SWITCHED-ACTIVE AND PASSIVE, HYBRID FILTER FOR RESONANCE-FREE TRANSIENT REGULATION

By

ATTMA SHARMA

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2006

This document is dedicated to my daughter, Nirvana Ardra who, from birth, brought heaven closer to earth.

ACKNOWLEDGMENTS

I express my sincere gratitude to my advisor Prof. Khai Ngo for his guidance over the years. His knowledge has been invaluable to me. I thank him for the tremendous kindness he has shown me on so many occasions. I wish to also thank my other committee members for investing their time in reviewing my work, offering new insights and making constructive recommendations. Their words of encouragement from time to time have meant a great deal to me.

In addition, I wish to thank my dear friend and co-worker, Xuelin Wu, for his considerable understanding of my work. Discussions with him have always been enjoyable and fruitful. He has shared references that have broadened my knowledge and helped to keep me on track. I also acknowledge the help of my other friends Yus Ko, Bharath Kannan, Shengwen Xu, Santanu Mishra, Alex Phipps, Ming He and C.G. Cho. Most notably, they have assisted me in using various software and laboratory equipment, finding parts, and loaning me tools. They have also been helpful in discussing ideas and issues.

The funding for this research was provided by CPES under an NSF grant. I'm grateful to these organizations' generous support. Furthermore, I owe Prof. Fred Lee thanks for including me on the CPES research team and for funding my portion of the project. I thank Dr. Ming Xu for sharing key references with me and I thank Julu Sun for allowing me the use of his VRM design in my simulations.

I also acknowledge the motivation, support and compassion of my family, most especially, my wife Vanessa Vidya, my parents and my brother, Avinash.

Finally, I thank God Almighty for bringing me this far at least with such good fortune.

TABLE OF CONTENTS

page

LIST OF TABLES

LIST OF FIGURES

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

SWITCHED-ACTIVE AND PASSIVE, HYBRID FILTER FOR RESONANCE-FREE TRANSIENT REGULATION

By

Attma Sharma

August, 2006

Chair: Khai D.T. Ngo
Cochair: William R. Eisenstadt
Major Department: Electrical and Computer Engineering

Passive filters in the form of cascaded stages of filter capacitors along the power supply path have thus far been used to keep the supply voltage at the point-of-load within specification—typically 100-150 mV of allowable droop under step loads of about 100 A at 1-10 A/ns. However, in faster, higher current, lower voltage microprocessors, for example, an all passive solution may not be realizable.

This dissertation describes a switched-active filter used as an auxiliary circuit to a passive filter in order to improve the dynamic load regulation of a power delivery system.

Useful, second-order models are derived for the multistage passive filter about the resonant peaks in impedance. Techniques to design a Switched-Active Filter based on this design-oriented model are described and the SAF is shown to complement the passive filter in order to minimize ringing in the general multistage power delivery system. Specific application of the resulting hybrid filter to a microprocessor power delivery system is described.

The SAF design comprises a switched-capacitor, pre-charged to the right voltage and energy to naturally inject or absorb the current needed to maintain regulation through pre-determined parasitic inductance and resistance. Once the parasitic inductance in the SAF path is determined from layout or loss constraints, analytic design equations yield the other SAF parameters which depend only on the passive filter, the ripple specification and the maximum load current. The design is flexible enough for the additional circuitry to be placed some distance away from the point-of-load and is therefore not constrained in the same way as the passive filter.

The design methodology is tested by means of scaled models of the microprocessor power delivery path. SAF experiments demonstrate a factor of five improvement in load regulation at acceptable losses, 0.1 μJ and 5 μJ with 20 nH and 200 nH DC path inductance respectively, per 1 A step load, even with 20% tolerance on the design parameters. Normalized power loss is about 0.4 W/MHzA. Simulations show feedback control is feasible while revealing some of the issues that need to be considered in practical implementation.

CHAPTER 1
INTRODUCTION

The dynamic load regulation specifications of power management systems are becoming increasingly stringent as electronic systems target higher speeds. In principle, passive filters may be designed to meet future requirements. However, in certain applications such as microprocessor power delivery, insufficient space for capacitors near the point-of-load (POL) and excessive series interconnect inductance will prevent an all-passive filter solution causing transient-induced ringing to exceed specifications. Described herein is a design methodology which uses a novel switched active filter to augment transient performance of a passive filter in order to efficiently achieve an almost resonance-free response. The analysis, design and experimental verification of the switched-active filter shall be the focus of the dissertation.

Although most of the application discussion will be geared towards future microprocessor power requirements, the techniques developed in this work can be applied to other situations where

- Parasitic path inductances limit available current slew rate.
- Intermittent load current steps occur.
- Sufficient bypass capacitors cannot be placed near the point-of-load (POL).
- A power efficient method of actively suppressing transients is required.

Examples and discussion of a few of the other possible applications will be mentioned later on in this work.

1.1 Review of Relevant Research

Given present layout and packaging technologies, future power delivery systems may require some form of active filter, which results in the system architecture shown in Figure 1.1. In addition to the main power path from the bulk voltage regulator (i.e., the source) to the load, there is an active filter block with the associated feedback control, power source and pre-charging mechanisms. Total loss will be the sum of the losses in the individual circuit blocks, not counting the load. Hence, the most efficient design that still meets load regulation specifications will be an optimization of all blocks taking into consideration any design-impact by one block on another.

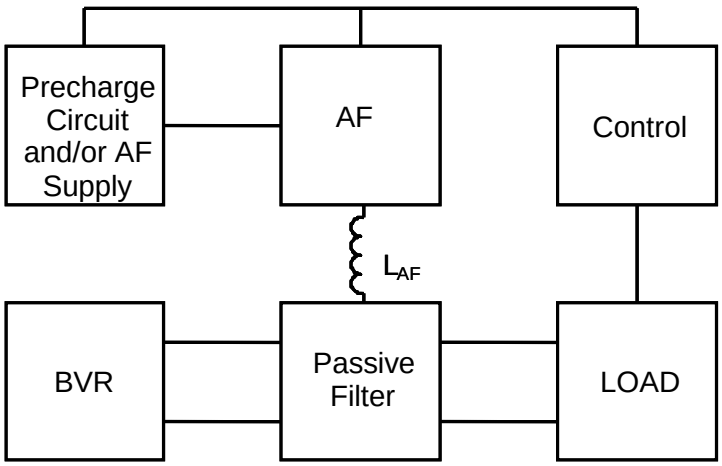

Figure 1.1. Future Power Delivery Architecture.

The required voltage supply, and hence the losses, for the active filter or AF is overwhelmingly a function of the parasitic path inductance, L_{AF}, from the AF power source to the injection point near the load. For a larger L_{AF}, a higher voltage is needed to achieve a given load current slew-rate or di/dt. For the same L_{AF}, losses in switched AF circuits may be lower than in linear AF's thanks to the switched-capacitors' voltage decay while it operates in pull-up mode, and thanks to energy storage while it operates in

the pull-down mode. The energy savings gained by using a switched AF may unwittingly be lost if the pre-charging circuit is poorly designed. In certain cases this may be acceptable, for the same reason that increased losses in the BVR may be acceptable: power dissipated further from the load (and hot-points in general) is less of a thermal concern. Still, optimizing for efficiency, resonant or switching converter techniques may be employed to make the loss in the AF supply and pre-charging circuit almost negligible. Sensing and signal processing in the control block should also consume negligible power.

