The Growing Importance of Inductance in Tantalum Capacitors

Erik Reed

KEMET Electronics Corporation, 2835 KEMET Way, Simpsonville, SC 29681
Phone: +1-864-963-6300, Fax: +1-864-228-4081
e-mail: erikreed@kemet.com

Abstract

Until recently, no one paid much attention to the equivalent series inductance (ESL) of tantalum capacitors other than to recognize that it was generally higher than that found in MLCCs. Tantalum capacitors were considered good for bulk charge storage at relatively low frequencies, but high frequency performance was thought to be the unquestioned domain of the MLCC. But that is now starting to change as power distribution networks (PDNs) for high-performance microprocessors, especially those used in mobile computing platforms, evolve.   In many instances it has been shown that a very few high-performance, low-inductance tantalum capacitors can outperform competing decoupling solutions that employ a larger number of conventional aluminum, tantalum, and ceramic capacitors.

With the advent of high-bandwidth voltage regulator modules (VRMs), PDNs with almost ideal impedance characteristics to frequencies above 10 MHz can be designed using very few capacitors, as long as those capacitors have moderately high capacitance, very low ESL, and predictable, low ESR. These characteristics are the key to the recent success of tantalum polymer capacitors in an application that was once dominated by MLCCs.

This paper briefly discusses the electrical theory behind the new success of tantalum polymer capacitors in mid-range decoupling applications. The physical design of these low ESL tantalum capacitors is described and some of the challenges and strategies of measuring low-inductance capacitors are discussed.

Introduction

Power distribution networks (PDNs) for microprocessors have the challenging task of delivering energy to a load whose current can change from a few amperes to over 100 amperes in nanoseconds. Moreover, since modern microprocessors operate at very low voltages (near 1 Vdc), the tolerance for voltage shifts due to large and rapid current change is very small (measured in tens of millivolts).

A circuit diagram of a typical microprocessor PDN appears in Figure 1. This figure is inspired by similar diagrams drawn by Larry Mosley of Intel Corporation[1].   The diagram deviates from convention by showing power flow from right to left instead of vice versa. This is done for the special purposes of showing the driving stimulus of the circuit on the left-hand side of the figure (µP load current change) and demonstrating the progressive response of the circuit on a conventional time scale with short response times on the left and longer response times toward the right side of the diagram.

During the first few instants after the load current changes, the required electrical charge is delivered by capacitance incorporated into the microprocessor die. Because this capacitance is relatively small, it can deliver the needed current for only a very short time before its voltage falls too much. Before this can happen, charge from capacitance on the package of the microprocessor arrives, relieving the burden on the die capacitance. As

time continues, charge arrives from the mid-frequency and bulk decoupling capacitors which are located on the motherboard close to the microprocessor. Eventually the current is supplied by the voltage regulator module (VRM).
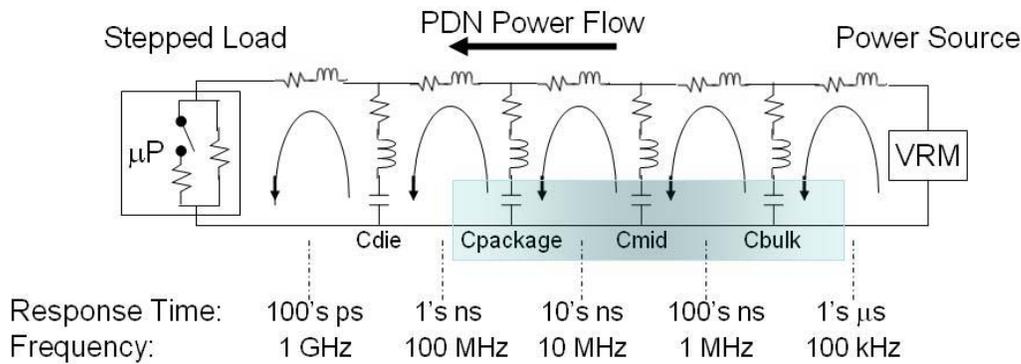

Fig. 1. Typical μP Power Distribution Network

The VRM is usually a DC-DC converter that steps down to about 1V a higher voltage (usually 12V) created by the computer's main AC to DC power supply. The VRM ultimately has the task of regulating the microprocessor's input voltage, but VRMs do not respond quickly to fast load changes. Thus it is the job of the decoupling capacitors in the PDN to hold the voltage steady until the VRM can sense the load change and respond appropriately.

