Emerging high frequency, high power applications need a rethink on capacitors to maintain power integrity

There is an alternative to using traditional multi-layer ceramic capacitors (MLCCs) to fulfill high frequency power demands in IoT devices, mobile phones and high performance computing applications. By **Mukund Krishna, Senior Manager, Product Marketing and Luca Vassalli, Customer Applications Engineering Director, Empower Semiconductor**

As one of the most fundamental components in electronics, capacitors are used in large numbers across a variety of designs to maintain power integrity. They are a key requirement for applications such as mobile phones, IoT devices and HPC (high performance computing) applications and address last mile power delivery challenges, such as high frequency power demands for the high performance processors. While traditional MLCCs have fulfilled the requirements thus far, stricter constraints on power density are challenging the continued usage of the existing model.

As system engineers look to deliver the promised and expected performance in smaller form factors, provision of the most efficient power de-coupling solution is a critical design consideration.

It is currently estimated that over one trillion capacitors are produced every year, of which 800 billion are surface mount MLCCs or chip capacitors.

These are used to address requirements across the complete power and voltage range in applications, including energy storage, filtering and decoupling of power rails to filter out unwanted ripple and noise. The exponential growth in the creation of HPC has been led by the rapid advent of artificial intelligence and machine learning (AI and ML).

The magnitude and frequency of instantaneous power demands are growing in steps of 100% for subsequent generations, placing stress on power integrity solutions.

When it comes to the latest data-intensive systems built around high performance, high speed processors and multiple power domains that operate with fast transients and low voltages, designers are finding a growing number of challenges with conventional MLCC capacitors.

These processors are increasingly used on highly dynamic workloads, such as running AI algorithms and neural network models for ML and inference. For such applications, the peak current swings become significant, with instantaneous peak processor currents of 800A to 1000A in tens of nanoseconds becoming the norm. This results in extremely challenging current transients (di/dt). These high-performance devices usually require multiple low voltage (0.4V DC to 1.0V DC) power rails and tight adherence to voltage regulation specifications, typically within ±1.0%.

Board-mounted switching DC/DC converters offer a viable method of provisioning high power direct to computational devices such as FPGAs, GPUs, and neural network processors (NPUs).

While the DC/DC converters mounted on the PCB provide adequate DC power to these workloads, their frequency of operation, and hence bandwidth (which is the ability to respond to ultra-fast current transients), is orders of magnitude lower than required. The sheer volume of such solutions render them being located at distances far enough away that any ability to service fast transients is rendered useless by the high impedance to the processor.

The electrical noise generated from transients, power supply ripple and other noise artefacts can significantly impact the performance of the computational ICs and other circuit functions. Signal integrity is tightly associated with power integrity in any complex application, and such artefacts can create 'ringing' oscillations across the whole system. Digital processors made with advanced process nodes such as 5nm have extremely tight tolerances on voltage supply to avoid brown-out at the lower end and over-voltage at the upper end. Analogue ICs used in data conversion signal chains are particularly vulnerable to PDN (power delivery network) noise, with its power supply rejection ratio as a critical indicator of susceptibility. As any analogue IC's datasheet will highlight, small variations of supply voltage can upset the function's operation. For example, the introduction of jitter on clock signals or the reduction of analogue conversion accuracy.

Decoupling power rails

Minimising transient, ripple and noise artefacts from the PDN in high frequency, high performance applications is paramount. To decoupe PDN noise artefacts, multiple capacitors, typically MLCCs of different values and case sizes, are placed across the supply rails. The aim is to provide a low impedance return path across a wide frequency range. To provide the most effective noise cancellation, the capacitors are placed closest to the noise sources and the power pins of sensitive ICs. Board layout will influence MLCC placement, but this situation is exacerbated by larger processor ICs requiring tens of capacitors. At higher switching and computational frequencies, PCB trace parasitics and the equivalent circuit characteristics of the MLCCs also become significant.

