28 May 2012

Dr Brent Layton
Chair
Electricity Authority
PO Box 10041
Wellington 6143

Dear Brent

SRC advice on security of supply

As part of its functions, the Security and Reliability Council (SRC) is tasked with providing the Electricity Authority (Authority) with independent advice on reliability of supply issues, including the performance of the electricity system, security of supply issues, system security assessments and security of supply forecasts. The purpose of this letter is to provide the SRC's advice to the Authority on current security of supply issues.

The SRC has been presented with, and discussed a range of information on security of supply matters including the policy framework, a security of supply event relevant to the electricity industry and an environmental scan looking at the future of natural gas, one of the power system's key fuel sources. Following these discussions the SRC has the following advice for the Authority:

1. Regulatory uncertainty is a significant security of supply risk

Regulatory uncertainty, across the wider legislative framework, can be a barrier to the long term investment that is necessary to ensure security of supply. This uncertainty is not limited to electricity sector legislation and regulation, but can include, for example, environmental legislation such as the Resource Management Act and the activities of the government-constituted Land and Water Forum.

2. Natural gas supply risks may not be well understood

While the system operator and thermal generators responded quickly and appropriately to the outage on the Maui gas pipeline in October 2011, an interruption to natural gas supply is a potential security of supply risk that may not be well understood by all industry participants.

3. Market signals do not always provide the correct incentives to participants in emergency situations

When the Maui gas pipeline outage occurred in October 2011, Genesis Energy took immediate action to increase coal-fired generation, despite market price signals not being of a level that would normally justify this action. Also, during these sorts of emergency situations the system operator can be left in a position of having to make real-time, security-based decisions on behalf of the industry. These can have significant commercial impacts, and there is a risk that they are judged inappropriately after the fact by others.

4. Under-frequency management project actions need to be progressed

The 13 December 2011 AUFLS event demonstrated that the short-term security of supply policy and operational approach is adequate, but that actions identified through the review of under-frequency arrangements need to be progressed.

5. **Improvements to security of supply information are supported**

The SRC supports the current review of the system operator's security of supply forecasting and information policy (SOSFIP) and the development of the hydro risk curves to improve industry participants' understanding of the potential supply risks.

The SRC will look further into the issue of gas supply risk and related market signals

The SRC considers that there is merit in looking further into the issue of gas supply risk and how the industry responds to gas supply issues. The secretariat has been asked to work together with relevant SRC members to establish an industry forum to discuss the events around the 2011 Maui pipeline outage to ensure there is a wider understanding of what occurred and how the industry did (and should) respond, and to report to the SRC on the results of that forum.

The SRC would be happy to meet with the Authority Board to discuss the advice in this letter further.

Yours sincerely

Dr Kevin Thompson
Chair
Security and Reliability Council

cc SRC members, Carl Hansen (Authority).

28 May 2012

Dr Brent Layton
Chair
Electricity Authority
PO Box 10041
Wellington 6143

Dear Brent

SRC advice on under-supply of AUFLS in the South Island

The Electricity Authority has requested the advice of the Security and Reliability Council (SRC) on the risks associated with the under-supply of automatic under-frequency load shedding (AUFLS) in the South Island. As there is currently no AUFLS system in operation at the Tiwai Point GXP, the Electricity Industry Participation Code (Code) requirements for AUFLS provision in the South Island are not being achieved and this may be increasing security of supply risk.

The SRC has received information on the issue from the secretariat, and also requested the secretariat obtain some additional information from the system operator on the actual costs associated with the measures the system operator has in place to mitigate this security risk.

The SRC considers that the non-supply of AUFLS at Tiwai Point is not a security risk as the system operator procures additional reserves to compensate for the deficit of AUFLS at Tiwai Point when extended contingent events are binding. The additional costs associated with the procurement of these additional reserves are not considered to be significant.

