THE RELIABILITY AND ACCURACY OF REMNANT LIFE PREDICTIONS IN HIGH PRESSURE STEAM PLANT

Ian Chambers
Safety and Reliability Division, Mott MacDonald

Studies have been carried out to show how failure probabilities can be calculated from plant measurements, to set confidence limits on predictions of time-to-failure for components under varying load or operating conditions.

The studies concentrate on creep damage in steam pipework, but the method developed can be generalised to a wide range of situations, because it relies only on general principles of measurement, rather than the specific technique or application. It uses numerical simulations to show how the reliability of time-to-failure predictions depends on the characteristics of measurements taken on plant, such as frequency of testing, accuracy of measurement technique, and number of samples taken at each inspection.

Keywords: Failure, probability, pipework, simulation, prediction

INTRODUCTION

This paper investigates what can be done to optimise predictions of the safe life of ageing plant, where the predictions are based on measurements characterising the current state of the plant and their historical trends. Taking a simple physical situation as a test case, it demonstrates how some simple simulation work can shed useful light on how best to concentrate finite resources on inspection using an approach that can be readily generalised to other forms of plant ageing.

The original work dealt specifically with high pressure, high temperature steam pipework, subjected to regular inspection, whose failure would cause significant financial loss to its operators and might pose a hazard to personnel. On the other hand, replacement of ageing pipework also involved a significant cost penalty, both in terms of the cost of the new component and in plant downtime. Over the years, a variety of techniques and design codes have been developed to give conservative predictions of a component's remaining life, but the simulation work described here goes beyond these to improve estimates of the actual remnant life, and so get maximum use (and therefore financial return) from it.

Such a situation is familiar to operators of many different types of plant, as they seek to optimise their inspection and maintenance strategies, and it leads to a number of questions concerning what can be done to improve the accuracy of the predictions of how much of a component's safe life still remains. Can the predictions be improved by making inspections more frequently? Would taking more measurements at each inspection have the same effect? What benefit would an improved measurement technique bring (to set against its development and implementation cost)? To try to answer these questions, safe life predictions have been simulated using a simple mathematical model, in which the frequency of inspection, number of measurements per inspection, and a range of other important parameters could be varied, and the effectiveness of the predictions could be quantified and compared. In this context, assessing effectiveness takes account of the undesirable instances where a component is marked for replacement prematurely, so incurring a cost penalty which would be avoidable with a more effective safe life prediction. The model is based on an idealised physical model of creep ageing of steam pipework, and the simulation work first investigates the effectiveness of different inspection and measurement strategies under ideal conditions of well-defined material properties and material behaviour, before going on to simulate uncertainties in both these parameters.

THE PHYSICAL MODEL: CREEP DAMAGE IN STEAM PIPEWORK

Creep is the gradual stretching of a piece of metal under the influence of an external load, even though the stress the load creates in the metal is below the material's yield stress. It is a relatively slow process, but the rate of growth is generally more pronounced at high temperatures and varies over the life of a component, as well as having its obvious dependence on the magnitude of the applied load [1].

Creep's variation over a component's lifetime is characterised by three distinct phases. The first, known as primary creep, occurs at the start of a component's life, is usually of relatively short duration and poses little risk of component failure, as the fresh metal of a new component is very likely to have high integrity, despite the fact that the strain rate is relatively high in this phase.

Following this, a component usually settles down into the secondary creep phase, a long period when the strain rate is considerably slower and almost constant for a given set of operating conditions, representing the great majority of the component's useful life. During this phase, however, component strain will be accumulating more rapidly at stress-raising features, via a variety of mechanisms, as well as slowly reducing the pipe's thickness and strength.

Eventually the accumulation of creep damage increases to such an extent that the component's ability to withstand the loads applied to it is significantly compromised, and the creep rate increases again. In this tertiary creep phase, the risk of the component failing in service is usually so high that it is deemed to have reached the end of its useful life. Inspections therefore have the aim of detecting the onset of tertiary creep, characterised by an acceleration of the creep growth rate after a long stable period of steady growth.

