

Electrochemical Capacitor Technology Basics for the Traditional Component Engineer

John R. Miller
JME, Inc.
17210 Parkland Drive
Shaker Heights, OH 44120
216-751-9537 phone/FAX
JMEcapacitor@att.net

Basic Technological Differences

Comparing electrolytic and electrochemical capacitors, the most obvious difference between the two is the very high capacitance of electrochemical capacitors, a capacitance often measured in units simply unheard for electrolytic capacitors. A second, and contrasting difference, is that electrochemical capacitors all have very low voltage, the highest-rated cell being at only 2.7 V. Electrolytic capacitors, on the other hand, may have much lower capacitance but operate at voltages ranging from only a few volts to more than six hundred. Stark differences in capacitance and voltage are what stand out immediately between the two technologies.

Furthering this comparison is helped by looking more deeply at the published specifications for the two technologies. Specifications for electrolytic capacitors will often include a working voltage, a maximum ripple current, an operating temperature range, a measure of capacitance (usually set at 120 Hz), and, quite often, an equivalent series resistance (ESR). In some cases a dissipation factor will appear, listed at, again, 120 Hz. Specifications for electrochemical capacitors will be markedly different, usually including a capacitance, an ESR value, a temperature range, sometimes the cycle life, and perhaps a leakage current, but nothing at all about either ripple current or dissipation factor.

Looking beyond specifications to the operational conditions of the two technologies, a still more striking difference appears in that while some electrolytic capacitors are able to operate effectively at 125 C, the maximum rated voltage for electrochemical capacitors will be at ~70 C. In regard to the low-temperature end, both technologies operate well at very low temperatures, some devices of course better than others, but all generally down to -40 C.

Differences like these have often made it difficult for component engineers to determine a proper use for the more recent electrochemical capacitors. The goal of this paper is to clarify electrochemical capacitor technology in such a way that it can be better, more appropriately, and most optimally exploited. To do this, we will attempt to put the two technologies on an equal footing in order to make the fundamental differences between them in both construction and operation as clear as possible. This will hopefully provide component

engineers familiar with traditional electrolytic capacitors a clearer path to optimally exploiting electrochemical capacitor technology.

Structural Fundamentals

Going back to basics, consider the classic parallel-plate capacitor, shown in Figure 1. A capacitor, as we see here, is a device that physically stores electrical charge by moving it from a first plate to a second plate, creating an electric field between the plates. It is between the plates that the energy is stored. The charges, of opposite signs but with identical magnitude on each plate, sit and attract each other, able to store energy with small leakage current virtually indefinitely. Because this charge separation is only from one plate to the other, cyclability represents no problem, any change involved being physical rather than chemical, restricted to moving charge from one plate to another. The phenomena related to other storage media like batteries are largely irrelevant here.




Figure 1: Simple capacitor formed by separating two parallel conductors by a distance d . A dielectric material is commonly placed between the conductors to increase the capacitance.

The capacitance of a device like this will be proportional to the plate area divided by the separation between the plates, times a dielectric constant, which latter is a property of the material between the two plates. This constant would be 1 for a vacuum and can be much larger for some materials, primarily polymers for which the constant might be in the range of 2–3. For some of the metal oxides like aluminum oxide, of course, the constant could be 8–10, while for tantalum oxides it might be 27. Very different values of the dielectric constant are clearly possible. Again, charge stored will be affected by three factors: plate separation, plate area, and the dielectric constant of the material between the plates.

Energy stored in a capacitor device is equal to the capacitance C divided by two, times the voltage V squared:

$$U = \frac{1}{2} C V^2 \quad (1)$$

This serves to remind us of the value of higher voltages: as the voltage increases, energy increases proportional to the voltage squared. A look at the energy density or energy-per-volume for such a parallel-plate capacitor shows that it is proportional to the voltage divided by the plate separation, that quantity then squared, times the dielectric constant:

$$U/Ad = 0.5 \epsilon \epsilon_0 (V/d)^2 \quad (2)$$

The maximum energy density will thus be achieved when the V/d value is at its maximum, and that value is really the breakdown strength of the dielectric, the material between the two plates. The metric for looking at the energy stored in a capacitor is precisely that breakdown strength squared times the dielectric constant.