In the first place, if a passive filter alone can meet load regulation specifications, the three additional active circuit blocks are not needed. A design method for passive filters to meet these specifications has been reported already [Wai01, Man04]. Unfortunately, the required inductance and capacitor values for these designs are not always realizable. Capacitors today have neither high enough capacitance density nor low enough parasitic series inductance to be used ubiquitously. Since the BVR is relatively slow to respond, the capacitors in the passive filter must be large (and hence may not fit in the available space, or, if available in small packages, the associated equivalent series inductance or esl may then be too large). The required capacitance may of course be reduced by increasing the bandwidth (and hence the switching frequency) of the BVR, but in doing so the BVR losses increase and efficiency goes down [Ren04]. With the trend towards higher load currents and lower supply voltages [Sta04], eventually it may become necessary to include the active filter, AF, together with its own power source and control circuitry.

Active filters, AF's, have been reported and continue to be a topic of research. Table 1.1 summarizes the key types of AF's.

Table 1.1. Sampled performance in the last seven years.

	References					
	[Bar05]	[Lim03]	[Wu01]	[Ang00]	[Amo99]	[Poo99]
Circuit	Sw. Reg.	Linear Regulators		Switched Capacitors		LAF
Output (V)	1.5	2.5	1-2	1.6	2.5	5
Load (A)	32	0.4	9	10	13.5	20
Load di/dt	< 0.1 A/ns			4A/ns	~0.15A/ns	< 0.1A/ns
Output ripple	14%	0.4%	1%	4%	2.4%	0.6-1%
Event freq.	50 KHz	< 10 KHz	< 1 MHz	50-70 MHz	< 50 KHz	100 Hz
Location	Off chip	On chip	Off chip	On chip	Off chip	
P_{AF}/P_{load} (%)	~10	32	>33	< 10	33	> 33
Input supply	5	3.3	4	1.6	5	12-18
GBW (MHz)	~0.1	0.03	250	> 200	N/A	3

The use of Linear Regulator techniques to implement the AF [Lim03, Wu01], is very inefficient, and excessive heat loss becomes a problem as load currents and transient event frequencies increase. An early topology (Figure 1.2) shows that the "active filter" block is made up of transconductance amplifiers where the transistors operate in the

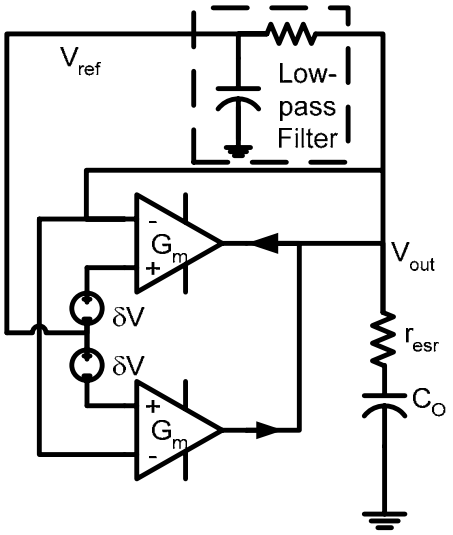

Figure 1.2. An example of a linear regulator, redrawn from: [Wu01].

linear mode. This topology is very lossy as high current flows through these regulating components across which there must be substantial voltage to maintain linear operation. Although the offset threshold voltages, δV , at the inputs of the transconductance amplifiers keep the circuit off when the point-of-load capacitor in the passive filter, C_0 , is sufficient to maintain regulation, for fast response and proper regulation, δV must be much smaller than the allowed ripple and in effect the linear regulator circuit is activated very easily for any sudden load change or transient event. As the rate of these transient events or transient event frequency goes up, the energy savings gained by this offset threshold technique is minimal and is merely an artifact of proper biasing of the two half-circuits. Linear regulator circuits, such as the one in Figure 1.3, with a class-A output stage avoid its use entirely. The linear regulator, in this case, is placed in series with the main supply bus and not in parallel with it. So even at DC the losses are high.

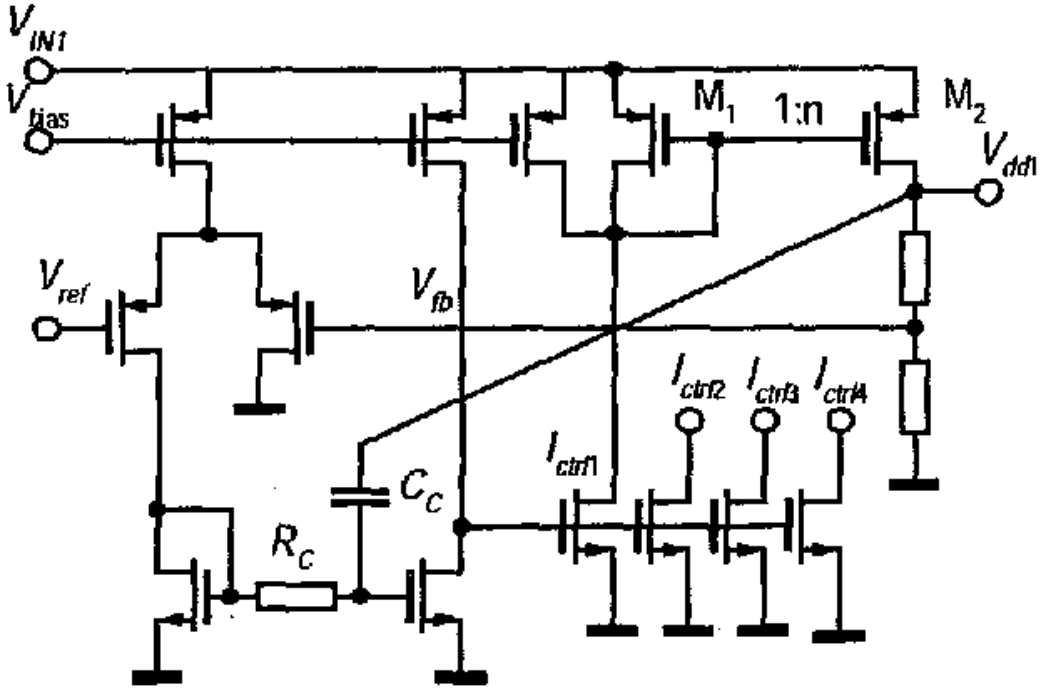

Figure 1.3. Another example of a linear regulator, source: [Lim03].

As summarized in Table 1.1, although the regulation performance in these types of circuits is very good, 0.4% [Lim03] to 1% [Wu01], they are inefficient. Power loss is comparable to the actual power delivered: P_{AF}/P_{load} is one third or more. In principle, the transconductance amplifier regulator [Wu01] placed in parallel with the supply bus would be more efficient but in experiments they appear just as lossy as the classical linear regulator which operates all the time.

By using a filter in the feedback loop to limit the operational range of a Linear Regulator, the resulting linear active filter [Poo99] can be made to operate only when needed for a short time after a sudden load change or transient event. One such topology is shown in Figure 1.4.