Very similar things happen when the microprocessor current rapidly falls from very high to very low levels. The major difference in this case is that the PDN's job now is to keep the voltage at the microprocessor from overshooting too much during the time it takes the VRM to realize that it doesn't need to supply such high load current anymore.

Winds of Change

Originally, package-level and mid-frequency decoupling were accomplished with multilayer ceramic capacitors (MLCCs), while bulk decoupling was accomplished with wet aluminum electrolytic capacitors. This was a natural division of duty because MLCCs had low equivalent series inductance and resistance (ESL and ESR), while wet aluminum electrolytic capacitors offered huge charge storage in modest volumes, but responded much more slowly because of higher ESL and ESR.

But as microprocessor voltage fell, required current rose, and systems became smaller while providing higher performance, the decoupling landscape shifted and new decoupling solutions appeared. The trend was toward lower ESR and higher capacitance in order to meet the increased performance requirements while attempting to avoid an explosion in the required number of capacitors.

For mid-frequency decoupling, MLCCs remain very popular because of their low ESL, low ESR, and acceptable cost. However, much higher-capacitance devices are now employed to limit growth in the sheer number of capacitors required. In spite of successes in miniaturization, these higher-capacitance MLCCs still tend to occupy more space in the circuit over time, and their cost is rising. Other challenges for high-capacitance MLCCs include lower than expected capacitance due to DC bias application, small AC signal levels, TCC effects at elevated temperature, and aging effects.

A significant change in desktop computer bulk decoupling has been the transition from wet-electrolyte wound-aluminum capacitors to wound-aluminum capacitors which use conductive polymer as a solid electrolyte. This

change dramatically reduced capacitor ESR.   Lower ESR allows acceptable performance with fewer capacitors while maintaining reasonable cost. In notebook computers where space is at a premium, tantalum polymer capacitors and stacked aluminum polymer capacitors are now commonly used for bulk decoupling because of their attractive combination of very low ESR, moderately high capacitance, volumetric efficiency, and low profile. But these capacitors do cost more than their wound-aluminum competitors.

New Low-Inductance Decoupling Solutions

Recently a completely new shift in thinking has appeared regarding bulk and mid-frequency microprocessor decoupling. New solutions are appearing which feature fewer, and sometimes smaller, capacitors.   The driver of this most recent paradigm shift is the desire to control cost, maintain high performance, and minimize circuit volume.

Two facilitators of this paradigm shift are the evolution of VRM technology and the development of very-low-ESL valve-metal capacitors made from aluminum and tantalum.   VRMs are now available with higher switching frequencies and wider control bandwidth, both of which reduce the required bulk capacitance in a PDN. Valve-metal capacitors with very-low-ESL and very-low-ESR can now match and sometimes exceed the mid-frequency decoupling performance of MLCCs while still satisfying the bulk-decoupling requirements. With the combination of new VRM technology and very-low-ESL valve-metal capacitors, it is now possible to replace parallel combinations of conventional aluminum, tantalum, and ceramic capacitors with just a few high-performance, low-ESL valve-metal capacitors while controlling costs and conserving board space.

So just what are these new valve-metal capacitors and how do they work with advanced VRM technology to provide a high-performance decoupling solution in much less space while maintaining reasonable cost?   Two broad categories of these capacitors are low-profile, very-low-ESL aluminum polymer capacitors and very-low-ESL tantalum polymer capacitors.   In the first category are aluminum polymer capacitors with configurations similar to items A and B in Figure 2.   In the second category are tantalum polymer capacitors with configurations similar to items C and D in Figure 2.
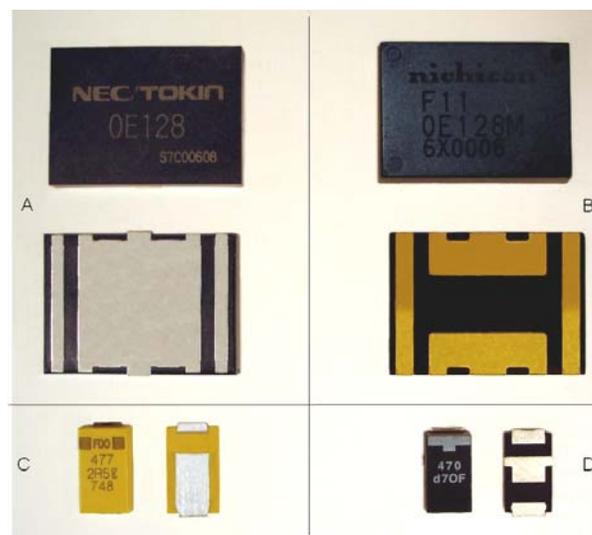

Figure 2. Very-low-ESL Aluminum and Tantalum Capacitors.