The SRC understands that a more appropriate and enduring solution for the provision of AUFLS at Tiwai Point will be in place by 1 January 2013, and that this is also the shortest practical time in which a permanent solution can be applied.

Yours sincerely

Kevin Thompson

Dr Kevin Thompson
Chair
Security and Reliability Council

cc SRC members, Carl Hansen (Authority).

3 August 2012

Security and Reliability Council

By Email

Dear Council members

WCM and WEM security of supply standards

This memo is to update the Security and Reliability Council (SRC) on the Authority's process for reviewing and updating the Winter Capacity Margin (WCM) and Winter Energy Margin (WEM) security of supply standards. The standards were discussed at the SRC's August 2011 meeting, where it was noted they would be reviewed by the Authority in 2012.

The function of these standards is to serve as points of reference in determining how likely it is that there will be efficient levels of generation and inter-island transmission available to meet demand in the next 5-10 years.

The capacity security of supply standard (WCM) is used in the process of assessing whether there will be an efficient level of peaking generation and transmission capacity for north transfer to meet peak demand.

The energy security of supply standards[1] (WEM) are used in the process of assessing whether there will be an efficient level of reserve generation and transmission capacity for south transfer to manage extended dry sequences.

The Authority has now completed its review of these standards (taking into account factors such as the imminent availability of HVDC Pole 3 and changes to the generation mix) and is proposing some changes. The Authority has published a consultation paper "Winter Energy and Capacity Security of Supply Standards" (http://www.ea.govt.nz/our-work/consultations/sos/winter-energy-capacity-security-supply-standards/) and has sought submissions from participants by 7 August 2012.

The consultation paper proposes to:

- amend Part 7 of the Code to change the security of supply standards; and
- submit to the system operator a proposal for a variation to the Security of Supply Forecasting and Information Policy (SOSFIP) to change the calculation of the WCM and WEM security metrics.

[1] There are separate energy margins for New Zealand and the South Island.

730929-1

Level 7, ASB Bank Tower, 2 Hunter Street, PO Box 10041, Wellington 6143, New Zealand TEL + 64 4 460 8860 FAX + 64 4 460 8879 WEB www.ea.govt.nz

The Authority Board has requested the SRC's comments on the proposed changes. The SRC secretariat will provide the SRC with a summary of the consultation paper and the industry submissions, in advance of the 29 August 2012 meeting.

Work on potential summer security standards is proceeding in parallel and an update will be provided to the SRC in due course (as requested at the August 2011 meeting).

Yours sincerely

Fraser Clark
SRC Secretariat

cc: Carl Hansen
Kieran Devine, Transpower NZ Limited

HYLAND MCQUEEN LIMITED
PO Box 1003
Dunedin 9054
New Zealand
ph +64 3 455 3583
web www.hylandmcqueen.co.nz
email info@hylandmcqueen.co.nz

31. July 2012

Dr Kevin Thompson,
Security and Reliability Council,
Electricity Authority.

Dear Dr Thompson,

I have been working for some time on a problem of using machine condition scores to determine a hazard of failure curve for that machine. I enclose a copy of a paper I recently presented at the EEA conference. I would like to take this work further and would like to make contact with any person or group that is involved in similar work.

If you know of such, I would be very grateful if you could send me any contacts.

Malcolm McQueen.

Malcolm A. C. McQueen

Malcolm McQueen Phd
Hyland McQueen Ltd
PO Box 1003
Dunedin 9054
New Zealand
Malcolm McQueen ph +64 3 477 5147
Pat Hyland ph +64 3 455 3583
fax +64 3 455 3557
web www.hylandmcqueen.co.nz

From Condition to Hazard

M McQueen
P Hyland
Hyland McQueen Limited.

Presenter: M. McQueen

EEA Conference & Exhibition 2012, 22 - 20 June, Auckland

Motivation

Prediction of the hazard of failure is important for financial modelling of an asset's value to an organisation. Towards this end measurement of the asset's condition is part of the routine operating procedure of many organisations' asset management practice.

