Bullet Proof Your DC Link Aluminum Electrolytic Capacitor Banks

Laird L Macomber
Cornell Dubilier
140 Technology Place
Liberty, SC 29657 USA
011-864-843-2277 ext 379
011-864-843-3800 fax

Sam G. Parler, Jr., P.E.
Cornell Dubilier
140 Technology Place
Liberty, SC 29657
011-864-843-2277 ext 224
011-864-605-2009 fax

Abstract

While screw terminal aluminum electrolytic capacitors are routinely used as bus capacitors in motor drives, UPS systems, welders and power supplies, design engineers are often concerned because aluminum electrolytic capacitors can wear out and can cause short service life. But in the large capacitor banks that are in high-power systems the problem is not wearout; the problem is short-circuit failure. The low milliohm source impedances delivered by multiple capacitors in parallel set the system up for short-circuit failure. When dc leakage momentarily increases at a partial discharge site in a capacitor's anode foil, the capacitor may not self heal as expected. The low source impedance may drive the current up and cause runaway local heating resulting in blowout-short-circuit failures.

The usual tool for assuring that aluminum electrolytic capacitors don't wearout and do deliver adequate service life is the load-life test, an accelerated life test at rated voltage, temperature and ripple current. Since the capacitors will operate at less than full rated conditions in the application, you can use the industry rule of thumb that for every 10 °C that you lower the temperature the life doubles, and surmise, for example, that surviving a 5000 hour 85 °C life test would demonstrate a greater than 40,000 hour life. That's about 5 year-service life at an ambient temperature of 55 °C.

But for a measure of the robustness needed to operate reliably in low-impedance capacitor banks you need a different tool; you need an overvoltage transient test. A method for assuring that the capacitors are resistant to partial discharge triggered short-circuit failure is embodied in an overvoltage surge test adapted from a test developed for capacitor banks in the traction drives of trains. This test pushes the capacitors past their anode-foil formation voltage into a semi-zener, flat-top high voltage mode with momentary, very-high leakage current and fixed voltage drop. It's a tough test and most capacitors fail in 20 to 100 shots.

This article explains the test and explores the performance of a new Cornell Dubilier capacitor design that extends the capacitors' capability

from a hundred to thousands of overvoltage transient surges. It is also the oxymoron of bus capacitors because it simultaneously delivers the previously mutually exclusive characteristics of ultra-low ESR and high-voltage transient robustness. Previous capacitors have not simultaneously been able to provide both high ripple current capability and high-voltage transient robustness.

Distinguishing Lifetime from Reliability

Capacitor lifetime is limited by electrochemical degradation that proceeds to wear out the capacitor in a fairly predictable fashion, accelerated by temperature and voltage stress. Along the way, however, random failures are bound to occur if the population is large enough. These are generally unrelated to wearout, and are instead linked to some latent weakness, usually in the paper, foil, or connections. These failures are most often short-circuits, and they may occur suddenly and without warning, although occasionally capacitors may begin to draw excessive leakage current and generate sufficient hydrogen gas pressure to rupture the safety vent then subsequently dry out and fail open circuit. During the manufacture of the capacitor, rated voltage and temperature are applied in the aging process. Capacitance, leakage current, and ESR are tested on a 100% basis. Usually we employ additional screening techniques to attempt to weed out infant mortalities. Such methods include burn-in, surge voltage test, and hot DC leakage test. At CDE at the present time, these screening methods may weed out an additional 0.1–0.5% of weak capacitors which may fail early in the field. The yield of large high-voltage capacitors has increased from 92% in 1990 to over 98% today. If the burn-in or other high--stress-screening processes are repeated, a small percentage (0.02–0.2%) may be expected to fail each time.

