Estimation of Power System Component Time-to-Failure Distributions with Limited Statistical Data

C.J. Dent, *Member, IEEE* and J.D. Gray

***Abstract*—We introduce to the power systems literature the statistical methods necessary for time-to-failure distribution estimation with a limited number of data points, and where few components have actually failed. These methods include both non-parametric estimation (i.e. no specific form of the distribution is assumed), and parametric maximum likelihood estimation (where a particular distribution family such as Normal or Weibull is assumed). Both of these are shown give results superior to methods previously suggested in the power systems literature. We explore the accuracy in estimation which might be achievable if the correct distribution form is known using simulated datasets. However, in realistic situations the only robust approach to extrapolating the time-to-failure distribution beyond the range of the data is to perform physical modelling of component ageing processes.**

***Index Terms*—Life estimation, Power system parameter estimation, Statistics**

I. Introduction

In many power systems worldwide, a substantial proportion of the network infrastructure was installed in the 1960s, and is thus approaching its design lifetime (typically 40 years) [1]. In order to plan replacement programmes, robust means of estimating the distribution of failure dates for classes of component are required.

This is complicated by the possibility that few or no components in a class have reached the typical failure age for that class, or indeed that few components may have failed at all. Data points typically are therefore of two types: failure times, and times at which observation ceased for a component with it still operating (the latter is known as censored data.) Standard elementary statistical techniques consider uncensored data, where actual failure times are available for all members of a sample. Specialised versions of these techniques are available for analysis of censored data; if censored data is available for a component the fact that it was still operating carries information, if not as much information as an actual failure time. Such specialised statistical techniques for handling censored data have received comparatively little attention in the power systems literature, while being an active subject of research in the general reliability community [2], [3].

Within power systems, [4] described an approach to estimating time-to-failure distributions with censored (and limited) data, by fitting Normal or Weibull distributions; [5] proposed a slight variant of this approach using a different estimate for the probability that a component fails by a given time, and [6] suggested using a Generalised Exponential model instead of Normal or Weibull. Relevant distribution estimation techniques drawn from the reliability theory literature have also been applied in the very different contexts of restoration times following hurricane damage [7], and residual deviations after subtraction of seasonal trends in spot prices [8] (in both cases however the data is uncensored).

This paper therefore introduces to the power systems literature the statistical techniques necessary for estimating time-to-failure distributions from censored data; Section II introduces the necessary terminology, and Section III reviews the statistical concepts required for distribution estimation. It will be shown that (contrary to the discussion in [4]) the powerful approach of maximum likelihood estimation provides an effective means of analysing censored data, and that alternative estimators are available which are in principle superior to those in [4]. These methods are demonstrated using real power system datasets in Section IV, and are shown to be superior to those methods previously proposed in the power systems literature. The question of what may be said robustly with limited failure data is discussed in Section V; this section presents new results on the robustness of mean lifetime estimates when few components have actually failed, in which fit are performed to data sampled from a known distribution. In realistic situations where correct form of fitting distribution is not known a priori, the only robust approach to extrapolating the time-to-failure distribution tvbeyond the range of the data is to perform physical modelling of component ageing processes. Finally, conclusions are presented in Section VI.

CJD was funded for this work through the EPSRC Supergen V, UK Energy Infrastructure (AMPerES) grant in collaboration with UK electricity network operators working under Ofgem's Innovation Funding Incentive scheme – full details on http://www.supergen-amperes.org/, and by the EPSRC and industry funded Supergen Flexnet consortium. JDG was supported by a bursary from the UK EPSRC.

C.J. Dent is with the School of Engineering and Computing Sciences, Durham University, South Road, Durham DH1 3LE, UK (Email: chris.dent@durham.ac.uk).

J.D. Gray is with the School of Engineering, University of Edinburgh, Edinburgh EH9 3JL, U.K. (Email: chris.dent@ed.ac.uk, James?)

II. Life Distributions and Associated Data

A. Definitions

It will be assumed throughout that the lifetimes of components in a class are independent and identically distributed random variables (a very comprehensive set of definitions and classifications of life distributions may be found in Chapter 9 of [9].) The *cumulative distribution function* (CDF) of this common probability distribution is then defined as

$$F_T(t) = p(T \leq t), \tag{1}$$

where T is the failure time for the component in question. The associated *probability density function* (PDF) is

$$f_T(t) = F_T'(t), \tag{2}$$

so that the probability of T lying in a given interval is the integral of the PDF over that interval. Any distribution for which negative values cannot occur, i.e. $F(0) = 0$, falls into the general class of *life distributions*; this non-negativity condition is clearly necessary for modelling lifetimes.

