John Field Consulting Pty Ltd

ABN 62 096 615 235

10 High St, Burnside SA 5066
Telephone: (08) 8332 5294
Facsimile: (08) 8332 1229
Mobile: 0409 097 586
Email: JohnField@ozemail.com.au

Defining Major Event Days

Report produced for

ETSA Utilities
GPO Box 777
ADELAIDE SA 5001

5 August 2008

Defining Major Event Days

EXECUTIVE SUMMARY

This report considers the method used to define Major Event Days by the IEEE Guide for Electric Power Distribution Reliability Indices, and its subsequent use by the Essential Services Commission of SA and the Australian Energy Regulator .

The validity of the method depends on the normality of the distribution of log(SAIDI) values. This condition is not satisfied for the three years of ETSA Utilities data available, and so the numbers of Major Event Days resulting from use of the method will not be valid.

Two modifications of the IEEE standard proposed by ETSA Utilities are examined, and a third modification is also examined. It is concluded that the most useful modification may be to use a rolling two-day period for the calculation of SAIDI values.

Concern is also expressed that the AER apparently fails to recognise that iteration may be necessary to remove previously unidentified extreme events from the estimation of a stable value of T_{MED}.

TABLE OF CONTENTS

1. The 2.5β method and the Normal distribution

The IEEE standard[1] uses the '2.5β method' to determine the number of Major Event Days. The rationale for this method is as follows.

Suppose we have some variable, X, which has a mean, α, and a standard deviation, β. We define a threshold, T, which is 2.5 standard deviations above the mean of the distribution, so

$$T = \alpha + 2.5\,\beta$$

We can do this for any variable with any distribution (provided α and β exist).

If we then assume that the variable X has some particular probability distribution, we can calculate the probability that we will observe values of X which lie beyond the threshold.

The IEEE Working Group used a Gaussian (or normal) distribution to do their calculations. They chose the 2.5 multiplier so that the probability of exceeding T was 0.000621, or 2.3 days/year. This choice was based on 'consensus reached among Distribution Design Working Group members on the appropriate number of days that should be classified as Major Event Days'[2], in other words, the Working Group members thought that 2.3 days excluded was 'about right': there was no scientific basis for this number.

The Working Group stated[3], without any justification, that SAIDI has a log-normal distribution. If this were so, then log(SAIDI) would have a normal distribution, and the probability calculations would apply to it.

The validity of the method, as used by the IEEE Standard, depends on the fact that the distribution of log(SAIDI) is normal. If it is not, the method is invalid, since the probability calculations are inappropriate.

The use of a mathematical probability distribution such as the normal or log-normal distribution to describe the actual distribution of some variable is a convenience to enable us to perform probability calculations about the variable. It is not some innate property of the variable, and is at best an approximation to the actual distribution.

In the current case, as the IEEE standard points out[4], SAIDI values experienced by a utility will be subject to variation because of geography (and consequent weather), system design, data issues and so on. These variations will have subtle influences on the actual distribution of the values, and we cannot guarantee that actual SAIDI values for all utilities will have the same distribution.

---

[1] IEEE Std 1366-2003, IEEE Guide for Electric Power Distribution Reliability Indices.
[2] IEEE Std 1366-2003, Section B.5.1, p32.
[3] IEEE Std 1366-2003, Section B.4.3, p30.
[4] IEEE Std 1366-2003, Section 6.2, p17.

2. ETSA Utilities data

ETSA Utilities supplied three years of daily unplanned SAIDI values for the period 1/7/2005 to 30/6/2008. This was the full dataset available from the ETSA Utilities newly-implemented OMS system. We use this to examine the assumption that ETSA Utilities SAIDI data is log-normally distributed.

We can use various methods to judge the normality of the log data.