Understanding Electrolytic Capacitors

Turning to the basic structure of an electrolytic capacitor, it is comprised of two series-connected capacitors, the two being electrically connected by a liquid electrolyte. The capacitor attached to the positive terminal of the device generally is made of aluminum foil that has been anodized to create a dielectric oxide on both its faces. The capacitor attached to the negative terminal of the device is also likely to be aluminum, except that the dielectric covering this foil is much thinner, causing it to have a much higher capacitance value. An “electrolyte,” defined as an ion conductor and electron insulator, is introduced in order to series connect these two capacitors. The dielectric covering the aluminum is generally quite fragile, and in order to make practical devices such an electrolyte is used. A genuinely practical electrolytic capacitor will thus be comprised of two capacitors in series. The capacitance C_T of two capacitors in series is calculated using the equation:

$$1/C_T = 1/C_+ + 1/C_- \quad (3)$$

where C_+ is the capacitance at the positive terminal and C_- is the capacitance at the negative terminal. If we make the negative terminal capacitance very, very large compared to the positive terminal capacitance, we will then realize essentially all the positive terminal capacitance, i.e. $C_T = C_+$. The positive terminal capacitor will have essentially all of the voltage standoff of the device, and the energy will be stored in the dielectric of this electrode. The negative terminal capacitor provides a second connection, allowing these series-connected capacitors to be charged and discharged.

Special design features are used to create electrolytic capacitors. Instead of using a planar surface, for instance, the surface is generally highly-etched aluminum. This provides a larger area for the oxide film to grow on, thus greatly increasing the plate area. Anodization is performed by applying a DC voltage to the aluminum foil, causing an oxide to grow on its two surfaces. In general ~ 1 nm of oxide will be grown per volt applied. Applying 100 V to the aluminum will eventually form an oxide dielectric of ~ 100 nm thickness on its surface. It is important to note, however, that this oxide will be of lower density than the aluminum metal. The volume of the structure as a whole will increase because of the density difference between the formed aluminum oxide and the aluminum metal.

For low-voltage anodic films, the aluminum foil can have many very-small-diameter pores, greatly expanding the surface area with a thin oxide layer covering this entire surface. For high-voltage anodic films, the etched surface will need a more open structure (with larger effective diameter pores) since the oxide is actually grown by consuming the aluminum and combining it with oxygen from the electrolyte, so much so that the metal/dielectric interface actually moves during the anodization process. There is also, as mentioned above, an increase of volume in the oxide over the metal, so that the pore diameter decreases as the anodic film grew. There will thus be a maximum voltage to use for any given porous structure in aluminum, as well as an optimal voltage.

In the electrolytic capacitor, the thickness of the oxide d on the surface will be proportional to the applied voltage V and further, that the capacitance C is proportional to reciprocal of the thickness of the dielectric. Combining these two equations, $d \sim V$ and $C \sim 1/d$, we find that the product $C \cdot V$, capacitance times voltage, will to first order be a constant. Because the CV product is approximately fixed for a specified electrolytic capacitor design, the voltage rating can be increased by growing the anodic oxide thicker, but not without decreasing capacitance.

The energy density of the capacitor, proportional to V^2 per equation 1, with $C \cdot V = \text{constant}$ means that the energy in an electrolytic capacitor will be proportional to the voltage, to first order. This relationship holds quite well for low voltage dielectrics. Since the ultimate applicable voltage is really the breakdown strength of the dielectric material, maximum energy density will be at the point of operation near the breakdown potential of that material. For electrolytic capacitors, one would expect 500–700 V devices, for instance, to have much higher energy density than 4–6 V ones constructed using the same physical design.

Again, in electrolytic capacitors the aluminum foil will generally not be just a planer surface but will have been etched to increase surface area, with the etching process dependent upon the operating voltage for which the device is being designed. Aluminum foil for low voltage devices will have much smaller pores etched into it, creating a highly etched, porous surface as shown in Figure 2.

For electrolytic capacitors, the capacitance of an etched and formed foil divided by the capacitance of the same size planer foil is referred to as “gain.” This ratio, the gain, can be above 200 in some cases, meaning that there may be more than 200 times higher capacitance in the etched foil than in the non-etched planer foil, the direct result of well-engineered etching of the material. This could also be described by saying that for every real cm^2 of foil, 200 cm^2 of real area has been created by the etching. Thus, the capacitance is increased by ~ 200 fold because of the etching.




