm which attempts to model this tendency to re-symmetrise whilst also being influenced by the most recent cycles has been presented and employed in Refs.[6,7]. However, this algorithm is not validated at this time. An alternative is to assume that every cycle is individually symmetrised. Though this is unphysical it produced very similar results to the interaction algorithm in Refs.[6,7]. It is *not* acceptable to model every cycle as if it strictly followed the previous cycle (i.e., such that the last stress-strain point of the previous cycle is the starting stress-strain point of the next cycle, with the hysteresis cycles being calculated using the R5V2/3 Appendix A7 procedure). This assumption would imply that the positions of all the stress-strain cycles throughout life were determined by the first cycle, perhaps 40 years earlier. This assumption is not recommended.

The handling of cycle interaction is currently unsatisfactory and must be addressed on a case by case basis.

10.3 Multiple Assessment Locations

Where there is more than one location on each component which might be limiting, such as illustrated by Figure 4, the assessment should address each location in parallel. The phrase "in parallel" means that the damage for all locations should be calculated at the same time (i.e., within the innermost loop

over cycles). This is because cracking will be conceded when damage reaches unity at any one of the locations. The trial can be terminated when the first location reaches a damage of one.

If, in contrast, the probabilistic code were run three times, once for each of the three assessment locations separately, the result would be an initiation probability separately for each location, $\widetilde{P}_J$. It would be an error to assume that the overall initiation probability is $\sum_J \widetilde{P}_J$. The probability that initiation occurs at any location (P, the desired quantity) cannot strictly be determined from the individual $\widetilde{P}_J$. In general the only certain statement is that $MAX\left(\widetilde{P}_J\right) \le P \le MAX\left(\sum_J \widetilde{P}_J, 1\right)$.

10.4 Treatment of a Population of Components

Probabilistic assessment can be carried out for a single component. The outcome would thus be the (estimate of the) probability that this specific component will initiate a crack. Hence, if N_f is the number of trials of this component which initiate a crack, out of a total of N_t trials of this component, then the estimated probability of crack initiation in this component is,

$$P = \frac{N_f}{N_t} \tag{10.1}$$

Alternatively, probabilistic assessment can be carried out for a population of components, e.g., all the superheater bifurcations in a given boiler or in a given reactor. If a population of components is to be addressed, the result will not be a probability of initiation, but rather the expected number of initiations or the frequency of initiation.

Suppose that P_i is the probability that the ith component initiates by the eol. Then $E = \sum_i P_i$ is the expectation value of the number of cracks by eol, which may be a fraction or may be greater than 1. If $E = \sum_i P_i$ is a ***small*** fraction then it approximates the probability that some component of the population cracks by eol. But in general $E = \sum_i P_i$ cannot be interpreted as a probability.

Alternatively, suppose that $\widetilde{P}_i$ is the probability that the ith component initiates *in* a given year (as opposed to "*by* a given year"). Then $\widetilde{E} = \sum_i \widetilde{P}_i$ is the expectation value of the number of cracks initiating in that year, i.e., the initiation rate. It may be a fraction or it may be greater than 1. If $\widetilde{E} = \sum_i \widetilde{P}_i$ is a ***small*** fraction then it approximates the probability that some component of the population cracks in the given year. But in general $\widetilde{E} = \sum_i \widetilde{P}_i$ cannot be interpreted as a probability.

The variation of $\widetilde{E}$ in future years can be of considerable interest as an indication of increasing risk (or not, as the case may be). In contrast, the end of life expectation value, E, does not provide this insight.

If all the components are identical in all respects, then all the P_i are equal and the simulation need only consider one component. The desired output is then $\sum_i P_i = NP_1$ where N is the number of components. However, if the components differ, perhaps in terms of stress or operating temperature or metal loss, for example, then there are two options,

(i) Carry out a probabilistic simulation of each individual component in the population, using component-specific data for each, hence calculating each individual P_i, in terms of the number of trials which initiate a crack in this component, N_{if}, out of a total of N_{it} trials of this component, so that $P_i = \frac{N_{if}}{N_{it}}$. The expectation value of the number of cracks is then found as indicated above, $E = \sum_i P_i$ where the subscript is summed over all the components in the desired population; OR,

(ii) Express the spread of data (say, operating temperature) between the components as a distributed variable so that one simulation addresses the whole population. If N_t such trials are carried out and these result in a total of N_f crack initiations then the expectation value of the number of cracks over the whole population is $E = \frac{N_f}{N_t}$.

In the case of just one component both the expressions for E in (i) and (ii) reduce to Equ.(10.1).

Option (i) has the advantage that individual component probabilities are determined, which may be helpful in identifying the specific items most at risk. However, obtaining convergent results for individual components will require a larger number of trails for the reactor as a whole (see §11.7). Hence (ii) may be more efficient if only reactor-wide (or boiler-wide) results are required.

Use of option (ii) may mean that some quantities involve two independent distributed variables. Take operating temperature as an example. Suppose that the temperature of any specified item is described by a mean T_m and a standard deviation σ_T. The latter may describe the aleatory uncertainty in temperature (e.g., due to measurement errors) or it may describe the epistemic uncertainty due to variation of temperature over time (or a combination of both). However, a different item may have a different best estimate temperature, so that T_m is itself subject to an epistemic distribution which describes the variation of temperature from item to item.

In some cases only a reactor-wide mean effective creep temperature (MECT) and its scatter may be available. If this has been derived on the basis of data covering a range of items and a range of times, then this MECT and its standard deviation implicitly cover all the above distributions (providing measurement errors are also incorporated). If this is all that is available, then option (ii) is the only possible option (i.e., the absence of component-specific data precludes component-specific assessment).

10.5 Perfect Correlation

The coding of partial correlation is discussed in §11.5 and Appendix A. However, perfect correlation ($C_{xy} = 1$), or perfect anti-correlation ($C_{xy} = -1$), is simply implemented by sampling both variables x and y using the same value of a single normalised error variable z. Note the crucial distinction between this and using different error variables, z_x and z_y, taken from the same distribution (see §7.4).

11 CODING ISSUES

11.1 Hardware Platforms

Whilst probabilistic R5V2/3 assessments are more demanding in terms of computer run times than many engineering calculations (other than large finite element analyses) good results have been

obtained using standard desktop PCs (see run time data in §11.9). Moreover, many PCs in the typical office stand idle much of the time so that use can be made of this readily available computing resource. This does not preclude the use of Unix or Linux workstations or specialist analysis boxes such as are used for big finite element analysis (FEA) jobs, but the use of these more specialist computing resources is not essential.

11.2 Software Platforms

In Refs[2-10] the software platform used was Excel with the bulk of the coding generally being implemented in a Visual Basic macro. Visual Basic is quick enough to be practicable. On the other hand, carrying out the coding within the spreadsheet is ***not*** recommended as this is far slower. This does not detract from the value of using spreadsheet(s) as the vehicle for input and output. However, calls to the spreadsheet(s) from Visual Basic during execution should be minimised since this is a relatively slow procedure. Hence, inputs from the spreadsheet(s) should be read only once and stored within the code. Any outputs to the spreadsheet during execution (e.g., per trial) should be kept to a minimum for good execution times, e.g., restricted to the percentage progress through the job, the number of initiations in the last trial and the running estimate of the expectation value of the number of initiations. The latter is valuable to monitor convergence in real time.

Another advantage of using Excel is that Latin hypercube proprietary add-ons, §11.3, are known to be available and known to work well.

Matlab would also provide a suitable vehicle for coding and has Latin hypercube facilities within it (e.g., the "lhsdesign" function).

Both Excel and Matlab have convenient features that facilitate pre- and post-processing, e.g., creating graphs and histograms, etc.

11.3 Proprietary Latin Hypercube Add-Ons

RISKAMP, Ref.[12], has been used successfully with an Excel implementation on a PC, e.g., Refs.[2,5,7]. @RISK, Ref.[14], is believed to have similar functionality although confirmation of this, in particular as regards its efficiency when coupled with Visual Basic coding in Excel, will require application experience (lacking within EDF Energy at the time of writing). @RISK is available within EDF Energy, whilst RISKAMP is not. However RISKAMP can be purchased for a single User/PC at very low cost.

11.4 Coding Efficiency

At the centre of the probabilistic simulation is the repeated construction of hysteresis loops and the associated creep and fatigue damage increments. The number of times this is done equals the number of trials (or bins, N_B) times the number of cycles over life, times the number of components assessed, times the number of assessment locations. This product can easily exceed 10^9 (for example, $N_B = 10{,}000$, 300 cycles, 300 components, 3 locations, giving a product 2.7 x 10^9). Consequently it is crucial that the hysteresis loop construction and creep and fatigue damage calculations be implemented efficiently. There are two areas where attention to efficient coding is particularly important,

- In constructing the hysteresis loops the Neuber construction is used repeatedly. This involves the inversion of the Ramberg-Osgood equation which must be done by numerical iteration. If code is written 'from scratch' to accomplish this, attention should be paid to the efficiency of the solution algorithm employed. Automatic solution procedures, such as the "Solve" facility in Excel, may not have the desired efficiency, depending upon the demands of the application. "Solve" can also fail to converge, resulting in errors that may be hard to detect.

- The evaluation of creep relaxation is usually carried out by integration of forward creep expressions. This requires numerical integration (time-stepping or the equivalent). Experience has shown that this can be the most significant determinant of run times. There is scope for the adoption of sophisticated algorithms here, including conditional use of approximate analytic expressions to reduce the number of integration steps required (see, for example, Refs.[5-7]). Care may be needed in the relaxation routine due to singular primary creep strain rates at zero datum time or strain.
- Efficiency issues related to input/output calls to Excel spreadsheet(s) have been discussed in §11.2.