In reliability analysis, it is often convenient to use the *survival distribution*

$$S_T(t) = p(T \geq t) = 1 - F_T(t) \tag{3}$$

instead of the CDF. The *hazard rate* (or failure rate, or 'force of mortality') $h(t)$ is then defined such that $h(t)\mathrm{d}t$ is the probability that the component fails in the next small interval $\mathrm{d}t$ given that it has survived to time t:

$$h(t) = \frac{p(T \in [t, t + \mathrm{d}t) | t \geq t)}{\mathrm{d}t} = \frac{f(t)}{S(t)}. \tag{4}$$

It thus gives the conditional probability of failure in the next small interval $\mathrm{d}t$ given that the component has already survived to age t. By integrating this expression, the relation

$$S(t) = \exp\left[-\int_0^t h(\tau)\mathrm{d}\tau\right] = \exp[-H(t)] \tag{5}$$

is obtained, where $H(t)$ is termed the *cumulative hazard function*.

B. Censored Data

1) Definition:

Lifetime data for a sample of components typically consists of an age for each component, plus a statement of whether this age is

- the age at which the component failed, or
- the age at which observation of the component ceased, with the component still functioning at that age.

In the latter case, the data is said to be *right-censored*[1]. While a full set of failure times from a sample makes statistical analysis straightforward, such uncensored data is usually only seen in lab experiments; it would almost be a contradiction in terms for a class of components in the field to be of interest, and for all of its members to have suffered end-of-life failures.

2) Importance of Independent Censoring:

The methods presented in this paper assume that the censoring processes are independent for different components. For the examples presented here, where censoring is due to a cutoff at the study date, this assumption is valid. However, if (for instance) the censoring is due to components being replaced for reasons other than failure, then this assumption of independent censoring might not be valid, and very careful treatment of the censoring would be required in the distribution estimation.

[1]Left-censoring, where it is known that some components failed before a given time, but the precise time is unknown, is also defined; however, left-censoring will not be relevant to the discussions in this paper.

C. Statistical Estimators

1) Definitions:

In many circumstances, information about probability distributions must be inferred from data. An *estimator* for a quantity associated with the distribution is a formula giving an estimate for the true value of that quantity based on the observed data.

As an example, a series of n coin tosses may be modelled by the Binomial distribution with n trials and success probability p. As estimator for p is

$$\hat{P} = N_h/n, \tag{6}$$

where N_h is the number of heads. As prior to an experiment N_h is unknown, this estimator may be regarded as a random variable. The expectation value of this estimator is then

$$\mu_{\hat{P}} = E[N_h]/n = p. \tag{7}$$

Estimators such as this one, where the expected value of the estimator is the true value, are called *unbiased*; lack of bias (at least in the limit of large sample sizes) is clearly a desirable property for an estimator.

It is also possible to calculate the variance of the estimator, which may be used to quantify the likely quality of the estimate made:

$$\sigma^2_{\hat{P}} = V[n_h]/n^2 = p(1-p)/n. \tag{8}$$

$\sigma_{\hat{p}}$ is then called the *standard error* of the estimator. Clearly in practical situations the true standard error remains unknown, but it may be estimated by inserting the estimate $\hat{p}$ into the formula:

$$\sigma^2_{\hat{P}} \simeq n_h(n - n_h)/n^3. \tag{9}$$

Clearly a small standard error is desirable. This estimator for binomial distribution probabilities will be revisited as a candidate estimator for survival probabilities in Section III-A3.

2) Maximum Likelihood Estimation:

Maximum likelihood estimation is a method of generating systematically estimators with good theoretical properties. The *likelihood function* is defined simply as the pdf (or probability mass function in the case of a discrete distribution), regarded as a function of the parameters instead of the data. As an example, for the Binomial distribution example above:

$$L(p; x) = \begin{pmatrix} n \\ n - x \end{pmatrix} p^x (1-p)^{n-x}. \tag{10}$$

Intuitively, it measures how well different values of p fit the data, and hence the value of p which maximises the likelihood given the data may be used as an estimator for p; this is known as the *Maximum Likelihood Estimator* (MLE). In practice, both to simplify the algebra, and to make numerical solutions better conditioned, it is usually easier to work with the log-likelihood function:

$$l(p; x) = \ln(L(p; x)). \tag{11}$$

It is straightforward to check that for a Binomial distribution, the MLE for the trial probability parameter is the same as the estimator in II-C1, which was derived by intuition as 'the obvious one'.