THE MATHEMATICAL MODEL: SIMULATION OF CREEP AND CREEP RATE

INSPECTION INTERVAL AND MEASUREMENT CHARACTERISTICS

Eqn. 1 shows the equation used to simulate the under-lying creep rate, R_{creep}, varying as a function of time, t,

$$R_{creep}(t) = R_{const}\left[1 + 0.1e^{\alpha(t-\tau)}\right] \qquad (1)$$

where R_{const} = a constant, which defines the secondary creep rate,
α = a constant defining how quickly the creep rate increases when tertiary creep sets in,
and τ = the time at which tertiary creep sets in.

(The equation, developed for this work, is a simple mimic of the typical strain rate behaviour, and is not intended as a model of the actual physical processes involved)

The process of taking measurements on creep-affected pipework was simulated by defining the pipe's actual creep strain history, based on Eqn.1. Measurements of the creep strain were simulated by generating a set of random numbers at regular intervals along the strain history (representing regular inspections). The measurements were taken from a population normally-distributed about the true underlying strain rate, whose standard deviation defined the accuracy with which a set of such measurements could be made. This allowed the effects of different inspection intervals, different levels of measurement accuracy, and different numbers of measurements per inspection to be investigated. In practice it is often found that when taking plant measurements some inspection staff produce much more tightly grouped sets of measurements than do others. This was simulated by allowing the standard deviation of the measurements to vary from inspection to inspection.

As shown in Fig. 1, the model produces an almost constant creep rate early on (representing secondary creep), followed by a very rapid increase, due to the exponential term in the bracket expression representing tertiary creep. In practice operating conditions will vary during a plant's life, subjecting the steam plant to a range of different temperatures and pressures, and so imposing a fluctuation on the creep rate. This was represented in some of the simulations by allowing the mean secondary creep rate to vary irregularly over time, partially masking the onset of tertiary creep. Note that creep strain is essentially the fractional change in the component's size, and is therefore dimensionless.

Life estimates of creep-affected pipework, are often made by using the creep rate in the following equation for the total component life [2],

$$t_{comp} = \frac{\varepsilon_R - 0.01}{\lambda \varepsilon'} \qquad (2)$$

where t_{comp} is the component's life (in hours),
ε' is the measured creep strain rate,
and ε_R and λ are material constants (the uniaxial ductility and a correction factor for multiaxial effects)

The expression is derived from the assumption that the creep strain at failure is the creep strain rate multiplied by the component's life. The constant 0.01 allows for the short primary creep phase when creep is relatively rapid. In practice, up-to-date creep rate data are used in the expression, which indicates the onset of tertiary creep when successive predicted values of t_{comp} show a trend to reduce markedly.

Results of the simulation work with Eqn.1 were used to investigate the accuracy and reliability of Eqn.2 under different inspection strategies.

FAILURE CRITERIA: DETECTING TERTIARY CREEP AND DEFINING END-OF-LIFE

Calculating the Creep Rate

Two pieces of information are required to establish whether creep has started to accelerate into the tertiary phase: a good estimate of the secondary creep rate (assumed constant for a given set of operating conditions), and an indication of the current trend.

The long-term average creep rate was calculated using a least squares fit, including all the data available up to the time of measurement (Fig. 2). So after two inspections the creep rate was simply based on two data points (the means of the measurements taken at each inspection). After three inspections, the rate was based on a least squares fit to three data points, and so on. As more data points were included in the fit, the fluctuations due to measurement inaccuracies tended to be smoothed out. This fitted long-term average was taken to be the best estimate of the secondary creep rate.

Least squares fitting was also used to calculate the current trend in the creep rate, following each inspection. However, in this case the number of previous data points included in the fitting process affects the method's ability to detect and follow trends. Fits based on small quantities of data are more susceptible to error and fluctuation than fits using many data points. When fitting to a small number of points each data point makes a relatively large contribution to the fit, and so points where the error is large tend to contaminate the overall result. Fits based on larger numbers of points fluctuate less, but when tertiary creep sets in, its presence is masked by the influence of earlier data points, from the phase of secondary creep. Therefore a longer interval elapses between the onset of tertiary creep and its unambiguous appearance in the fitted creep rate calculations.