11.5 Implementation of correlation

Perfect correlation or anti-correlation is discussed in §10.5.

The use of proprietary codes to implement the Latin hypercube (§11.3) has the added benefit that these codes allow an arbitrary correlation, C_{xy}, between two variables to be implemented simply by inputting this correlation coefficient. If the analyst is writing his own code, a method for implementing an arbitrary correlation between two variables is defined in Appendix A. This Appendix also specifies the procedure for implementing correlation between more than two variables.

11.6 Coding the Latin Hypercube

An example of coding the Latin hypercube in Visual Basic is provided in Appendix B.

11.7 Numbers of Trials Required

The number of trials required shall be determined by satisfaction of the convergence criteria (§11.8).

The number of trials required will depend upon the initiation probability or rate being estimated. The smaller this is, the more trials will be required.

Fewer trials will be required to determine the eol initiation probability (or eol number of initiations) compared with the number of trials required to resolve each yearly initiation rate. The convergence tests shall be applied to the quantity which is required. Hence, if resolution of the individual yearly rates is required, the convergence tests shall be applied to each yearly rate separately. However, if only the eol result is required (equivalent to the average yearly rate) then the convergence tests need be applied to this eol quantity alone. Consequently, significantly more trials will be required to resolve rates for individual years separately. However this may be desirable if a worsening trend of increasing annual rates is anticipated.

Fewer trials will be required to determine the number of initiations across a population of components than to determine the initiation probability for each component individually. (Recall that “trial” is used in this report to mean the assessment of a single component with a single set of randomly sampled inputs. Hence systematically assessing all components from a population of N_C components with a single set of randomly sampled inputs is N_C trials).

The most efficient means of evaluating the number of initiations across a population of components will depend upon how the initiation probability varies across the population. If the individual component initiation probabilities are all of comparable order of magnitude then simulating the whole population in a single computer run is appropriate. However, if a sub-population has substantially higher initiation probabilities than the rest, it will be most efficient to carry out runs for the two sub-populations separately. In this situation it may be acceptable for one (but only one) of the sub-populations to fail the convergence tests if enough has been done to demonstrate that this sub-population contributes negligibly to the total incidence of cracking.

The acceptable initiation probability, or rate, will vary depending upon,

- the nuclear safety duty of the component;
- the strength of other legs in the safety case; and,
- the implications of cracking for operability/lifetime.

For example, consider a population of 2000 components in the "Frequent" tolerability of failure category. It is desired to demonstrate that the *average* initiation rate over the whole population does not exceed 0.2 per year, applicable over the next 10 years. This is equivalent to 2 cracks in 2000 components over the required 10 year period, i.e., a fraction 10^{-3}. Hence, a first estimate is that 10,000 trials may be adequate (i.e., ten times the reciprocal of 10^{-3}). This must, however, be confirmed as adequate via the convergence tests. If so, computer runs might take perhaps ~14 minutes (see §11.9).

On the other hand, it may be desired to establish for a single "High Integrity" component that the eol cracking probability does not exceed 10^{-6}. A first estimate is that this would require 10 million trials to establish or refute. This might take the order of 10 days to run on a standard PC (if the example applications of §11.9 are indicative). Moreover, this level of reliability calls into question the appropriateness of the tails of the distributions assumed (see also §6 and §11).

11.8 Convergence Tests

It is essential to confirm that a sufficient number of trials has been used in a given run to obtain convergence of the predicted probability, or rate, of cracking. The two tests defined in §11.8.1 and §11.8.2 must be carried out.

The tests must be carried out on the output quantities which are to be reported. If only eol (or averaged-over-years) quantities are required, then the tests can be confined to these. If only whole-reactor, or whole-boiler, quantities are required, then the tests can be confined to these. Conversely, if yearly rates are required then the tests shall be applied individually to each yearly rate. If component-specific initiation probabilities are required then the tests shall be applied to individual components.

The convergence tests ensure only that a sufficient number of trials have been used. They do not guarantee accuracy.

11.8.1 Real-Time Monitoring

The first test is to monitor in real time how the current estimate of the required output quantity develops whilst the simulation is running. An example of this is shown as Figure 6. A sufficient number of trials shall be used to ensure that towards the end of the simulation the results, as further trials are completed, vary by no more than ±10% about a converged value. This test shall be carried out for every reported computer run (though the convergence data need not be reported for more than one or two example runs).

11.8.2 Repeat Identical Runs

The second test involves running a simulation with unchanged inputs at least three times. Results will vary, of course, due to random sampling. A sufficient number of trials shall be used to ensure that the range of results for all the output quantities of interest is within ±10% of the corresponding average. This test need only be carried out for one set of inputs to establish the adequacy of the number of trials.

11.9 Run Times

The following typical run times have been found for an implementation of the procedure in Excel using Visual Basic for the bulk of the coding. Standard office desk-top PCs at Barnwood were used. Trials relate to one component assessed over the whole of life.

- Based on Ref.[6], with 261 cycles over life, 3 assessment locations per component, and incorporating complex time-varying oxidation metal loss calculations and time-varying flow restrictions: 0.33 seconds per trial for single processor machines and half this for dual core processors (or 0.15 seconds per trial for a core i5 processor).
- Based on Ref.[9], with 385 cycles over life and 3 assessment locations per component: 0.085 seconds per trial using a quad core i5-2400 processor at 3.1 GHz (now in standard desktop PCs).
- Based on the last example, 10,000 trials would take 14 minutes whereas 10,000,000 trials would take 10 days.

12 DOES THE TAIL WAG THE DOG?

12.1 When Does the Assumed pdf Introduce Errors?

The pdf chosen will affect the result. However, if the trials which result in cracking only sample values within the range of the available data then the form of the chosen pdf is unlikely to affect the results greatly if alternative pdfs all provide reasonable fits. The chosen pdf is more likely to make a substantial difference to the outcome if the trials which result in cracking sample values beyond the range of the available data. In this case, since alternative pdfs may have significantly different tails, they may produce very different results. In this case it will be necessary to attempt to justify a given form of tail (and hence pdf) on physical grounds. If there is no reason to prefer one pdf to another, and they produce very different results, then the calculated initiation rates must be declared correspondingly uncertain. The bottom line is that the probabilistic approach is not magical and cannot produce knowledge from nothing.

However, it should not be assumed that an initiation probability of (say) 10^{-4} cannot be justified because it requires an unattainable level of reliability (i.e., individual quantities will not generally have data out to the 99.99% confidence level). This is because, roughly speaking, a probability of 10^{-4} can be achieved by two variables each sampled at a probability of 10^{-2}, or by four variables each sampled at a probability of 10^{-1}. Moreover, there are many more combinations of the latter type than single variables sampled at a probability of 10^{-4}. Consequently outcome probabilities of $\sim 10^{-4}$ may be achieved robustly, without great reliance on the tail of individual distributions. For this reason the large number of distributed variables in a typical R5V2/3 assessment is actually advantageous.

Nevertheless, at a sufficiently small output probability, dependence upon the tails of the pdfs will ultimately arise.

12.2 Higher Nuclear Safety Duty Applications

There is no principle which prevents the probabilistic approach being used for the higher nuclear safety duty items. However there are two practical issues which may make it more problematical in practice, both a result of the very small probabilities which must be demonstrated. The first is that small probabilities require a correspondingly large number of trials to calculate (e.g., a probability of 10^{-6} may require 10^{7} trials and hence 10 days computing time). This is merely an inconvenience rather than an insurmountable problem. The second problem is more fundamental. This is that the very small required probabilities may result in a dependence on the tails of the pdfs. For the reasons discussed in §6 and §12.1 this would have implications for the veracity of the results. Where the initiation probabilities are so small that there is a potential for a sensitive dependence on the tails, this should be tested by examining the single-variable probabilities which contribute to trials which initiate.

13 EXAMPLE MATERIALS DATA: 316H STAINLESS STEEL

It is beyond the scope of this report to discuss specific inputs to a probabilistic R5V2/3 assessment. However there will be some generic features across many such assessments, for example the issue of the distributions used for key materials properties. In the case of 316H stainless steel and compatible weld consumables operating in the creep regime but not exceeding about 580°C, Ref.[15] has provided advice for most of the required materials property distributions. Note however that Ref.[15] used component-specific, certificate-based tensile data which would need to be replaced for the application in mind. Further note that Ref.[15] was intended for application to thin components where the weld/HAZ properties could reasonably be assumed equal to those of the parent. For thick section weld/HAZ this may not be appropriate.

13.1 Reset of Primary Creep

If creep relaxation is evaluated by integration of forward creep, a rule must be assumed for the carry-over of creep hardening from one cycle to the next. The options are, (i) to accumulate creep strain between cycles, known as "continuous hardening", or, (ii) to reset the effective creep strain to zero at the start of each dwell, known as "primary reset". The latter will result in substantially larger creep strains, and hence greater creep damage. The current advice in deterministic R5V2/3 assessments of 316H, originating from Ref.[16], is to use continuous hardening at and below 550°C but to adopt primary reset above 550°C. However, more creep-fatigue test data were analysed in Ref.[15] which led to a modified recommendation *for probabilistic R5V2/3 assessments*. This is,

(a) Adopt primary reset at all temperatures so long as the last unloading prior to the creep dwell caused a reverse plastic strain of 0.01% or more (otherwise use continuous hardening);

(b) Ameliorate the creep strain resulting from (a) by multiplying by the factor "zeta" (ζ), whose distribution of values is defined in Ref.[15];

This advice, originating from Ref.[15], is recommended for probabilistic R5V2/3 assessments as being more appropriate in this context than the advice of Ref.[16]. It is likely that this issue will be examined more closely over the next three years under work programmes now being devised.

