Capacitors with Fast Switching Require Distributed Models

C. R. Sullivan
A. M. Kern

Found in *IEEE Power Electronics Specialists Conferenc,* June 2001, pp. .

Capacitors With Fast Current Switching Require Distributed Models

Charles R. Sullivan and Alexandra M. Kern
Dartmouth College, Hanover, NH, USA, http://engineering.dartmouth.edu/inductor

Abstract—**A transmission line model for capacitors is proposed and developed. The model relates electrical parameters to geometrical parameters, leading to simple design rules for power applications. In contrast, design rules based on lumped RLC models are shown to lead to design errors. Experimental results match a simulation based on the distributed model whereas a lumped model gives erroneous predictions.**

I. Introduction

Conventionally, capacitors used in power applications are modeled as lumped RLC (resistor-inductor-capacitor) networks [1]. The inductance (or ESL, for equivalent series inductance) includes inductance of the leads and inductance of the current path through the capacitor itself. When the lead inductance dominates, the lumped model is accurate. However, when the inductances of the paths through the capacitor itself dominate, an accurate model must be based on distributed capacitance and inductance. Increasingly, in applications where low inductance is needed, packages with little or no lead inductance are used, and a distributed model becomes necessary if accuracy is needed. Capacitors with low-inductance packages include leadless multilayer ceramic (MLC) chip capacitors, similar stacked-film capacitors, and some wound film capacitors, although the lead inductance is significant in most wound-film capacitors.

The interconnect inductance for a capacitor is often also significant, and when it dominates, a lumped RLC model for the combination of a capacitor and its interconnect path can be accurate. However configurations with low interconnect inductance are increasingly common. Such configurations include low-inductance laminated bus bars, mounting of capacitors directly to IGBTs or other semiconductor switches, and the careful low-inductance PC board layout needed for power delivery in high-current low-voltage digital systems. Thus, the inductance of the capacitor itself often dominates, and a true distributed model is needed. For cylindrical capacitors with coaxial terminations a rigorous distributed model was developed in the frequency domain [2]. The model in [2] has been further developed and applied in [3], [4]. In this paper we develop a similar model for rectangular geometries, but emphasizing analysis in the time domain. In addition to being applicable to rectangular MLC and stacked-film capacitors, the rectangular model is a good approximation for wrapped film capacitors without coaxial terminations, to which the analysis in [2], [4] does not apply. The rectangular analysis is also much simpler than the cylindrical analysis. A simple model combined with time-domain analysis facilitates simple design calculations and intuitive understanding of capacitor behavior with switched waveforms possible.

Other previous high-frequency distributed models of capacitors have been based on measurements of the capacitor in a 50 Ω system [5], [6]. For power and bypass applications, this is not helpful, as the impedance desired is many orders of magnitude lower, and 50 Ω measurements would need an impractical level of precision to be useful.

II. Extended Lumped Model

Consider modeling a capacitor as shown in Fig. 1 a). We wish to include, in particular, stray inductance effects that become important at high frequency. Note that inductance cannot be defined except for a path forming a complete loop. Thus, we consider the impedance at the terminals to the left, assuming the cap is mounted on a two-layer board. This is problematic, because the impedance of the cap now depends on the particular board characteristics. So we choose to define the impedance of the cap as the *difference* between the impedance of the configuration in Fig. 1 a) and that of the same board with the cap replaced by a short, shown in Fig. 1 b). The model of the whole setup will be the model of the cap in series with the impedance of the board and vias in Fig. 1 b). With a low-impedance capacitor at high frequencies, the board and via impedance is usually primarily an inductance.

To model the configuration in Fig. 1 a), we draw a capacitor for each pair of plates. For each path, we also add inductance. In particular, the inductance of a path that goes across at a height x is

$$L = \mu_0 \frac{x\ell}{w} \tag{1}$$

where μ_0 is the permeability of free space, ℓ is the length and w is the width (into the page in Fig. 1a), and we assume $w >> x$. Note that in some configurations, this last condition does not hold. In those cases, the conclusions developed here will be qualitatively, but not quantitatively, correct. The simplification is still useful because the simple results lend intuition, and because a correction based on a single measurement can be used to extend the results to other situations.