As high frequency processing

applications become the norm, several factors will limit MLCCs' abilities for specific applications. There are several parasitic factors that will cause a capacitor's impedance to change across frequency. The equivalent circuit of a capacitor includes a series resistor (ESR) and series inductance (ESL). For example, the metal electrodes and end caps contribute to an equivalent series inductance or ESL that impacts the capacitor's resonant frequency. All things being equal, the lower the ESL, the higher the resonant frequency. Above its resonant frequency, a capacitor's effective impedance becomes inductive in nature, i.e., increases with frequency. There is an imperative to keep ESL as low as possible for capacitors used in high performance, high frequency applications.

Another critical issue is capacitor de-rating. Key de-rating factors that impact MLCCs are voltage, temperature and age.

An average MLCC will see its capacitance value reduce as the DC bias voltage increases. Capacitance also reduces as the temperature increases, with the degree of change dictated by EIA (American Electronic Industries Alliance) code used to classify the temperature co-efficient. AC or DC bias voltage has an impact on ageing characteristics that produce a decrease in capacitance over time due to changes in the dielectric's crystal structure. Increasing the number of MLCCs to address these characteristics ensures that the necessary decoupling capacitance is provided over the product's lifetime and across the entire anticipated range of operating conditions. Accommodating the many MLCCs, however, effects the product's mechanical design attributes, from form factor and power density to PCB layout flexibility, reliability, and cost. In general, the more capacitors that are deployed, the further away the capacitor network is likely to be from the processor, increasing the series inductance and introducing further opportunities for parasitics, especially at high frequencies. The product's calculated reliability metric will also be lower because the overall component count negatively influences reliability.

Differences between MLCCs and silicon capacitors

MLCCs are formed by alternating plates of metal (electrodes) with a dielectric material in between. The long plates connect to the terminals on either side providing contacts to the outside world. Since parasitic inductance is proportional to length of the metal path that current or charge must travel along, the long electrodes within MLCCs lead to inherently higher ESL.

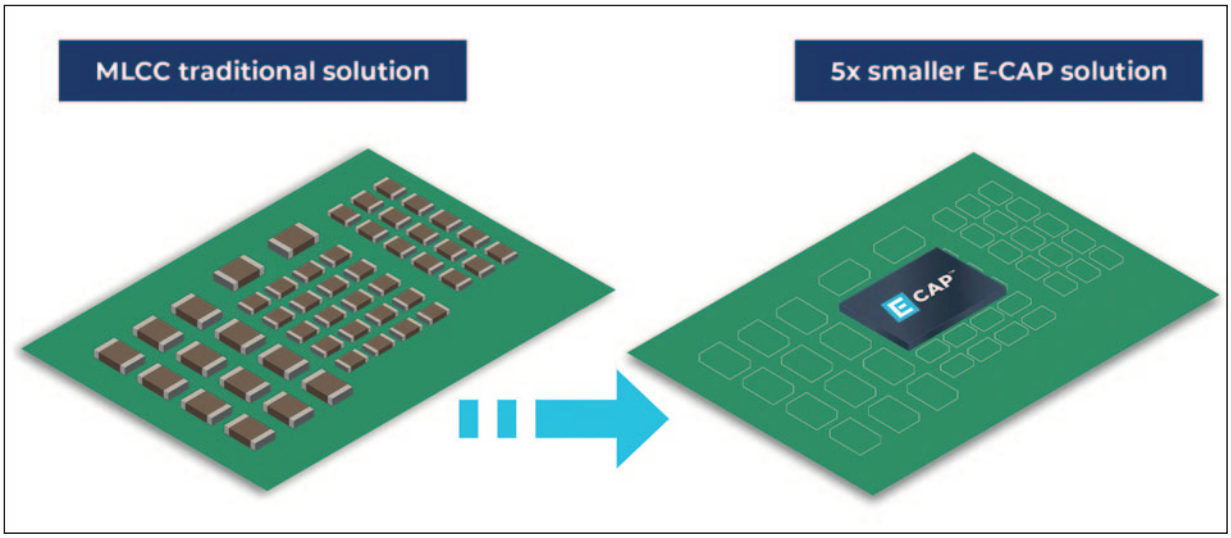

Silicon capacitors, like E-CAP, can improve layout density compared to standard MLCCs.