Defining End-of-Life

The artificial data were used to investigate how easily the onset of tertiary creep could be resolved from the general scatter of measurements. Scatter tends to mask trends in the data, and to introduce spurious fluctuations into the calculated creep rate, which can give the false appearance of tertiary creep. Studies were carried out to establish the effect of measurement accuracy, inspection interval, etc. on the detectability of tertiary creep.

THE MOST IMPORTANT MEASUREMENT CHARACTERISTICS

Frequent measurements ought to indicate when failure is imminent, and prevent in-service failures, but this has to be set against the cost and operational penalties of carrying out inspections, and the possibility that measurement errors will lead to the component being replaced prematurely. For an optimum balance to be struck, it is necessary to know how the probability of failure varies with inspection interval.

Factors which affect the measurement accuracy also affect the predictions of failure probability, as does the criterion chosen to indicate that failure is imminent. The importance of different measurement characteristics was investigated using the artificial data described previously. In each case, a series of twenty simulations were produced, to examine how detection of tertiary creep depends on each of the measurement characteristics.

A BASE CASE FOR COMPARISONS

First of all a set of artificial measurements were generated that were typical of the measurements taken in practice from plain pipework, using the characteristics shown in Table 1.

An example of a set of measurements based on this definition is shown in Fig. 1 and Fig. 2 shows the corresponding creep rates, along with the best estimate secondary creep rate and the criterion used to detect tertiary creep. In this case the measured creep rate was based on just the two most recent sets of measurements, so that increasing creep rate can be detected with minimum delay. Even so, tertiary creep was not detected until three inspection intervals after its onset. During this time the actual creep rate had increased nearly seven-fold, so there is a clear risk that the component would fail before tertiary creep could be detected.

Table 1 Typical Data Describing Creep of Plain Pipework, and Measurements Made During the Pipe's Life: *The inspection interval has been normalised to 1.0 to ease comparisons with other inspection intervals. The natural scatter of the data sometimes leads to negative measured creep rates, but these were excluded from creep rate calculations, as would probably be done in practice.*

Characteristic	**Value**
True secondary creep strain rate	4×10^{-4} per inspection interval
Inspection interval	Normalised to 1
Maximum allowed standard deviation of measurements	10^{-3}
No. of measurements per inspection	4
Censor out negative creep rates?	Yes
Criterion for onset of tertiary creep	Doubling of creep rate

There is also one spurious indication of tertiary creep in Fig. 2, where the criterion is satisfied early on, although tertiary creep has not actually begun. In practice, however, indications of failure so early in life can usually be recognised as spurious and discarded. Therefore, in this case the indication was ignored. In general, any spurious indications occurring in the first quarter of life were ignored on the same grounds.

Out of twenty simulations, all sharing the measurement characteristics in the table above, all but one resulted in detection of tertiary creep within three inspection intervals of its start. The probability of detection within three intervals is therefore 95%. However, thirteen simulations resulted in spurious detection, so the reliable detection of genuine tertiary creep is offset by 65% probability of premature retirement from service.

SENSITIVITY STUDIES

Inspection Interval

More frequent inspections might be expected to improve the reliability and promptness with which tertiary creep can be detected, but in fact detection of accelerating creep was found to occur at the same time as in the base case. Furthermore there was considerable scatter in the measurements which would probably have led to spurious indications of tertiary creep. The data would have to be subjected to careful interpretation to determine whether the increase in creep rate were genuine or just an artefact of the measurement inaccuracies.

On average, halving the inspection interval did lead to earlier detection of tertiary creep, but spurious detection becomes so frequent as to mask the true result. In practice the criterion for tertiary creep detection would have to be relaxed, further inspections would be necessary, and overall there would be little or no benefit from the extra effort expended. Relaxing the detection criterion from a factor of two increase in creep rate to a factor of four restored the frequency of spurious detection of tertiary creep to a value comparable to that found in the base case, but the time taken to detect tertiary creep deteriorated back to the base case value.