We proceed with the modeling process by adding an inductance of

$$\Delta L = \mu_0 \frac{\Delta x \ell}{2w} \tag{2}$$

where Δx is the vertical distance between the center of one plate and the center of the next, on each side for each plate. This results in the model in Fig. 1 c).

In Fig. 1 d), the model is simply redrawn a bit more neatly. In Fig. 1 e), it is transformed into a model with exactly the

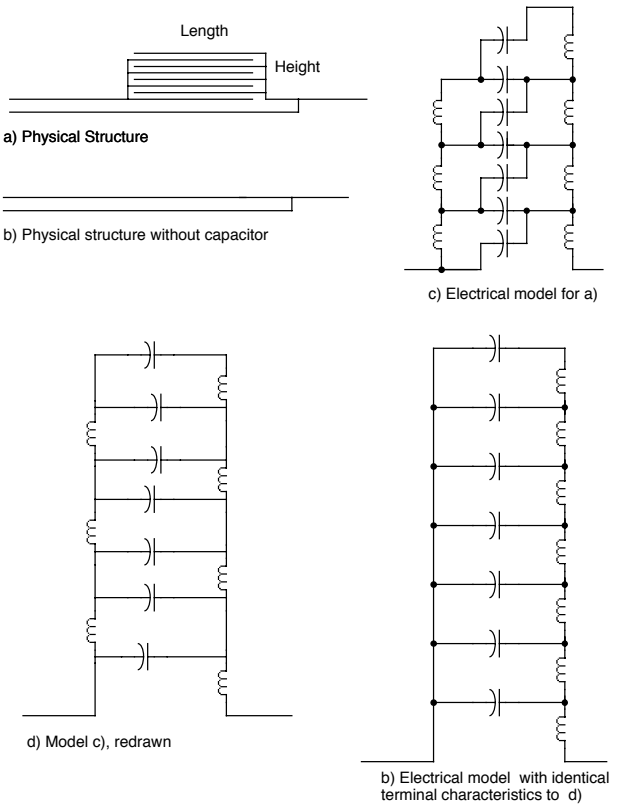

Fig. 1. Capacitor models

same terminal impedance, but a less confusing structure. Note that this is the lumped equivalent of a transmission line. This inspires modeling the capacitor as a transmission line.

III. Transmission Line Model

We use the model Fig. 1 e) to develop a transmission line model. It might seem that we need to know details such as the number of plates and their spacing to calculate model parameters such as ΔL. But actually, when we model it as a transmission line, we are making the continuous approximation that there are so many plates that they don't need to be modeled individually [2]. Certainly, with modern capacitor technology the number of plates is typically very large, and so a continuous approximation is good. So what we really need is the capacitance and inductance *per unit length* (actually, height, in this orientation). Capacitance per unit height is just the manufacturer's rated capacitance divided by the height. Inductance per unit length is just (2) divided by x, or

$$L' = \mu_0 \frac{\ell}{w} \tag{3}$$

If we want to design a system, i.e. chose a capacitor or set of capacitors, a useful way to describe the transmission line is in terms of characteristic impedance and phase velocity. Or, better, characteristic impedance and the round-trip time for a wave traveling up the transmission line and back down. The characteristic impedance can be expressed as:

$$Z_0 = \sqrt{\frac{L'}{C'}} = \sqrt{\frac{\mu_0 \ell / w}{C/h}} \tag{4}$$

For a particular type of capacitor that has a given voltage rating, dielectric material and thickness, and plate thickness, we can also describe the capacitance in terms of a constant capacitance per unit volume, C_v. In terms of this, we can rewrite (4) as

$$Z_0 = \sqrt{\frac{\mu_0 \ell / w}{C_v w \ell}} = \frac{1}{w}\sqrt{\frac{\mu_0}{C_v}} \tag{5}$$

Thus we see that for a given capacitor technology, a given requirement for characteristic impedance determines the width required. For example, for a 22 μF 6.3 V X5R dielectric capacitor, C_v is about 1000 F/m^2 (i.e., 1 μF/mm^3). For 1 mm width, Z_0 is about 35 mΩ. For a 1 V power bypass application, 35 mΩ is near the upper limit one might like for 1 A of load current. Thus, we can conclude that in such an application, with that capacitor technology, we need 1 mm of capacitor width for each amp of load current, independent of the height and length.