Bank Considerations

In the field, it is often the case that several or many capacitors are connected in series and in parallel. Due to adverse effects of fuses (resistance, cost, size and inductance), it is usually only practical to do this in a manner that unfortunately creates the situation that when one capacitor fails short-circuit, the bank and system cease to function. It is the case, therefore, that the failure rate of the bank is approximately equal to the number of capacitors in the bank multiplied by the failure rate of each capacitor. It is generally true that smaller capacitors (CV rating) are more reliable than larger capacitors. But when a large CV is needed, the highest reliability is usually achieved by using a smaller number of physically larger capacitors. This is true for two reasons. First, the lower number of terminals, welds, connections, potential failure sites of the fewer, large capacitors is an inherent reliability advantage. Second, large capacitors are generally designed, constructed, and screened differently from small capacitors to withstand the application of a large amount of energy connected with miniscule source impedance, which is the situation of a capacitor bank. That said, the capacitor design engineer can tailor the design of a large or small capacitor to the application of his capacitor in a large capacitor bank. But the larger capacitor has an inherent advantage in this regard. Wearout lifetime ratings for large and small capacitors are about the same. So, in summary, a price is paid for high-energy banks with regards to reliability, but the wearout lifetime is essentially energy-inelastic.

Reliability Modeling

Even when manufacturing processes and materials and screening methods are state-of-the-art, random failures will still occur in the field. Classical methods to predict reliability of wet aluminum electrolytic capacitors are MIL-HDBK-217 and Bellcore. We are familiar with MIL-HDBK-217 and use it often, but we consider the magnitudes of the FIT rates it predicts at moderate temperatures to be obsolete, as are life models of aluminum electrolytic capacitors from the era when the handbook was compiled and written. Interestingly, the factors of improvement in life and reliability are about the same. The MIL Handbook basically starts from the maximum rated conditions then states that the incidence of random failures diminishes by half for each 20% drop in applied DC voltage and for each 20 ºC drop in core temperature. What we have found is that the predicted failure rate at elevated temperature for small capacitors is somewhat high, usually by a factor of 2–10, but for large capacitors the MIL Handbook is too large by more than a factor of 10. This is in part because of incorrect scaling factors. For example, the failure rate of a capacitor is in actuality related strongly to its energy storage, which is proportional to the square of the rated voltage, and yet in the MIL Handbook there is no reliability factor for the rated voltage, only for the capacitance. This capacitance only factor ($C^{0.18}$) is not correct because it ignores the effect of voltage rating.

We have also found that reliability approximately doubles every 10 ºC, not every 20 ºC. We do not disagree with the voltage derating factor, but we choose to use a cube-law instead.

What all this means is that for a typical 45–65 ºC application of large electrolytics, the MIL Handbook may indicate a failure rate that is over 100 times larger than for actual CDE capacitors. We believe the reasons are possibly due to excessive conservatism in the military handbook’s temperature and voltage factors, along with real improvements in the reliability of our aluminum electrolytic capacitor performance over the past 30 years, associated with some of the same advances in lifetime performance that we have already discussed. One might argue that production yield is related to field reliability, and in that case we note particularly the yield improvements in large capacitors (greater than 300 joules) with voltage ratings of 400 Vdc and higher.

Reliability Data

We rely on customer field failure rate data because from an economic and practical standpoint, it would be impossible to acquire sufficient up-to-date reliability information from our laboratory to assess our present levels of reliability performance in the typical derated conditions that our customers' equipment experiences in the field. This is because billions of unit-hours would be needed, requiring about a million units and a staff of 50 people. Instead, we have had detailed discussions with some of our major customers regarding their applications (bank quantity, temperature, ripple current, applied voltage) and field-return history. In some cases, it was impossible to tell if some of the failures were caused by a capacitor failing, or whether there was another cause, such as an IGBT shorting, causing the capacitor to be destroyed. In those cases, we assumed 50% liability for purposes of reliability estimation. This effort resulted in our acquiring tens of billions of unit-hours data in a variety of applications to give us a good baseline confidence in the 45–75 ºC core temperature range where we most needed data. Most of the field data experienced FIT rates of 0.5–20 from 2–3 inch diameter high-voltage capacitors in multiple-capacitor (2–24 caps) banks. We combined this wealth of field data with observations from our lab regarding the effects of capacitor size, design criteria (basically rated life and temperature), core temperature, and voltage derating. Our combined QA and Engineering labs contain about 60 ovens and 170 power supplies, of which usually 80% or more are in use. We also took into account some published prior studies and the small number of published reliability models for aluminum electrolytic capacitors. We used all of these sources of information together to develop an empirical best-estimate reliability model of our present capacitor reliability.