Figure 2: Scanning electron microscope view of cross section of aluminum foil that has been etched on both faces (with expanded view on right). This creates two high-surface area porous surfaces connected to a solid center. The surfaces are then anodized to fully cover the aluminum with capacitor dielectric.

The operating voltage for an electrolytic capacitor is generally somewhat less than the anodization voltage. For devices now available commercially, this means anodizing the foil at perhaps 400 V to create an electrolytic capacitor with a rated working voltage of 300 V. Structurally the device consists of two foils with two paper separators, wound together. One foil is the positive electrode connected to the positive capacitor C_+ and the other is the negative electrode connected to the negative capacitor C_- . The separator between the anodized foils is wetted with an electrolyte that series connects the two capacitors. This results in very large area plates, particularly in large-format devices the size of soft drink cans. Tightly wound, with the separators plus foils perhaps only $250\ \mu\text{m}$ thick, a considerable amount of formed aluminum can be packed into a package. The size of a typical 1 F, 15 V rated aluminum electrolytic capacitor is a right cylinder 7.5 cm in diameter by 22 cm tall.

Understanding Electrochemical Capacitors

In an electrochemical capacitor what corresponds to the “plate” in an electrolytic capacitor is actually created naturally at the interface of an electrolyte and an electrode. To illustrate how charge is stored this way, consider a very basic electrochemical capacitor constructed by putting two electrodes into an electrolyte, using, for the sake of concreteness, the example of two carbon rods immersed in salt water as shown in Figure 3.




Figure 3: Electrochemical capacitor formed by placing two electron conductors in an electrolyte and applying a voltage (left). Charge separation occurs at the solid-liquid interface of both electrodes and persists after the voltage source is removed (right), creating two capacitors in series.

As voltage is applied between the two carbon rods, charge flows from one electrode to the other, and a balancing charge accumulates at the surface of the rod on the electrolyte side. This is an equal but opposite charge in the rod that is attracted to its surface. There is no electron transfer, provided that the applied voltage is below the breakdown potential of the electrolyte.

It is in fact the breakdown potential of the electrolyte that limits the voltage that can be applied between the two electrodes. For an aqueous device, this might be 1.2 V, while for some non-aqueous electrolytes it might be >3.0 V. The breakdown potential of the electrolyte itself constitutes a critical and inherent limitation on the operating voltage of electrochemical capacitors. It is this that makes it necessary, in order to create devices for practical use, to series connect electrochemical capacitor cells in the same way battery cells are series connected in order to create higher voltage operation. Thus, if 2.5 V is the working voltage of a cell, it will take ten such in series to create a 25 V capacitor. As it happens, such higher-voltage electrochemical capacitors are generally created by series connecting individual cells externally, i.e. assembled as a module.

When the switch for the right-hand beaker shown in Figure 3 is opened, charge will persist on each electrode. In fact, two capacitors in series have actually been created. Details of the charge storage with an aqueous electrolyte are shown in the electric double layer model of Figure 4. The thickness of the dielectric at the interface is of very small dimensions indeed:




Figure 4: Model of a charged electric double layer at a solid interface for an aqueous electrolyte. The thickness of this region is typically less than 1 nm, thus creating a very small “plate” separation.

As illustrated, the ions in the aqueous electrolyte are attracted to the surface of the solid but are prevented by a solvent layer from reaching it to effect an electron transfer. This layer would be constituted by the water molecules up against the surface, and the distance of closest approach by the ion would be on the order of ~ 1 nm. Charge accumulates on the solid surface, attracted by an equal but opposite charge in the electrode itself. The two charges together constitute the so-called the “electric double layer,” one layer being in the electrolyte and the other being just below the surface of the solid. Separation is on the order of $1 \text{ nm} = 10 \text{ \AA}$, and the area for a planer surface would $1 \text{ cm}^2/\text{cm}^2$ for a gain of 1, using familiar terms from electrolytic capacitor technology. In the aqueous system illustrated here there would be a dielectric in the electric double layer with a dielectric constant of ~ 6 . This is much lower than the bulk value, which for water would be ~ 80 , due to orientation effects. The final result is electrostatic charge storage in the electric double layer that has a plate separation of ~ 1 nm with a dielectric constant of ~ 6 .