The velocity is

$$v = \frac{1}{\sqrt{L'C'}} = \frac{1}{\sqrt{\mu_0 \frac{\ell}{w} C_v w \ell}} = \frac{1}{\ell}\frac{1}{\sqrt{\mu_0 C_v}}. \tag{6}$$

As discussed above, velocity is not directly of interest. More important is the round trip time, $T = 2h/v$. Substituting, we find

$$T = 2h\ell\sqrt{\mu_0 C_v}. \tag{7}$$

For the round trip time, we see that width does not matter, but increasing length or height will directly contribute to increasing round trip time—length and height are in this sense interchangeable. However, returning to Fig. 1 a) and b), we can note the series inductance due to the board increases as we increase the length. We can also note that we might, in some cases, also deviate from the condition of $w >> h$ if we use height rather than width to increase round trip time. In this case, L' would be nonuniform, decreasing at the top of the cap. This would result in dispersion, which may be beneficial. For the same 6.3V X5R technology discussed above, we can compute round trip time as

$$T = \ell h (71 \text{ ns/mm}^2) \tag{8}$$

For example, a 3 mm long, 3 mm high cap has a round trip time of about two-thirds of a microsecond.

When the assumption $w >> h$ does not hold, (7) and (5) cannot be expected to be accurate. In this case, either T or

Z_0 can be measured experimentally, and the other calculated from the relationship

$$2Z_0C = T. \tag{9}$$

Alternatively, more extensive electromagnetic modeling could be used.

IV. Application to a circuit

Consider a capacitor with a step change in current. This could be a bypass capacitor with a stepped load and no power source connected, or a model for a short time period near a switch transition in one of many power converter circuits. The distributed model predicts a waveform like that shown in Fig 2. As the round trip time approaches zero, the steps merge into a ramp; the behavior of an ideal lumped capacitor.

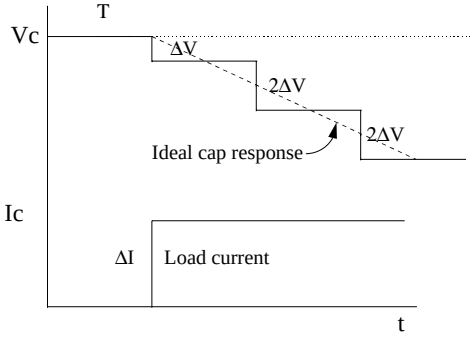

Fig. 2. Behavior with current step.

Note that this model shows behavior that is preferable to that of the lumped RLC model, in that even with an abrupt load step (infinite di/dt), there is no voltage spike, but just a step of magnitude $\Delta I_{load}Z_0$. One way to design the system is to size the width in order to make this step the maximum size tolerable by the system, and to make the height and/or length large enough such that the round-trip time is adequate for the power supply and/or interconnects to the power supply and/or additional bypass capacitance to respond and supply the current.

These design guidelines, as well as the predicted behavior, are quite different from what would be predicted by the conventional lumped RLC model. With a lumped model, if one is concerned about the voltage step at the current transition, one wants to minimize the inductance, which can be achieved by minimizing the height of the capacitor. However, the distributed model shows that increasing the height of the cap only affects the round trip time, and does not affect the height of the initial voltage step. Thus a thicker capacitor does no harm, contrary to the predictions of the lumped model.

A conventional lumped model also leads to the conclusion that additional smaller capacitors with lower inductance may be needed in parallel with a higher-inductance capacitor. However, if lead inductance and interconnect inductance can be kept low enough, additional parallel capacitors are not necessary—a single multilayer capacitor comprises multiple parallel plates, and the plates closest to the board have very low inductance. So a low-inductance capacitor is already present, connected in parallel. Additional parallel capacitors are not needed, and may even make matters worse by complicating the wiring and making it difficult to achieve low inductance interconnects. Effective strategies for coordinating impedance of multiple parellel decoupling capacitors have been developed based on lumped RLC models [7]. It is not clear what effect the transmission-line effects have on the net impedance in such a scheme, but in any case a better understanding of distributed capacitor behavior shows that simpler schemes may be at least equally effective. Note that this conclusion assumes that the capacitor terminations do not introduce additional unnecessary inductance. As discussed below, this is not always the case, and good design of capacitor terminations is essential.