Scalability issues

What we found with regard to performance of various capacitor types and sizes was that:

1. Premium grade capacitors are more reliable,
2. Equal energy storage results in equal reliability.
3. For a given capacitance, lower voltage ratings are more reliable,
4. For a given voltage, lower capacitance is more reliable,
5. For a given CV rating, lower voltage is more reliable.
6. For a given size, lower voltage ratings are more reliable because capacitor volume is nearly proportional to $CV^{1.5}$.
7. As we said for a given bank stored energy, fewer capacitors give higher reliability,

8. The rate of random (pre-wearout) failures of a bank of capacitors is equal to the number of capacitors times the failure rate of a single capacitor.
9. Failure rate doubles every 10 °C hotter the capacitor core is operated, and
10. Failure rate is proportional to the cube of the ratio of the applied to rated voltage.

CDE's Reliability Model

We have developed a semi-empirical model for the FIT rate λ for pre-wearout random failures that satisfies the 10 criteria outlined in the preceding section.

$$\lambda = 400NVa^3C^{½}2^{(Tc\text{-}Tm)/10} ⁄ (L_BVr^2)$$

λ	failure rate, FIT, ppm/kh
N	number of capacitors
Va	applied voltage, V
C	nominal capacitance, μF
Tc	actual core temperature, °C
Tm	maximum core temperature, °C
Lb	base life at full rated conditions, h
Vr	rated voltage, V

This basically means that for a given capacitor family (type, such as CDE 550C), temperature, and voltage derating, the failure rate is proportional to number of capacitors N, the product of the rated voltage Vr and the square root of the capacitance C. All of these variables have been covered in the life equation except for the nominal capacitance C (μF) and the number of capacitors in the bank, N. The FIT rate λ can then be used as the parameter in the Exponential Distribution, which we will discuss in the next section.

Reliability Model for Pre-wearout Mortality

The manner in which the occurrence of failures is distributed in time determines the appropriate failure distribution. There are many standard distributions, some of which are just special cases of others. The two distributions we will be using are the Exponential Distribution with the λ parameter just discussed for the random failures and the Normal Distribution for wearout failures. The shape of a probability distributions depends on whether one is looking at the normalized instantaneous failure rate of the initial population, known as the Probability Density Function (PDF) f(t), the total area under which is normalized to equal unity; the proportion of units from the initial population that have failed, which is the integral from time t=0 to the plotted time coordinate of the PDF, in other words the area under the PDF curve, known as the Cumulative Distribution Function (CDF) F(t); or the instantaneous failure rate of the surviving population, known as the Hazard Function (HF) h(t). For any distribution,

$$h(t) \equiv f(t) / [1-F(t)]$$

The nature of the Exponential Distribution is that its hazard rate has a constant value of λ. This means that the initial failure rate is λ and even after units fail, the failure rate of the remaining units retains the same λ value. Thus we have for the pre-wearout failures a probability density function

$$f_{exp}(t) = \lambda e^{-\lambda t}$$

and a cumulative distribution function

$$F_{exp}(t) = 1 - e^{-\lambda t}$$

leading to a hazard function

$$h_{exp}(t) = \lambda = \text{FIT Rate (ppm/kh)}$$

Examples of FIT rate and MTBF

The Mean Time Between Failures is the reciprocal of the FIT rate,

$$\text{MTBF} = 1/\lambda = 1/\text{FIT (Gh)}$$

and it is straightforward to deduce its units and meaning. See table on next page for a summary of common units. The units of FIT are ppm/kh, which is failures per billion unit hours, or nanofailures per unit hour. When we reciprocate this unit, we obtain the units of gigahours. The main way to keep the dimensions straight is to use the units if they are given, and if they are given in FIT, remember that the FIT is expressed in ppm/kh so its reciprocal is billions of unit hours.