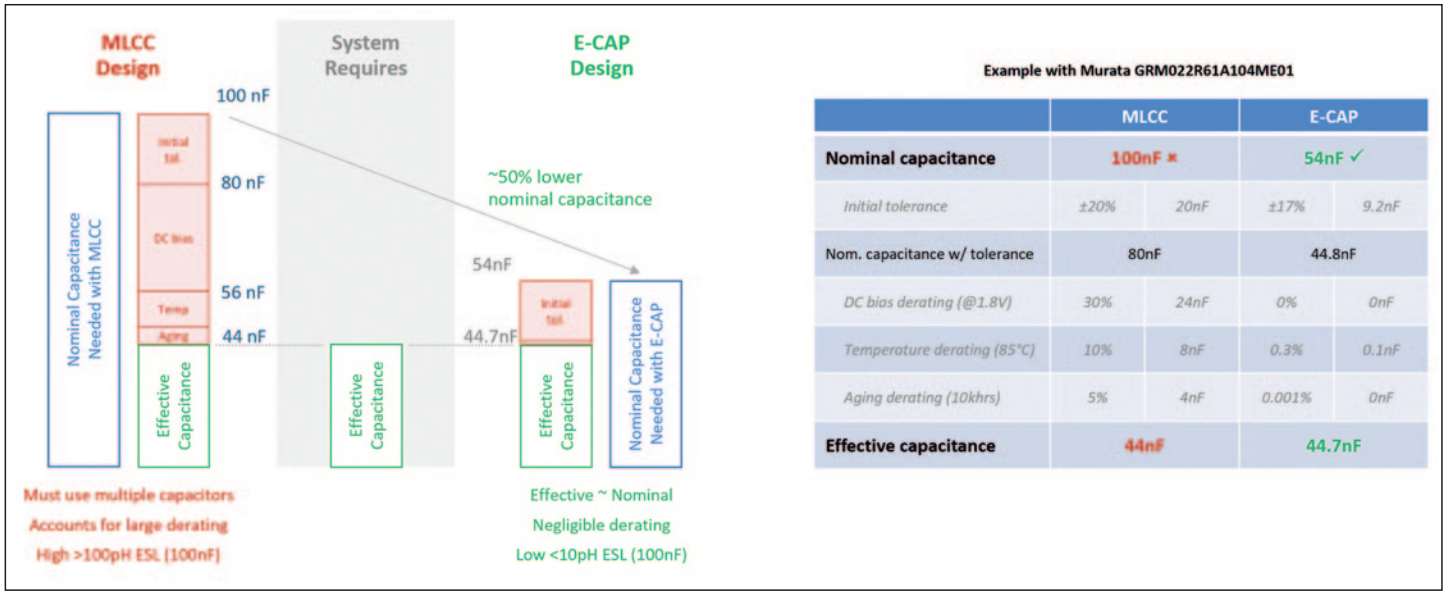

Example with Murata GRM022R61A104ME01

	MLCC		E-CAP	
Nominal capacitance	100nF ✖		54nF ✓	
Initial tolerance	*±20%*	*20nF*	*±17%*	*9.2nF*
Nom. capacitance w/ tolerance	80nF		44.8nF	
DC bias derating (@1.8V)	*30%*	*24nF*	*0%*	*0nF*
Temperature derating (85°C)	*10%*	*8nF*	*0.3%*	*0.1nF*
Aging derating (10khrs)	*5%*	*4nF*	*0.001%*	*0nF*
Effective capacitance	44nF		44.7nF	

Figure 2: **Comparison of MLCC nominal capacitance values to achieve a desired effective capacitance goal. (Source: Empower)**

Silicon capacitors are relatively new, created by vertical trenches in silicon using an ultra-fine lithography process. This allows many capacitors to be formed in a very small space where electrodes are orders of magnitude shorter than MLCCs, providing the same order of magnitude reduction in ESL. A large number of such capacitor cells (100s or 1000s) are connected in parallel to form a single capacitor further reducing effective ESL. Standard metal layers available in silicon-based semiconductor processes can connect to the electrodes anywhere on the silicon die, enabling termination flexibility and performance. Using silicon-based semiconductor processes also provides an inherent stability against variations in voltage, temperature and ageing that MLCCs struggle with, resulting in a much more stable and reliable product.