MLEs have a number of useful theoretical properties. In the limit as the sample size tends to infinity, the MLE is unbiased and Normally distributed, with the variance of the MLE given by minus the reciprocal of the second derivative of the log-likelihood (for the Binomial example, the variance thus derived is the same as that derived by less formal means above.) These results generalise to estimates for distributions with more than one parameter; the estimator is then multivariate-Normal distributed, with covariance matrix equal to minus the inverse of the matrix of second partial derivatives of the log-likelihood function. For an accessible and detailed statement of the properties of MLEs see Section 2.6.3 of [10]; a more formal treatment may be found in [11].

This property that the MLE is asymptotically Normally distributed may be used for large sample sizes to specify confidence intervals for the estimates found. For the small-sample examples presented here, this useful property does not apply, and hence the standard error gives a less specific quantification of the uncertainty in the estimates.

III. Estimating the Survival Distribution

Estimation methods for the survival distribution fall into two broad classes:

- *Non-parametric Estimation.* No assumption is made a priori as to the shape of the survival distribution.
- *Parametric estimation.* The survival distribution is assumed to lie in a family of distributions (such as Weibull, Exponential or Normal), each of whose members is defined by a set of parameters. The estimation process for the survival distribution then consists of finding values for those parameters for which the model survival function is (in some defined sense) the best fit to the observed data.

A. Non-parametric Estimation

Non-parametric estimation provides an estimate of the survival function $S(t)$ for each time t within the range of the data. For a readable explanation of how to apply these methods, see chapter 4 of [12], and for an approachable statement of the proofs, see Section 2.3 of [13]. A more rigorous set of proofs may be found in [14].

Suppose that events (failures or censorings) occur at n distinct times, that d_i is the number of failures at the i^{th} time, and that Y_i is the number of components which are at risk at time t_i (i.e. which have not failed or been censored before t_i).

1) Kaplan-Meier Estimator:

The most commonly used non-parametric estimator for the survival function at time t is by Kaplan and Meier [15]:

$$\hat{S}(t) = \begin{cases} 1 & t < t_1 \\ \prod_{i|t_i \le t} \left(1 - \frac{d_i}{Y_i}\right) & t_1 \le t \end{cases} \quad (12)$$

The estimated survival function is thus a step function, with downward steps occurring at each observed failure time. The logic of the proof is that

- In any interval containing a single event time t_i, d_i/Y_i provides an estimate of the probability of a component failing during the interval conditioned on it working at the start of the interval.
- $(1 - d_i/Y_i)$ thus provides an estimate for the probability that a component is working at the end of the interval, provided that it was working at the start.
- (12) follows from combining these conditional probabilities, and taking the limit of small intervals.

The variance of this estimator may be estimated from Greenwood's formula:

$$V[\hat{S}(t)] = \hat{S}(t)^2 \sum_{i|t_i \le t} \frac{d_i}{Y_i(Y_i - d_i)}. \quad (13)$$

2) Nelson-Altschuler Estimator:

An alternative estimator for the survival function is the Nelson-Altschuler estimator [16]. This directly estimates the cumulative hazard as

$$\hat{H}(t) = \begin{cases} 0 & t < t_1 \\ \sum_{i|t_i \le t} \frac{d_i}{Y_i} & t_1 \le t \end{cases}, \quad (14)$$

with estimated variance

$$V[\hat{H}(t)] = \sum_{i|t_i \le t} \frac{d_i}{Y_i^2}. \quad (15)$$

$S(t)$ may then be estimated using (5). The derivation may be found in Section 2.3.2 of [13]; using a similar philosophy to that for Kaplan-Meier, it assumes that for each interval the number of failures is a Poisson process, and estimates that rate of this process based on the observed data. [12] (page 86) states that Nelson-Altschuler has better performance with small sample sizes.

3) Failure Ratio:

[4] proposed a non-parametric estimation method for the survival function (which was subsequently used to estimate parameters in Normal and Weibull fits). Component ages were recorded rounded to the nearest year t. The estimator suggested for the cumulative distribution function of the time-to-failure distribution is then

$$\hat{F}_t = \sum_{t' \le t} \frac{d_t}{Y_t}. \quad (16)$$

This may immediately be rejected as a reasonable estimator for the cumulative distribution function, as it is precisely the Nelson-Altschuler estimator for the cumulative hazard function, and hence is not bounded above by 1. As an example, consider for example a system of two components which fail in years 5 and 10; after 10 years, the estimated value for the CDF would then be 1.5.