Failure Rate and MTBF Unit Conversions

λ FIT (ppm/kh)	λ $f / 10^6 h$	λ (%/kh)	$1/\lambda$ MTBF (kh)
0.1	0.0001	0.00001	10,000,000
1	0.001	0.0001	1,000,000
10	0.01	0.001	100,000
100	0.1	0.01	10,000
1,000	1	0.1	1,000
10,000	10	1	100

The Large Capacitor Bank Exception

While the above failure rate model accurately predicts the failure rate for most of Cornell Dubilier's customers' aluminum electrolytic capacitor banks, it fails to accurately predict failure rate for large capacitor banks and in smaller capacitor banks intended for high-peak current applications like welders without especially designed capacitors. As an illustration the model predicts a failure rate of 3.2% per year for twenty CDE Type

DCMC 4700-µF, 400-V capacitors operating at a 65 °C core temperature with 400 V applied. However, the actual customer was seeing capacitor bank failures of almost 10% per year, all of them blowout short circuit failures. The customer's laminated bus structure connected the capacitors in parallel creating a source impedance capable of delivering thousands of amps of peak current. We discovered that when a small burst of dc leakage current tries to heal a weakness in the anodic oxide, the low source impedance can deliver a runaway surge; local heating can drive the current up and cause a blowout short circuit typically at the beginning edge of the anode foil or other spot with poor anodic oxide.

The corrective action turned out to be multifaceted. To assure continuous coverage of anodic oxide the orientation of the capacitors during aging was adjusted to assure that electrolyte contacted all of the cut edges of the anode. The aging voltage was permitted to drift up to well over the rated voltage. A special separator pad between the anode and cathode foils was developed that restricted the peak current without significantly increasing the capacitor's ESR. And the electrolyte and anode formation voltage were optimized to assure overvoltage transient withstanding capability.

When an overvoltage approaches the anodic formation voltage, the dc leakage current begins to increase rapidly. The resulting voltage division with the source impedance causes the applied voltage to flat top like a zener diode. This mechanism can be used to prevent scintillation, the electrical breakdown and sparking of the electrolyte. To prevent scintillation we require that the electrolyte's scintillation voltage be at a higher voltage than the flattop voltage.

By changing the electrolyte and paper pad and changing the aging process we were able to reduce the example customer's capacitor failure rate from nearly 10% per year to less than 1%. And to assure control of the process and continuing low failure rate we implemented a 100% Hot Surge Voltage screen and an Overvoltage Transient sample test.

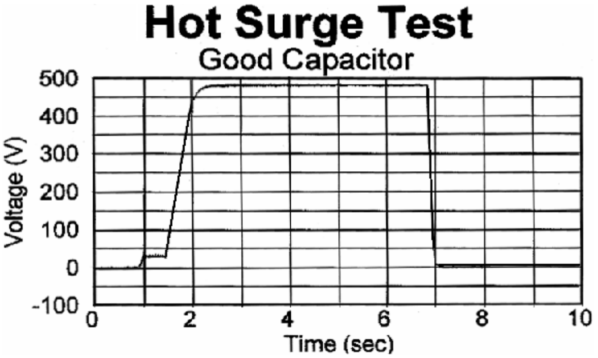

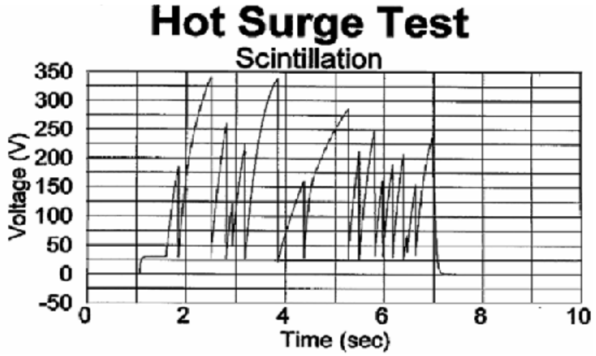

Hot Surge Test

The Hot Surge detects weaknesses in the anodic oxide. Parts are stabilized at 85 °C in an oven, removed from the oven and connected to a 3 amperes constant current power supply. This provides a voltage ramp that is monitored to detect any downward change in the voltage as this would indicate a partial discharge therefore an oxide fault. We credit the our steady implementation of capacitor design and process improvements allowing higher yields on this test as the principal reason that the production yield of our large high-voltage capacitors has increased from 92% in 1990 to over 98% today.