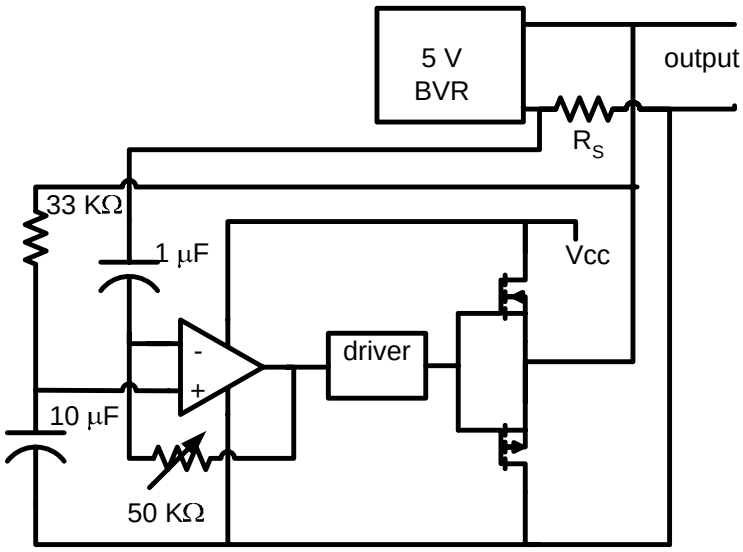

Figure 1.4. Example of a linear active filter, redrawn from: [Poo99].

The 33 KΩ and 10 μF extracts the reference voltage. The voltage across R_S which is proportional to the load current is fed to the opamp circuit. The opamp with the 1 μF, and 50 KΩ potentiometer in a negative feedback connection is configured as a differentiator to sense di/dt of the load current. (In reality this circuit forms the

compensation block seen in Figure 1.1 and is a bandpass filter due to the opamp gain roll-off at the dominant and the 2nd most dominant poles.) For a load transient, the circuit senses the di/dt and the driver operates the push-pull output stage to source or sink the required current to maintain regulation. At DC the circuit is biased so that the driver sets the output transistors to be just barely off. These linear active filter circuits can handle higher load current and higher transient event frequencies than the Linear Regulator approach, but are still lossy. Furthermore, the response is unpredictable if driven into saturation mode by high di/dt's. Note that the additional power source, shown as one block in Figure 1.1, is shown explicitly in Figure 1.4 as Vcc. It is also understood to be present in Figure 1.3 as well although it is not shown.

Switched circuits can be used as active filters. Figure 1.5 shows the topology of a single-shot prototype [Amo99] which demonstrated improved load regulation on load step-up but the design is cumbersome, using two separate branches for current injection and qualitatively derived design equations. The required injection current is first obtained

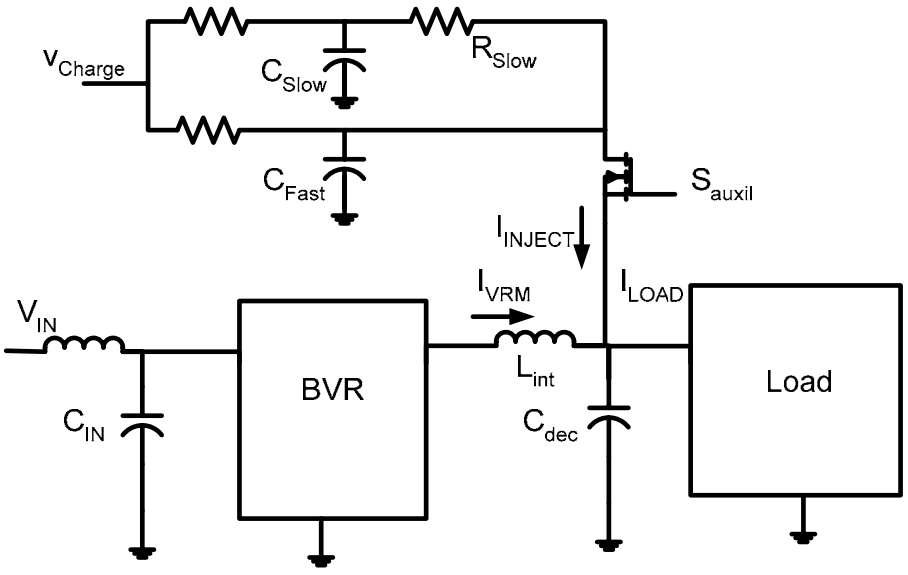

Figure 1.5. An example of a switched active filter, redrawn from: [Amo99].

through simulation for the worst case step load; then the resistor and capacitor networks are chosen such that the discharge closely matches the required injection current. No explicit design equations are given. The capacitors and extra charging voltage are such that there is sufficient energy to regulate the transient. For the slow network the idea of matching time constants is used as a guide in choosing components. The design of the fast network is mostly done by trial and error. The parasitic inductance in the network is not considered in the design phase. Although it is a switching scheme, the data listed in Table 1.1 indicate that the relative efficiency, 33%, at moderate current, 13.5 A, and transient event frequency, less than 50 KHz, is no better than the linear regulator included in the survey. This suggests that the experimental designs were not optimum.

An integrated switched-capacitor AF has been reported by Sun Microsystems Inc. [Ang00]. Although high di/dt, 4 A/ns, and transient event frequency, 70 MHz, have been achieved, this performance is largely due to the process technology which allowed for total on-die capacitance greater than 1.6 uF and low-esl. The key features of the topology (Figure 1.6) are the switched capacitor charge pump circuit controlled by an average voltage tracking loop. On a pull-up load transient, C1 and C2 are connected in series, ideally doubling the effective voltage which then quickly discharges its energy into the load, thereby providing pull-up regulation. By the end of the discharge time, the voltage across the series combination settles to V_{dd}-V_{ss}. If a pull-down transient comes, capacitors are switched into a parallel bank at half the total series voltage, ideally half of V_{dd}, a lower voltage which helps pull-down regulation as the capacitors absorb the excess energy. However, a close study of the experimental data revealed the switch-capacitor circuit itself only afforded small improvement relative to a static 1.6 uF on-die capacitor.

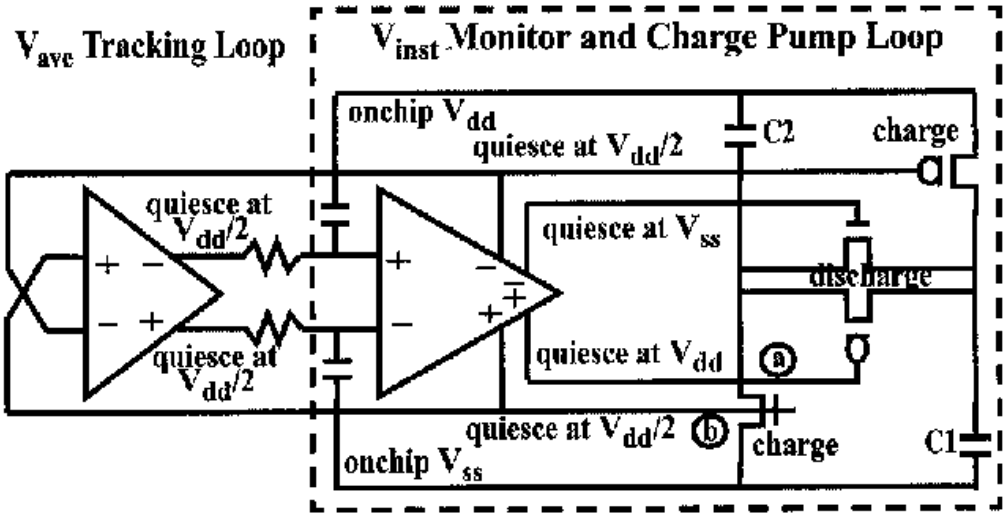

Figure 1.6. Another example of a switched active filter, source: [Ang00].

Furthermore, since parasitic inductances and resistances were ignored in the analysis, the circuit concept fails if applied to more general situations off-chip. The IC implementation did show that if care is taken, and if layout constraints allow, the active filter may run off of the main supply bus but there is no guarantee of this in general. Since the circuit merely switches capacitor bank configurations, there is never a stiff voltage source injecting current from the AF directly into the load. This makes the system naturally more stable and easier to control because there is less chance for excessive energy to be injected into the system. Furthermore, $P_{AF}/P_{load} < 10\%$ is very good.