However, the condition measurements alone cannot provide such prediction. They only provide a point estimate of the most likely time of failure. Additional data and a methodology incorporating some model of the asset is required.

Markov process models have been proposed that statistically describe an asset's failure modes but are not of practical use to industry because the required knowlege of the machine's internal states is insufficient.

This paper describes a model whose parameters can be calibrated against commonly available data and so provides a practical way of estimating an asset's hazard of failure. The method is applicable to a wide range of machine types. It has been successfully applied to hydro and thermal generators, and power transformers. It is sufficiently general to be applied to other asset classes as well.

The Problem

In addition to the condition measurement data, population survival data for the machine type is often available. These two sources of information may be combined using a machine model to provide an estimator of a machine's future hazard of failure.

There are multiple parts to this problem; we consider it to be structured as follows:

1. A set of condition measurements is made on a particular machine at some time.
2. These condition measurements are combined to provide a condition score which is an estimate of the condition of the machine at that time; its condition value.
3. The condition value for the machine is combined with population survival estimates for that type of machine to yield a hazard over time function for the particular machine.

We principally consider step 3 in this process but we must first consider the notion of what the "condition" of a machine means.

Interpretation of Condition

Engineers, familiar with some machine, recognise certain physical attributes as indicative of the machine's deterioration. Measurements of such physical attributes can be used as indicators of condition.

It may be that a single physical characteristic is sufficient to fully characterise the condition but, more often, a collection of measurements is considered necessary and a scoring system is used to calculate a condition score. This scoring system is developed by engineers experienced with machines of this type and is, implicitly, a predictor of the age at failure. In this case, the condition score can be interpreted as the expected fraction of life remaining of the machine. Alternatively the condition scoring system may directly provide an estimate of the machine's remaining life in which case its condition is taken as the fraction of life remaining at the time of the condition measurement.

For the present considerations we will assume that it is the case that the condition linearly decreases from a value of one at the start of the machine's life to zero at failure.

Prediction of hazard

With the above interpretation, a condition value for a machine implies its expected age at failure. It does not give a prediction of the time of actual failure. Actual failure is most likely to occur somewhere near this time but may occur before or after. For the purposes of financial modelling it is the hazard of failure that is important and the condition alone does not give this.

Further information relevant to the hazard of failure is available from population survival curves of similar machines. While there are often problems with the homogeneity of the information derived from this source, it does provide direct evidence of actual failures of the machines of interest and so contains information on the hazard function of the machines making up the population. This survival curve is usually well represented by a Weibull curve which is described by two parameters.

From these two sources of information we derive an estimator of the hazard function for a particular machine.

The derivation is carried out using a model of the machine then fitting the survival curve derived for a population of models to the observed survival curve of the population of machines to determine the parameters of the model. Then the model is run with these parameters using the observed machine condition as input to give a hazard function for that machine.

Attention should be paid to what is actually meant by the term 'failure'. In practice 'failure' is often considered as being 'end of life' which is usually considered as occurring when the condition of the machine drops below some minimum acceptable value. The criterion of failure should be used consistently throughout the method and the hazard function interpreted accordingly.

Typical data that we may have for input to the method is shown in Figure 1.