14 REQUIRED OUTPUTS

It is recommended that the following items be reported,

- The inputs as would be required for a deterministic R5V2/3 assessment, but also including the pdfs and their defining parameters (e.g., mean and standard deviation, or the equivalents).
- Any correlations assumed between input variables.
- The type of implementation, i.e., what form of coding was used, whether proprietary software was used, etc. At least some of the input/code files should be retained in the electronic Task File (though perhaps not all since they may be large).
- The outcome of the convergence tests defined in §11.8.1 and §11.8.2, justifying the adequacy of the number of trials used in the production runs.
- The hardware platform used and typical run times.
- The distribution (e.g., in histogram form) of the R5V2/3 creep and fatigue damages, D_c and D_f, for trials which initiate a crack (optionally also for trials which do not crack).
- The distribution (e.g., in histogram form) of the following two fractions,
  - "frac1": If the start-of-dwell stress evaluated using the R5V2/3 hysteresis cycle construction procedure is less than the rupture reference stress (used in R5V2/3 as the steady state stress), then the rupture reference stress must be used as the start-of-dwell stress. The definition of

frac1 is the fraction of dwells for which the start-of-dwell stress is set equal to the rupture reference stress;

- "frac2": the fraction of dwells which use continuous hardening to evaluate creep relaxation (as discussed in §13.1).

It can be enlightening to find these two fractions separately for those trials which predict cracking and those which do not.

- The desired output quantities will vary according to the application but will include at least one of the following,
  - The eol number of cracks predicted across some defined population of components;
  - The eol probability of cracking for certain specified individual components;
  - The graph against time of the expected numbers of cracks initiating per year across some defined population of components;
  - The graphs against time of the probability of cracking per year for certain specified individual components.
- For selected input variables, evaluate and report the correlation between these variables and the incidence of cracking. This is very valuable to establish the most significant factors which result in cracks initiating.
- The distribution (e.g., in histogram form) of selected parameters of interest in the particular application. Examples might include,
  - Histograms of thickness for applications with metal losses incorporated (perhaps restricted to cases which initiate);
  - Histograms of creep ductility restricted to cases which initiate cracks (which will therefore differ from the input ductility distribution);
  - Histograms of some relevant stress, perhaps restricted to cases which initiate cracks;
  - Histograms of operating temperature, perhaps restricted to cases which initiate cracks;

15 VERIFICATION

Perhaps the main disadvantage to deploying Monte Carlo probabilistic assessment is the difficulty presented to the independent Verifier. Independence of Verification to the procedure of Ref.[17] requires that the Verifier use an independent code. An exception to this is made for finite element analyses where it is acknowledged that repeating an FE analysis with an alternative code would be prohibitively costly in time and money (for QA Grades 2-4). However, in this case "reasonableness checks" are offered in lieu of the use of an independent code. Moreover, the FE code used must be recognised as robust and reliable with a track record of previous applications, preferably a proprietary code (see Ref.[17]). Neither of these precautions is available for an R5V2/3 Monte Carlo assessment. There is no general purpose probabilistic R5V2/3 code which has achieved the status of being "recognised as robust and reliable". Indeed, there is no such code at all at the present time. It is important that the Verifier carries out whatever "reasonableness checks" are feasible. However, the nature of a probabilistic assessment is, unfortunately, that it is all but impossible to gain sufficient confidence without reproducing the Monte Carlo simulation independently.

Consequently, the advice is that the Verification should involve the use of an independently produced code by the independent Verifier so as to independently derive the reported initiation probabilities or rates. Consequently the Verification may nearly double the number of mandays required for the task.

Unlike deterministic assessments where such reproduced results might be expected to be accurate, in the case of probabilistic assessments exact reproduction of results will never be possible if only

because of the random sampling involved. In practice the differences obtained from independent codes will generally be greater than is explicable on the basis of different random samples alone. Experience with independent coding produced in parallel by different engineers has shown that there is a myriad of places for differences of detail to arise (compare Ref.[4] with Ref[5], and Ref[6] with Ref.[7]). It is not possible to be definitive about the level of agreement which might be considered satisfactory since this will vary with the application. However, it is suggested that agreement to within a factor of 2 or 3 is probably adequate in most applications (i.e., when the absolute initiation probabilities are small).

16 CONCLUSION

A procedure has been defined for the application of Monte Carlo probabilistic simulation in creep-fatigue crack initiation assessments to R5 Volume 2/3, Ref.[23]. This report does not address the 'precursor' assessments required by R5 Volume 2/3 but only the crack initiation assessment. The report addresses only probabilistic aspects, not the deterministic methodology to be used as the core of the probabilistic model, which is assumed to be that of R5 Volume 2/3. One exception relates to the adoption of primary creep reset versus continuous hardening in estimating creep relaxation (specific for 316H parent stainless steel). Other limitations are as specified in §2.

17 RECOMMENDATION

To: Group Head, Assessment Technology

The probabilistic procedure defined herein should be considered for inclusion within R5 Volume 2/3.

18 REFERENCES

[1] D.P.Kroese, T.Taimre and Z.I.Botev, "Handbook of Monte Carlo Methods", Wiley Series in Probability and Statistics, Wiley-Blackwell, April 2011.

[2] P.J.Holt, "Hinkley Point B and Hunterston B Power Stations: Probabilistic Predictions of Creep/Fatigue Crack Growth and Eddy Current Inspection Results for Superheater Bifurcations for Full Power and Reduced Power Operations", E/REP/BBAB/0009/AGR/08, Rev.001, August 2010.

[3] P.J.Holt and R.A.W.Bradford, "Application of probabilistic modelling to the lifetime management of nuclear boilers in the creep regime: Part 1", International Journal of Pressure Vessels and Piping 95 (2012) 48-55.

[4] R.A.W.Bradford, "HYA/HAR: Specification and User Guide for BIFLIFE-RB: A Verification Facility for a BIFLIFE Probabilistic Creep Rupture Assessment of the Hartlepool Boilers' Superheater Bifurcation Inlets", E/REP/BBAB/0019/HAR/11, Rev.000, December 2012.

[5] P.J.Holt, "HYA/HAR - BIFLIFE: A Monte Carlo Probabilistic Creep Rupture and Plastic Collapse Assessment of the HYA/HAR Boilers' Superheater Bifurcation Inlets with Application to Hartlepool Reactor 2", E/REP/BBAB/0021/HAR/11, Rev.000, December 2011.

[6] R.A.W. Bradford, "Heysham I / Hartlepool Power Stations: BIFINIT-RB - A Probabilistic Creep-Fatigue Crack Initiation Program for the HYA/HAR Boilers' Superheater Bifurcations; Theory, User Guide and Application to HAR Reactor 2", E/REP/BBAB/0023/AGR/12, Rev.000, December 2012.

[7] P.J.Holt, "Heysham I / Hartlepool Power Stations: BIFINIT: A Monte Carlo Creep-Fatigue Crack Initiation Assessment of the HYA/HAR Boilers' Superheater Bifurcation Inlets with Application to Hartlepool Reactor 2", E/REP/BBAB/0025/HAR/12, Rev.000, December 2012.

[8] R.A.W.Bradford and P.J.Holt , "Application of Probabilistic Modelling to the Lifetime Management of Nuclear Boilers in the Creep Regime: Part 2", International Journal of Pressure Vessels and Piping, Vol.111-112 (2013) 232-245. .

[9] R.A.W.Bradford, "Hinkley Point B and Hunterston B: Probabilistic Creep-Fatigue Crack Initiation Assessment of the Superheater Bifurcations 0° Crotch at Full Power Including Allowance for a Carburised Surface", E/REP/BBAB/0027/AGR/13, Rev.000, August 2013.

[10] T.Siddall, "DNB Power Station: Deterministic Critical Defect Size and Probabilistic Structural Integrity Assessment of Fuel Plug Unit Partially Penetrated Welds", E/REP/BBJB/0181/DNB/12, Rev.001, May 2013.

[11] M.R.Schroeder, "Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise", W.H.Freeman & Co (1991). See also A.Chechkin, "A Short Introduction to the Theory of Lévy Flights", http://www.maths.qmul.ac.uk/~klages/bee_wshop/bbees_chechkin.pdf in *Understanding the Dynamics of Foraging Bumblebees*, a one-day workshop at Queen Mary University of London, 10 July 2009.

[12] RiskAMP: A Monte Carlo Add-In for Excel, Structured Data LLC, New York, http://www.riskamp.com/

[13] R.A.W. Bradford, "Analysis of Fatigue Cycles for Superheater Bifurcations and Tailpipe-Pintle Welds, Hinkley Point B and Hunterston B Power Stations", E/EAN/BBJB/0101/AGR/07, Rev.000, December 2007.

[14] @RISK, Risk analysis software for Monte Carlo simulation, Palisade Corporation, http://www.palisade.com/risk/

[15] R.A.W.Bradford and P.J.Holt, "Materials Data Inputs to Probabilistic R5V2/3 Creep-Fatigue Crack Initiation Assessments of HYA/HAR Superheater Bifurcations and Adjacent Tubing in 316H Parent and HAZ Materials", E/REP/BBAB/0022/AGR/12, Rev.000, November 2012.