There is an immediate density improvement for silicon capacitors compared to MLCCs for the same effective capacitance and decoupling requirements. Empower's E-CAP silicon capacitor portfolio, for example, features a single capacitor of 220nF that fits standard 0201 footprints, up to an array of 17 single capacitors (4,800nF) in a single, low profile, surface mount, chip scale package (Figure 1).

The combined impact of an MLCC's ageing and de-rating factors is illustrated in Figure 2. The table in Figure 2 showcases a 54nF E-CAP compared to an example 100nF MLCC. The effective capacitance achieves the required 44nF, but the E-CAP achieves this with only the initial tolerance specification added. It is also more resistant to ageing. In this example, double the number of MLCCs would be required, effecting board space and layout.

Impedance versus frequency

Figure 3 illustrates how the low ESL attributes of an E-CAP offer high frequency impedance characteristics above 50MHz. For example, an average MLCC may have an ESL of 200pH compared to an E-CAP's 15pH. Even a commonly used network of MLCCs with reducing values and case sizes features more than two times the impedance at many 100s of MHz than a solution using E-CAPs which uses 40% fewer components.

An examination of the physical placement of capacitors is warranted to ensure their effective utilisation. When series inductances in the order of pH start to make an impact on the de-coupling capability, the inductance of the traces connecting the capacitors to their de-coupling points start to feature the same order of magnitude as the capacitor's ESL. Placing such capacitors too far away will render them useless and the benefits versus MLCCs may not be apparent.

High performance processors are typically mounted as die on silicon substrates that are then moulded over to form a package. The packages house multiple die, i.e., the processor, memory, communications chips, as well fan out the fine pitch of the die (e.g., 150µm) made in deep sub-micron processes to a manageable I/O pitch for PCB mounting (e.g., 500µm). There are multiple levels of