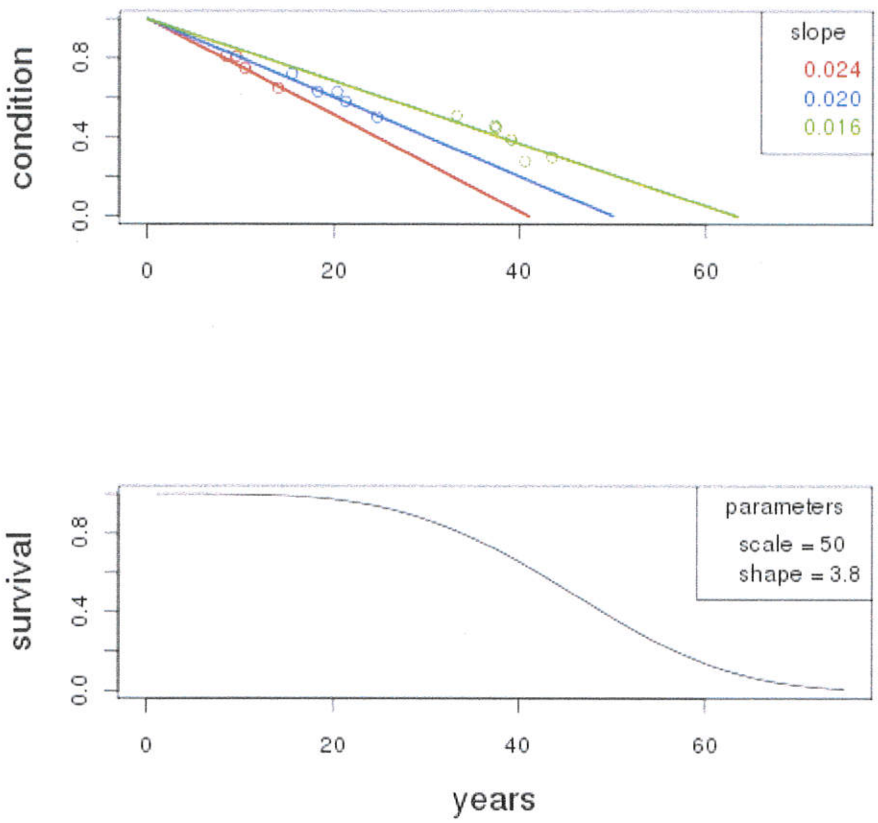

Figure 1.

The top panel shows what might typically be measured for three different machines. The open circles show condition scores, calculated from sets of condition measurements. The bottom panel shows the survival curve for the population of machines of which the individual machines are members.

Machine Model

A major problem confronted in developing the model is that the number of parameters required should not exceed the number of measurements available. This constrains the type of model. It excludes detailed machine models such as those based on Markov process because although, the processes of physical degradation are known, the parameters governing these for any particular machine are not known in sufficient detail.

We develop a simple model of the machine that captures the essence of the problem. This model comprises a deterministic and a random component.

It is assumed that stresses are placed on the machine during its operation. When new, these stresses are insufficient to cause the machine to fail in normal operation but, as the machine ages, its ability to resist the stress decreases. We refer to this ability to resist strength as the machine's strength. Eventually a time is reached when the strength becomes insufficient to

resist the instantaneous stress placed upon the machine, at which point the machine fails.
The model describes the machine by a single variable which characterises its strength and a random variable characterising the stress. The strength and stress variables are composites of all the physical, electrical, and chemical stresses and strengths of the machine components that are critical in determining the failure of the machine. The strength is normalised to start at one and decreases linearly during operation to become zero. Thus the strength can be considered as some quantity that is used up during operation and corresponds to some combination of quantities such as bearing thickness and insulation strength. It must be noted that the machine will fail before its strength reaches zero.
The model simulates the operation of the machine by calculating the strength at successive time-steps and checking whether the maximum stress in the preceding time-step has not exceeded its strength. The values of maximum stress is taken from an exponential distribution.
Thus the model describing a specific machine has two parameters: the slope of the strength and the exponent of the stress distribution.