[5] presented an alternative estimator for use in the same framework (which will be referred to as the 'failure ratio' estimator). Defining Y'_t as the number of components operating in year t *plus* the number which had failed in previous years, and d'_t as the number of failures in years up to and including t, this estimator is:

$$\hat{S}_t = 1 - \frac{d'_t}{Y'_t}, \quad (17)$$

i.e. it is based on estimating directly the proportion of components surviving to time t. It is essentially the binomial probability estimator of Section II-C, and has variance

$$\sigma^2_{\hat{S}(t)} \simeq \frac{d'_t(Y'_t - d'_t)}{(Y'_t)^3} = (\hat{S}(t))^2 \frac{Y'_t - d'_t}{Y'_t d'_t}. \quad (18)$$

We have found no reference to this estimator in books on reliability such as those cited in this paper, despite it being perhaps the most intuitively obvious. One possible theoretical explanation may be seen by comparing the variance with that of Kaplan-Meier, at times for which the cumulative number of failures is much less than the number surviving (i.e. $d'_t \ll Y_t, Y'_t$). The variance of the failure ratio estimator in [5] is then approximately $(\hat{S})^2/d'_t$, whereas that for Kaplan-Meier is approximately $(\hat{S})^2 d'_i/Y_t^2$; by the assumption of a small proportion of failures, the latter is clearly smaller. Also, it is doubtful whether the estimate of the survivor function should decrease as age increases merely because of further censoring times (and the consequent increase in the proportion of failed components), as opposed to further failure times.

B. Parametric Estimation

1) Maximum Likelihood Estimation with Censored Data: [4] proposed an approach to fitting Normal or Weibull distributions to censored data, based on a least-squares fit to a non-parametric estimate for the CDF/survivor function. This will be described in Section III-B4, after more formally justified methods, based on maximum likelihood estimation, have been introduced.

Contrary to the claim in [4] that maximum likelihood estimation cannot incorporate censored data in the likelihood function, an MLE may be constructed as follows [10]. Let δ_i equal 1 if the time t_i for component i is a failure, and 0 if it is a censoring time. Suppose that the distribution for T, the time-to-failure for a single component, has a distribution parametrised by β. The likelihood function for a sample of n components is then

$$L(\beta;\mathbf{t}) = \prod_{i=1}^{n} \left[f_T(t_i;\beta)\right]^{\delta_i} \left[1 - S_T(t_i;\beta)\right]^{1-\delta_i}. \tag{19}$$

This is the product of the pdfs for the components which have failed, and the survival probabilities for those which have not. Any multiplicative constants, which depend on neither $\mathbf{t}$ or β, may be neglected, as they do not affect the estimation process.

2) Weibull Distribution:
The Weibull distribution has survival function

$$S_T(t) = \exp(-(t/\theta)^{\kappa}), \tag{20}$$

where θ and κ are the scale and shape parameters respectively. It is commonly used in survival analysis due to the form of its hazard function:

$$h(t) = \frac{\kappa}{\theta}\left(\frac{t}{\theta}\right)^{\kappa-1}, \tag{21}$$

which is respectively increasing / decreasing / constant in time for ($\kappa > 1$) / ($\kappa < 1$) / ($\kappa = 1$).

The maximum likelihood estimator is found using standard calculus by differentiating the log-likelihood with respect to θ and κ, and setting the resulting expressions equal to zero. Following some algebraic manipulation, an equation for the shape parameter estimator $\hat{\kappa}$ may be derived [10]:

$$\frac{\sum_i t_i^{\hat{\kappa}} \ln(t_i)}{\sum_i t_i^{\hat{\kappa}}} - \frac{1}{\hat{\kappa}} - \frac{\sum_i \delta_i \ln(t_i)}{\delta} = 0, \tag{22}$$

where $\delta = \sum_{i=1}^{n} \delta_i$, the number of failed components. The estimator for the scale parameter is then

$$\hat{\theta} = \left(\frac{\sum_i t_i^{\hat{\kappa}}}{\delta}\right)^{1/\hat{\kappa}}. \tag{23}$$

3) Normal Distribution:
The Normal distribution has pdf

$$f_T(t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right), \tag{24}$$

where the parameters μ and σ are respectively the mean and standard deviation (SD) of the distribution. There is no closed-form expression for the survival function; the log-likelihood must therefore be maximised numerically. For the results here, the nonlinear optimisation problem was specified in the AIMMS optimisation modelling environment, and passed to the CONOPT nonlinear programming solver for solution [17].

4) Least Squares Fits: [4] proposed a method for fitting a Normal or Weibull distribution to censored data by making a least-squares fit to a non-parametric estimate for the survival distribution. As discussed in Section III-A3, the specific non-parametric estimator proposed in [4] can result in probabilities greater than one being obtained, and therefore the alternative estimate in [5] (based on the proportion of failed components) is the preferred option from those in the power systems literature.

[4] suggested adding two extra years' data to the least-squares calculation, beyond the estimates for the survival distribution at the actual observed fallure times:

- A value of 0.999 for the survival distribution in the year before the first failure.
- The estimate for the survival distribution at the last censoring time is the same as that at the last failure time.