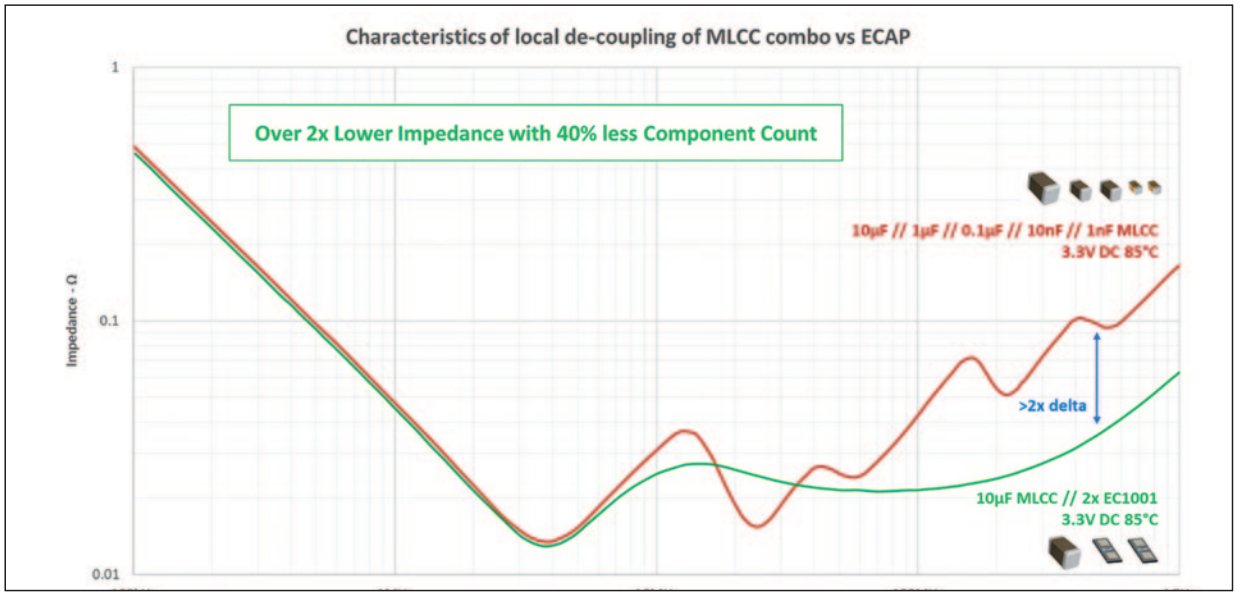

Figure 3: Impedance versus frequency curves comparison of MLCCs and E-CAPs. (Source: Empower)

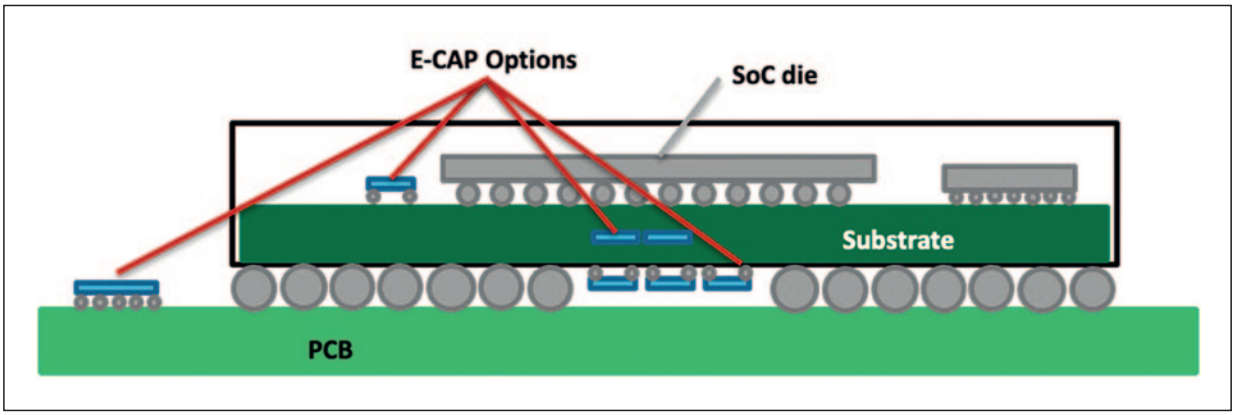

Figure 4: E-CAP mounting locations.

layers separating the processor die pins and the PCB (see figure 4).

A de-coupling capacitor located on the PCB is thus separated by impedance contributed by the many steps to reach the processor die pins.

De-coupling capacitors on the processor die are limited in capacitance and are meant to de-couple the transients with signatures exceeding 1GHz. Capacitors placed on the PCB will see at least two levels of RL impedances to get to the pins of the die. Since E-CAPs demonstrate low ESL and superior de-coupling capability up to many 100s of MHz, their ideal location should be on the SoC package, so that they see minimal impedance to the processor die pins.

Silicon capacitors can be realised in profiles down to 50µm and therefore capable of being located within the processor package in tight spaces such as the ball-height of the package, or even embedded within a package substrate which would provide the lowest possible impedance to the processor pins (apart from the on-chip capacitance themselves).

High frequency integrated voltage regulators

Another notable trend for DC/DC converters used in high performance applications are integrated voltage regulators (IVRs). IVRs leverage high switching frequency techniques in DC/DC converters, enabling orders of magnitude higher bandwidth, and fast response to load steps with minimal droop and recovery. Often neglected at low frequencies, the ESR and ESL of the output capacitor become critical design elements for IVRs that operate at high frequencies, in order to minimise noise and ripple.

IVRs are meant for PoL (point of load) power, targeting output voltages from 0.4V to 2.0V. Figure 5 illustrates the difference in ripple of solutions that use MLCC vs E-CAPs on the output ripple signature of a DC/DC converter operating at 10MHz. In addition, the ripple waveform is smoother, resulting in fewer harmonics limiting the EMI signature of the design.