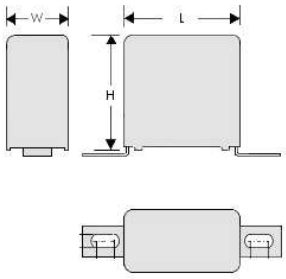

Fig. 3. Film capacitor A.

V. Experimental Verification

Three capacitor configurations were tested by discharging through a resistor switched by a set of parallel MOSFETs, and monitoring their terminal voltages. The circuit was constructed on a double-sided copper-clad 50μm thick polyimide substrate for low inductance, with the main current path across the top, returning in parallel on the bottom. Four S08 packages containing eight IRF7103 MOSFETs were connected in parallel to minimize stray inductance in the MOSFETs. The inductance of the loop including the board, resistors, and MOSFETs, with the capacitor replaced by a copper foil short was measured as 1.4 nH. The capacitor was charged by an external voltage source through a relatively large resistor, and the MOSFETs were periodically pulsed on in order to allow observation of the capacitor voltage waveform during the transient. The data was saved with a digital oscilloscope, and later plotted along with simulation data.

A. Film capacitor A

The first test specimen to be discussed was a 1.5 μF wound metalized polypropylene film capacitor formed into a plastic

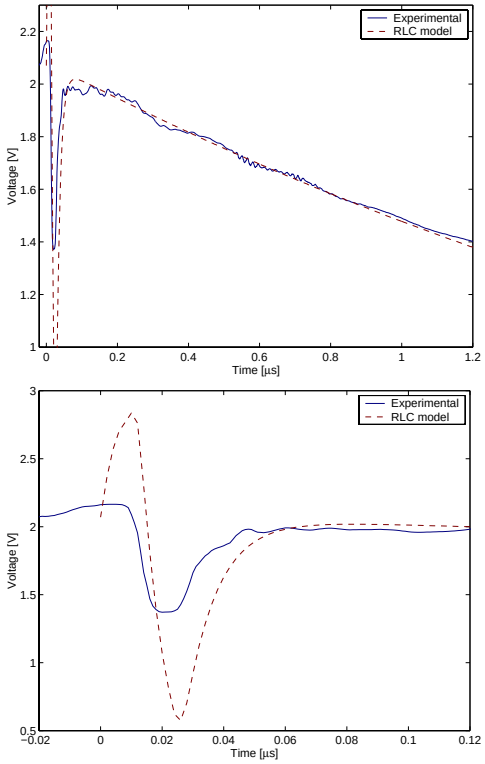

Fig. 4. Comparison of measured voltage waveform and simulated waveform using a lumped RLC model for film capacitor A. The lower plot is the same data as the upper plot, but with the scale chosen to show the initial transient.

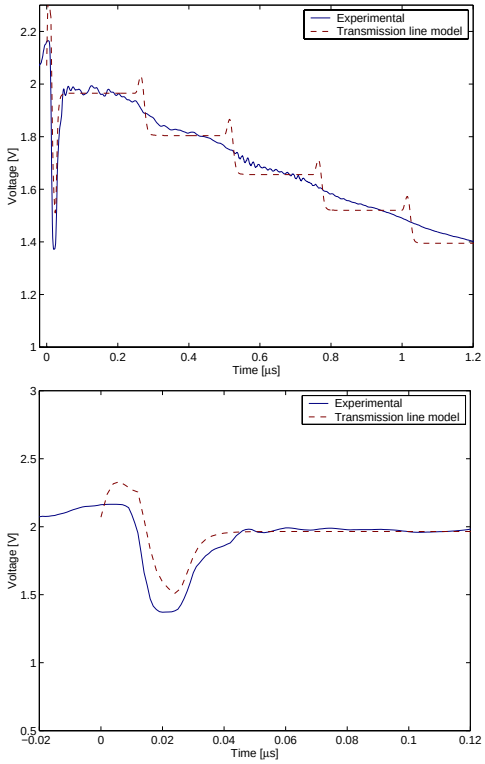

Fig. 5. Comparison of actual of measured voltage waveform for film capacitor A with a simulated waveform a transmission line model, with additional stray inductance representing the capacitor leads. The lower plot is the same data as the upper plot, but with the scale chosen to show the initial transient.

box, with metal tab leads as shown in Fig. 3. Dimensional and electrical details of all the capacitors are summarized in Table I. Fig. 4 shows the experimental waveforms, compared to a lumped RLC model based on the measured capacitance and ESR, and an 18 nH inductance to match the measured self-resonant frequency. The most serious error in the model is the initial spike which is more than twice the size of the spike in the experimental measurement. Both the measurement and the simulation show an upward spike first, as a result of charge coupled through the MOSFET gate-to-drain capacitance, although this is again much more pronounced in the simulation with the lumped model. After the initial spike, the measurement shows a regular ripple, slightly tending towards the steps shown in Fig 2. The RLC model does not capture any of this behavior.