[16] M P O'Donnell, "Use of forward creep deformation to estimate the cyclic stress relaxation behaviour in Type 316 stainless steel at temperatures from 550°C to 650°C", E/REP/ATEC/0027/AGR/01, Rev.000, October 2001.

[17] W.J.J. Vorster, "Verification and Technical Review of Calculations and Assessments Carried Out By, or on Behalf of, Structural Integrity Branch", E/PROC/ENG/BI/048, Rev.002, March 2013.

[18] M.J.Chevalier, "The Reliability of Degrading Structural Systems Operating at High Temperature", EngD thesis, University of Bristol, March 2013.

[19] M.Abramowitz and I.A.Stegun (editors), "Handbook of Mathematical Functions", Dover (1965).

[20] C.Austin, "Ks Values for CMV Steels", AMEC report AMEC12416/R003.

[21] M P O'Donnell, "Proposed Changes to R5 Volume 2/3 Appendix A4, Treatment of Weldments", British Energy Generation Ltd Report, E/REP/BDBB/0067/GEN/05, Rev.000 (2005).

[22] R.A.W.Bradford, "AGR Power Stations' Boiler and Reactor Internals: Ground Rules for Creep and Fatigue Component Life Assessment (CLA)", BEG/SPEC/ENG/CTS/209, Rev.001, July 2013.

[23] "Assessment Procedure for the High Temperature Response of Structures", R5 Issue 3, Revision 1, May 2012, specifically Volume 2/3, "Creep-Fatigue Crack Initiation Procedure for Defect-Free Structures".

19 TABLES

Table 1: Confidence Levels for the Normal Distribution

z	$P_{cum}\left(\langle x\rangle - z\sigma_x\right) = \widetilde{P}_{cum}\left(\langle x\rangle + z\sigma_x\right)$	Confidence Level*
0	0.5	50%
1	0.1587	84.13%
1.2816	0.1	90%
1.6449	0.05	95%
2	0.0227	97.73%
2.3263	0.01	99%
3	0.00135	99.865%
4	0.000032	99.9968%

*The confidence level is defined by the single-sided value of $1 - P_{cum}\left(\langle x\rangle - z\sigma_x\right) = 1 - \widetilde{P}_{cum}\left(\langle x\rangle + z\sigma_x\right)$. It differs from the confidence interval which is double-sided. For example, the 95% confidence interval would have $P_{cum}\left(\langle x\rangle - z\sigma_x\right) = \widetilde{P}_{cum}\left(\langle x\rangle + z\sigma_x\right) = 0.025$ and hence $z = 1.96$.

Table 2: Numbers of Trials (Bins) and Corresponding Number of Standard Deviations

Number of Trials (Bins)	Number of Standard Deviations (± this number about the mean)	Probability of being outside the simulated range per variable*
10	1.75	0.079
30	2.23	0.026
100	2.67	0.0077
300	3.02	0.0025
1000	3.37	7.6×10^{-4}
3,000	3.66	2.5×10^{-4}
10,000	3.96	7.5×10^{-5}
30,000	4.21	2.5×10^{-5}
100,000	4.48	7.5×10^{-6}
300,000	4.71	2.5×10^{-6}
1,000,000	4.95	7.5×10^{-7}
10,000,000	5.38	7.5×10^{-8}

*i.e., 1 – the confidence interval

20 FIGURES

Figure 1 **Possible Flow Diagram for the Probabilistic Procedure**
Details will vary according to the application. Inputs Decisions Outputs

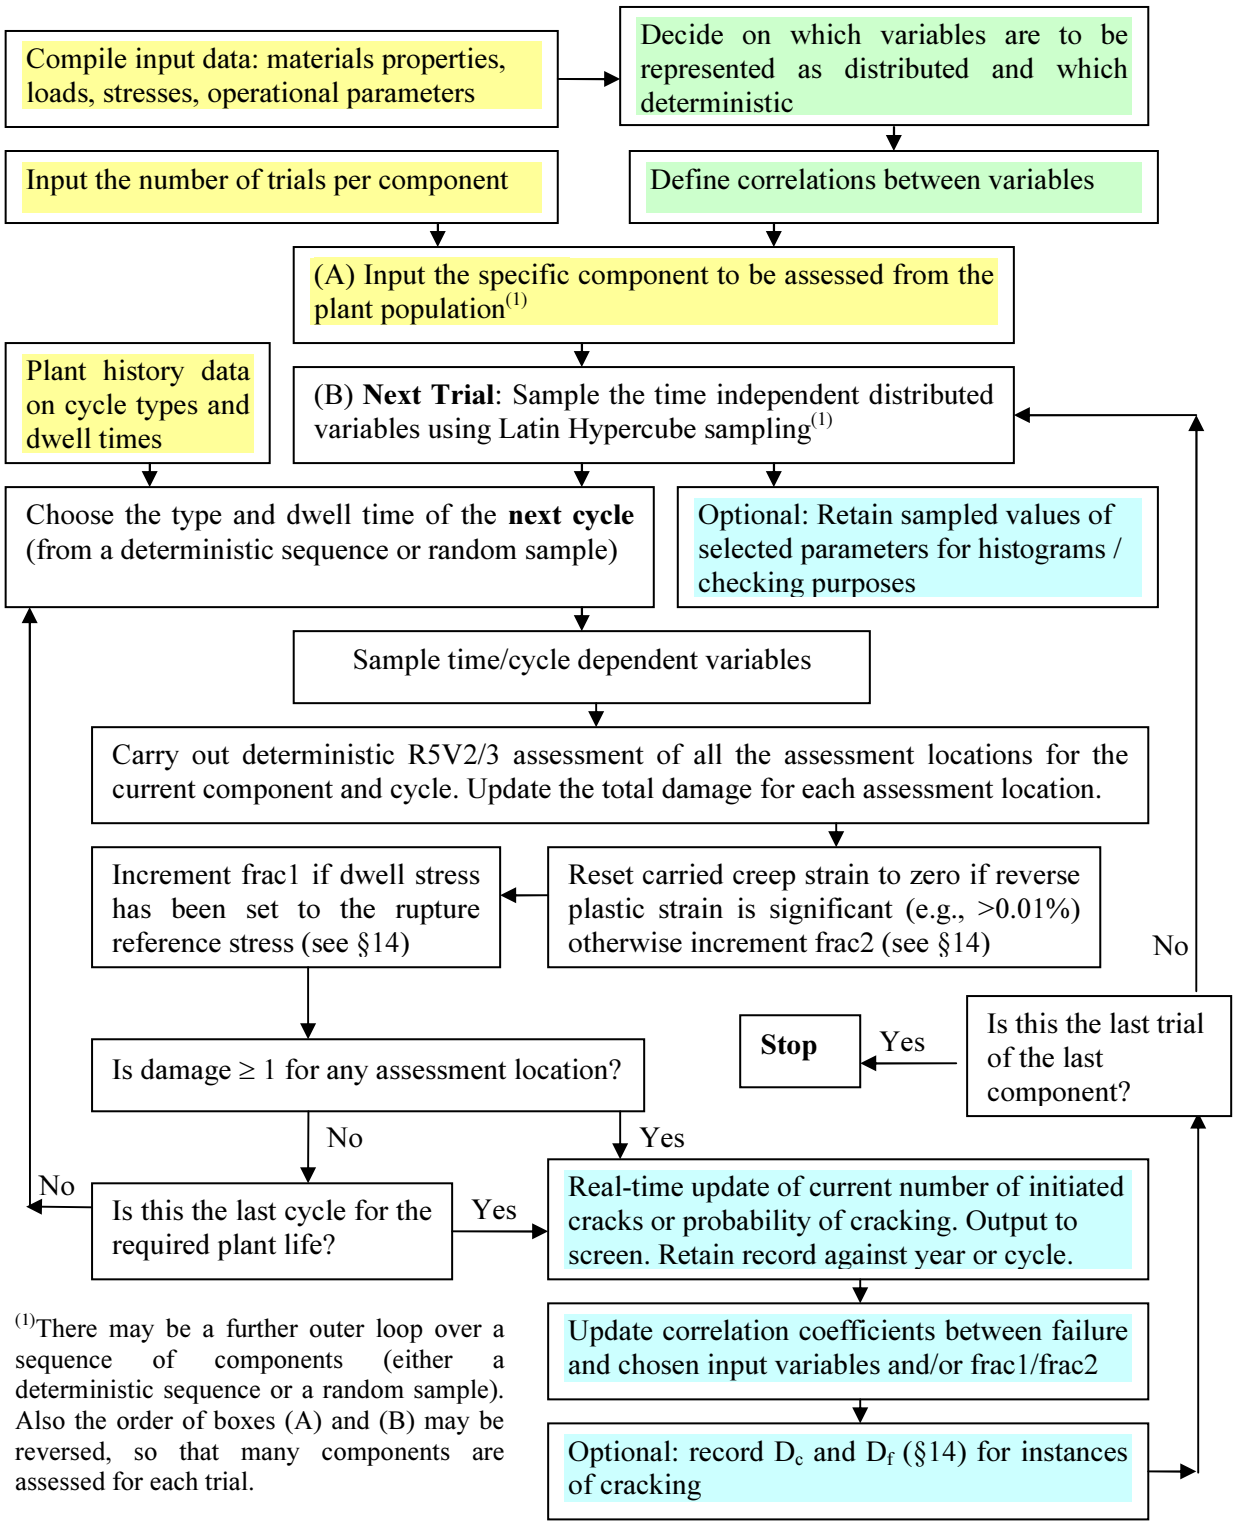

(1)There may be a further outer loop over a sequence of components (either a deterministic sequence or a random sample). Also the order of boxes (A) and (B) may be reversed, so that many components are assessed for each trial.