Each of these capacitors competes for the designer's attention with various desirable performance characteristics: very-low-ESL, very-low-ESR, moderately high capacitance, small footprint, low-profile, and modest cost. However, within in each category, performance is similar and competition for even better

performance is fierce.   A few comments are now offered on the roles of evolving VRM technology and these very-low-ESL valve-metal capacitors in the recent decoupling-solution paradigm shift.

Modern VRMs and Very-Low-ESL Capacitors

As suggested above, wider control bandwidth of modern VRMs reduces the need for bulk capacitance in the decoupling solution. High capacitance allows the decoupling solution to provide low and constant-impedance power to the microprocessor, even at "low" frequencies. Not long ago, conventional VRMs were unable to respond to load changes faster than a few kHz. So the decoupling solution had to provide charge/power to the microprocessor during the time required by the VRM to sense and respond to faster load changes.

In those older designs, good voltage stability required a lot of bulk capacitance. But now, advanced VRM designs can track load changes as fast as 100 kHz. This reduces the bulk capacitance requirement from as high as tens of thousands of microfarads (uF) to as low as a few hundreds of microfarads. This reduction in required bulk capacitance nibbles away at one of the advantages of the wound aluminum polymer capacitors – cheap bulk capacitance.

In the mid-frequency regime, MLCCs were once the unquestioned capacitor of choice because they have relatively low inductance, good capacitance retention at high frequencies, and low cost. These characteristics allow MLCCs to maintain constant microprocessor supply voltage as the processor's load current shifts at frequencies as high as 10-15 MHz.

Conventional MLCC do not have dramatically low ESL, but they do have lower ESL (~ 0.5-1.0 nH) than conventional aluminum and tantalum capacitors (~ 1-10 nH). Moreover, they can be connected in parallel in high volume to provide very-low ESL at reasonable cost if you have the board space. But now, valve-metal capacitors are available that have even lower ESL (~ 25-100 pH) than conventional MLCCs, and very few (perhaps only one) devices are needed to provide stable power delivery in the mid-frequency decoupling range.

With high-performance VRM technology limiting the need for huge bulk capacitance and with very-low-ESL capacitors able to handle the mid-frequency range, it is now possible for a very few, or perhaps only one, valve-metal capacitor(s) to handle the entire bulk- and mid-frequency decoupling job. So where space is at a premium, low-profile aluminum polymer and tantalum polymer capacitors are starting to displace wound-aluminum bulk capacitors and many MLCCs, while saving space and controlling costs. This is the essence of the microprocessor decoupling paradigm shift.

In Search of the Ideal Solution

It was said earlier that low-profile, very-low-ESL aluminum and tantalum polymer capacitors compete for the designer's attention on the basis of ESL, ESR, capacitance, size, and cost. However, as is usually the case, the exact mix of these attributes varies from device to device.   So it is important to understand important trade-offs among the devices as one searches for the optimal solution to his own design needs.

Some of these attributes are fairly straightforward to measure and compare. These include physical dimensions, price, and lower-frequency bulk capacitance.   A little more difficult to measure, but not excessively so, are milliohm-level ESR which has now been successfully measured at 100 kHz for years and capacitance stability at high frequency which usually involves algebraically removing the effects of self inductance from high-frequency capacitance measurements in the frequency-domain or observation of short-term, time-dependent charging rates in the time-domain.

Detailed descriptions of milliohm-ESR and high-frequency capacitance roll-off measurements will not be covered here, but the origin and measurement of ESL will now be briefly discussed because the topic has not received much attention for valve-metal capacitors.

Origin and Measurement of Inductance

To develop low-ESL capacitors or to facilitate fair comparison of competing low-ESL devices, one needs to understand how ESL occurs in devices and circuits, and also have a means to quantify (measure) it. The origin of inductance in electrical circuits will be briefly described with the aid of a simple circuit diagram. Then with the introduction of a few equations, the most common frequency-domain measurement strategy will be discussed.