Overvoltage Transient Test

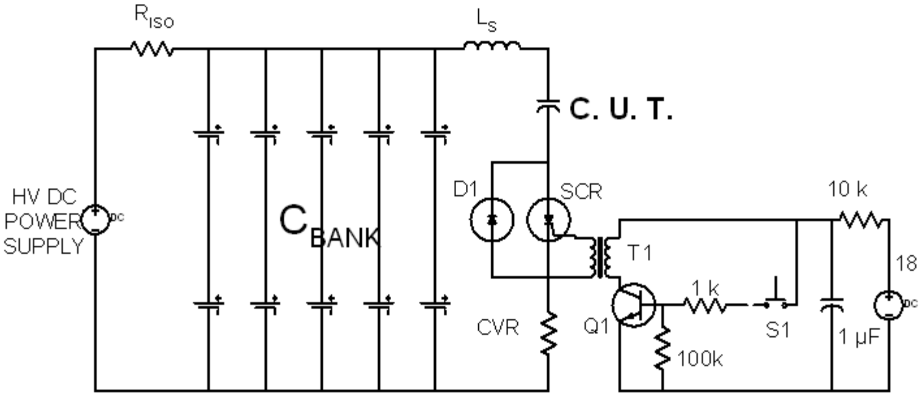

The elements of the overvoltage transient test are a capacitor bank charged to a voltage above the formation voltage of the capacitors to be tested, an inductor acting as a current source and a thyristor switch to dump the energy into the capacitor under test. The test was originally developed to demonstrate that the capacitor banks in the traction drives of electric trains could withstand the overvoltage surge that would occur if one of the cars were disconnected from the power rail and the rail inductance were to drive the total rail current into the remaining cars.

The test has proven effective in demonstrating the robustness of power bus capacitors and their ability to withstand transients in service. Where the first capacitor we subjected to this test in the early 1990s only survived 2 shots, Cornell Dubilier's standard production bus capacitors typically withstand more than 50 shots. In the case of the capacitor develop for the mentioned customer that was optimized for the test, we have not been able to get it to fail. It has been subjected to tens of thousands of shots, and it is still good.