Figure 2a: Illustration of Normal and Lognormal Distributions for Coefficient of Variation 0.1

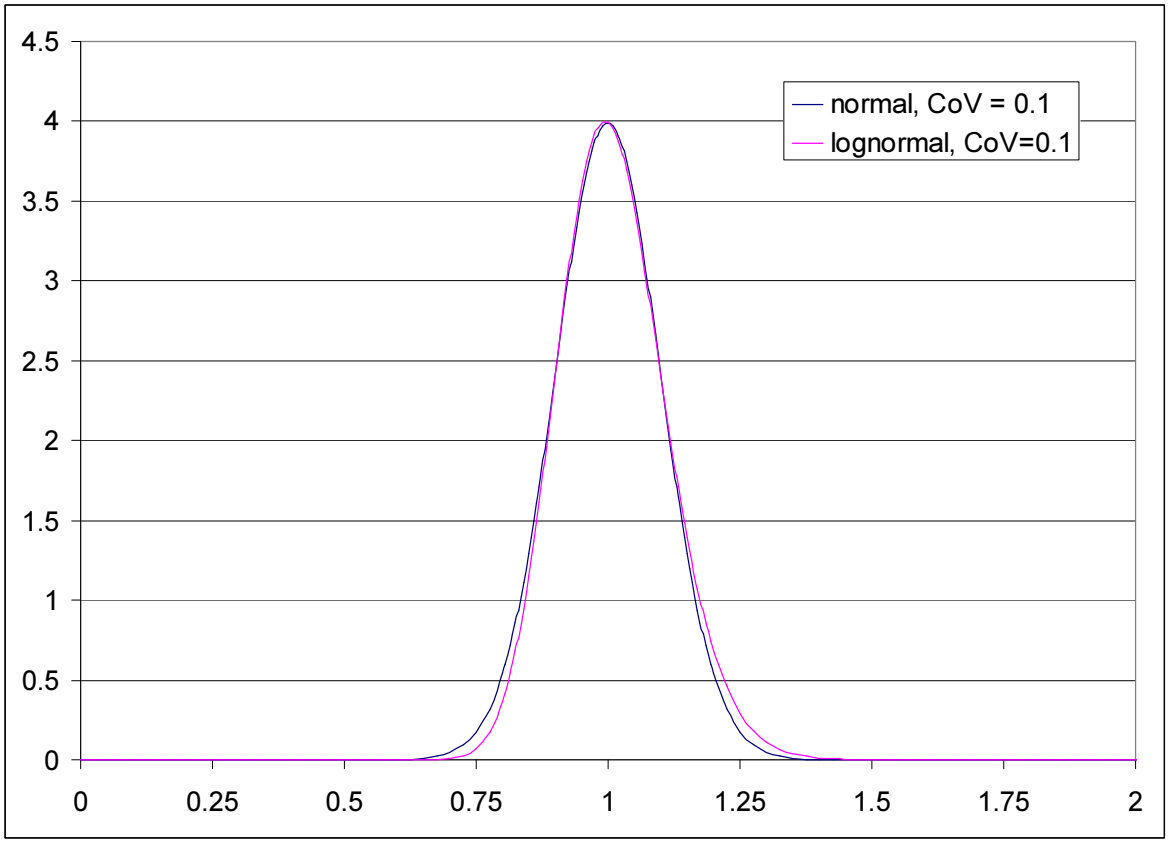

Figure 2b: Illustration of Normal and Lognormal Distributions for Coefficient of Variation 0.2

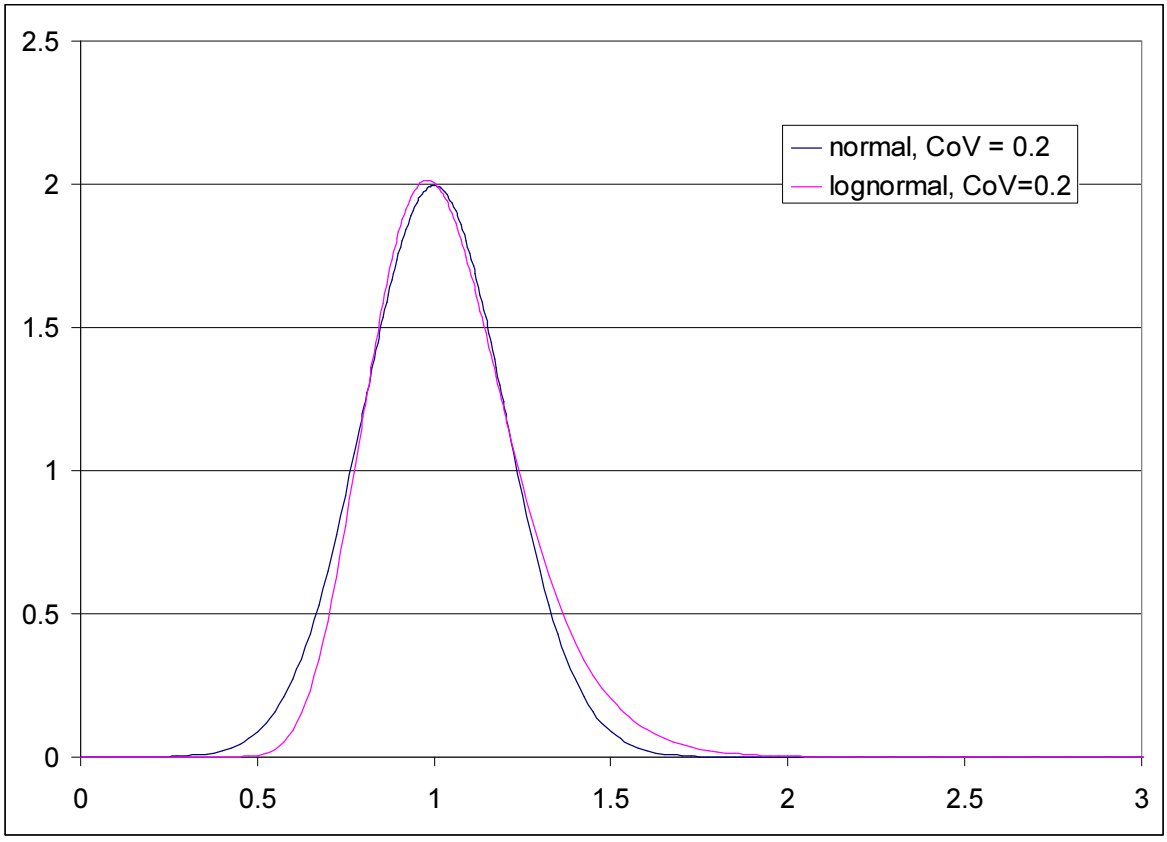

Figure 2c: Illustration of Normal and Lognormal Distributions for Coefficient of Variation 0.4

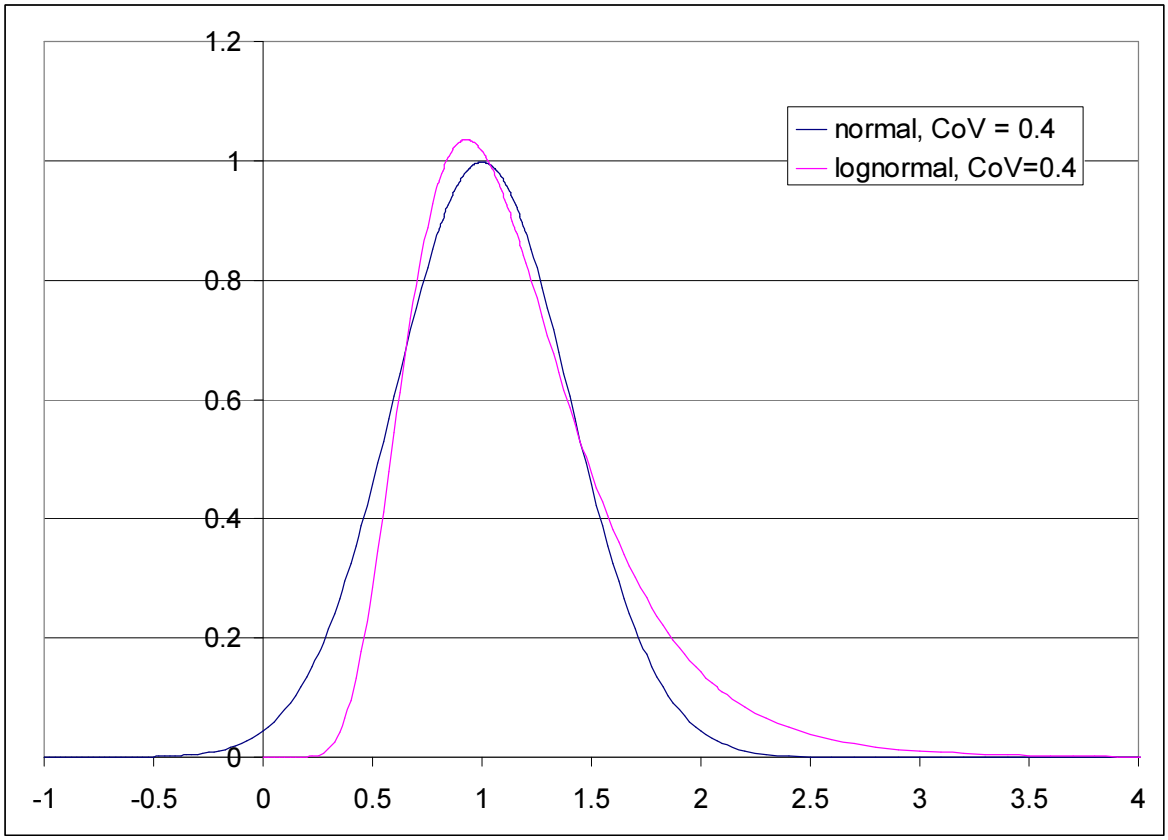

Figure 2d: Illustration of Normal and Lognormal Distributions for Coefficient of Variation 1.0