- If the data is log-normally distributed, then we would expect the mean and median of the log data to be similar. The mean is -1.551 while the median is -1.440, evidence that the data is skewed.
- We can calculate the skewness and kurtosis for log(SAIDI). The skewness is a measure of the symmetry of the distribution, and kurtosis is a measure of whether the distribution is peaked or flat relative to the normal distribution. For the normal distribution we would expect both to be zero. For this data, skewness = -0.321 with a 95% confidence interval of (-0.466 to -0.176). The kurtosis is 0.604 with a 95% confidence interval of (0.314 to 0.894). Neither confidence interval includes zero, and we conclude that the distribution differs from a normal distribution. The distribution is skewed to the left (ie the left hand tail is long relative to the right hand tail) and the distribution is more peaked than a normal distribution.
- We also use the Anderson-Darling test to test for normality[5]. This test is one of the most powerful for testing for departures from normality. It is based on the empirical cumulative distribution function of the data, and tests how similar this is to the cumulative distribution function for a normal distribution. It tests for all sorts of departures from normality, but puts emphasis on the tails of the distribution. The usual statistical practice is to reject the hypothesis that the data come from a normal distribution if the significance probability is less than 0.05; for the ETSA Utilities data, the test gives a significance probability of P=0.0006; that is, there is a chance of only 6 in 10,000 that the log(SAIDI) data come from a normal distribution.

We conclude that the distribution of log(SAIDI) is significantly different from the normal distribution. Hence the results of the 2.5β method are invalid for this data.

However, as a comparison to what follows, we use the data to calculate T_{MED}. We initially start with all the data, including previous extreme events. We find $\alpha = -1.551$, $\beta = 1.353$, $T_{MED} = \exp(\alpha + 2.5\beta) = 6.248$. There are three days over the three years with SAIDI values greater than this. This is rather less than the 2.3 days per year expected by the IEEE standard.

As the AER's Service Target Performance Incentive Scheme points out[6], it is not clear in the IEEE standard how previous extreme events are to be excluded from the data. The AER suggests iterating to find T_{MED}, but only for one cycle. It is not clear why iteration should not continue until a stable value is reached. This is not an issue here, but it is in later cases. So we exclude the three SAIDI values identified above

[5] Stephens, M. A. (1974). EDF Statistics for Goodness of Fit and Some Comparisons, *Journal of the American Statistical Association*, **69**, pp730-737.
[6] AER, June 2008: Electricity distribution network service providers – Service target performance incentive scheme, Appendix D, p30.

and repeat the calculation; we find that TMED has reduced slightly, but the same three events are still excluded. The table shows summary results.

Table 1: Calculation of MEDs with IEEE method

Iteration	T_{MED}	No. days excluded	AD test of remainder
1	6.248	3	P=0.00030
2	6.019	3	P=0.00030

It is clear from the Anderson-Darling (AD) test results that eliminating the three days with log(SAIDI) > T_{MED} has not affected the (lack of) normality of the data.

3. Suggested ETSA Utilities modifications

ETSA Utilities has made two suggestions for modifications to the IEEE definition of a Major Event Day[7]. These are

- Applying the 2.5β method to two consecutive rolling calendar days; and
- Excluding all days outside the range $\alpha \pm 2.5\ \beta$.

We examine each of these in turn.

3.1 Two consecutive rolling calendar days

The IEEE standard uses a 24 hour window in which to calculate SAIDI, but this really is an arbitrary decision, and no reasons are given for it. ETSA Utilities argues that since the vast majority of SA's significant weather events are spread over two or more days, using a longer period than 24 hours to calculate individual SAIDI values may compensate for this factor. SAIDI for a two day period is the sum of the two days' SAIDI values, assuming the total number of customers remains the same for the two days.

- The mean of the rolling two-day SAIDI data is -0.575, while the median is -0.562, so this option has reduced the skewness to a large extent.
- The skewness is -0.006 with a 95% confidence interval of (-0.080 to 0.068). This interval includes zero, and so the skewness is not significantly different from that of a normal distribution. Kurtosis is 0.336 with a 95% confidence interval of (0.188 to 0.484), so the distribution is still more peaked than a normal distribution.
- An Anderson-Darling test for normality gives a probability of P=0.007. That is, there is still evidence of non-normality in the data, but the evidence is less strong than for the log of the daily SAIDI values.

We show the results of using this option in Table 2.

Table 2: Calculation of MEDs using two day rolling period

Iteration	T_{MED}	No. days excluded	AD test of remainder
1	7.600	9	P=0.030
2	6.887	9	P=0.030

The Anderson-Darling (AD) test result shows that there is still a significant difference from normality at each step of the iteration process, although excluding the MEDs has helped to bring the distribution closer to normality.

[7] ETSA Utilities: Submission to ESCoSA's Draft Decision: South Australian Jurisdictional Service Standards to apply to ETSA Utilities in the 2010-2015 Regulatory Period.