In practical electrochemical capacitors there is no planer surface, but rather a high-surface-area electrode, activated carbon being the most popular material for that electrode. The surface area of an activated carbon will depend upon how the carbon has been activated, but it can in fact be greater than $1500 \text{ m}^2/\text{g}$. Electrochemical capacitor activated carbon is a very highly engineered material, its pore-size distribution, which needs to be matched to the ion size of the electrolyte, being a most important feature. There needs to be a good relationship between these two feature sizes in the initial design so that electric double layer charge storage is set up over as much of the available surface area of the carbon as possible. Generally, however, more surface area is not something to strive for without recognizing that at some point a higher surface area will mean smaller pore sizes that, owing to their diminished physical size, will not be accessible to the ions. A device constructed with such a

material can have a very much larger surface area than one with a planer surface could possibly have. With a plate separation of 1 nm and a surface area potentially as great as 10^3 m²/g, very substantial capacitances of up to hundreds of Farads per gram of activated carbon material become possible.

In practice electrochemical capacitor technology, just like electrolytic capacitor technology, requires two electrodes, as shown in Figure 3. These will generally be identical, activated carbon of the same mass, and thus the device capacitance will be approximately one-half the capacitance of one of the electrodes. Devices intended for actual use do not, of course, leave spaces between the electrodes as the Figure 3 illustration depicts, but rather activated carbon particles with a binder are pasted directly onto the surface of a lightly-etched aluminum foil. The thickness of the activated carbon coating the aluminum can be perhaps 100 μ m.

Using terms usually applied to aluminum electrolytic capacitors, this activated carbon coating can produce a capacitance of ~ 0.5 F/cm², in contrast to the electric double layer capacitance of a planer surface, which typically has ~ 10 μ F/cm². Thus the gain in this case would be 50,000. This is, of course, starkly different from the ~ 200 gain possible for an etched aluminum low-voltage electrolytic capacitor. Thicker carbon electrodes, for instance 1 mm thick, will obviously produce even higher gains, and this does appear in some device designs.

The characteristic differences in gain and dielectric thickness just described accounts for enormously different capacitance values possible between electrolytic and electrochemical capacitors. The fact that electrochemical capacitors have two capacitor elements in series and that they are nearly identical in capacitance value prevents the device from operating at voltages greater than the breakdown potential of the electrolyte. Electrolytic capacitors, on the other hand, have two capacitor elements in series with vastly different values, causing essentially all of the voltage in a device to be on the one electrode that is covered by the oxide dielectric. This accounts for the differences in operating voltage of a device. And finally, the difference in the dielectric material structure accounts for differences in operating temperature. The typical electrochemical capacitor has a liquid dielectric that experiences ever increasing instability as its temperature approaches the boiling point. That is why electrochemical capacitors have ratings of ~ 70 C or less. In sharp contrast the dielectric of aluminum electrolytic capacitors is a solid metal oxide that is totally stable to much higher temperatures, which permits them to operate reliably to temperatures of 125 C or higher.

Points of Comparison

One way to characterize a capacitor is by its natural response time, which for an ideal series-RC circuit, is the mathematical product of R times C. For electrolytic capacitors the RC product is generally ≤ 1 ms. This means that, assuming an ideal circuit, during that response time a shorted capacitor will discharge to $1 - 1/e = 0.632$ times its original voltage. If, for purposes of concreteness, we consider an electrolytic capacitor with RC=0.3 ms charged to 100 V, once it has been dead shorted its voltage will drop to 63.2 V during 0.3 ms, 63.2% of

its original value, which very much illustrates the dynamics of an electrolytic capacitor as a storage medium. Using equation 1, we see that the capacitor delivered $(100^2 - 63.2^2)/100^2 = 0.601$ of its stored energy. Thus the capacitor, after being shorted, discharges ~60% of its energy during its response time. It cannot discharge any faster.

Electrochemical capacitors, on the other hand, are not well represented by a single response time. They have a natural response characteristic of a multiple-time-constants device, which has to do with the high-surface-area carbon that is used to create the electrodes. Because of this, the charge stored within the device has been essentially distributed over a very broad area, too wide for all that stored charge to be accessed in the same time. Nevertheless, the RC-product can be used to approximate the response time of an electrochemical capacitor and it generally is on the order of ~1 s for many commercial products. Thus, it will take at least ~1 s to charge or discharge an electrochemical capacitor, in comparison to an electrolytic capacitor, which requires only ~1 ms.