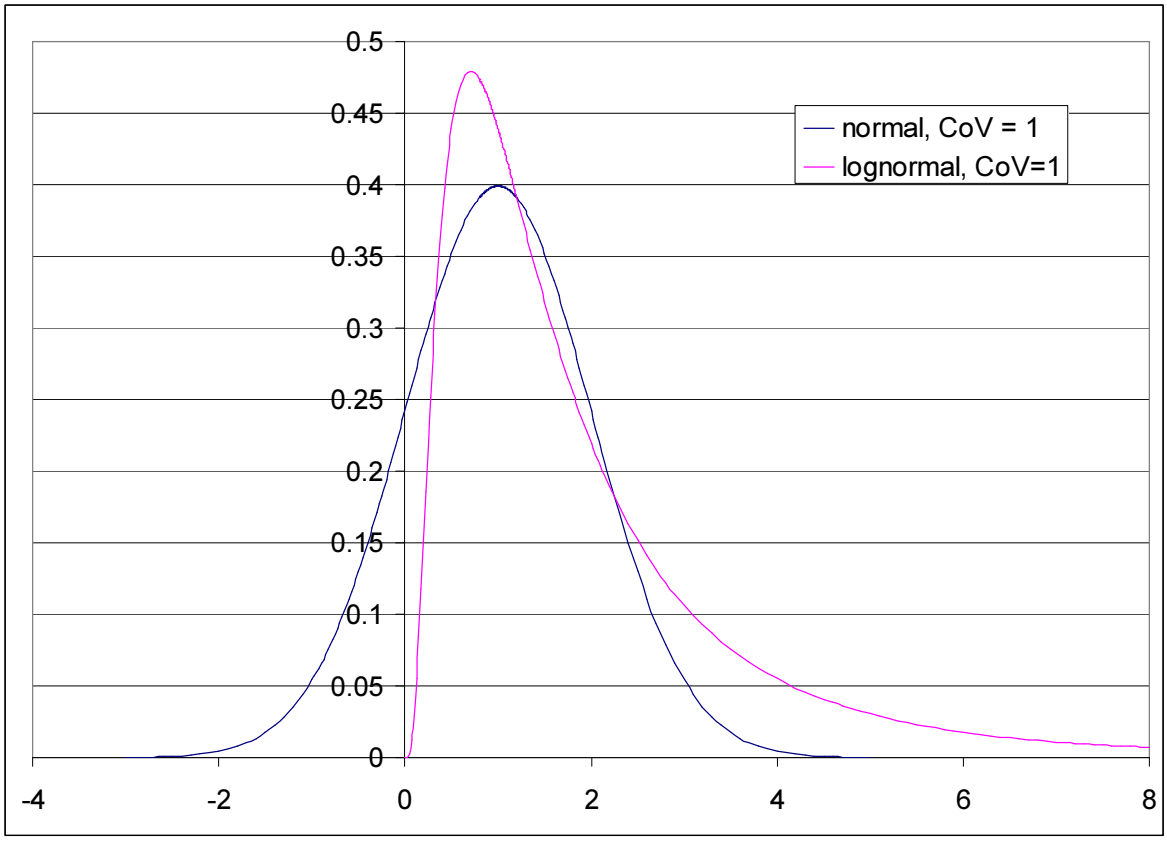

Figure 3: Illustrative Histogram of Input Data (Standby Temperatures)

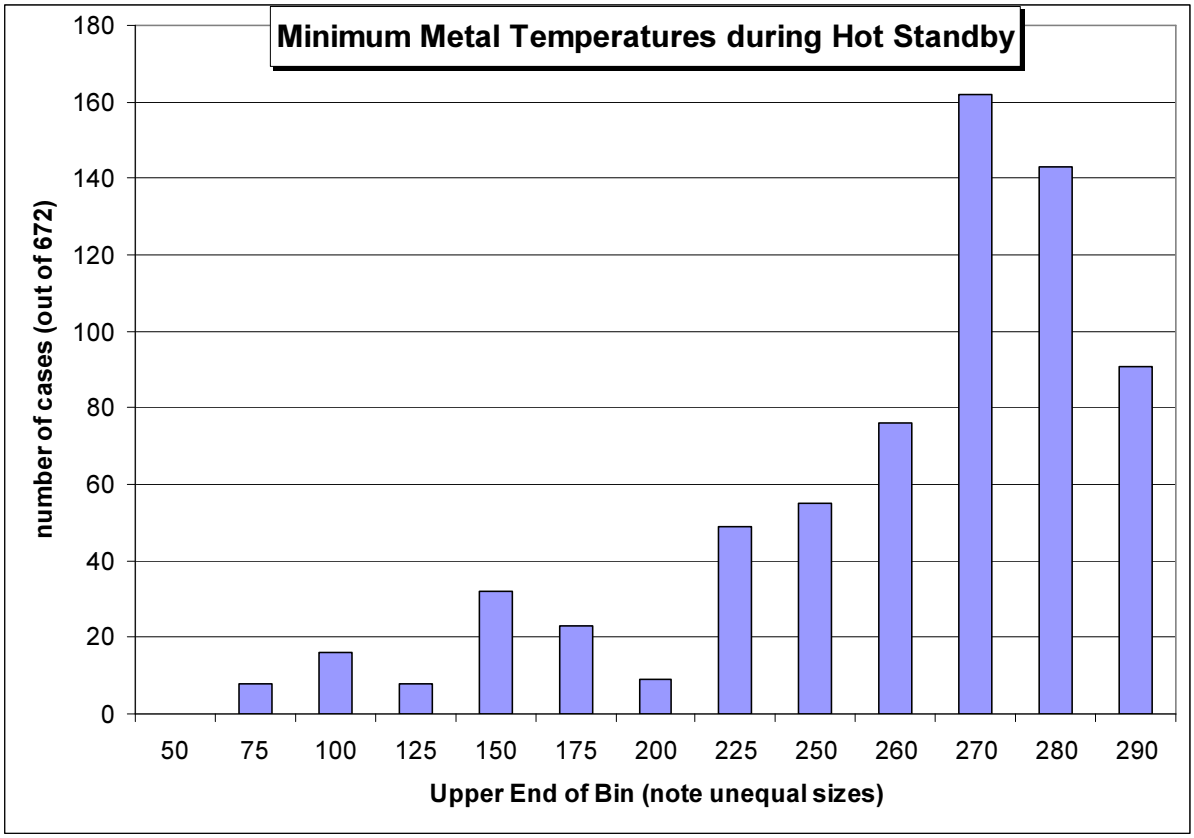

Figure 4: Illustrative Structure to be Assessed: Three Assessment Locations

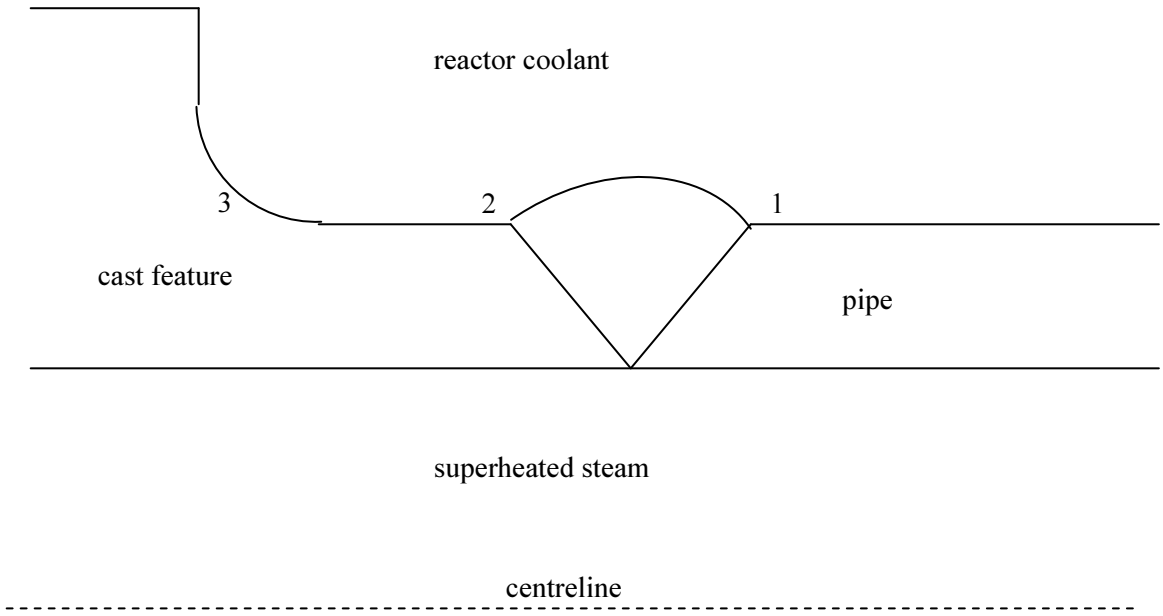

Figure 5: Illustration of a Latin Hypercube: $N_p = 2$ Variables; $N_B = 5$ Bins
- showing how a Latin hypercube may be generated by permutation of the bin numbers (row and column numbers in 2D).