Measurement Accuracy

Accuracy was represented in the artificially-generated data by the measurements' standard deviation; small standard deviation corresponded to high accuracy. The measurements were allowed to fall on a normal distribution, whose mean was the true creep strain at that time, so the mean of a large number of measurements would tend towards the true strain value.

The base case, in Fig. 1, used a standard deviation on the creep strain measurements of up to$1x10^{-3}$, which is a fifth of the strain at the onset of tertiary creep. To see what effect improved accuracy would have, the allowed standard deviation of the measurements was reduced to $3x10^{-4}$. Improving the accuracy of the measurements decreased the scatter, and made the onset of tertiary creep easier to detect. In this case, the failure criterion could be tightened, allowing detection of tertiary creep two inspections after its onset, compared with three in the base case.

Number of Measurements at Each Inspection

The base case assumed that four measurements were taken at each inspection, their mean being the best estimate of the creep at that time. A set of twenty simulations were carried out assuming twelve measurements per inspection. This is an increase of a factor of three compared with the base case, and allows direct comparison with the results of the previous section, where the measurement accuracy was improved by a factor of three. As might be expected, tertiary creep was detected with similar promptness and reliability to the case with improved accuracy.

Overall, spurious detection was much more common in this case than the previous one, where the accuracy of the measurements was improved. Nonetheless, there is a clear benefit from increasing the number of measurements, compared with the base case.

THE EFFECTS OF PROCESSING THE MEASUREMENT DATA

USING MORE DATA POINTS TO IDENTIFY THE TREND

As well as the characteristics of the measurements, detecting trends in the data and determining end-of-life depend on how the data are processed. Investigating the effect of inspection interval showed spurious detection became much more likely if the interval were short. However, this was based on a very simple calculation of the creep rate, using just the most recent measurement and the one before it. Smoother trends result from using more data points to determine the recent behaviour, at the penalty of longer delays in spotting the onset of tertiary creep.

To see how the additional data point affects the reliability of the measurements, a set of twenty simulations were carried out. Compared with two-point fitting, the results showed a small improvement. They were also a little better than those produced by the base case, but the overall difference was not very great, and may well not justify the extra expense and effort involved.

QUANTIFYING THE UNCERTAINTY

Uncertainty in the Measure Creep rate

Since in practice it is not possible to determine the under-lying creep rate exactly, a least squares fit through all the measured data must be used as the best estimate of the true rate. Clearly, this changes over the life of the component, as more and more inspections are carried out, and an increasing number of data points go into the least squares fit.

Quantifying the uncertainty associated with this procedure involves quantifying the degree of scatter about the least squares fit. The simplest way of doing this is to calculate the difference between each creep rate measurement and the long-term best estimate. A reasonable approximation of the best estimate's accuracy after n measurements can then be found from the root mean square of all these individual differences, that is,

$$\delta_{RMS}(n) = \sqrt{\frac{\sum_{i=1}^{n}\left[\varepsilon'(i) - \varepsilon''(i)\right]^2}{n}} \qquad (3)$$

where $\delta_{RMS}(n)$ is the root mean square error after n measurements,
$\varepsilon'(i)$ is the best estimate of the long-term creep rate after i measurements,
and $\varepsilon''(i)$ is the best estimate of the instantaneous creep rate at the i^{th} measurement.

This expression can then be used as an estimate of the range of values for creep rate. After n inspections, the long-term creep rate is estimated to be $\varepsilon'(n) \pm \delta_{RMS}(n)$, which can serve as a means of calculating the probability of obtaining a creep rate within that range. In the examples given above, where the measurements were assumed to be normally-distributed about the mean, so it would be reasonable to assume the creep rate is normally-distributed about $\varepsilon'(n)$, with standard deviation $\delta_{RMS}(n)$.