No formal justification for these additional data points is given in [4]. Moreover, these additional data points might be considered unnatural when compared with realistic shapes for the true underlying survival distribution; the first implies looking for a fitted parametric survival distibution with a sudden step from 1 at the first failure time, and the second implies looking for a survival function which is constant between the last failure and censoring times. For the example presented here we therefore perform the least squares fit to the non-parametric estimate using points at the observed failure times only.

IV. Estimated Survival Distributions: Results

A. Data Used

The methods described above will be demonstrated on two datasets previously used in the power system literature, to facilitate comparison with previously published approaches.

1) Canadian Reactor Data: Data for a sample of 100 500 kV reactors from the British Columbia Hydro network is given in [4]. These were commissioned between 1969 and 1996, and just 4 had reached an end-of-life state by the study (i.e. censoring) year 2000. The 'failed' components were actually retired because field inspections deemed their condition to be

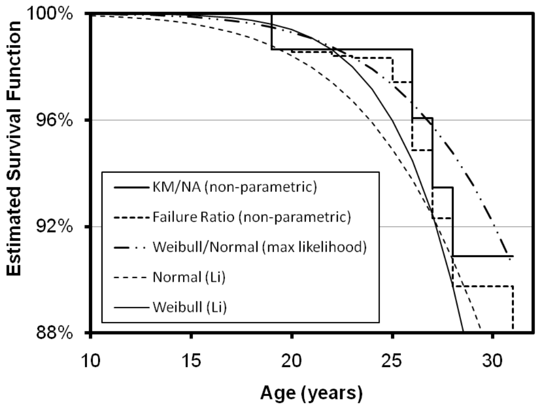

Fig. 1. Parametric and non-parametric estimates for the survival function based on the data in [4].

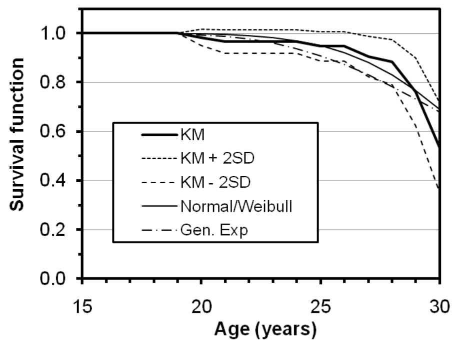

Fig. 2. Parametric and non-parametric estimates for the survival function based on the data in [6].

sufficiently poor; Depending on how these inspections were carried out, the assumption of independent identical time-to-failure distributions might or might not be valid; the methods given here are applicable as long as the criteria for retirement were consistent across all inspections.

2) Mexican Transformer Data: Data for a sample of 100 13.8 kV distribution transformers from a Mexican utility is given in [6]. These were commissioned between 1973 and 2006, and 16 had failed by the study year (2008). No detail is given as to whether the failures were true ends-of-life, or whether they were replaced following inspections while still operating.

B. Results: Reactor Data

Results for the data in [4] using the various parametric and non-parametric estimators described above are shown in Fig. 1. The Kaplan-Meier and Nelson-Altschuler non-parametric estimators give almost identical results, and hence are not plotted separately. The survival disribution estimates based on the failure ratio are rather more pessimistic. The maximum-likelihood Normal and Weibull estimates are not plotted separately as once more the curves are almost indistinguishable on this scale. Normal and Weibull fits using the least-squares method in Section III-B4 are also shown.

The KM and NA estimators, the non-parametric failure ratio, and the two maximum likelihood estimates all give consistent-looking results. Of the non-parametric estimators, the KM and NA are to be preferred to the failure ratio due to their stronger formal justification, see Section III-A3.

The Normal and Weibull least-squares fits based on [4] and [5] are somewhat more pessimistic; using the non-parametric estimators, which are inherently robust within their scope, as a baseline, this version of the least-squares estimators appears to give poor results. The maximum likelihood estimators both give better agreement with the non-parametric estimators, and moreover include in a formally justified and systematic way the available information from both failure and censoring times.

C. Results: Transformer Data

1) Survival Function:

Results for the transformer data in [6], using the various parametric and non-parametric estimators described above, are shown in Fig. 2. The Kaplan-Meier and Nelson-Altschuler non-parametric estimators again give almost identical results, and hence are not plotted separately. Similarly, the maximum-likelihood Normal and Weibull estimates are not plotted separately as once more they are almost indistinguishable on this scale.