On the other hand, for comparable relative efficiency, the use of an additional switching DC-DC converter in parallel with the main VRM [Bar05] requires more careful control (both linear and non-linear control must be used). The diagram of the topology is shown in Figure 1.7. The lower switching regulator is designed to be the main power supply, with linear feedback control and low ripple. Consequently it is slow. The output voltage sense signal, however, not only goes to the linear feedback control circuit but also to a digital threshold logic circuit block that switches very fast to control the upper buck

regulator. This regulator is designed to provide high di/dt's without regard for switching ripple. The inductance in the output filter of the upper switching regulator is about 10 times lower than the equivalent component in the main regulator, 0.7 μH versus 7 μH. For an input voltage of 5 V and nominal output voltage of 1.5 V, maximum positive di/dt available from the fast regulator is only $(5-1.5)/0.7 = 5\,\text{A}/\mu\text{s}$ and negative di/dt is even lower. In principle a larger input voltage and a smaller inductance would give higher di/dt's. However, with this method, the designer runs into trouble trying to find sufficiently fast components far sooner than if using any of the other approaches previously discussed. This is because, in addition to precise switch-on, very fast and precise switch-off is necessary. This method, therefore, is better suited for slower, off-chip implementation.

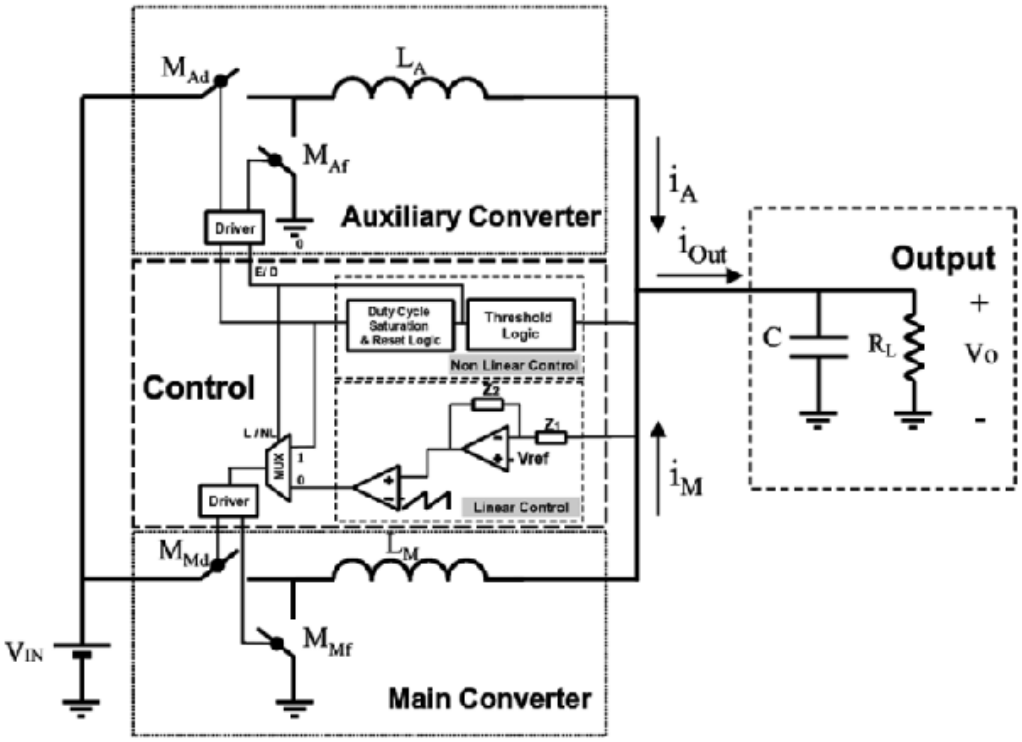

Figure 1.7. Example of an auxiliary switching regulator used in fast transient regulation, source: [Bar05].

1.2 Research Goal and Contributions

Thus far, the problem that has been largely ignored is two-fold. The impact of the passive filter on the design of the active filter and the impact the design of the active Filter, in turn, has on power loss have not been considered in depth. Furthermore, open loop or saturation mode performance has not been studied in detail even though high di/dt transients usually drive the AF into this mode of operation initially.

Therefore, the goal of the dissertation is to show the coordinated design of a wwitched-active and passive hybrid filter topology that achieves dynamic load regulation when driven in the open loop mode. First the design equations for the switched-active filter or SAF are derived based on the equivalent passive filter components. These equations and the corresponding design methodology are then verified experimentally.

The open-loop SAF circuit topology is essentially the equivalent saturation mode circuit in the other AF topologies previously described. So the SAF design equations lend additional insight into the design and performance analysis of these other AF topologies. More importantly, the SAF may be used with feedback control or advance warning signals in place of the other AF topologies for dynamic load regulation of power supplies with relative benefits as summarized in Table 1.2.

Despite the excellent line and load regulation performance associated with linear regulators, they are very lossy- especially when designed for continuous operation. The linear regulator response is unpredictable when driven into saturation- it may incur too much overshoot or undershoot, and may even oscillate. The SAF offers the potential for lower loss and is naturally designed for a stable open loop response. The SAF does share some common design concerns with linear regulators that are off at steady state, namely, the threshold triggering mechanism and the choice of sufficiently fast components.

Table 1.2. Relative advantages of the SAF over other AF topologies.

Topology	Features	Relative Advantages of the SAF
Linear regulator	Stiff extra supply, wideband	Less loss, works when driven into saturation, pre-charged capacitor serves as a soft supply
Linear active filter	Stiff extra supply, narrow bandwidth	Slightly less loss, works when driven into saturation, pre-charged capacitor serves as a soft supply (which affords better stability)
Switched capacitor AF	Pre-charges and switches capacitor-bank configuration	Simpler design, parasitic inductance included in the analysis and design, well-defined design methodology, may be implemented on-chip as well as off-chip
Switched regulator	Stiff extra supply, parallel buck converter topology	pre-charged capacitor serves as a soft supply, sharp turn-off not needed, simple linear control feasible (usually), works with higher di/dt's

Linear active filters can be low loss and are comparable to the SAF in this respect. However, since it operates off of a stiff power source, instead of a pre-charged capacitor with a smoothly decaying discharge curve, the SAF may be designed for slightly lower loss. Also due to the "soft" supply formed by the pre-charge capacitor, the SAF is again likely to be more stable.

The SAF and by extension, the hybrid filter, is far simpler to design and apply practically than the other switched capacitor techniques surveyed. This is expected to remain true even with the inclusion of the feedback control and pre-charge circuits for the SAF based on promising simulations of the envisioned schemes. However, it is not expected that the SAF will be as efficient as some of the other switched capacitor schemes that are designed to be ideally lossless. One should note however that any practical implementation of such ideally lossless switching schemes is very likely to suffer from excessive ringing and may prove even more difficult to control and stabilize.

This is particularly applicable to the switched regulator approach. While high efficiencies are possible, the design criteria placed on the control circuit are very

demanding. Given present technology, it is doubtful that this approach is capable of handling di/dt's much greater than 25 A/μs.

The SAF will be presented as a flexible auxiliary circuit that may function in a hybrid filter to provide dynamic load regulation for a short time after a sudden load transient. Furthermore, the SAF will be discussed as a building block which may be used in other power delivery applications such as EMI filtering and switched-coil damping.