In Figure 3 is shown a sketch of a circuit board with two mounting pads and a short loop of wire that connects one pad to the other. It is known that the loop of wire has inductance because when an AC (alternating current) current flows in the wire, a small voltage appears across the ends of the wire and the observed voltage is out of phase with the applied current by roughly 90 degrees. This effect occurs even in super-conducting wires, so we know that it is not caused by wire resistance. Also, the voltage grows as the frequency is increased, which is exactly how inductors behave. So we are confident that the wire is acting as an inductor.
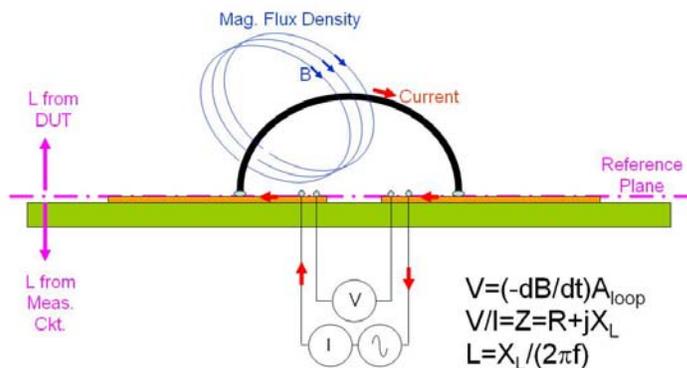

Figure 3. Circuit Board with Attached Loop of Wire to Demonstrate Principles of Inductance Measurement.

The current in the wire creates a time-varying magnetic field around the wire according to the "right-hand rule" of physics. Some of the resulting magnetic flux "couples" with the loop of wire and generates a voltage across the wire in much the same way that voltage is created in a car's alternator by the spinning magnetic rotor. It is the time-varying nature of the magnetic field that generates the voltage.

The precise amount of generated voltage can be predicted by the equation $V=(-dB/dt)A_{loop}$, where B is the magnetic flux density at right angles to the surface of the loop. The equation says that the generated voltage depends directly on the area of the loop and the amount of time-varying magnetic flux density. While the flux density is proportional to the current in the wire, the final magnitude also depends on the shape of the current path. Wider current paths make lower flux density than narrow paths.

Now that we have an AC current and a resulting AC voltage for the loop of wire, it is straightforward to calculate the wire's inductance using a few simple electrical formulas. The ratio of the voltage to the current is known as the impedance of the wire, which is given by the equation $Z=V/I=R+jX_L$. The resistance (R) and inductive reactance (X_L) can be found from Z by means of the formulas $R=Z\cos\theta$ and $X_L=Z\sin\theta$ if the phase angle θ between the voltage and current is known. The inductance is then calculated using the formula $L=X_L/(2\pi f)$. Alternatively, if the frequency of the current source is reasonably high, we can approximate X_L as having the same magnitude as Z and make a pretty good estimate of L without actually knowing the precise phase angle.

Almost all impedance analyzers measure inductance by variations of this simple method. It is only the desire for higher precision, speed, and frequency range that makes modern analyzers so expensive.

There is one more related concept that deserves mention here. Not only is there inductance in the loop of wire in Figure 3, but there is also inductance in the measurement system. So after the total inductance is measured,

it's not always easy to decide how much of the inductance should be blamed on the device we are trying to measure and how much should be blamed on the measurement circuit.

There is magnetic flux generated by both the wire we want to measure and the wire that connects the current source to it. Magnetic flux from both wires generates voltage in the bigger loop formed by both the wire we want to measure and the wire that connects the wire we want to measure to the voltmeter. So we always get an answer that is too big (too much flux and too much loop). By convention, we try to sort this out by defining a plane of reference in the circuit. We blame what happens on the measurement side of the reference plane on the measurement circuit and the rest on the device under test.

To find the inductance of the measurement circuit, we replace the device under test with a short circuit and make a measurement. The short is removed and we then connect and measure the device under test. The difference of these two measured inductances is assigned to the device under test. This is rarely a very precise result because there is mutual coupling between the two sides of the reference plane that depends on the geometry of the circuit on each side. The best we can do to optimize measurement accuracy for small inductances is to minimize the inductance of the measurement circuitry so that what we subtract from the total measurement is much smaller than what we are trying to measure.

Low Impedance Measurement Fixturing

In the process of building and measuring low-ESL tantalum capacitors, it was quickly apparent that once the device inductance fell below roughly 300 pH, it was very hard to measure further improvements. This was because there was excessive inductance in the connections between the impedance analyzer and the capacitors in the existing measurement fixtures. Mutual coupling between the measurement circuit inductance and the inductance of the device under test was limiting measurement resolution and defied the simplistic "zero-short" correction technique described earlier.

While we reached this conclusion independently, others have experienced similar difficulty and have documented their experiences[2,3]. Our solution to this problem bears remarkable similarity to that of Larry Smith[3] who was measuring the inductance of MLCCs.