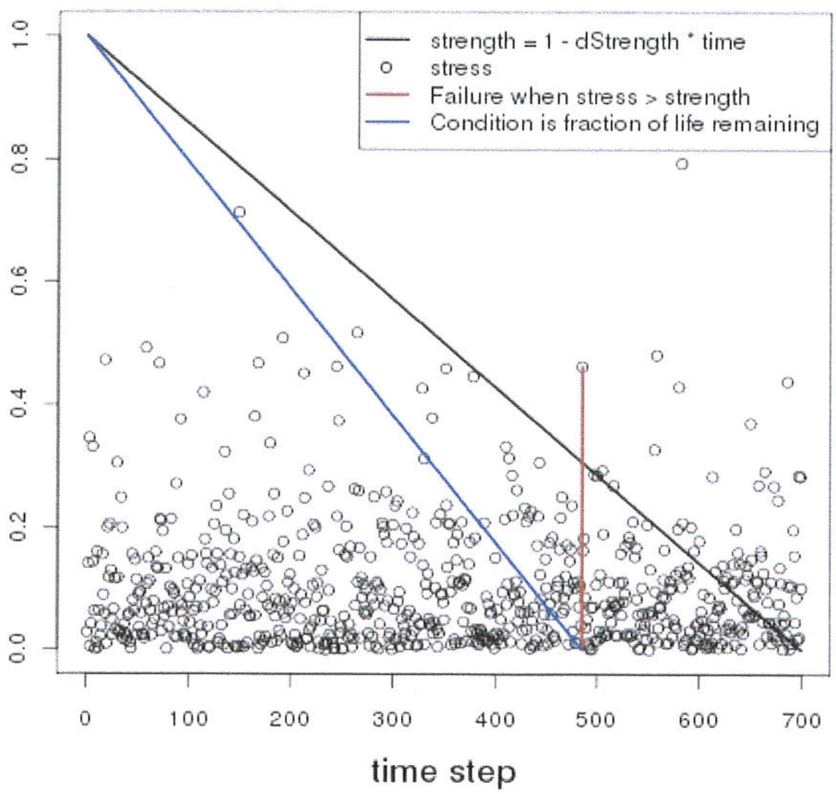

Figure 2 illustrates this model.

Figure 2.

Example Machine

We illustrate with an example of the application of the model. A problem with which we are all familiar is estimating the range of a motor vehicle setting out on a trip with a full tank of fuel. How far will it go? We consider fuel consumption/km to be constant. We also consider there to be some slosh of the fuel in the tank caused by the roughness of the road. The vehicle will stop when there is, instantaneously, no fuel reaching the engine and this will occur some time before the fuel tank is completely empty. The fuel gauge registers the strength (in the above sense) of the vehicle but, by itself, it does not allow an accurate estimate of the range. We must also rely on our experience of the vehicle to estimate its remaining range. The stress on the machine (in the above sense) is created by the sloshing of the fuel in the tank.