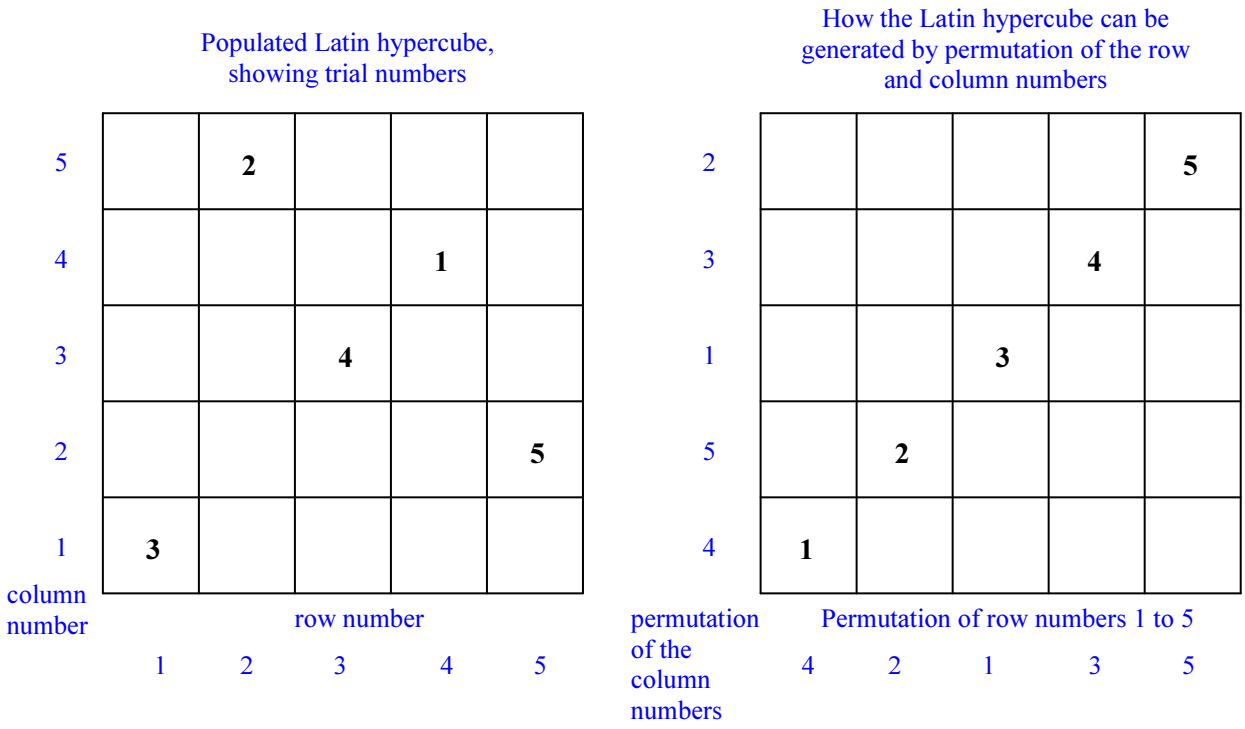

Figure 6: Illustration of Convergence of Initiation Rate

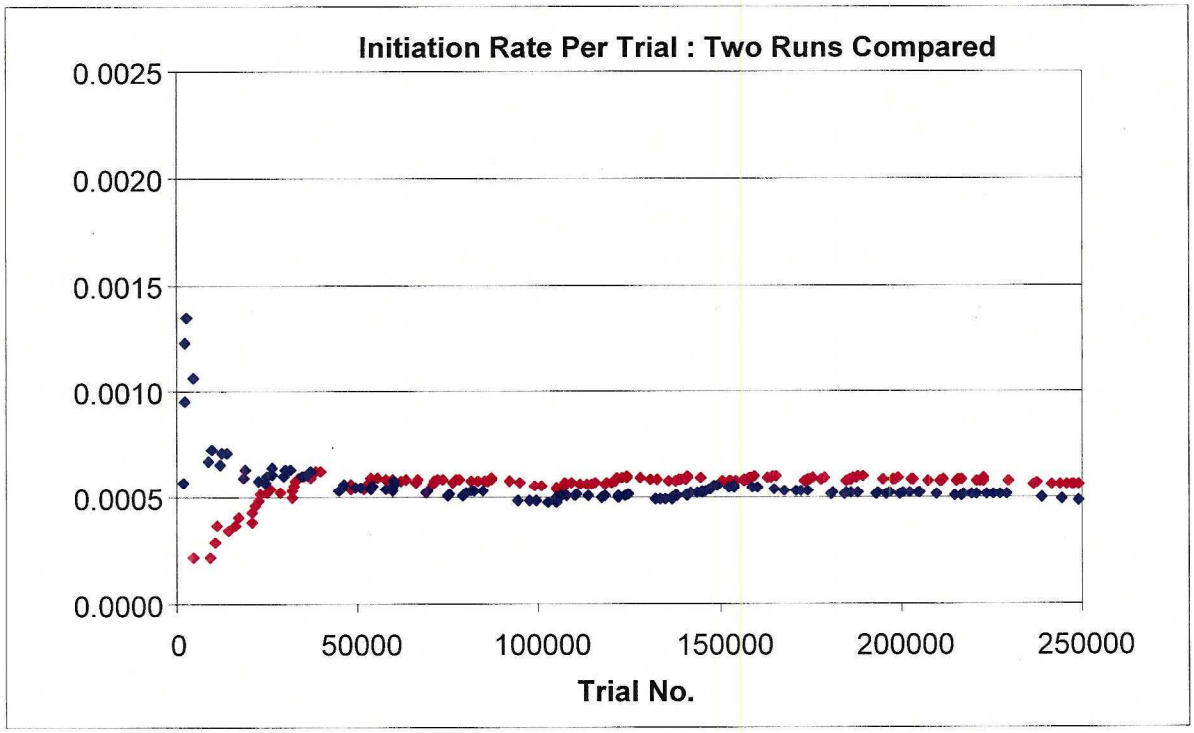

Appendix A - Implementing Correlation

Suppose you have two distributed variables, x and y. They are assumed to have been put in standard form, with zero mean and normalised to unit variance. How can correlation between x and y with a given correlation coefficient, C_{xy}, be implemented? The method for accomplishing this is stated in §A.1 and the proof given in §A.2. For Users writing their own code, the case of correlation between two variables is easily implemented.

More generally you may have a number of variables, $x, y, z, w...$, which are all mutually correlated. How is this multivariable correlation imposed in a Monte Carlo simulation? The method is given in §A.3. The imposition of multivariable correlation requires Cholesky decomposition of the correlation matrix. Coding Cholesky decomposition 'from scratch' is roughly as difficult as matrix inversion, so is feasible but not recommended. Users wishing to impose multivariable correlations may opt to use proprietary software (e.g., Refs.[12,14]) which have such facilities in-built. Alternatively, using the more mathematical software platforms, such as Matlab, which includes Cholesky decomposition facilities, may provide the best of both worlds in terms of flexibility and ease of use.

A.1 ALGORITHM FOR TWO VARIABLES

The contention is that correlation between two variables can be imposed by using a third variable, ξ, also with zero mean and unit variance, then sampling x and ξ independently (i.e., uncorrelated) and setting y to be,

$$y = C_{xy}x + \sqrt{1 - C_{xy}^2} \cdot \xi \tag{A.1}$$

Thus, if correlation were perfect ($C_{xy} = 1$) then (A.1) would reduce to $y = x$ as required, and perfect inverse correlation ($C_{xy} = -1$) would reduce to $y = -x$ as required. Conversely, if there were no correlation between x and y ($C_{xy} \to 0$) then (A.1) becomes $y = \xi$, i.e., a random variable completely independent of x, again as required. The proof that Equ.(A.1) imposes the desired correlation C_{xy} between x and y in the general case is given below.