Our strategy was to minimize the short-circuit impedance of the measurement fixture by employing small-diameter (0.047") coaxial cable for both the drive (current) and sense (voltage) circuits, and by making the connections as close to the device under test as possible. Also, the center conductor of each cable is deflected toward the "reference plane" where it leaves the shield. This overall geometric configuration reduces the escape of magnetic flux from the drive part of the measurement circuit that can couple with both the sense circuit and the inductance of the device under test. Even better results can be obtained with smaller-diameter cable (0.020"), but practicality suffers.
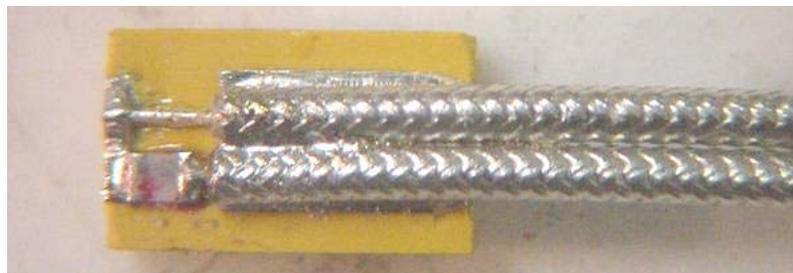

Figure 4. T528Z Low-ESL Tantalum Capacitor Connected to Measurement Circuit via Low-Z Fixturing.

A capacitor fixtured by the improved technique appears in Figure 4. One difference between Smith's approach and our approach is the inclusion of a small 50Ω SMT resistor in the drive circuit. The purpose of this resistor

is to provide a 50 Ω termination for the 50 Ω "drive" coaxial cable. This prevents transmission-line reflections from the end of an improperly-terminated transmission line for measured devices whose impedance is much less than 50 Ω at high frequencies.

Such reflections can significantly perturb impedance measurement accuracy, especially at frequencies where the line is a multiple of 0.25 λ in electrical length. However, for maximum accuracy the presence of the resistor must be accommodated in the impedance calculations. When the impedance of the measured device grows beyond roughly 1 Ω at high frequencies, these calculations become more difficult. But useful decoupling devices generally have much lower impedance than 1 Ω at the highest frequencies of interest.

When such fixturing is applied to a conductive copper sheet to simulate a "zero-short," the residual inductance is rough 30-35 pH. So the mutual-coupling error in inductance measurements employing this technique probably does not exceed this value and is likely smaller. Device inductances well below 100 pH have been resolved with this technique.

Low-ESL Capacitor Designs

There is much in common between low-ESL capacitors and the loop of wire we are trying to measure in Figure 3. In capacitors, as in the loop of wire, both loop area and the amount of magnetic flux that couples with the loop area affect the device's inductance.   These are the keys to designing low-ESL capacitors. Another similarity is that it is not possible to completely separate the ESL of the capacitor from the inductance of the circuit in which it operates. This means that high performance circuits demand both low-ESL capacitors and carefully designed, low-inductance circuits.

X-ray photographs of three very-low-ESL tantalum capacitor designs appear in Figure 5. If the loop of wire in Figure 3 were replaced by one of these capacitors, we could readily identify a current path and a current loop.
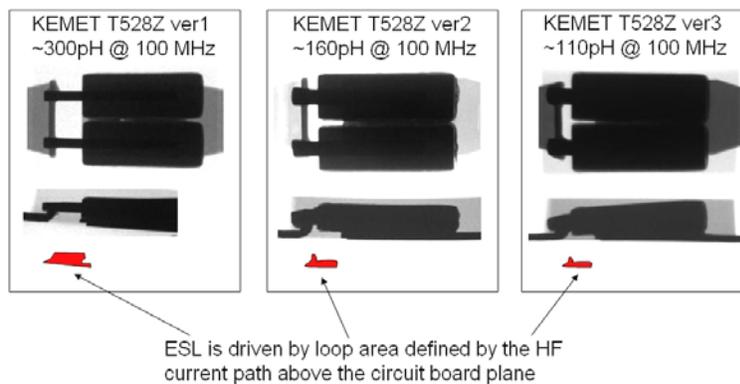

Figure 5. Construction Details of Very-Low-ESL Tantalum Capacitors.

It is the nature of electric circuits that high-frequency current is inclined to take the shortest path available along the surface of the loop formed by the device and the connecting circuitry. Accordingly, the loop defined by the current path of the capacitors of Figure 5 and the reference plane of Figure 3 is identified and highlighted by the outlined surface that appears below the respective x-ray pictures in Figure 5.   The smaller this outlined loop area, the lower the resulting inductance. These tantalum capacitors have much lower inductance (not more than 0.3 nH) than conventional tantalum capacitors whose ESL usually falls between 1 and 4 nH.