The Normal and Weibull parametric maximum-likelihood estimators again give results which are consistent with the non-parametric Kaplan-Meier and Nelson-Altschuler results. In this case, a 2-standard deviation error bar for the K-M estimator is also displayed. For this small sample size the asymptotically Normal distributed property of the MLE does not apply (this is demonstrated by the 2 SD error bar including probabilities above 1), but neverless the SD may be used to give an order-of-magnitude estimate for the estimation error.

The generalised exponential distribution fit from [6], based on the least squares approach in [4], is also displayed. Once more, this is less consistent with the inherently robust non-parametric estimators than the Normal and Weibull MLEs; this, combined with its lack of sound formal justification, demonstrates that the MLEs should be preferred.

2) Cumulative Hazard Function:

The Nelson-Altschuler estimate for the cumulative hazard function is plotted in Fig. 3. While it contains the same information as the survival function, this plot reveals transparently important information on the classification of failures suffered.

The cumulative hazard rate is increasing superlinearly with age over the whole range of ages observed. This implies that its derivative, the hazard function, is increasing with age, and hence that the failures observed are age-related. If the transformers were suffering random failures, caused by

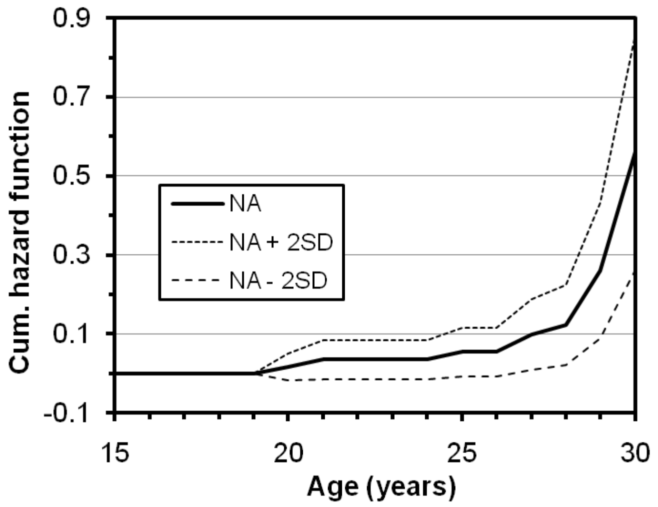

Fig. 3. Nelson-Altschuler estimate for the cumulative hazard function using the Mexican transformer data, and 2 standard deviation error bar.

external effects such as extreme short-term overloading, then the hazard rate would be constant with time.

V. Estimation with Limited Failure Data

A. *Parametric Versus Non-Parametric Estimation*

As mentioned previously, non-parametric estimators are inherently robust within their scope; they essentially derive an empirical observed hazard rate for the sample considered, and do not make any a priori assumptions about the shape of the distribution.

However, as observed above, the common non-parametric estimators produce step functions for the survivor and cumulative hazard functions, whereas the actual underlying functions are typically smooth. Provided the parametrised family of distributions used is sufficiently flexible in shape, parametric estimation can provide a degree of smoothing which produces a more realistic result. The price paid, however, is that if a poor choice of parametric distribution is made, then the resulting distribution estimate will not be realistic at all.

B. *Non-robustness of Extrapolation*

1) Reasons for Non-Robustness:

A pure statistical data-based approach, such as those discussed in this paper, cannot in general robustly say anything about the survival distribution outside the range of component ages used in the estimation process; the only exception to this is when there is good reason to assume that the same member of a particular parametric family of distributions applies across a broader range of ages. There might well be no reason to suppose that the same trend in the hazard rate extends across a wider range of ages, nor even that the same failure mechanisms remain relevant (in particular, there might be a transition from random failures being dominant, to aging-related failures becoming more significant.)

The motivation behind [4]–[6] was to estimate mean times-to-failure based on samples where a small proportion of components have failed; this implicitly means that most of the data is from the low-lifetime tail of the time-to-failure distribution, and hence is not directly relevant to lifetimes around the mean.

The data in [4], where just 4 components have failed, is simply too limited for a meaningful parametric estimate with two fitting parameters. In both cases, the above arguments about the dangers of extrapolating parametric fits outside the range of the data apply, and hence the estimates for mean lifetimes cannot be regarded as being in any way robust.

2) Inaccuracy of Results: Simulation Results:

To illustrate these points regarding the difficulties of statistical fits with data in only the extreme early tail, 1000 sets of 100 component lifetimes for the Canadian reactors have been simulated from a Normal distribution with mean 43.1 years and standard deviation 9.3 years, and from a Normal distribution with mean 32.2 years and standard deviation 4.4 years for the Mexican transformers (these are the Normal means and SDs obtained using the MLEs described earlier.) The commissioning years for each component, and the study (censoring) years, are the same as in the original papers. Fig. 4 shows a scatter plot of means of the Normal and Weibull distributions, fitted using maximum likelihood estimation to the 1000 sets of simulated data for each of the sets of components.