A simulation using a transmission line model is compared in Fig. 5. Because this capacitor does not come close to satisfying $w >> h$, (7) and (5) were not expected to be accurate, and so the round-trip time was determined from the experimental waveform and used with (9) to find $Z_0 = 88$ mΩ. The experimental waveform also shows a significant transient spike, unlike the ideal waveform in Fig. 2. This is because of lead inductance; although the geometry in Fig. 3 could potentially be made with very low lead inductance, the capacitor is actually assembled by placing the wrapped film into the top of the plastic box, and filling the extra space in the bottom with epoxy. All of this extra space adds lead inductance. The simulation used an additional 5 nH inductance in series with the transmission line to model this lead inductance, as well as a series resistance to model capacitor ESR. The simulation matches the initial transient well. Both voltages then stay approximately constant for about 200 to 300 ns. The subsequent steps are then much more pronounced in the transmission line model than in the actual capacitor voltage. The difference can be explained by damping and dispersion in the non-ideal transmission line formed by the actual capacitor, but not modeled in the simulation. The ideal transmission line model also differs in that it reproduces the initial transient in the form of a spike that recurs at each step.

B. Film capacitor B

The lead inductance in film capacitor A is only one third the ESL calculated from the resonant frequency, and the size and width of the voltage spike are much smaller than would be produced by the full ESL. However, it is still much larger than it needs to be, and the initial spike could be further reduced with lower lead inductance. To demonstrate this, we modi-

TABLE I
TESTED CAPACITORS

		Film A	Film B	Single MLC	Stacked MLC
ℓ	mm	50	27.4	6	3.2
w	mm	24	23.5	5.3	1.4
h	mm	40	13	2.7	13 (stack)
C	μF	1.5	10	100	26.4 (stack)
C_v	F/m^2	0.031	1.2	1160	453
Z_0, (5)	mΩ	264		6.2	37.6
Z_0, exp.	mΩ	88			17.4
T, (7)	ns	793		1240	2000
T, exp.	ns	250			9200
f_0	kHz	986		400	544
ESL (from f_0)	nH	18		1.68	2.24

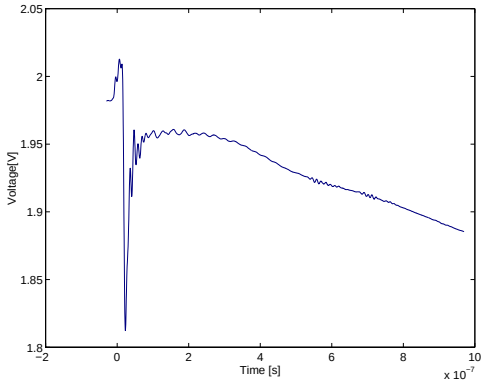

Fig. 6. Measured waveform on film capacitor B.

fied a conventional wire-lead metalized polyester film capacitor by filing off the epoxy coating on the end and soldering copper-foil leads to the end metalization. The resulting low-lead-inductance capacitor (film capacitor B) was tested in the same circuit as film capacitor A. The waveforms show a significantly reduced initial voltage spike of only 170 mV, as shown in Fig. 6. The expected lead inductance can be calculated as

$$\Delta L_{lead} = \mu_0 \frac{\Delta x \ell}{w} \tag{10}$$

where Δx is the thickness of the epoxy coating on the bottom of the capacitor of approximately 1 mm. The lead inductance thus calculated is 1.47 nH. A simulation with a transmission-line model and 1.47 nH lead inductance gave 183 mV spike. Thus, it appears that the initial spike is in fact determined predominantly if not exclusively by the lead inductance and not by the internal inductance of the capacitor.