For the sake of accuracy, it must be stated at this point that the effective loop area of the device is somewhat larger at lower frequencies where the current path is not completely confined to the perimeter of the identified loop area in Figure 5. The consequence of this expanded current loop is higher inductance at medium

frequencies than at high frequencies. Thus there is not a single "correct" value of inductance for a given capacitor geometry, but rather a range of frequency-dependent inductance. Moreover, the device's inductance is also sensitive to the inductance of nearby circuitry as previously discussed. Nevertheless, practical efforts for reduction of capacitor ESL are still focused on reducing the loop area identified in Figure 5.

Measuring Decoupling Effectiveness

Capacitor ESL is only one contributor to the overall performance of a decoupling solution. Other factors that affect performance include the layout of the circuit (primarily circuit inductance), variation of capacitor impedance over the desired frequency range, and interactions (resonances) among capacitors in a multi-capacitor solution. These effects and interactions are frequently hard to predict and it is useful to have a measurement method to quantify the overall relative performance of various bulk and mid-frequency decoupling solutions.

An idealized test circuit was developed to evaluate the decoupling effectiveness of various competing decoupling solutions. The intent of this circuit is to realistically reproduce the bulk and mid-frequency decoupling environment of a microprocessor while avoiding the complexity and uniqueness of a specific computer circuit board.

The strategy was to fabricate a standard test board to which one or more capacitors could be mounted. This test board would be connected to a single high-performance VRM on one end and to a high-performance (t_r = 18 ns) electronic load on the other. The electronic load would then be rapidly switched from low current to high current while the stability of the voltage is monitored with a high-speed oscilloscope.
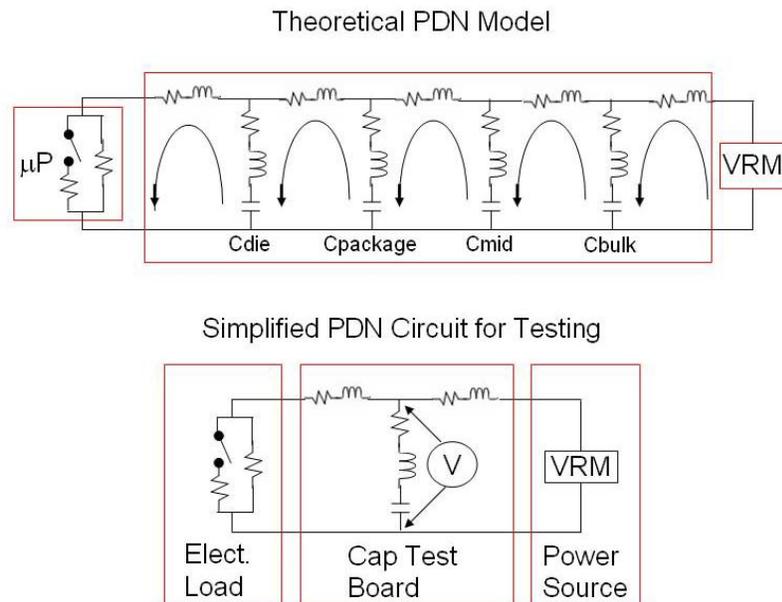

Figure 6. Theoretical Power Distribution Network for Microprocessor Decoupling and Simplified Circuit for Laboratory Decoupling Measurements.