A.2 PROOF OF TWO VARIABLE ALGORITHM

The pdf of x is written $P(x)$ and the pdf of ξ is denoted $\tilde{P}(\xi)$. We must first check that (A.1) is consistent with y having a mean of zero and a variance of unity. This is proved as follows…

$$\begin{aligned}\langle y \rangle &= \iint y P(x) dx \tilde{P}(\xi) d\xi = \iint \left(C_{xy}x + \sqrt{1 - C_{xy}^2} \cdot \xi \right) P(x) dx \tilde{P}(\xi) d\xi \\ &= \left[C_{xy} \int x P(x) dx + \sqrt{1 - C_{xy}^2} \int \xi \tilde{P}(\xi) d\xi \right] = 0\end{aligned} \tag{A.2}$$

$$\begin{aligned}\sigma_y^2 &= \iint \left(y - \langle y \rangle \right)^2 P(x) dx \tilde{P}(\xi) d\xi = \iint \left(C_{xy}x + \sqrt{1 - C_{xy}^2} \cdot \xi \right)^2 P(x) dx \tilde{P}(\xi) d\xi \\ &= \left[C_{xy}^2 \int x^2 P(x) dx + \left(1 - C_{xy}^2\right) \int \xi^2 \tilde{P}(\xi) d\xi + 2C_{xy} \sqrt{1 - C_{xy}^2} \int x \xi P(x) \tilde{P}(\xi) dx d\xi \right] \\ &= \left[C_{xy}^2 \sigma_x^2 + \left(1 - C_{xy}^2\right) \sigma_\xi^2 + 0 \right] = 1\end{aligned} \tag{A.3}$$

where we have used $\int xP(x)dx = \int \xi\widetilde{P}(\xi)d\xi = 0$ (i.e., zero means), and $\int P(x)dx = \int \widetilde{P}(\xi)d\xi = 1$ (total probability is unity), and $\int x^2 P(x)dx = \sigma_x^2 = 1$ and $\int \xi^2 P(\xi)d\xi = \sigma_\xi^2 = 1$ (unit variance), and $\int x\xi P(x)\widetilde{P}(\xi)dxd\xi = 0$ (i.e., x and ξ are uncorrelated). The correlation coefficient between x and y is,

$$\begin{aligned}\frac{\iint xyP(x)dx\widetilde{P}(\xi)d\xi}{\sigma_x\sigma_y} &= \iint x\left(C_{xy}x + \sqrt{1-C_{xy}^2}\cdot\xi\right)P(x)dx\widetilde{P}(\xi)d\xi \\ &= C_{xy}\iint x^2P(x)dx\widetilde{P}(\xi)d\xi = C_{xy}\end{aligned} \tag{A.4}$$

as required. Note that this proof applies for any probability density functions, $P(x)$ and $\widetilde{P}(\xi)$, provided only that the standard deviation exists, i.e., any "with-σ" pdfs (see §6). However the nature of the pdf's assumed for x and ξ will impose a particular pdf on y. If $P(x)$ and $\widetilde{P}(\xi)$ are both normal distributions then y will also be normally distributed. Consult specialist texts such as Ref.[1] for the implications of this method of imposing correlation with non-normal distributions.

A.3 ALGORITHM FOR MULTIPLE VARIABLES

If there are N variables which are all correlated, the pair-wise correlation coefficients can be arranged in a matrix called the "correlation matrix", (C). The element C_{ij} of this matrix is the correlation coefficient between the i[th] and the j[th] variables. Hence, the diagonal elements of (C) are all unity and the matrix is real and symmetric. A restricted form of the Cholesky decomposition theorem states that any Hermetian, positive-definite matrix, H, can be written in a unique way as the absolute matrix square of a lower triangular matrix with real, positive diagonal elements, i.e., $H = LL^+$ where $^+$ denotes the complex conjugate transpose in the general case. Real symmetric matrices are a special case of Hermetian positive-definite matrices, and hence any correlation matrix can be written in the form LL^T where ^{T} denotes the transpose and L has zeros above the main diagonal. A 3×3 example illustrates this,

$$(C) = \begin{pmatrix} 1 & 0.5 & 0.3 \\ 0.5 & 1 & 0.7 \\ 0.3 & 0.7 & 1 \end{pmatrix} = LL^T \quad \text{where,} \quad L = \begin{pmatrix} 1 & 0 & 0 \\ 0.5 & 0.866 & 0 \\ 0.3 & 0.635 & 0.712 \end{pmatrix} \tag{A.5}$$

Again we assume that all variables have been put in standard form, with zero means and unit variances. The variables which are to be correlated are $x, y, z, \ldots$. To impose this correlation, start with uncorrelated variables $\xi_1, \xi_2, \xi_3, \ldots\xi_N$, also in standard form. If L is the Cholesky "square root" of the correlation matrix, i.e., $(C) = LL^T$, then the desired correlated variables are obtained from the independently randomly sampled $\xi_1, \xi_2, \xi_3, \ldots\xi_N$, written as a column vector $\overline{\xi}$, as follows,

$$\begin{pmatrix} x \\ y \\ z \\ etc \end{pmatrix} = L\overline{\xi} \tag{A.6}$$

The two-variable method of §A.1 is just (A.6) in the case of a 2×2 correlation matrix because in that case we have,

$$(C)=\begin{pmatrix} 1 & C_{xy} \\ C_{xy} & 1 \end{pmatrix}=LL^T \qquad \text{where,} \quad L=\begin{pmatrix} 1 & 0 \\ C_{xy} & \sqrt{1-C_{xy}^2} \end{pmatrix} \tag{A.7}$$

so that (A.6) gives $x=\xi_1$ and $y=C_{xy}\xi_1+\sqrt{1-C_{xy}^2}\cdot\xi_2$, in agreement with (A.1). For the above 3×3 example we get,

$$\begin{aligned} x &= \xi_1 \\ y &= 0.5\xi_1 + 0.866\xi_2 \\ z &= 0.3\xi_1 + 0.635\xi_2 + 0.712\xi_3 \end{aligned} \tag{A.8}$$

The general case is now clear and very simple to implement providing that the Cholesky decomposition can be carried out. By construction it is clear that the correlation coefficients amongst the variables $x, y, z, ...$ will be as given by (C). Note that to use this method in practice the variables $x, y, z....$ resulting from (A.8) will be in standard form and must be converted to the desired physical variables by multiplying by the standard deviation and adding the mean. For this reason the method can only be used for "with-σ" pdfs. Another practical issue is that the User must specify a correlation matrix, (C), which is mathematically possible, a significant constraint on the possibilities for (C). Provided that (C) is real symmetric and positive definite, the above construction guarantees that it is a valid correlation matrix, because L exists. Consequently, the test of a real symmetric matrix being a valid correlation matrix is that it is also positive definite. This is equivalent to all its eigenvalues being positive. The User should therefore test that this is the case before proceeding with the analysis

Appendix B - Algorithm for Generating a Latin Hypercube

Let Nb be the number of bins and Nv the number of distributed variables. Hence, the hypercube is of dimension Nv and each of its Nv sides is divided into Nb bins. The array LHC(c,v) defines the Latin hypercube. The possible values taken by LHC are the bin numbers, [1, Nb]. The first index, c, in LHC(c,v) is a sequential identifier, from 1 to Nb, of the occupied cells (which represent the Nb trials). The second index, v, is the variable number, from 1 to Nv. Thus, occupied cell No.1 (i.e., trial No.1) puts variable 1 in bin LHC(1,1), variable 2 in bin LHC(1,2), the variable 3 in bin LHC(1,3), etc.

For every variable, v, the Nb numbers LHC(c,v), for c from 1 to Nb, is a permutation of the integers from 1 to Nb, i.e., they are all different. Hence, LHC is defined by setting each of its Nv columns to a permutation of [1,Nb]. Note that it is acceptable for different variables, v, to be assigned the same bin value in a given occupied cell. Different variables might even have the same permutation of bin values (i.e., the same bin for every occupied cell). This is an unlikely but acceptable occurrence. If all variables have the same permutation then the occupied cells of the Latin hypercube are the principal diagonal (regardless of the permutation), though this will not happen by chance for sensible numbers of bins.

It is assumed that there is available a routine “RandomPerm” which will provide a random permutation (array “perm”) of the integers 1 to Nb. The following Visual Basic code will generate a random Latin hypercube,

```
Dim LHC(Nb, Nv), perm(Nb)
For v = 1 To Nv
Call RandomPerm(perm, Nb)
For c = 1 To Nb
LHC(c, v) = perm(c)
Next c
Next v
```

“RandomPerm” can be coded in Visual Basic as follows,

```
Sub RandomPerm(perm, Nb)
Randomize
For i = 1 To Nb
perm(i) = i
Next i
For i = 1 To 3 * Nb
j = Int(1 + Nb * Rnd())
10 k = Int(1 + Nb * Rnd())
If k = j Then GoTo 10
keep = perm(j)
perm(j) = perm(k)
perm(k) = keep
Next i
End Sub
```

This coding for “RandomPerm” may not be terribly efficient and you may devise better. However, note that the efficiency of the coding for generating the Latin hypercube is quite unimportant since it is done only once.

21 DISTRIBUTION LIST

M Jee	AA00819
C D Hamm	AUTSD39
EDPJB01	EDPJB01
D W Dean	EDDWD01
M W Spindler	EDMWS01
S L Spindler	ED00748
M P O'Donnell	EDMXO01
Y Lei	ED00353
D J Tipping	EDDJT01
L Allport	PT00062
S Thozhur	ED01886
C J Aird	ED03171
M Chevalier	ED01589
R A W Bradford	EDRXB02
P M Cairns	EDPMC01
G Rayner	ED00739
T Siddall	ED01625
R H Priest	EDRHP01
C J Kirk	EDCIK01
D R McLennon	ED00359
R.Griffiths	AA02771
J.G.Johns	ED00153
M.Schmitz	ED01148
L.Worthy	ED01534
T.Cronin	BBEAY85
D.Hughes	HU03349
P.Rose	ED03273
J.Tao	ED03082
A.Kidd	ED01015
D.John	ED03531
G.Marshall	EDGWM01
D.Knowles	ED03101
C M Davies	Imperial College, catrin.davies@imperial.ac.uk
D J Smith	Bristol University, david.smith@bristol.ac.uk
R A Ainsworth	Manchester University, robert.ainsworth@manchester.ac.uk
M Martin	Rolls Royce, michael.martin@rolls-royce.com
N G Smith	AMEC, godfrey.smith@amec.com
S K Bate	AMEC, steve.bate@amec.com
W M Payten	ANSTO, wmp@ansto.gov.au
I Bretherton	Office for Nuclear Regulation, ian.bretherton@hse.gsi.gov.uk