Uncertainty in the Component Life Prediction

Measured creep rates are used in Eqn. 2 to predict component life. Having established the uncertainty in the creep rate, this can be fed into Eqn. 2 to estimate the uncertainty in the component life prediction, which can be shown as error bars on a graph of predicted component life, Fig. 3.

It is a relatively simple matter to apply the same approach to other uncertainties that arise in Eqn. 2, such as the range of possible values for the material constants. The end result is a clear visual indication of the reliability of all the life estimates that have been made (from the length of the error bars), and note how the error bars shrink and life predictions depart from the long term average as the end of the component's life approaches.

DISCUSSION

On the basis of the simulations presented here, the most effective means of improving predictions of in-service failure is to improve the accuracy of the technique used to characterise the damage state of the plant. In many cases this will be easier said than done, but the simulations give a basis for determining whether the costs of developing or procuring a more accurate system would be cost-beneficial, compared with the outage and inspection time involved with some of the other alternatives.

This message is reinforced by comparing two of the other studies: it was considerably more beneficial to collect a relatively large number of measurements at a few, widely-spaced inspections, than to take frequent inspections consisting of a small number of measurements each. Using this knowledge, the most cost-effective inspection strategies for a piece of plant can be drawn up.

An important point to bring out from the study is that the principles can be applied to a much wider range of situations than the simple case of creep in steam pipework. As long as a reasonable physical model of the underlying damage process is available, mathematical simulations of the measurement process can be carried out, to guide and inform the development of inspection strategies. The simulations were performed with readily available and inexpensive mathematical modelling software, and provide considerable flexibility, such as allowing for operator-to-operator variations, and comparing different means of fitting and analysing the data once it has been collected. The main difficulty in setting up the simulations is to ensure that all the important factors affecting accuracy have been taken into account. To some extent this can be eased during development by carrying out some scoping calculations with the model, to rank the factors according to their importance, which helps determine the appropriate level of detail in modelling them.

CONCLUSIONS

The main feature of the results of this study into the effectiveness of different measurement strategies is that a small number of good quality, reliable measurements may well prove more useful than a large group of poor ones. This was evident from the relatively small improvement obtained from reducing the inspection interval and the relatively high benefit from improving the accuracy of the measurement technique.

The study also makes evident the need for a pre-defined inspection strategy, based on a knowledge of how inspection data can be best used, which will allow operators to get the most from their inspection efforts.

REFERENCES

1. Iron and Steel Standards Policy Committee ISE/73
   Guidance on Methodology for the Assessment of Stress-Rupture Data Part 1: Procedure for the Derivation of Strength Values
   PD 6605 : 1997 (Final Draft)

2. Clews J A, Barraclough D R, Browne R J, Cane B, Hepworth J K, Phillips J, and Plastow B, 1985,
   Creeplife Assessment of Boiler Pressure Parts Using Strain Data
   CEGB