A simplified schematic diagram of the circuit appears in Figure 6 while photographs of the VRM, electronic load, and two test boards appear in Figure 7. Observe that the two test boards are the same size and are configured similarly, even though the tested decoupling solutions are quite different (combination of MLCCs and conventional tantalum capacitors versus a very-low-ESL aluminum solution). A photograph of the fully-assembled circuit appears in Figure 8.
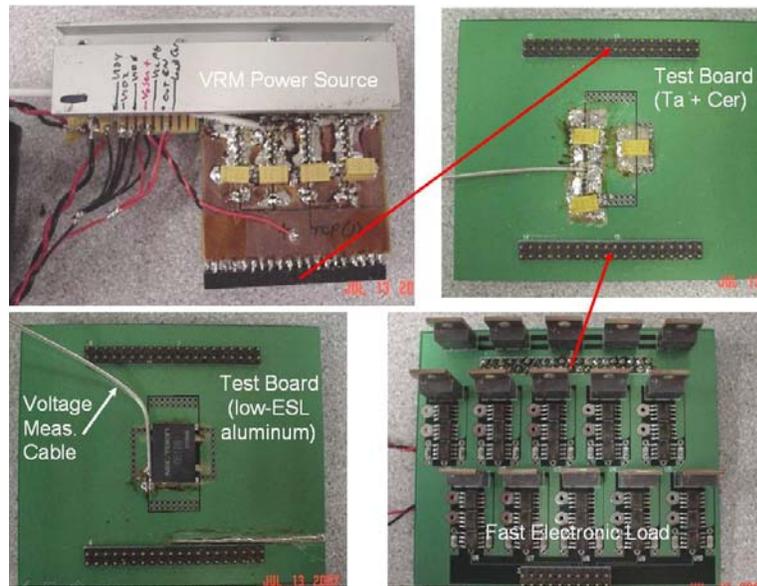
Figure 7. Photographs of VRM, Test Boards, and Electronic Load Board for Laboratory Decoupling Measurements.
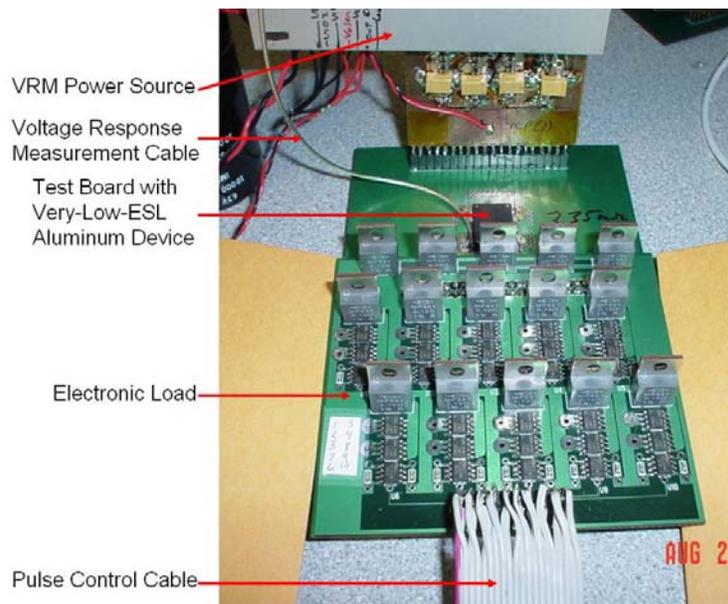

Figure 8. Photograph of Assembled Apparatus for Laboratory Decoupling Measurements.

A graph of the response of 3 combinations of conventional low-ESR tantalum capacitors and MLCCs appears in Figure 9 for a load-current step of approximately 80 A. The ideal desired response is a voltage drop of about 0.1 volts with no undershoot or ringing. This is consistent with a constant power system impedance of 1.25 milliohms. Unfortunately, because of interaction between the tantalum capacitors and the MLCCs, there is both undershoot and ringing. For the combinations shown (a small subset of the combinations explored), duration of ringing could be traded off against depth of undershoot, but neither could be effectively eliminated.
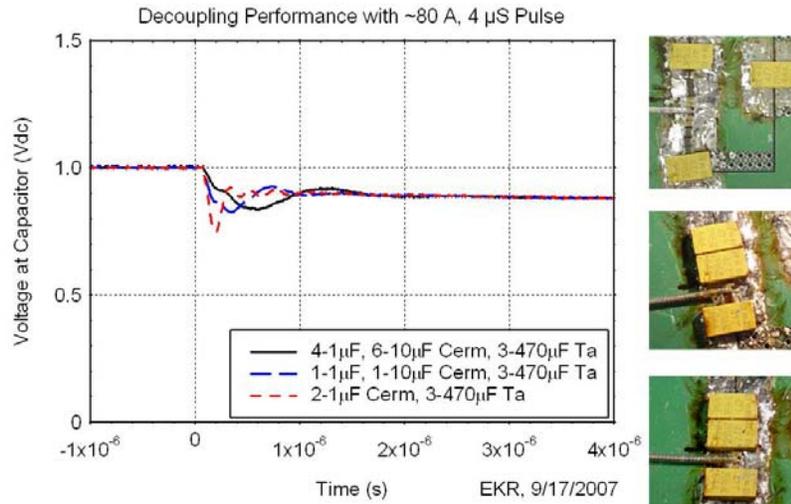
Figure 9. Decoupling Response of Three Configurations of Conventional Tantalum Polymer Capacitors and MLCCs.