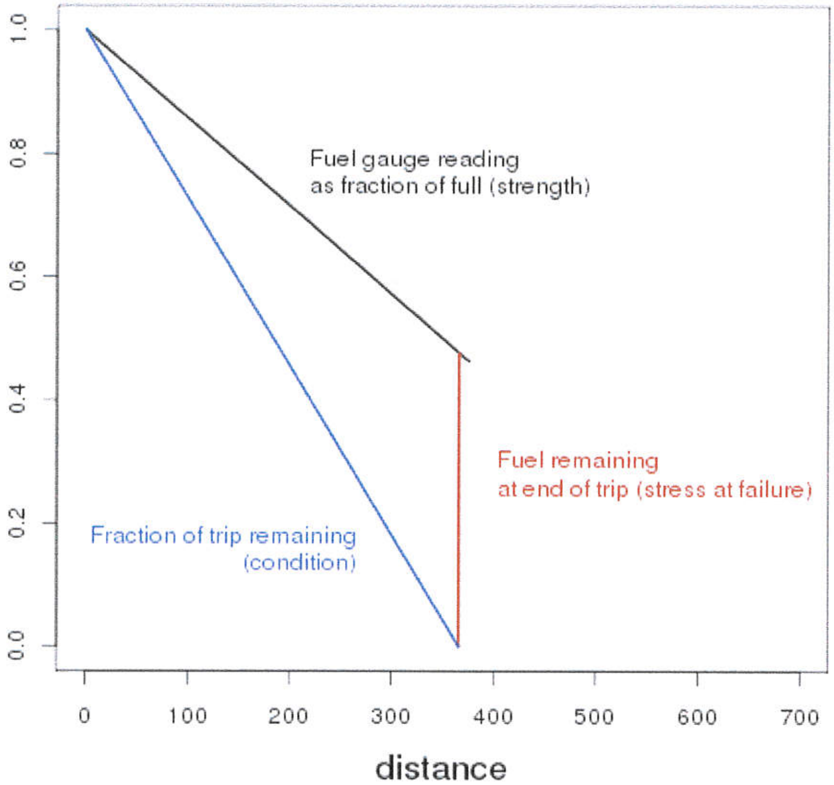

This is illustrated in Figure 3.

Figure 3

Deriving the Parameters of the Model

The two parameters of this model are not directly given by the data available. We determine them by matching population survival curves. One calculated from a population of machine models; the other from the known survival curve for the population of machines of which the one being considered is an example. The relationship between slope of strength and the slope of the condition for the machine being considered can then be determined. This process requires assumptions regarding the statistical distributions involved. Due to the lack of data available to this process the distributions are constrained to single parameter distributions; a Rayleigh distribution is assumed which is also broadly justifiable on physical grounds. This approach still requires too many parameters for a solution. A further assumption is necessary which is as follows:

The survival curve of the population of machines is influenced by two factors. Firstly, individual machines may deteriorate more quickly than their design dictates due to manufacturing defects and, second, the environments in which they operate vary. On the basis that these machines have been manufactured for a sufficiently long time for the manufacturers to optimise their manufacturing processes, the assumption made is that these factors contribute equally to a decrease in age at failure.

With this assumption there is sufficient data to determine the parameters of the model and the hazard function of the model can be calculated.

Using the above method, the hazard functions for the three example machines (red, blue, green) given in the fig 1 are presented in figure 4.