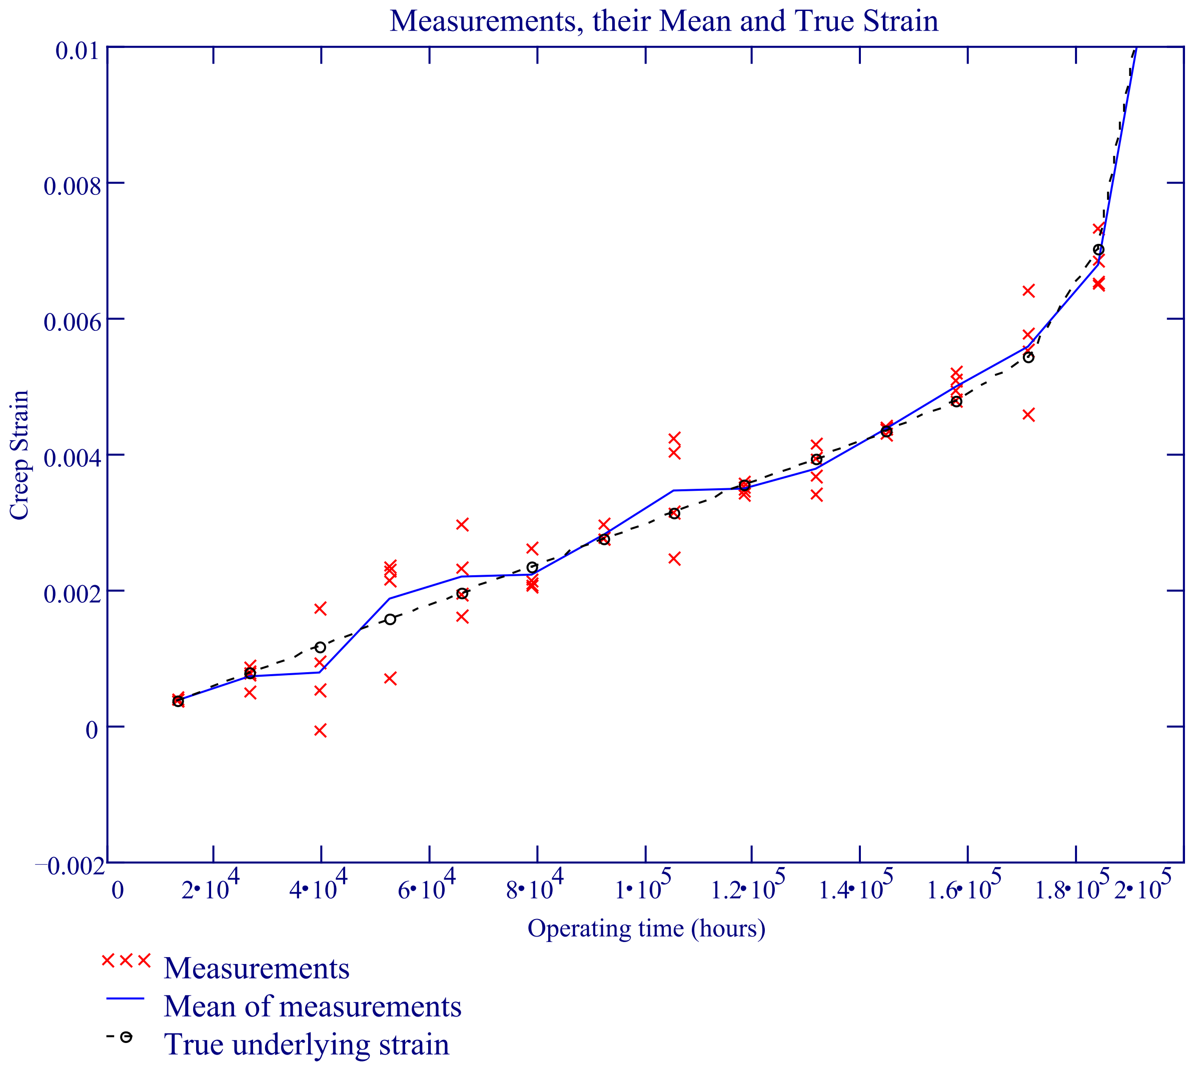

Fig. 1 Example of Artificial Measurements Generated from Base Case Data: *Initially the up-turn in the creep rate late in life is masked by scatter in the data. If strains of the order of 0.005 are sufficient to cause failure, there is significant probability of failure between the last two inspections.*