Growth in demand

Over the last decade, the demand for high power applications has grown considerably. Use cases such as compute-intensive server boards used for cloud computing, machine learning deployments, the rise of electric vehicles (EVs) and the need for fast, energy-efficient charging stations are just some examples. There is also the need for ever-faster computing and highly dynamic computational workloads in high frequency, high bandwidth systems and in power conversion applications, high frequency switching topologies yield fewer energy losses while reducing the size of the critical supporting inductors and capacitors. With these game-changing technological innovations, power rail decoupling is even more critical than before. The availability of silicon-based capacitors significantly aids the development of these high frequency, high power applications.

www.empowersemi.com

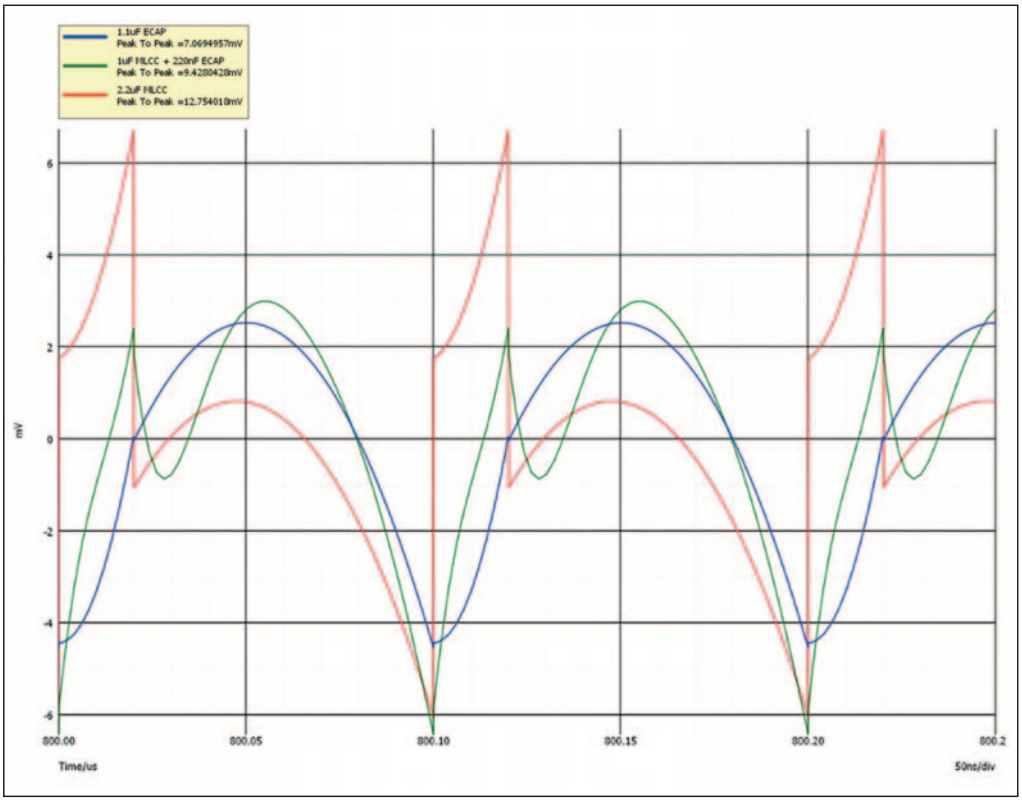

Figure 5: MLCC / E-CAP output ripple comparison at 10MHz, showing the MLCC at 12mV and E-CAP at 7mV.

Parameter	Standard MLCC	ECAP
DC bias de-rating	*44% @ 3 V*	*None*
Temp de-rating	*-11% up to 85°C*	*Negligible - ~0.3% (measured in ppm / K) – equivalent to C0G*
Aging	*~5-10% / 10k hrs*	*<0.001% / 10k hrs*
ESL	*>100pH (100nF)*	*<10pH (100nF)*

Table 1: Summary of technical differences between a conventional MLCC and a silicon E-CAP.