3) Reactor Data:

The estimates for the mean lifetime obtained by fitting a Normal distribution typically range from 32 to 50 years, with some outliers up to 60 years; the mean of these estimates is 38.4 years and the SD 3.8 years. Fitting Weibull distributions to the same Normal-generated data, however, gives a much greater spread; the mean estimate is 46.7 years, the SD is 23.0, and there are outliers up to 350 years. The mean number of failures across the 1000 samples is 6.8, which is greater than the 4 in the original data.

It must be emphasised that the Normal fits represent an idealised situation in which the form of distribution from which the data is sampled is assumed known. The Weibull fits demonstrate just how non-robust extrapolation outside the range of the data can be if the correct form of the distribution is not known, even where no new failure mechanisms become relevant at typical lifetimes. The inevitable conclusion is that estimating mean lifetimes on the basis of such limited data cannot be a robust practice.

4) Transformer Data:

The estimates for the mean lifetime obtained by fitting a Normal distribution in this case are much more tightly clustered, having mean 32.2 years and SD 1.1 years. The Weibull-estimated mean lifetimes show similar performance (mean 31.8 years and SD 1.0 years). The mean number of failures across the 1000 samples is 17.6, which is slightly greater than the 16 in the original data.

This demonstrated that if one can sensibly assume that the trend in hazard rate seen in the Mexican transformer data remains valid up to the mean lifetime, then quite a robust estimate could be obtained from the original transformer dataset. However, it is difficult to envisage circumstances under which this assumption can be deemed reasonable a priori, and hence the range of estimates seen in the scatter plot

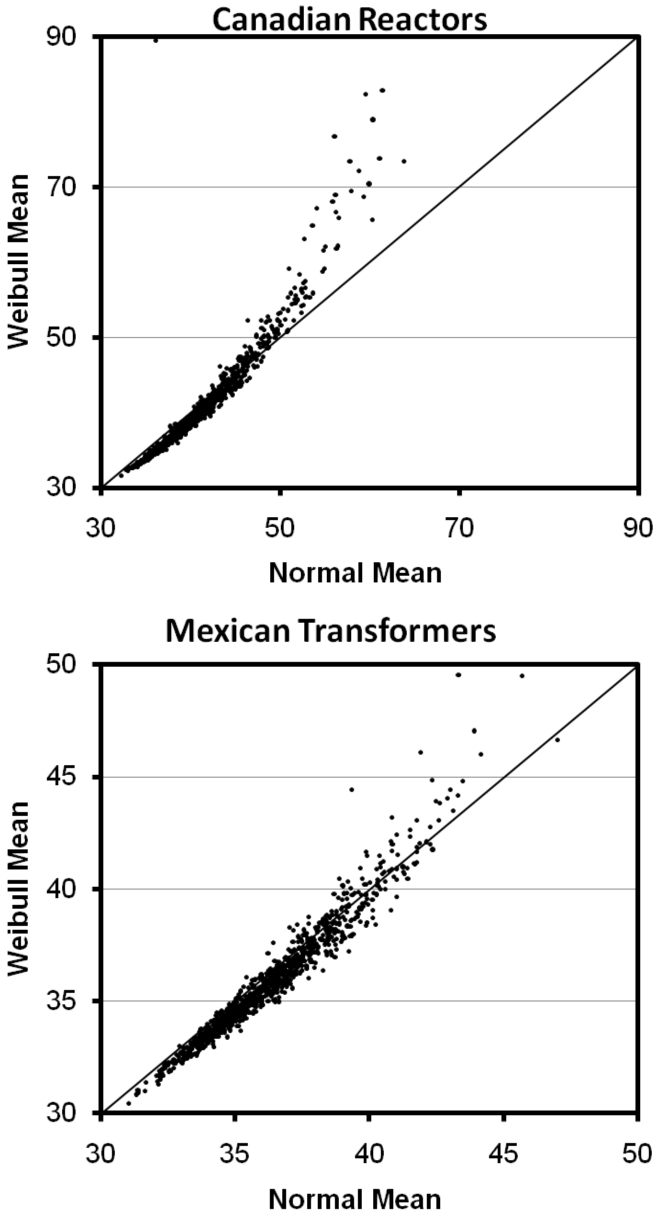

Fig. 4. Scatter plot of means from Normal or Weibull fits to simulated reactor and transformer data.

gives a highly optimistic view of the method's robustness.

C. Physical Modelling

Where the data does not exist to make robust projections of typical failure times by pure statistical methods (i.e. based on data alone), then such projections might still be made using physical modelling of aging processes.