This serves to explain quite clearly why electrochemical capacitors will not filter 60 Hz ac ripple voltage, one of the major and obvious differences between the two technologies in regard to potential applications. Electrolytic capacitors are used primarily for filtering on a DC bus, to reduce the ripple voltage involved when ac power is rectified. This is the reason why an allowed ripple current is so often listed for electrolytic capacitors in their product literature. Traditional component engineers, for whom the similarities between these two technologies may not be so apparent, can easily find themselves frustrated by the fact that electrochemical capacitors refuse to filter. Concluding that “This dog won’t hunt,” however, should not be the end of the matter for them.

More closely examined, are these differences quite what they seem to be in regard to the applications for which each technology is best suited. Consider in this connection two capacitors, one electrolytic and the other electrochemical, as shown in Figure 5. These devices are about the same physical size but with vastly different ratings, 350 F versus 0.047 F, and 2.7 V versus 25 V.



Figure 5: Photograph of a 350 F, 2.5 V, 3.2-cm-diameter electrochemical capacitor (left) and a 0.047 F, 25 V aluminum electrolytic capacitor (right). The devices are approximately the same physical size but the electrochemical capacitor can store ~75 times more energy than the electrolytic capacitor.

Consider the impedance plots of these two capacitors, shown in a complex-plane representation in Figure 6, and in a Bode representation in Figure 7. First note that the electrochemical capacitor has about one-half of the equivalent series resistance of the electrolytic capacitor, about $3\text{ m}\Omega$ versus about $7\text{ m}\Omega$ for the electrolytic capacitor.

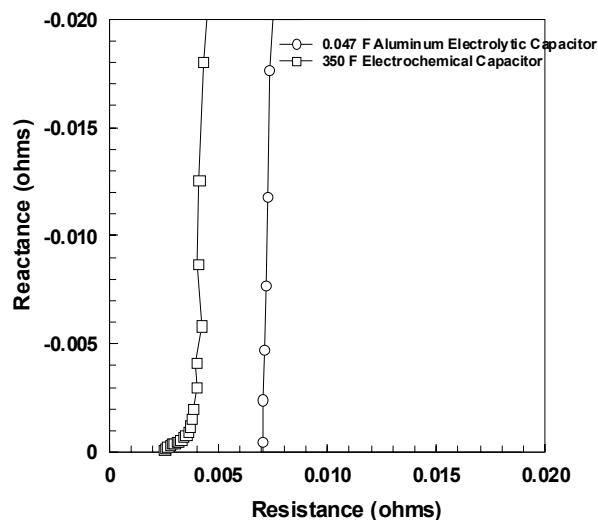


Figure 6: Complex-plane plot of the impedance of the two capacitors shown in Figure 5. Note that the equivalent series resistance of the electrochemical capacitor is about one-half the value of the electrolytic capacitor. The $\sim 45^\circ$ intersection of the electrochemical capacitor data with the real axis represents classical porous electrode behavior and is due to the use of high-surface-area electrode material.

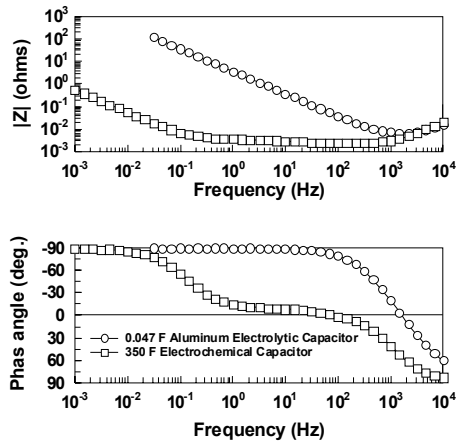


Figure 7: Bode plot of the impedance data from the two capacitors shown in Figure 5. The upper figure is the magnitude of the impedance versus frequency, and the lower figure is phase angle versus frequency. Note that both have -90° phase angles at low frequencies, and inductive behavior (positive phase angle) at high frequencies.

The electrolytic capacitor has a self-resonance frequency (phase angle of 0°) of ~ 2 kHz, while the electrochemical capacitor value is well below 100 Hz. Clearly despite this difference there are still some similarities. These two different types of capacitors demonstrate similar behavior at very low frequencies; both have phase angles of approximately -90° . The electrolytic capacitor reaches this value at about 100 Hz while for the electrochemical capacitor it is at about 0.01 Hz. So the two devices show similar behavior except that the electrochemical capacitor's phase angle is shifted to the left to a location that is 10,000-times lower in frequency than the electrolytic capacitor.