1.3 Practical Basis for the Experimental Setup

Equipment availability and cost constraints aside, the fact that the development of the SAF was primarily intended for application to future microprocessor power delivery systems, which themselves are yet to be fully developed and disclosed, posed a peculiar obstacle in doing the related experiments. Presently, only a passive filter similar to the ladder network of Figure 1.8 is used to provide the necessary load regulation. Parameter values for realistic models are given in the first two data columns of Table 1.3.

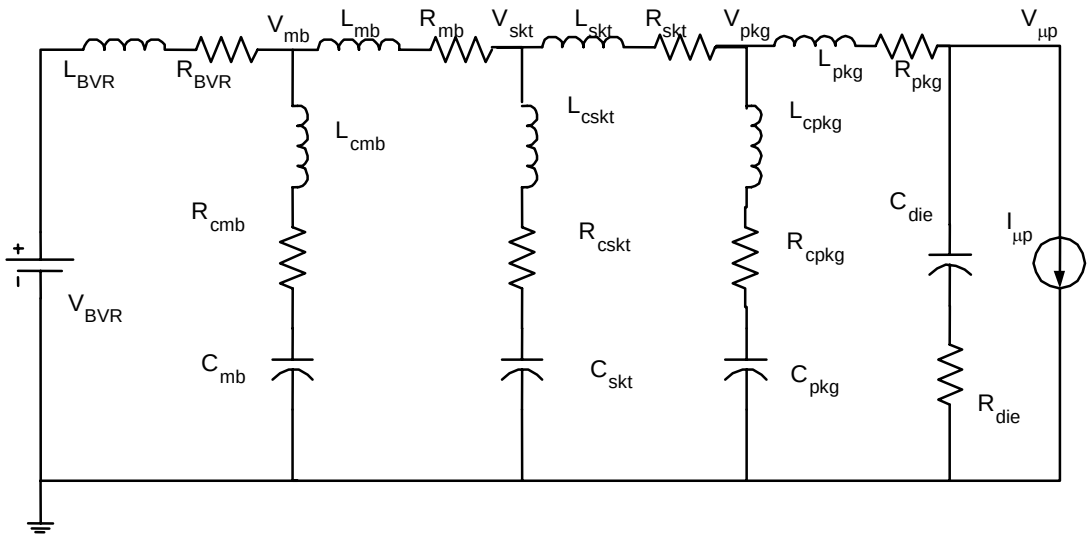

Figure 1.8. Model of the microprocessor power delivery path.

Table 1.3. Power Delivery Path Model Parameters.

Model Element	Model Parameters	1998 Model [Luo02] 478 pin PGA	2004 Model 775 pin LGA	Magnitude (Impedance) Scaling 100x	Frequency Down-shifting 10x
C_{mb}	C_{mb} [uF]	5 600	6 700	67	670
	L_{cmb} [pH]	800	700	70 nH	700 nH
	R_{cmb} [mΩ]	1	0.6	60	60
L_{mb}	L_{mb} [pH]	21	100	10 nH	100 nH
	R_{lmb} [mΩ]	0.1	0.2	20	20
C_{skt}	C_{skt} [uF]	240	850	8.5	85
	L_{cskt} [pH]	340	140	14 nH	140 nH
	R_{cskt} [mΩ]	0.17	0.2	20	20
L_{skt}	L_{skt} [pH]	120	40	4 nH	40 nH
	R_{lskt} [mΩ]	1.1	0.6	60	60
C_{pkg}	C_{pkg} [uF]	26	23	0.23	2.3
	L_{cpkg} [pH]	4.6	1	0.1 nH	1 nH
	R_{cpkg} [mΩ]	0.54	~0.02	2	2
L_{pkg}	L_{pkg} [pH]	6	~3	0.3 nH	3 nH
	R_{lpkg} [mΩ]	.03	~0.015	1.5	1.5
Micro-processor	C_{die} [uF]	0.53	~1	0.01	0.1
	R_{die} [mΩ]	0.1	~0.7	10	10
	L_{die}	Assumed to be negligible			
	I_{up}	~100 A, ± 5-100 A/ns		~1 A, 0.05-1 A/ns	

The 2004 Model data was obtained through private communications with the Center for Power Electronics Systems, CPES, headquartered at Virginia Tech. It was taken as the best available power delivery filter to date using only passive elements. Maintaining a desired ripple specification of about 100 mV, the model was scaled down in magnitude 100 times to allow for testing on the order of 1 A at 0.05-1 A/ns. This effectively scales all impedances by 100. All resistances and inductances scaled by 100x while the capacitors scaled 1/100x. In order to make measurements of inductance and time with acceptable accuracy, using the available equipment, it was further necessary (in some cases) to frequency shift the model 10x to a lower range. Keeping all resistances fixed, all capacitors and inductors in the model were scaled up 10x.

As a side-note, the wide range in di/dt stemmed from the difference in values that are merely desired or claimed to values that are experimentally verifiable. Currents slew-rates of up to 100 A/ns have been predicted and are certainly desired. However, the best certifiable electronic load (manufactured by Chroma Inc.) is only capable of 1 A/ns and the most aggressive claim by any electronic load manufacturer is 3.5 A/ns. Hence 5 A/ns was aggressively taken as the best di/dt that could be measured.

The scaled model parameters of the last two columns in Table 1.3 served as the basis for the experimental work. The stages most frequently used were at the package (subscript "pkg") and the socket (subscript "skt") levels.

1.4 Content Organization

The remaining chapters delve into design theory, experimental verification, applications and a discussion of issues related to the SAF.

The SAF topology is described and the design methodology developed in Chapter 2. Design equations for the SAF are derived, normalized and then used to identify key design tradeoffs in the same chapter. In Chapter 3 the SAF theory is verified experimentally by means of scaled test circuits comprising printed circuit boards, discrete devices and custom IC's. Other issues relevant to practical application such as pre-charging and control are also discussed. Chapter 4 delves into applications of the SAF, focusing on its use in microprocessor power delivery systems. Ideas for future work are presented that offer viable solutions when final stage capacitor esl becomes significant and the lumped element models breakdown. Further applications to EMI filter systems, and active damping of switched coils are suggested before concluding in Chapter 5.

CHAPTER 2
DESIGN THEORY AND ANALYSIS

2.1 Equivalent Passive Filter Model and SAF Topology

Since the SAF discussed herein is intended to supplement an imperfect EAVP passive filter, the latter is first reviewed so that the issues can be identified and the approximations made. Later it will be shown that the equivalent single stage model that shall be derived for the EAVP passive filter is also applicable to multistage passive filters in general.

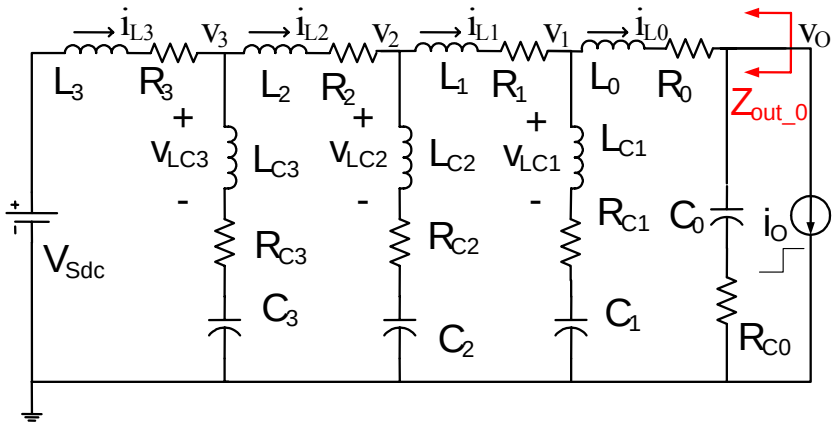

Figure 2.1. Multistage passive EAVP filter model.