C. Large multilayer ceramic capacitor

A 100 μF multilayer ceramic (MLC) capacitor was tested in the same low-inductance discharge circuit, but using only the on-resistance of the FETs and no explicit resistor in the circuit, in order to obtain higher currents and larger voltage deviation with this low-impedance capacitor. We estimate that the resulting derivative of current was greater than 1 A/ns, limited

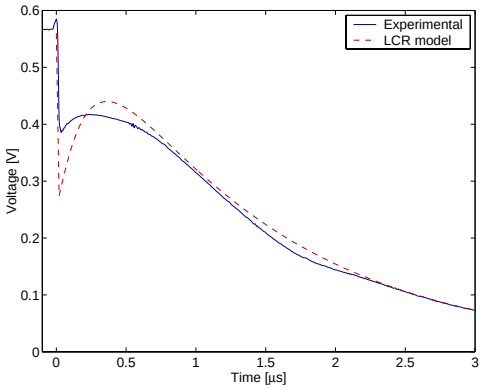

Fig. 7. Comparison of actual measured 100 μF MLC capacitor voltage waveform and simulated waveform using a lumped RLC model.

primarily by the inductance of the paralleled SO8 MOSFET packages.

Fig. 7 compares the measured voltage waveform to the simulated waveform using the lumped RLC model, with inductance based on the measured self-resonant frequency of 400 kHz. Here the spike predicted by the lumped model is almost entirely absent in the measured data, showing that the contribution of ESL in the lumped model is more misleading than helpful. The experiment also shows some non-uniformity in the voltage decay, tending toward stepped behavior, that is not captured at all by the lumped model.

In this capacitor, w is about twice h. This falls short of the $w >> h$ criterion for the simple inductance formula (1) to be highly accurate, but it is close enough for (1) to be a useful approximation. Thus, in Fig. 8 we compare the measurements to a simulation based upon parameters calculated from the geometry using (5) and (7). This capacitor presents another interesting opportunity for verification of the transmission line behavior: the voltage can be probed at both the top and bottom of the capacitor, with the top corresponding to the open-circuited far end of the transmission line. The qualitative behavior matches the expectations the simulation: the voltages at the top and bottom alternately step down, with the step in the top voltage happening at time halfway between the steps in the bottom voltage. Again, however, attenuation and dispersion in the transmission line substantially smooth out the steps, an effect that is of course completely absent in the simulated ideal transmission line.

An additional interesting unexplained effect is the immediate jump in the voltage at the top of the capacitor, well before a wave propagating through the capacitor would be expected to reach the top. Initially we suspected a measurement problem, but the effect is repeatable and is significantly larger than any noise pickup we were able to observe in our test set-up. It may be due to a decrease in dielectric constant of the X5R ceramic material at high frequencies, allowing a high-frequency edge to propagate through the capacitor at much higher velocity than the wave that arrives half a microsecond later.

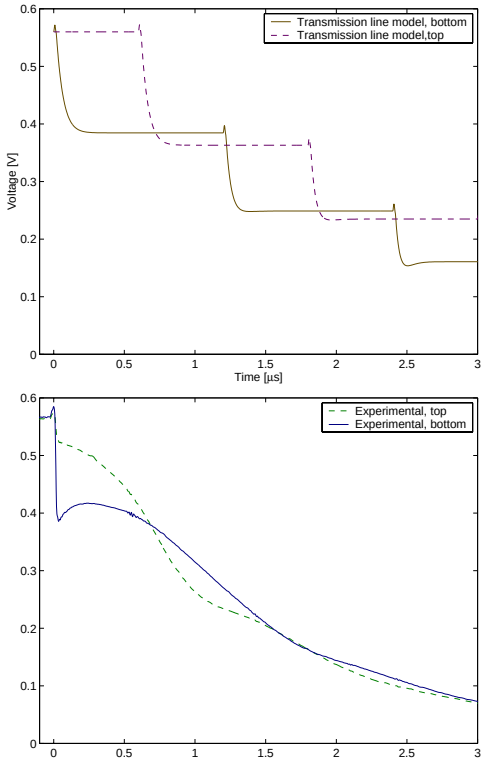

Fig. 8. Comparison of actual measured 100 μF MLC capacitor voltage waveforms (top plot) and simulated waveforms using a transmission line model (lower plot). Each shows both the voltage at the bottom where the capacitor is soldered to the board and at the top, 3 mm above the board.