Relating this difference in impedance characteristics to a practical application, in a hybrid vehicle with a braking time of ten seconds, charging the capacitor would occur over roughly a ten-second time. If the vehicle were then accelerated in a similar time, the capacitor would be discharged over another ten seconds, the period totaling some twenty seconds. Under repetitive start/stop operation this charge/discharge cycle roughly corresponds to a frequency of $1/20 \text{ s} = 0.05 \text{ Hz}$. Thus, important operational frequencies in this application are obviously below 0.1 Hz, where the phase angle does approach -90° as shown in Figure 7, i.e., where the electrochemical capacitor shows nearly ideal capacitive behavior. In this hybrid vehicle example, the electrochemical capacitor is in effect filtering voltage ripple in the vehicle's electrical system exactly like an electrolytic capacitor filters voltage ripple in rectified 60 Hz power in a power supply, the only difference of note being the "ripple" frequency.

Equivalent circuit models can be derived using complex impedance data. The electrolytic capacitor is well represented by a series-RC circuit model with element values as shown in Figure 8. The electrochemical capacitor is poorly represented by such a model--a more accurate representation can be obtained by using a three-time-constant truncated ladder network, as shown on the right in Figure 8. Such a multiple-time-constant model is needed because of the distributed resistance and distributed capacitance brought about from the use of high surface area electrode materials.




Figure 8: Equivalent circuit models for the two capacitors shown in Figure 5. The series-RC model (left) well represents the aluminum electrolytic capacitor. Notice that the RC product is ~ 0.3 ms. The right circuit is a three-time-constant model derived for the 350 F electrochemical capacitor. The response times of this equivalent circuit model are ~ 0.25 s, ~ 0.34 s, and ~ 0.87 s.

Discussion

What I have attempted here is to describe how electrolytic and electrochemical capacitors operate in order to point out differences and similarities between the two technologies. Both physically store charge, and both have relatively high capacitance compared to that obtainable on a planer surface. The electrochemical capacitor electrode gain is generally much much higher than that of an electrolytic capacitor. This really is what explains why their frequency responses are so different, why electrochemical capacitor performance is so “sluggish” compared to that of an electrolytic capacitor. Added to this are energy density differences that arise directly from the thickness of the dielectric and its surface area. The electrochemical capacitor provides very high capacitance, generally 100-times or even higher capacitance per unit volume than available from electrolytic capacitors of comparable voltage.

Indicated earlier as well, there are major differences in operating temperatures to be considered. Electrochemical capacitors store charge in the electric double layer, and that electric double layer will have leakage current. There is a finite probability of electron transfer across the double layer, a probability affected by the absolute temperature of the device: the higher the temperature, the higher the probability of an electron transfer. For most electrochemical capacitor devices, the logarithm of the leakage current is proportional to the applied voltage. There is thus a limit on the voltage at which an electrochemical capacitor can be operated. This will generally be well below the breakdown potential of the electrolyte, due to the leakage current limitations. This can be quite a practical matter, in that electron transfer often results in gas generation, a potentially life-

limiting event in a sealed device. Accordingly, electrochemical capacitors will generally have some published upper limit on their operating voltage, usually derived from life considerations.

Because of their solid aluminum oxide dielectric, electrolytic capacitors generally do not have leakage current issues. So the temperature limitations that differentiate the two technologies are explained simply by breakdown differences between liquid and solid dielectrics.

It may also be useful to remember that electrolytic capacitors come in a substantial range of different designs. Some are designed for high volume efficiency like in photo-flash applications, while others designed to effectively shed heat are directed specifically at high ripple current filtering applications.

The same considerable range of optimizations made in electrolytic capacitors has likewise been developed for electrochemical capacitors. Some devices have been optimized for very low-rate applications and so have very minimal leakage current. Others have been optimized for high cycle efficiency, with correspondingly lower energy density, not unlike the high ripple current electrolytic capacitors.

Such a range of optimizations makes it necessary to look as carefully as possible not only at the specifications for each product but also at how it is constructed. Devices intended to handle high ripple current will generally have a form factor allowing heat to be dissipated much more efficiently. Electrochemical capacitors may sometimes be structured with terminals on each end while others have both terminals on only one end, each configuration optimized for different uses.