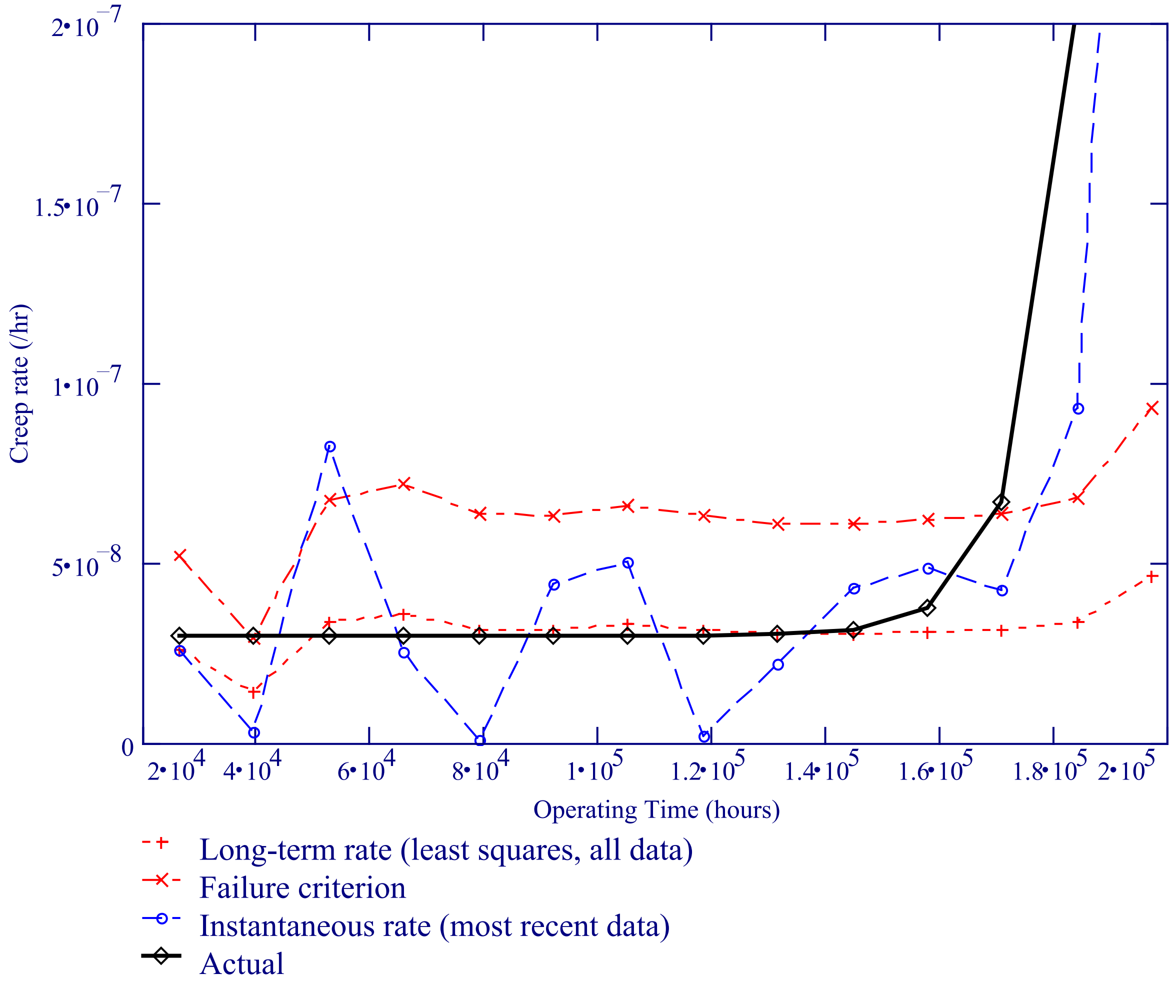

Fig. 2 The Actual and Measured Creep Rates from the Base Case: *The measured rate follows the actual rate with only a short time lag, although there is considerable fluctuation due to scatter in the data. The two near-horizontal lines represent the best estimate of the secondary creep rate (lower line) and the criterion chosen to remove the component from service, i.e. creep rate accelerates to double the secondary creep rate.*