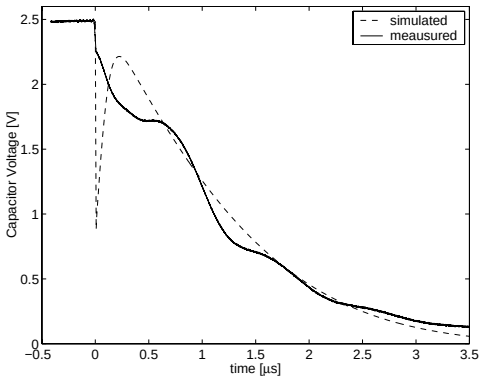

Fig. 9. Comparison of measured voltage waveform on stacked MLC capacitors and simulated waveform using a lumped RLC model.

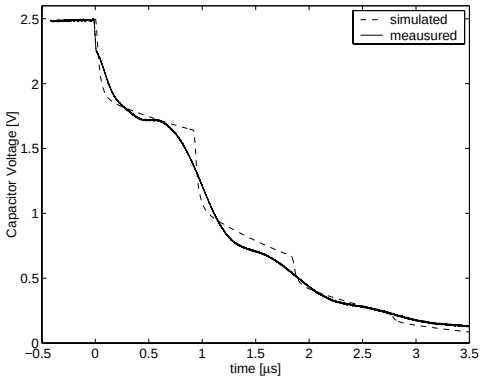

Fig. 10. Comparison of measured voltage waveform on stacked MLC capacitors and simulated waveform using a lossy transmission line model.

D. Stacked multilayer ceramic capacitors

In the previous experiments, the stepped waveforms produced by the ideal transmission line model were much more abrupt than the measured waveforms. In this experiment, a lossy transmission line simulation was compared to experimental data. This is experiment was performed before the ones reported above, when the low-inductance test set-up was not yet available, and before we had obtained samples of large 100 μF capacitors. In order to get a round-trip-time long enough to be accurately measurable, we stacked smaller MLC capacitors to increase h. In order to minimize stray inductance, they were soldered directly to the leads of a TO-220 package MOSFET.

A set of eight 3.3 μF MLC capacitors was used to obtain a total capacitance of 26.4 μF. Figs. 10 and 9 compare the experimental voltage response of a the capacitor to simulated response, using either a lumped RLC model or a transmission line model. The simulation includes non-ideal switching of the MOSFET, 5 nH of inductance in the MOSFET package, and capacitor ESR modeled as series resistance in a lossy transmission line. The use of a lossy transmission line explains the deviation from the ideal steps in the previous transmission line models. The transmission-line simulation in Fig. 10 matches the experiment quite well, but the lumped RLC model simulated in Fig. 9 deviates substantially from the actual measured behavior, predicting a large spike that does not occur in practice. These results confirm that the conventional model is misleading, and that the transmission line model can be simulate actual behavior much more accurately, especially when the transmission line parameters are adjust to include loss and dispersion effects. Note that in this configuration, $h > w$, and so the transmission line parameters calculated from geometry are not valid, and values calculated from time-domain waveforms were used instead.

E. Discussion

The experimental results consistently show that the conventional lumped LRC model does not accurately predict the voltage waveform with a fast current step, particularly when peak voltage deviations are important. The discrepancy between actual capacitor behavior in the time domain and the spike produced by the RLC model has been noted previously, for example in [3]. In [3], the time-domain measurements were discarded because of the lack of an inductive spike, and frequency-domain measurements were used instead. However, the behavior in the time domain corresponds to the ac-

tual performance in a typical power circuit, and should not be ignored.

The ideal transmission line model exaggerates the definition of the steps in the voltage waveform, particularly after more than one round-trip time. If one simply evaluates rms error between the simulated waveform and the actual waveform, one might think that the ideal transmission line is a poor model, and that more accurate transmission line models, as begun in section V-D, are necessary before transmission line models can be useful to the designer. We intend to work on such improved models, but we believe that the simple ideal transmission line model will often be more useful to the designer. It provides a conceptual model of what behavioral features to expect, and relates those features directly to geometry. In particular, the height and timing of the first and second steps are probably the most important features. The height of the first step is given by $Z_0 \Delta I$ for the height of the first step, and twice that for the second step.