Conclusion

This paper has been intended to further understanding of electrochemical capacitor technology by evoking those aspects of this technology most likely to be familiar to traditional component engineers. Hopefully such an enhanced familiarity will enable greater confidence in turning to electrochemical capacitor technology whenever it may make sense to do so. There are applications where each of these technologies shines, for very different reasons, and the task of the component engineer is to select the most optimum technology for any given system being designed.

At the very least, I hope this paper puts us well beyond a situation that memorably occurred at one of this organization's earlier meetings, when someone in the audience loudly, repeatedly, and even rudely, continued to correct another presenter to the tune of "It's microfarads Lady!" Times have changed. We're talking farads and kilofarads today, get used to it!

Further Reading on Electrochemical Capacitor Technology

Barnett, B.M., Dowgiallo, E., Halpert, G., Matsuda, Y. and Z-i. Takehara (eds.) (1993). *Proceedings of the Symposium on New Sealed Rechargeable Batteries and Supercapacitors*. **93-23**, Pennington, NJ: The Electrochemical Society.

Brodd, R.J., Doughty, D.H., Kim, J.H., Morita, M., Naoi, K., Nagasubramanian, G., and C. Nanjundiah (eds.) (2002). *Proceedings of the International Symposium of Electrochemical Capacitor and Hybrid Power Sources*. **2002-7**, Pennington, NJ: The Electrochemical Society.

Conway, B.E. (1999). *Electrochemical Supercapacitors: Scientific Fundamentals And Technological Applications*. New York: Kluwer Academic.

De Levie, R. (1963). On porous electrodes in electrolyte solutions. *Electrochimica Acta* **8**, 751-780.

Delnick, F.M., and M. Tomkiewicz (eds.) (1996). *Proceedings of the Symposium on Electrochemical Capacitors*. **95-29**, Pennington, NJ: The Electrochemical Society.

Delnick, F.M., Ingersoll, D., Andrieu, X., and K. Naoi (eds.) (1997). *Proceedings of the Symposium on Electrochemical Capacitors II*. **96-25**, Pennington, NJ: The Electrochemical Society.

Kinoshita, K. (1988). *Carbon: Electrochemical and Physicochemical Properties*. New York: John Wiley and Sons.

Korzeniewski, C. and B.E. Conway (eds.) (1997). *Proceedings of the Symposium on the Electrochemical Double Layer*. **97-17**, Pennington, NJ: The Electrochemical Society.

Kamath, H. and Miller, J. R., (2002). "Electrochemical Capacitor Energy Storage" In *Handbook of Energy Storage for Transmission or Distribution Applications*, 1st Ed, Electric Power Research Institute (EPRI), Palo Alto, CA.

Linden, D. and Reddy, T.B. (2002). *Handbook of Batteries*. 3rd ed. New York: McGraw-Hill.

Marincic, N. (organizer) (2003-2008). *Proceedings of International Seminar on Double Layer Capacitors and Hybrid Energy Storage Devices*. (6 vols.) Bonita Springs, FL: Redox Engineering.

Miller, J.R. (chair) (2003-2007). *Proceedings of Advanced Capacitor World Summit* (5 vols.) Portland, Maine: Intertech.

Nazri, G.A., Takeuchi, E., Koetz, R., and B. Scrosati (eds.) (2003). *Proceedings of the International Symposium on Batteries and Supercapacitors*. **2001-21**, Pennington, NJ: The Electrochemical Society.

Raistrick, I.D. (1992). Electrochemical Capacitors. In McHardy, J. and Ludwig, F. (eds.) *Electrochemistry of semiconductors and electronics: processes and devices*, pp 297 – 355. Park Ridge, NJ: Noyes Publications.

Sato, Y. (chair) (2003). *Proceedings of the 2003 International Conference on Advanced Capacitors*, Kyoto, Japan: Electrochemical Society of Japan.

Simon, P., Brodd, R.J., Abraham, K.M., Kim, K.B., Morita, M., Naoi, K., Park, S-G., Srinivassan, V., Sugimoto, W., Zaghbi, K. *Electrochemical Capacitors and Hybrid Power Batteries 2008*, ECS Transactions 16-1, Pennington, NJ: The Electrochemical Society.

Wolsky, S. and N. Marincic (organizer) (1991-2002). *Proceedings of International Seminar on Double Layer Capacitors and Related Devices*. (12 vols.) Boca Raton, FL: Florida Educational Seminars.