Shown in Figure 2.1 is an EAVP passive filter connected between a well-regulated low-frequency bus V_{Sdc} and a microprocessor load represented by i_O, C_0, and R_{C0}. Parasitic inductance associated with C_0, is assumed to be zero, while i_O is assumed to be an ideal step for simplicity. Real effects of non-zero esl and finite di/dt are discussed later on. If the filter components are designed appropriately [Wai01], Z_{out_0} and the impedance seen into each filter section equal a real value R_{out_n}. As the reactive components become progressively smaller toward the high-speed end, they are more difficult to realize, i.e.,

the EAVP design constraints tend to be violated at or near the load, resulting in a resonant spike in the impedance.

Let the filter section responsible for the resonance peak in the output impedance be the section containing C_n, R_{Cn}, L_{Cn}, R_n, L_n, L_{C_n+1}, R_{C_n+1}, and C_{n+1}. (The reader could assume n = 1 to facilitate understanding.) This section is excited by i_{L_n-1} at node v_n, and by i_{L_n+1} at node v_{n+1}. To see how the section could be reduced to a simple circuit that facilitates SAF design (Figure 2.3, eventually), the key voltages and currents on the section are plotted in Figure 2.2.

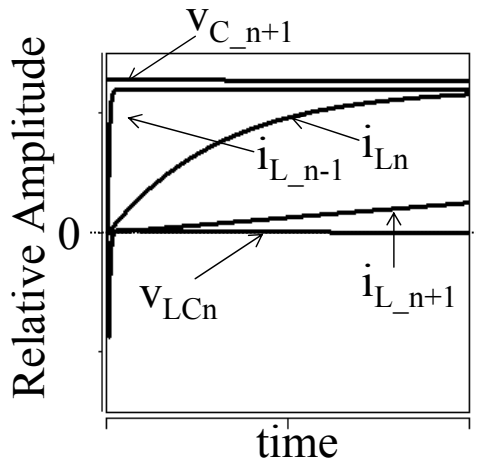

Figure 2.2. Typical waveforms of the passive EAVP filter.

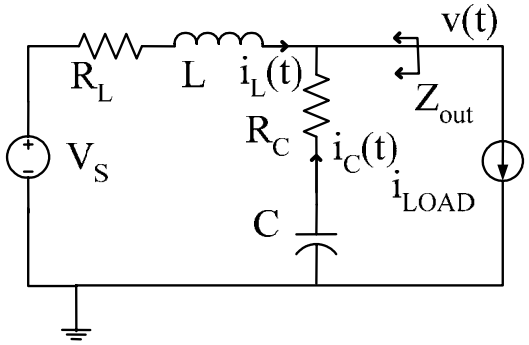

Figure 2.3. Equivalent single-stage passive filter model.

Consider first the current waveforms. Since i_{L_n-1} rises to its steady state much quicker than i_{Ln}, i_{L_n-1} can be approximated as a step current source in the analysis of i_{Ln}. Thus, the inductive current i_{L_n-1} exciting node v_n in Figure 2.1 is represented as the current source i_{LOAD} exciting node *v(t)* in Figure 2.3. On the other hand, since i_{L_n+1} rises to its steady state much slower than i_{Ln}, it can be approximated as an open circuit while the SAF operates. Thus, i_{L_n+1} is not seen in Figure 2.3; the series combination of L_n and L_{C_n+1} is represented as L, and the series combination of R_n and R_{C_n+1} by R. Consider now v_{LCn} and v_{C_n+1} waveforms. Since $v_{LCn} \sim 0$ (except for a very narrow initial spike) while i_{L_n-1} is approximated as a constant current, L_{Cn} is approximated as a short, leaving C_n as C and R_{Cn} as R_C in Figure 2.3. Chapter 4 will address design options to deal with L_{Cn} when the magnitude of the spike induced by it poses a problem.

Since v_{C_n+1} is essentially constant, C_{n+1} is represented as the stiff source V_S in Figure 2.3.

The model of Figure 2.3 is also applicable to the general multistage passive filter which exhibits multiple resonances as in Figure 2.4 for example. After obtaining the equivalent impedance at node n, Z_{out_n}, by simulation or impedance analyzer measurement, each resonant peak in the impedance can be fitted with a second-order impedance curve corresponding to the model of Figure 2.3.

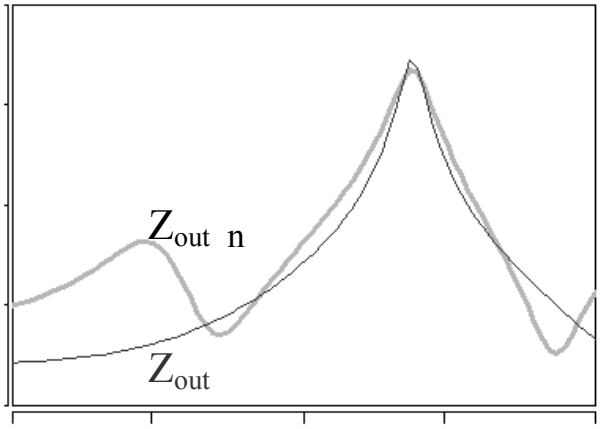

Figure 2.4. General multistage passive filter impedance and second order fit to dominant resonance peak.

Given the simplified, yet generalized, section model Figure 2.3, resonance-free Z_{out} (= R_{out}) and transient response are achieved when the EAVP design constraints $R_C = R = \sqrt{L/C} = R_{out} = \Delta V/\Delta I$ are satisfied [Wai01], where ΔI is the step load current and ΔV the step output voltage.

Power is normally supplied from source, V_S through the "DC" or low frequency path modeled as L in series with R_L. High frequency components of the power to the load flow mainly from the point of load filter capacitance represented by C in series with some effective esr, R_C . Altogether, L, R_L, C, and R_C form a passive filter.

When these passive filter parameters cannot be designed, practically, to meet the load regulation specification, the SAF (highlighted in Figure 2.5) comprising a switch, a capacitor C_{AF} with initial voltage V_{EXT}, an inductor L_{AF}, and a resistor R_{AF} may be used to achieve the desired response by injecting some additional current i_{AF} (t) for a short period after a sudden transient.

The SAF is a narrow-band filter that is activated by an event scheduler to deliver/absorb a short current pulse, and then remains off. A simple pre-charger using V_{EXT} and a resistor is shown to establish V_{AF}.

The value of V_{AF} is set to overcome inductance L_{AF} to provide the initial high di/dt. C_{AF} is chosen to store just the required amount of charge needed for the voltage to settle to steady state with a time constant dependent on R_{AF} after which the main DC path takes over. The SAF parameters are dependent on the passive filter parameters hence the hybrid nature of the overall filter operation. If the timing of the transient event is known in advance, as is often the case in digital signal processing applications, the SAF may be programmed to switch in time with the transient edge. Otherwise fast sensing and feedback with acceptable delay may be used to control the SAF. In this work, however, open loop control is assumed. Methods to deal with significant esl, the design of more efficient pre-charge topologies and feedback control feasibility for arbitrary load profile are future topics to be addressed later.