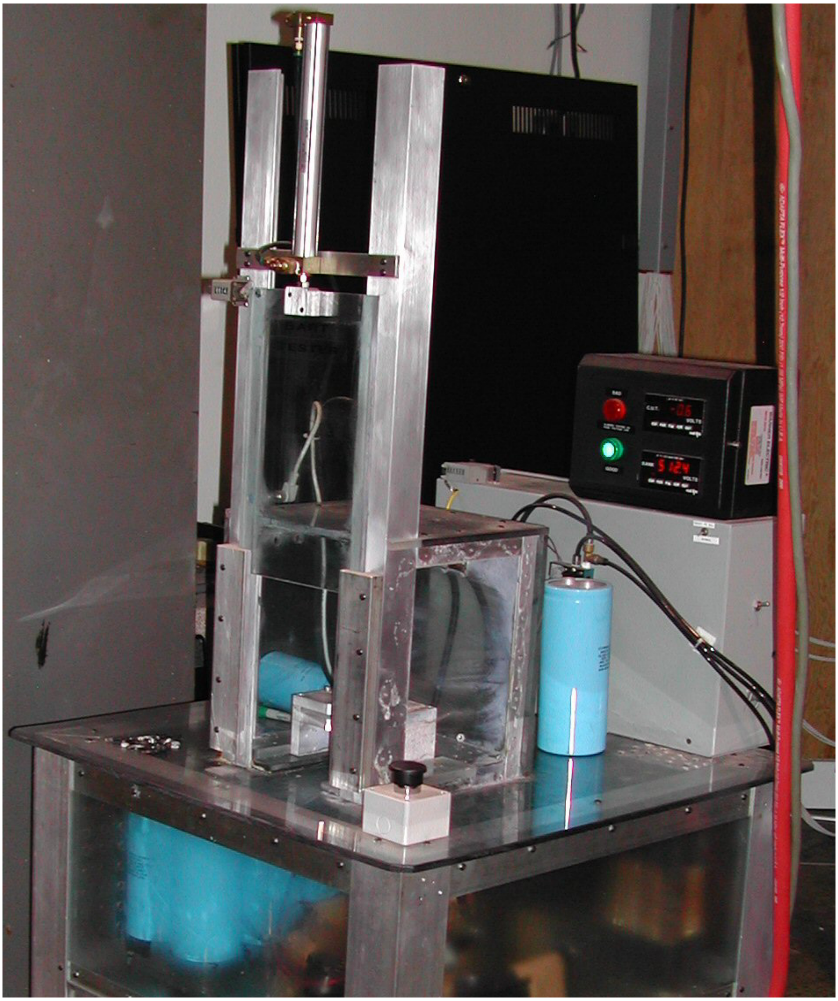

This is the production test fixture for the Transient Overvoltage Test. The capacitor bank is on the shelf below the test fixture. The capacitor under test's voltage is monitored to assure that it charges with each shot and has not failed short, and a counter keeps track of the number of shots.

Overvoltage Transient Test, Successful Waveform

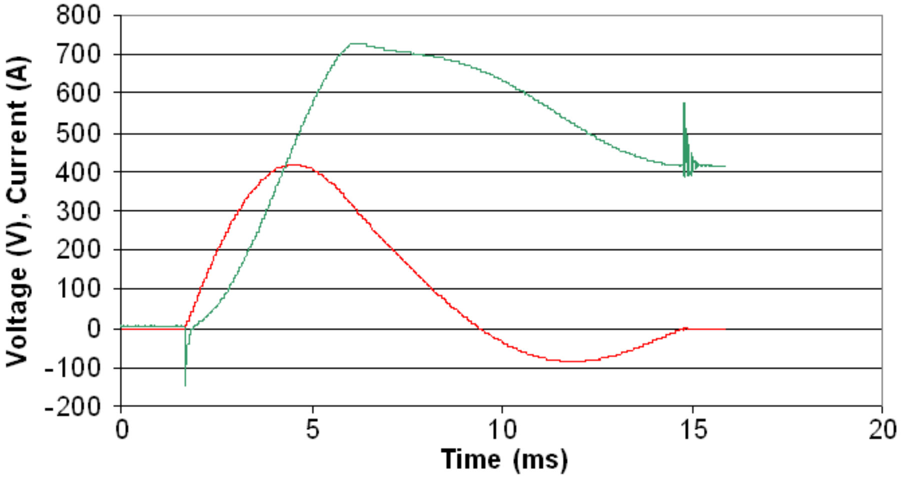

Overvoltage Transient Test, Failed Waveform

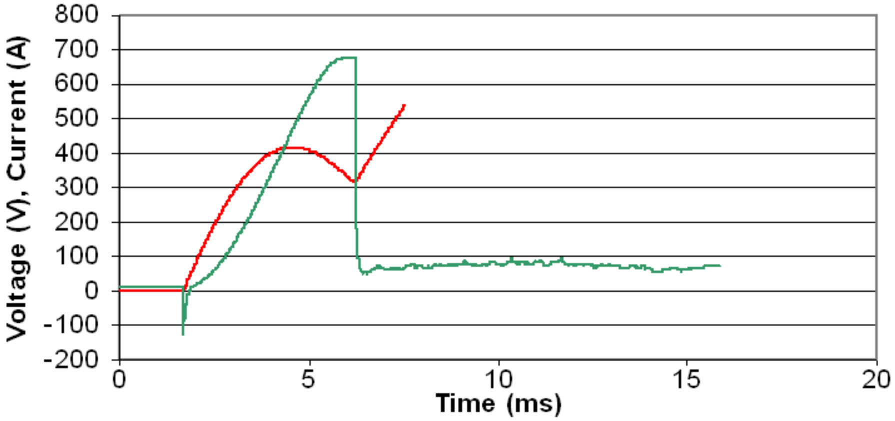

Conclusion

Overvoltage transient testing is an effective way to show the efficacy of aluminum electrolytic bus capacitors for use with very low source impedances. Adequate anode foil formation voltage and the right paper separators and electrolyte can achieve excellent results and greatly reduce the likelihood of short-circuit failure’.