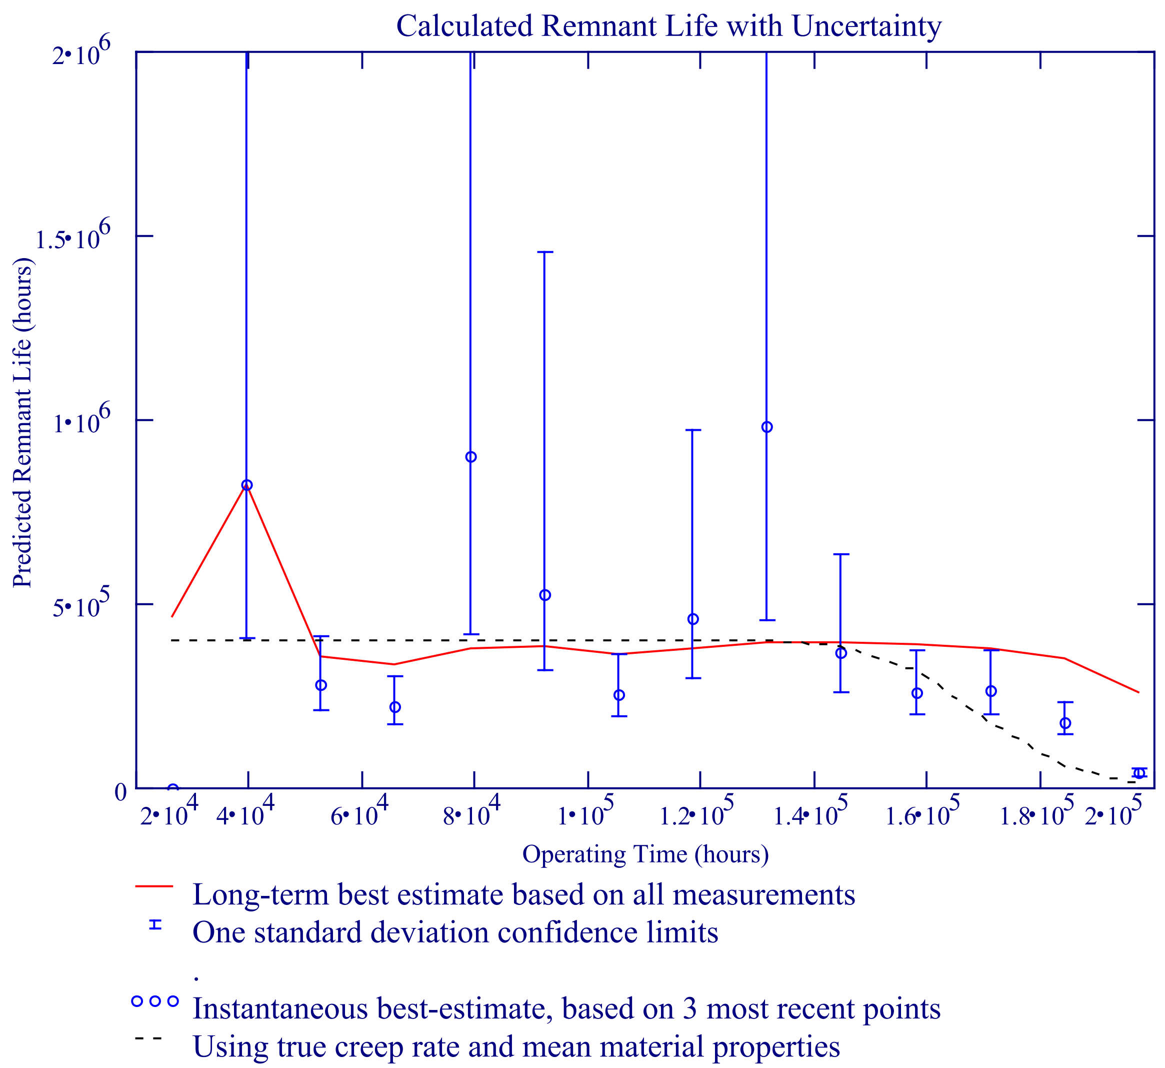

Fig. 3 Predicted Component Life, with One Standard Deviation Confidence Limits, from the Base Case: *The confidence limits are based on the scatter of the creep rate data, assuming that material properties are fixed. The best-estimate and standard deviation were based on a least squares fit using the three most-recent inspection, which was found to be a good compromise between accuracy and number of inspections required to detect imminent failure.*