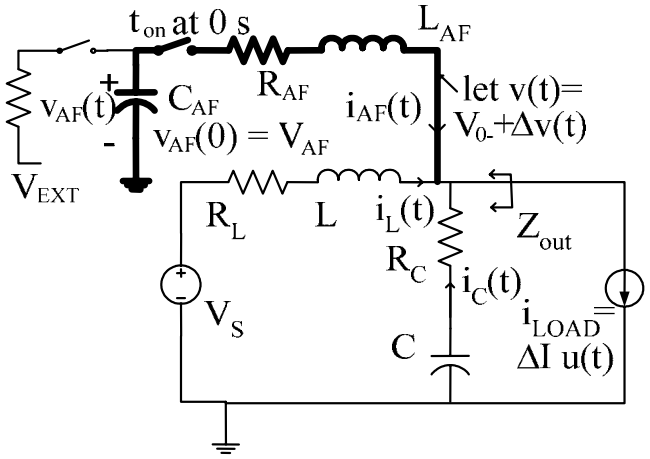

Figure 2.5. Simplified diagram of the hybrid filter. The switched active filter is highlighted in bold print.

Extremely detailed and rigorous analysis of a 4th order network may prove unnecessarily tedious. Insights drawn from the analysis of two special cases shall yield useful design equations, a general design methodology and a qualitative understanding of the circuit behavior in all cases.

In the first analysis, the ability to damp the passive filter with precise series resistors is assumed. This is the best case design situation. The EAVP [Wai01] or purely resistive output impedance response is imposed as a specific condition on the design. Next, the worst case analysis is done by assuming a lossless passive filter which exhibits maximum ringing. Finally, it will be demonstrated that the design equations converge in the most practical cases. Loss calculations are shown to be applicable in both cases. By comparing the voltage responses for both EAVP and non-EAVP designs, general circuit design and response of the passive and switched-active hybrid filter will be discussed along with examples simulated in Synopsys Saber.

2.2 SAF Design for the Damped Passive Filter (EAVP)

Filter design for purely resistive $Z_{out} = R$, commonly termed EAVP [Wai01], ensures that the response to a step load current, $i_{LOAD}(t) = \Delta I \cdot u(t)$, is a step output voltage: $v(t) = V_{0-} + \Delta V \cdot u(t)$, where

$$\Delta V = -\Delta I \cdot R \tag{2.1}$$

and V_{EXT} is the unit step function. Just after the step load, $V_{0+} = V_{0-} + \Delta V$.

To generalize the design of the SAF, the following normalized variables are defined:

$$Q = \frac{\sqrt{L/C}}{2R};\ K_L = \frac{L_{AF}}{L};\ K_C = \frac{C_{AF}}{C};\ K_R = \frac{R_{AF}}{R};\ K_V = -\frac{V_{AF} - V_{0+}}{\Delta V} \tag{2.2}$$

Imposing the desired step voltage response to the specified step in load current, in the Laplace frequency domain, the required transient current that must be injected by the active filter in order to achieve perfectly regulated response appears as

$$\begin{aligned}\mathbf{I_{AF}}(\mathbf{s}) &= \frac{C_{AF}\left(V_{AF}-V_{0+}\right)}{1+sR_{AF}C_{AF}+s^2L_{AF}C_{AF}} = \mathbf{I_{LOAD}}-\left(\mathbf{I_L}+\mathbf{I_C}\right)\\ &= \frac{\Delta I}{s}+\frac{\Delta V}{s}\left[\frac{1}{R_L+sL}+\frac{sC}{1+sR_CC}\right] = \frac{\Delta I}{s}-\frac{\Delta I\cdot R}{s}\left[\frac{1/R_L}{1+sL/R_L}+\frac{sC}{1+sR_CC}\right]\\ &= \frac{\Delta I}{s}\frac{1-R/R_L+s\left(L/R_L-\left(R+R_C/R_L-R_C\right)C\right)+s^2LC\left(R_C/R_L-R/R_L\right)}{1+s\left(R_CC+L/R_L\right)+s^2LC\,R_C/R_L}\end{aligned} \tag{2.3}$$

Note, the network is linear in each mode of the switch. Therefore, the DC components can be set to zero in deriving (2.3) in order to extract just the transient response. By matching the corresponding terms in (2.3), the design equations for the SAF can be derived as follows:

$$R_C = R_L = -\Delta V/\Delta I = R \tag{2.4}$$

$$L_{AF}C_{AF} = LC \Rightarrow K_L = 1/K_C \tag{2.5}$$

$$R_{AF}C_{AF} = L/R + RC \Rightarrow K_R = \left(4Q^2+1\right)/K_C \tag{2.6}$$

$$C_{AF}\left(V_{AF}-V_{0+}\right) = \Delta I\left(L/R - RC\right) \Rightarrow K_V = \left(4Q^2-1\right)/K_C \tag{2.7}$$

Note the corresponding quality-factor of the SAF network,

$$Q_{AF} = \frac{\sqrt{L_{AF}/C_{AF}}}{R_{AF}} = Q\frac{2R}{\sqrt{L/C}}\frac{\sqrt{L_{AF}/C_{AF}}}{R_{AF}} = \frac{2Q}{1+4Q^2} \tag{2.8}$$

is always less than 0.5, or non-resonant.

Equation (2.4) is a requirement already derived in [Wai01]. It states that, for the desired ideal EAVP response, the esr values in the model of the damped passive filter must be equal to the desired output resistance. However, it shall be shown that the rest of

the design equations are still useful in improving regulation even if (2.4) is not satisfied exactly.

The characteristic frequency of the SAF must equal that of the passive filter, (2.5). The time constant of SAF is equal to the sum of the time constants in the "DC" and "high-frequency" branches of the equivalent passive filter, (2.6). For Q > 2, K_V is essentially K_R, and R_{AF} scales with V_{AF} via (2.6) and (2.7). Note that (2.7) shows V_{AF} is proportional to ΔI and the level of deviation from the passive EAVP condition (which requires the time constants in the "DC" and "high-frequency" branches to be equal). Furthermore, (2.7) reveals the design dependence on Q, the quality factor of the passive filter alone when (2.4) is satisfied. If $Q = 0.5$, the EAVP design constraint is met and no SAF is needed. If $Q < 0.5$, the passive filter keeps the output voltage excursion within ripple specification, but exhibits an overdamped settling response. The SAF might be used to take advantage of the quick settling of the resonance-free transient. The SAF is most useful when $Q > 0.5$; without it, there may be excessive ringing.

Furthermore, for $Q > 1.25$, $Q_{AF} < 0.35$, $R_{AF} \to \dfrac{V_{AF} - V_{0+}}{\Delta I}$, $L_{AF} << R_{AF}^2 C_{AF}$, response is insensitive to L_{AF}, and the SAF behaves as an RC circuit with the time constant, $\tau_{AF} \approx R_{AF} C_{AF} = RC + L/R = (1 + 4Q^2)RC$. The current peaks almost instantly, then settles to zero within $5\tau_{AF}$. Note that the settling time of the SAF depends only on the passive filter parameters.

With four SAF variables in three equations, there is one degree of freedom in the SAF design. The set of allowable designs depends on layout and loss constraints.