3.2 Trimming both ends

The IEEE standard excludes events with duration of less than five minutes. The AER proposes to exclude only events less than one minute. Thus under the AER proposal there may be more extreme events at the low SAIDI end of the SAIDI distribution than were originally intended by the IEEE standard.

The normal distribution is symmetric, and so the shape of the extreme upper tail is dependent on the shape of the extreme lower tail. If we add extra values to the extreme lower tail, this will change the shape of the extreme upper tail of the fitted normal distribution. By including SAIDI values of less than 5 minutes duration, the upper tail will also be affected: the standard deviation of the distribution will increase, and so will the MED threshold, thus eliminating fewer major events.

The IEEE standard omits days with zero SAIDI, so there is an argument for omitting days with very small SAIDI as well. Omitting SAIDI values less than exp(α - 2.5 β) sets a definite limit.

The results of fitting this option to the data are shown in Table 3. In this case there are two values of TMED, one for each tail of the distribution.

Table 3: Calculation of MEDs by trimming both ends of SAIDI distribution

Iteration	T_{MED}	No. days excluded (low SAIDI, high SAIDI)	AD test of remainder
1	0.0072, 6.248	18 (15, 3)	P=0.074
2	0.0096, 5.136	31 (19, 12)	P=0.009
3	0.011, 4.529	38 (26, 12)	P=0.0013
4	0.011, 4.243	41 (22, 19)	P=0.0006
5	0.011, 4.137	42 (23, 19)	P=0.0006
6	0.011, 4.118	42 (23, 19)	P=0.0006

The problem of the number of iterations needed to reach stable values of T_{MED} is clearly shown. However, stopping iteration before a stable solution is reached will have a large effect on the values of T_{MED} to be used in the next regulatory period. Of the 42 values excluded, 23 are in the lower tail and 19 in the upper tail.

The mean of the non-excluded log(SAIDI) values is -1.527, and the median is -1.436, so this option has not had a large effect on the skewness of the distribution by these measures. The skewness is -0.134 with a 95% confidence interval of (-0.210 to -0.059), and the kurtosis is -0.499 with a 95% confidence interval of (-0.650 to -0.348). So removing values at both ends of the distribution of log(SAIDI) leaves the distribution still with a longer tail on the left (ie for low SAIDI values) but this time the distribution is more flattened than a normal distribution. The Anderson-Darling test confirms that the distribution is still non-normal.

4. An alternative modification

The reason for transforming the SAIDI values by taking logs is to try and arrive at a normal distribution of values, so that the 2.5β method can be validly applied. As we have seen this has been unsuccessful thus far.

An alternative method is to use a different transformation which better converts the distribution to normal.

A possibility is the Box-Cox transformation[8], which for a variable X is defined as

$$X^{(\lambda)} = (X^{\lambda} - 1) / \lambda \text{ for } \lambda \neq 1, \text{ and } X^{(\lambda)} = \log(X) \text{ for } \lambda = 0.$$

We use the data to estimate λ. If we do this, we find $\lambda = 0.0678$. Denoting the transformed data as $SAIDI^{(\lambda)}$ we calculate

$$SAIDI^{(\lambda)} = (SAIDI^{0.0678} - 1) / 0.0678$$

Calculation of λ and $SAIDI^{(\lambda)}$ are not difficult; standard programs exist to calculate λ, and then calculation of $SAIDI^{(\lambda)}$ is straightforward. In this instance we used the (public-domain) software package R[9].

The mean of the transformed data is -1.417 and the median is -1.372. The skewness is not significantly different from zero: 0.010, 95% confidence interval (-0.063 to 0.084). The kurtosis shows the distribution is still slightly peaked compared to a normal distribution: 0.329, 95% confidence interval (0.181 to 0.477). The Anderson-Darling test of this data however shows that the distribution is not significantly different from a normal distribution (P=0.153).

Using this transformation, as previously we define α and β as the mean and standard deviation of the transformed data, so

$$\alpha = \text{mean}(SAIDI^{(\lambda)}) \text{ and } \beta = \text{sd}(SAIDI^{(\lambda)})$$

$$\text{and then } T_{MED} = \alpha + 2.5\,\beta$$

Then any day where $SAIDI^{(\lambda)} > T_{MED}$ is defined as a Major Event Day.