More detailed models are unlikely to be necessary or useful for typical conceptual design or hand calculations for design. However, they can be very valuable for improving the accuracy of simulations. For this purpose, frequency-domain measurements may be performed, and lumped models developed by curve fitting can provide arbitrarily good accuracy. For some geometries, such models can be developed analytically [4]. The possibility of obtaining arbitrarily good accuracy from models fit to measurement data is attractive, but an important caution is that the accuracy is still limited by the accuracy of the measurement. Typical impedance measurement jigs for surface-mount components have stray inductance on the order of 100 nH, which makes resolving capacitor inductances on the order of 1 nH very difficult. In addition, the mounting configuration including the location of the current return path can significantly affect affect the field and current distributions in the component under test. Thus, for a measurement to be relevant to the application, it should duplicate the positioning of the component relative to the current return path.

More accurate transmission line models will need to take into account the actual loss and dispersion mechanisms in a capacitor. Standard lossy transmission line models are based on resistance in series with the inductance, and in parallel with the capacitance. The series resistance modeled corresponds, for a capacitor, to the resistance of the end contacts in the direction h. This is typically a negligible resistance, with more significant resistance corresponding to the resistance of the plates and the dielectric loss. This would be better modeled as resistance in series with the capacitance of the transmission line. Although further improving transmission line models may not be essential to accurate simulations (because higher-order lumped models based on curve fits to measured data can work well), and may not lead to models simple enough for use in everyday design, it will enhance our understanding of the phenomena leading to the observed time-domain and frequency-domain behavior. It may also lead to improved capacitor designs, particularly for applications in which time-domain step response is important.

VI. Conclusion

A simple distributed model for a multilayer capacitor proves easy to use in design work, and captures important aspects of experimentally observed waveforms better than traditional series RLC models do.

Traditional RLC models indicate that increased capacitor height increases ESL, and thus increases the inductive voltage spike to be expected during abrupt current changes. The distributed model shows that this is incorrect. The height of the bottom of the capacitor above the board contributes to lead inductance and directly contributes to the voltage spike, but the overall height of the capacitor only increases the round-trip time before the voltage steps down further and does not affect the initial voltage spike. Thus, the lead configuration is of much more importance than the overall capacitor size.

An important question is when to use which model: simple ideal capacitance, lumped ESR-C, lumped RLC, higher-order lumped models, or transmission line models. We believe that whenever design for good time-domain performance is of interest, and current transition times are faster than the period of the capacitor self-resonant frequency, the transmission line model is the best choice—other models can lead to poor design decisions. When accurate modeling is of more interest than understanding how to design for high performance, lumped models can be preferable. If and only if interconnect and lead inductance dominate over capacitor ESL, a lumped RLC model can be adequate. Otherwise a higher-order lumped model can be used for accurate simulations, although accurately determining parameters for a higher-order model may be difficult.

References

[1] Philip T. Krein, *Elements of Power Electronics,* Oxford University Press, 1998.

[2] Ch. Joubert, G. Rojat, and A. Beroual, "Magnetic field and current distribution in metallized capacitors", *Journal of Applied Physics*, vol. 76, pp. 5288, 1994.

[3] S. Siami, N. Daude, Ch. Joubert, and P. Merle, "Minimization of the stray inductance in metalized capacitors: Connections and winding geometry dependence", *The European Physics Journal Applied Physics*, vol. 4, pp. 37–43, 1998.

[4] S. Siami, Ch. Joubert, N. Daude, P. Ropa, and C. Glaize, "High frequency model for power electronics capacitors", *IEEE Transactions on Power Electronics*, vol. 16, no. 2, pp. 157–166, 2001.

[5] Ingalls M and Kent G, "Monolithic capacitors as transmission lines", *IEEE Transactions on Microwave Theory and Techniques*, vol. MTT-35, no. 11, pp. 964–70, 1987.

[6] de Vreede L C N, de Kok M, van Dam C, and Tauritz J L, "A high frequency model based on the physical structure of the ceramic multilayer capacitor", *IEEE Transactions on Microwave Theory and Techniques*, vol. 40, no. 7, pp. 1584–7, 1992.

[7] A. Waizman and Chee-Yee Chung, "Extended adaptive voltage positioning (EAVP)", in *IEEE 9th Topical Meeting on Electrical Performance of Electronic Packaging*, 2000, pp. 65–8.