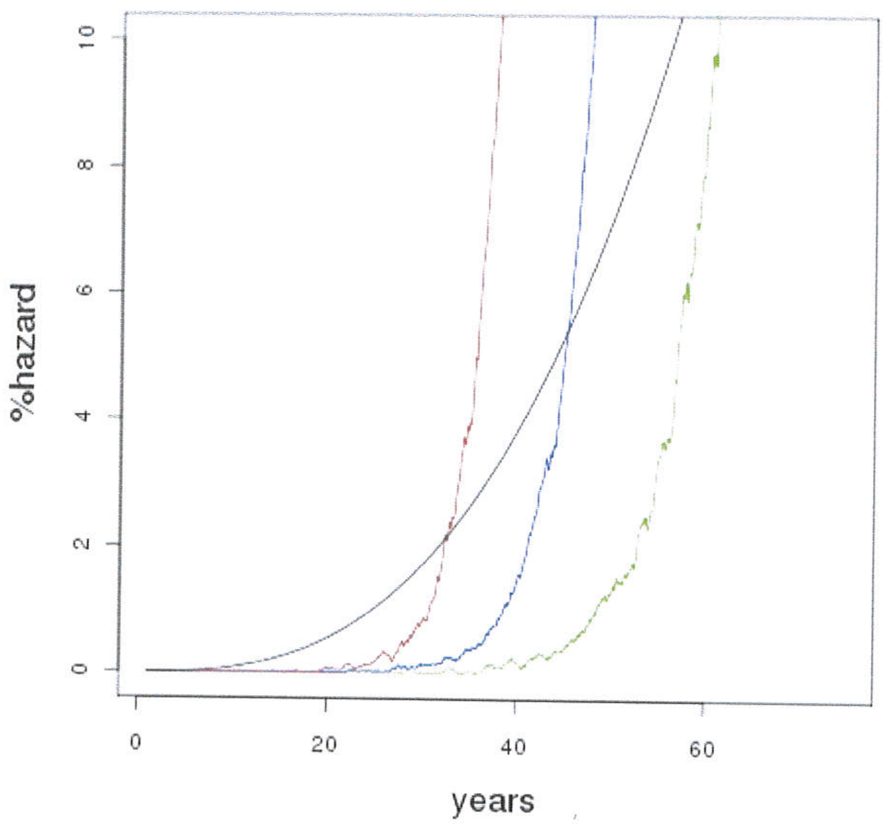

Figure 4

Also shown, in black, is the hazard curve of the machine population. In the absence of condition information, this curve is the best estimate of the machine's hazard function.

Assumptions of the Method

The method described is based on a number of assumptions which affect the estimation of the hazard function.

- While the assumptions of the stress and population distributions do affect the calculated hazard function, calculation using other distributions indicate that the hazard function is not unduly affected.

- The definition of condition is fundamental to this method. There seems to be no universal definition of this quantity. In general usage, 'condition' has a very wide range of application. Nevertheless these different usages have a commonality that allows comparison to carry meaning. A more general definition than that used in this paper is to define condition as a fraction of ideal utility. If the utility of machines is defined as their output over their lifetime, then the definition used in this paper follows.
- No data is available to verify that there is an equal partition of the effects of stress and machine variability but is supported by general economic argument.

These assumptions are unavoidable. They are required by the lack of data that characterises the problem considered.
The hazard function calculated from the model is an estimator of the hazard function of the machine.

Errors

Above we have considered the condition of the machine as the predictor of the time of failure. The scoring system provides an estimate of this quantity from condition measurements. Both these steps are subject to errors and contribute to uncertainty in the estimate of the hazard function.
The condition measurements will be subject to measurement errors. Good practice will try to minimise these but even if they could be reduced to zero, uncertainties also arise from the scoring system used to derive the condition score. The scoring system will depend on some concept of normal operation of the machine and some concept of its normal ageing processes. These are judged on the basis of a population of machines and so characterise a typical machine. The scoring system applied to any particular machine will inevitably yield a condition estimate that deviates from that of such typical machine population depending on its specific operation history. Care must be taken as to what is considered normal operation. It should consider what aspects of maintenance and refurbishment are included as part of normal operation and also the degree of the utilisation and the effects of startup and shutdown.

Fitting condition line

The input to the current method is the slope of the machine condition given condition score estimates. Generally, successive condition sores will be estimated and so, in effect, it seems the problem is one of fitting some straight line to these values. It is tempting to use a least-squares fit to this line but this would be in error as successive condition score measurements are not independent. The process of determining a condition score is a diagnostic one where each estimate is a refinement of the previous one as more information on the machine's deterioration becomes apparent. The engineer should fit a line depending on which condition scores, in his judgement, most accurately indicate the machine's condition. This is likely to be heavily weighted towards the last score because at the time that was determined the most experience with the machine has been gained.

Summary

A method has been described for determining an estimator of the hazard function for a machine.
It uses a simple machine model comprising a deterministic and a random component calibrated using condition measurements and a machine population survival curve.
It provides a consistent approach and so allows reliable comparison between different machines.
For further details on the method, please contact the author.

30 August 2012

Hyland McQueen Ltd.
PO Box 1003
Dunedin 9054

Attention: Malcolm McQueen

Dear Malcolm

Use of machine condition scores to determine hazard of failure curves

Thank you for your letter of 31 July 2012 discussing your work on the problem of using machine condition scores to determine a hazard of failure curve for that machine. Detailed engineering investigations of this type are outside the remit of the Security and Reliability Council (SRC), so we are unable to assist you in this instance (the role of the SRC is to provide independent advice to the Electricity Authority on the performance of the electricity system and the system operator; and reliability of supply issues). We wish you every success in finding partners or other people or groups involved in similar work.

Yours sincerely

Dr. Kevin Thompson
Chair
Security and Reliability Council