An example of this approach is given in [18]. Data is presented on failures of transmission transformers in Great Britain. Here, a very small proportion of the transformers have actually failed, and moreover a plot of the empirical hazard rate (which is approximately constant with time) reveals that these are essentially random failures, and are thus irrelevant to the risk of ageing failures which motivates major asset renewal programmes.

[18] therefore proposes using data on the condition of insulation in transformers which have been scrapped for other reasons, to project when those transformers' insulation would have reached a state which made replacement necessary. If the number of scrapped transformers is sufficiently large, and they are considered representative of the entire population, then this provides a means of estimating the distribution of times for the entire population at which insulation breakdown makes replacement necessary.

VI. Conclusions

We have introduced to the power systems literature the statistical methods necessary for time-to-failure distribution estimation with a limited number of data points, and where few components have actually failed. These methods include both non-parametric estimation, and parametric maximum likelihood estimation; both of these are shown to give superior results to methods previously suggested in the power systems literature.

The accuracy in estimation which might be achievable has been explored using simulated datasets. This suggests that if the correct form of the time-to-failure distribution is known, then estimating mean times-to-failure on the basis of a limited number of early failures might give acceptably low error. However, this is not a situation which is realised in practice. In reality, there is usually no reason to suppose that the form of the survival distribution at typical lifetimes is the same as that in the extreme early life, and hence the only robust means of extrapolating beyond the range of the data is physical modelling of the ageing processes.

Acknowledgements

The authors acknowledge valuable discussions with F. Coolen, Q. Zhong, Z. Wang, S. Blake, J.W. Bialek and their colleagues in the Supergen-AMPerES consortium.

References

[1] "Special Issue on 'Our Aging Power Systems: Infrastructure and Life Extension Issues'," IEEE Power and Energy Magazine, vol. 4, no. 3, May/June 2006.

[2] T. A. Maturi, P. Coolen-Schrijner, and F. Coolen, "Early Termination of Experiments in Nonparametric Predictive Comparisons," in *International Workshop on Applied Probability*, 2008.

[3] F. Coolen, "On Probabilistic Safety Assessment in Case of Zero Failures," *Journal of Risk and Reliability*, vol. 220, pp. 105–114, 2006.

[4] W. Li, "Evaluating Mean Life of Power System Equipment With Limited End-of-Life Failure Data," *IEEE Trans. Power Syst.*, vol. 19, 2004.

[5] N. Enkhmunkh, G. W. Kim, K.-J. Hwang, and S.-H. Hyun, "A parameter estimation of weibull distribution for reliability assessment with limited failure data," in *International Forum on Strategic Technology*, 2007, pp. 39–42.

[6] J. E. Cota-Felix, F. Rivas-Davalos, and S. Maximov, "A new method to evaluate mean life of power system equipment," in *Cired*, 2009, p. Paper 1011.

[7] H. Liu, R. A. Davidson, and T. V. Apanasovich, "Statistical Forecasting of Electric Power Restoration Times in Hurricanes and Ice Storms," *IEEE Trans. Power Syst.*, vol. 22, no. 4, 2007.

[8] F. Olsina and C. Weber, "Stochastic Simulation of Spot Power Prices by Spectral Representation," *IEEE Trans. Power Syst.*, vol. 24, no. 4, 2009.

[9] H. Pham, Ed., *Handbook of Reliability Engineering*. Springer, 2003.

[10] J. I. Ansell and M. J. Phillips, *Practical Methods for Reliability Data Analysis*. Oxford Science Publications/Clarendon Press, 1994.

[11] J. E. Freund, *Statistics*. Oxford Science Publications/Clarendon Press, 1994.
[12] J. P. Klein and M. L. Moeschberger, *Survival Analysis: Techniques for Censored and Truncated Data*. Springer, 1997.
[13] P. Hougaard, *Analysis of Multivariate Survival Data*. Springer, 2000.
[14] M. Crowder, *Classical Competing Risks*. Chapman & Hall/CRC, 2001.
[15] E. L. Kaplan and P. Meier, "Non-parametric estimation from incomplete observations," *J. Am. Stat. Assoc.*, vol. 53, 1958.
[16] W. Nelson, "Hazard plotting for incomplete failure data," *J. Qual. Tech.*, vol. 1.
[17] J. Bisschop and M. Roelofs, *AIMMS - The User's Guide*, Paragon Decision Technology, 2006.
[18] P. Jarman, Z. Wang, Q. Zhong, and T. Ishak, "End-of-Life Modelling for Power Transformers in Aged Power System Networks," in *6th Cigr´ Southern Africa Regional Conference*, 2009, p. Paper C105.