The SAIDI threshold value equivalent to T_{MED} is found by inverting the transformation, thus:

$$SAIDI_{MED} = (\lambda\, SAIDI^{(\lambda)} + 1)^{1/\lambda}$$

We iterate as previously to find a stable value of T_{MED}. This involves recalculating λ and $SAIDI^{(\lambda)}$ at each stage. Details are given in the table below.

---

[8] Box, GEP and Cox DR (1964) An analysis of transformations (with discussion). *Journal of the Royal Statistical Society B* **26**, pp211-252.

[9] R Development Core Team (2008). *R: A language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria, http://www.R-project.org

Table 4: Calculation of MEDs using Box-Cox transformation

Iteration	λ	T_{MED}	$SAIDI_{MED}$	No. days excluded	AD test of remainder
0	0.068				P=0.153
1	0.124	1.607	4.330	17	P=0.815
2	0.149	1.241	3.121	25	P=0.314
3	0.157	1.093	2.742	28	P=0.183
4	0.157	1.048	2.639	28	P=0.183

Note that at no stage is the non-excluded data significantly different from normal.

The value of T_{MED} and λ for use in the following period would be 1.048 and 0.157.

This transformation has the advantage that the normality of the data is retained, and so the methodology is detecting Major Event Days with the correct probability.

5. Discussion

The IEEE method is clearly inappropriate for the ETSA Utilities data. By any measure the log(SAIDI) data for ETSA Utilities, at least as represented by the three years available, is non-normal. This means that the IEEE method will produce probabilities, and hence numbers of Major Event Days per year, which are incorrect.

Three variations of the IEEE method have been examined. A summary of the effectiveness of these methods in detecting MEDs is shown in Table 5 below.

Table 5: Numbers of MEDs by the various methods

Method	Total MEDs in upper tail	MEDs per year
IEEE	3	1.0
Two consecutive days	9	3.0
Trimming both ends	19	6.3
Box-Cox transformation	28	9.3

The effect on total SAIDI is shown below

Table 6: SAIDI (in minutes) for each year by the various methods

Method	2005/06	2006/07	2007/08	Total
Raw data	200.1	197.1	138.7	535.9
IEEE	184.5	197.1	132.2	513.8
Two consecutive days	168.2	181.7	138.7	488.6
Trimming both ends	152.1	155.6	121.8	429.4
Box-Cox transformation	135.2	145.2	118.3	398.6

Of these three variations, the Box-Cox transformation produces data whose distribution is consistently not different from a normal distribution, and hence complies with the assumptions of the IEEE standard. Yet it produces many more MEDs per year (9.3) than the 2.3 days per year envisaged by the IEEE Working Group. The IEEE states[10] 'The b multiplier of 2.5 was chosen because, in theory, it would classify 2.3 days per year as major events. If significantly more days than this are identified, they represent events that have occurred outside the random process that is assumed to control distribution system reliability'. It is quite possible that the data which the IEEE Working Group used was not quite log-normally distributed, despite their assertions that it was – they appear to have done no testing of the assumption.

In this case, we may be better choosing a method which produces a value somewhere near the 2.3 days/year. Using a period of two consecutive days to define SAIDI, rather than the single day used by the IEEE may be preferable, as it has a physical

[10] IEEE Std 1366-2003, Section B.1, p26

justification in terms of the time taken for significant weather events to cross South Australia.

We note that other variations of the IEEE method are also possible: for instance, censoring the lower half of the distribution of log(SAIDI) values and fitting the normal distribution to the upper half only, or combining some of the features of the variations investigated above. We have not investigated these however.

We also are concerned that the steps of the 2.5β method outlined by the AER[11] fail to recognise that iteration will be necessary to get a stable T_{MED} value. This matter should be discussed with the AER.

This analysis has been done using three years' data, the only data available from ETSA Utilities' OMS system. It would, of course, be possible to resample this data, allowing for seasonality, to produce an extra two years' data to achieve the IEEE desired run length of five years. However this will not take into account variations in weather patterns other than those observed in the last three years. It is felt that this would be a serious shortcoming of the approach, and so it has not been attempted.

[11] AER, June 2008: Electricity distribution network service providers – Service target performance incentive scheme, Appendix D, p30.