Figure 10 compares the performance of a very-low-ESL aluminum polymer device and two T528Z low inductance tantalum polymer devices manufactured by KEMET. It is clear that either of these configurations provides performance that is superior to the combinations of conventional tantalum capacitors and MLCCs seen in Figure 9. However, the two T528Z KEMET devices occupy less area than the very-low-ESL aluminum device which could be an advantage in some designs. The T528Z devices also provide standard EIA 7343 footprint.
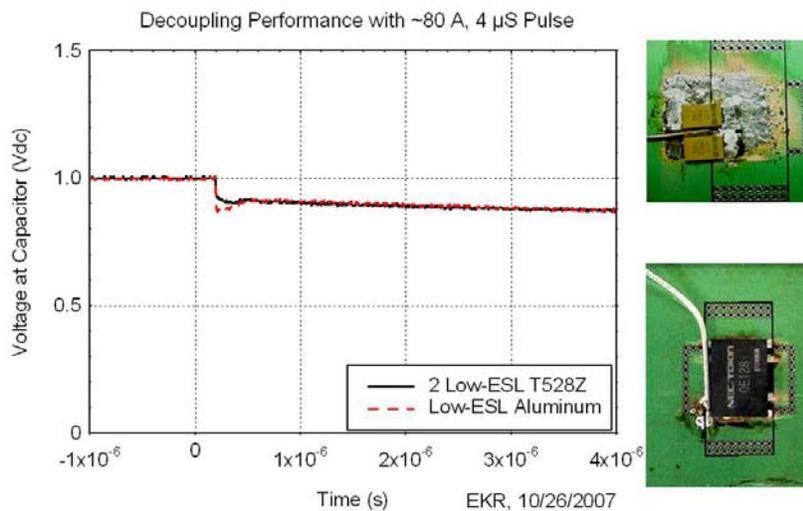

Figure 10. Decoupling Performance of Two KEMET T528Z Devices and One Very-Low-ESL Aluminum Device.

It is clear from this testing that very-low-ESL aluminum and tantalum capacitors can be used for decoupling solutions that require fewer capacitors, can consume less board space, and frequently have better performance than traditional combinations of conventional valve-metal capacitors and MLCCs. Growth of adoption of these high-performance decoupling solutions will now largely depend on favorable combinations of cost, size, and ease of implementation.

Summary and Conclusion

Power distribution networks for high-performance microprocessors are described in the introduction to this paper. Rapid and large load current changes are normal and getting worse while absolute requirements for voltage regulation become tighter with falling processor voltage. These trends have forced significant performance improvement in capacitors, but these improvements have not been good enough. New efforts are now underway to further improve performance while controlling cost and reducing required space, especially in mobile systems.

These forces have created an opportunity for a new approach to microprocessor decoupling. Improved voltage regulator module performance and development of very-low-ESL valve-metal capacitors have generated decoupling solutions that use far fewer capacitors and occupy less circuit volume. Also these new solutions displace some MLCCs in an application that MLCCs have dominated for many years.

Tantalum polymer capacitors are playing a significant role in this paradigm shift because of their very low ESR, high volumetric efficiency of capacitance, and newly-developed, very-low ESL. These capacitors can now be made with ESL that is a small percentage of previous devices' ESL and development efforts continue to drive down ESL even further.

But reduction of ESL requires understanding of what causes ESL, so that practical strategies can be formulated to reduce it. Also required is the ability to measure these improvements and to validate the performance of low-ESL devices in the target application. Each of these subjects was discussed in the context of developing and characterizing low-ESL tantalum polymer capacitors.

Problems with ESL measurement are discussed and a new fixturing method is introduced that improves our ability to quantify ESL reductions. Also, validation techniques are described that demonstrate the improved performance of low-ESL valve-metal capacitors under simulated decoupling conditions.

Evolution of VRM technology continues. As the switching frequency and control bandwidth of VRMs increase, even less bulk capacitance will be required. If the control bandwidth eventually reaches the 10-15 MHz range and if techniques are developed to locate the VRM even closer to the microprocessor to minimize circuit inductance, it is conceivable that no bulk and mid-frequency decoupling solution will be needed at all, since these duties could be handled entirely by the VRM. Such a dramatic shift is not expected overnight, but the trends in this direction have already been established.

References

[1] L. Mosley, *"Capacitor Impedance Needs for Future Microprocessors,"* CARTS USA 2006, p. 194.

[2] Z. Yang and S. Camerlo, *"Inductance of Bypass Capacitors, How to Define, How to Measure, How to Simulate,"* DesignCon 2005, TecForum TF7.

[3] L. Smith, *"MLC Capacitor Parameters for Accurate Simulation Model,"* DesignCon 2005, TecForum TF7.