2.3 SAF Design for the Lossless Passive Filter (non-EAVP)

With $R_L = R_C = 0$, $V_{0+} = V_{0-} = V_S$ and EAVP is no longer applicable. Instead, the ripple is now a general function of time, $\Delta v(t)$, which is described in the Laplace domain by

$$\mathbf{\Delta V(s)} = \frac{-\Delta I \cdot L}{1 + s^2 LC} \frac{1 + \left(R_{AF} - \frac{V_{AF} - V_{0-}}{\Delta I} \right) C_{AF} \cdot s + L_{AF} C_{AF} \cdot s^2}{1 + s R_{AF} C_{AF} + s^2 LC} \tag{2.9}$$

Typically for large ΔI, the expressions in terms of SAF parameters in the numerator and the denominator approximately cancel leaving $-\Delta I \cdot L \big/ \left(1 + s^2 LC\right)$, the inverse of a sum of a constant and a square in s, which appears as prolonged sinusoidal oscillation in the time domain. With care, however, the SAF can be designed to cancel this factor that is primarily responsible for the ringing.

Keeping

$$L_{AF} C_{AF} = LC \Rightarrow K_L = 1 \big/ K_C \tag{2.5}$$

while letting

$$R_{AF} = \frac{V_{AF} - V_{0-}}{\Delta I} = \frac{V_{AF} - V_{0+}}{\Delta I} \tag{2.10}$$

equation (2.9) simplifies to a standard second order expression:

$$\mathbf{\Delta V(s)} = \frac{-\Delta I \cdot L}{1 + s R_{AF} C_{AF} + s^2 LC} \tag{2.11}$$

Keeping (2.2) while further defining

$$\omega_0 = 1 \big/ \sqrt{LC} \tag{2.12}$$

$$Q_{AF} = \frac{\sqrt{L_{AF}/C_{AF}}}{R_{AF}} = 2Q\frac{K_L}{K_R} \tag{2.13}$$

$$a = -\frac{\omega_0}{2Q_{AF}}\left(1+\sqrt{1-4Q_{AF}^2}\right) \tag{2.14}$$

$$b = -\frac{\omega_0}{2Q_{AF}}\left(1-\sqrt{1-4Q_{AF}^2}\right) \tag{2.15}$$

one obtains,

$$\Delta v(t) = -\Delta I\,\omega_0^2 L\frac{e^{bt}-e^{at}}{b-a} \tag{2.16}$$

for which the extremum occurs at

$$t_{\max} = \frac{\ln(a/b)}{b-a} \tag{2.17}$$

and is given by

$$\Delta V_{\max} = \Delta I\,\omega_0^2 L\frac{e^{\frac{a}{b-a}\ln\frac{a}{b}}}{b} \tag{2.18}$$

For the highly overdamped SAF, $Q_{AF} << 0.5$, ($Q_{AF} < 0.35$ is better than 12% accurate)

$$\begin{aligned}\Delta v(t) &\approx -\Delta I\,\omega_0 L Q_{AF}\exp\left(-Q_{AF}\omega_0 t\right)\\ &= -\frac{\Delta I\cdot L}{RC}\frac{K_L}{K_R}\exp\left(-\frac{K_L}{K_R}\frac{t}{RC}\right)\end{aligned} \tag{2.19}$$

but with

$$t_{\max} \approx \frac{-2Q_{AF}\ln Q_{AF}}{\omega_0} = -2\frac{L}{R}\frac{K_L}{K_R}\ln Q_{AF} \tag{2.20}$$

Using (2.1) and (2.2) in (2.19), the extremum is approximately given by

$$\Delta V_{\max} \approx -Q_{AF}\Delta I\omega_0 L\cdot\frac{2R}{2R} = -2QQ_{AF}\Delta I\cdot R$$
$$= 2QQ_{AF}\Delta V = 4Q^2\frac{K_L}{K_R}\Delta V;\ Q_{AF} < 0.35 \tag{2.21}$$

The expression for $\Delta V_{\max}$ is put in terms of ΔV so that comparison can later be made with the SAF design equations for the damped passive filter.

Approaching the critically damped case, $Q_{AF} \sim 0.5$,

$$\Delta v(t) \to -\Delta I\omega_0^2 Lte^{-\frac{\omega_0}{2Q_{AF}}t} \tag{2.22}$$

$$t_{\max} \to 2Q_{AF}/\omega_0 \tag{2.23}$$

and

$$\Delta V_{\max} \approx -2Q_{AF}\Delta I\omega_0 L/e = -4QQ_{AF}\Delta I\cdot R/e = \frac{2Q\Delta V}{e};\ 0.35 < Q_{AF} < 0.5 \tag{2.24}$$

Now, in terms of ΔV, the peak ripple of the lossless passive filter alone is

$$\left|\Delta V_{pk}\right| = \left|\Delta I\omega_0 L\right| = \left|\Delta I\sqrt{L/C}\right| = \left|\Delta I\cdot R\cdot\frac{1}{R}\sqrt{L/C}\right| = \left|2Q\Delta V\right| \tag{2.25}$$

Thus for a desired ΔV ($2Q$ improvement), simply set $\Delta V_{\max} = \Delta V$; and (2.21) implies

$$Q_{AF} = \frac{1}{2Q};\ Q_{AF} < 0.35 \text{ or } K_R = 4Q^2K_L;\ Q > 1.5 \tag{2.26}$$

It should be noted that, as implied by (2.24)

$$Q_{AF} \approx \frac{1}{2} \Rightarrow \Delta V_{\max} \approx \Delta V_{pk}/e \tag{2.27}$$

In other words, the use of the SAF automatically provides a factor of e or better improvement in ripple.

For L, C, ΔV and ΔI specified, the SAF design equations for the lossless passive filter are (2.5), (2.10), (2.13) and either (2.26) or (2.27). For practical purposes, (2.27)

may be ignored since (2.26) yields a conservative design when $Q_{AF} \sim 0.5$. If this is done, the SAF design equations can be summarized as

$$\Delta V = -\Delta I \cdot R \tag{2.1}$$

$$Q = \frac{\sqrt{L/C}}{2R};\ K_L = \frac{L_{AF}}{L};\ K_C = \frac{C_{AF}}{C};\ K_R = \frac{R_{AF}}{R};\ K_V = -\frac{V_{AF} - V_{0+}}{\Delta V} \tag{2.2}$$

$$L_{AF}C_{AF} = LC \Rightarrow K_L = 1/K_C \tag{2.5}$$

$$Q_{AF} = \frac{1}{2Q} \tag{2.26}$$

$$R_{AF} = \frac{\sqrt{L_{AF}/C_{AF}}}{Q_{AF}} \Rightarrow K_R = \frac{4Q^2}{K_C} \tag{2.28}$$

$$V_{AF} - V_{0+} = R_{AF}\Delta I \Rightarrow K_V = K_R \tag{2.29}$$

Now, high-Q passive filters are approximately lossless. Hence the SAF design of the damped passive filter converges to the lossless case as the quality factor, Q, increases.

2.4 Loss Calculations

The energy lost in the SAF in Figure 2.5 is the same as that lost in a series RLC circuit which discharges to zero from zero initial current and ($V_{AF} - V_{0+}$) initial voltage:

$$E_{loss} = \frac{1}{2}C_{AF}\left(V_{AF} - V_{0+}\right)^2 \tag{2.30}$$

and is valid for both the damped and lossless cases. For the damped (EAVP) case,

$$E_{loss} \approx \frac{C \cdot \Delta V^2}{2K_C}\left(4Q^2 - 1\right)^2 = \frac{L \cdot \Delta I^2}{2K_C} \cdot \frac{\left(4Q^2 - 1\right)^2}{4Q^2} = \frac{L_{AF} \cdot \Delta I^2}{2} \cdot \frac{\left(4Q^2 - 1\right)^2}{4Q^2} \tag{2.31}$$

whereas for the lossless passive filter case,

$$E_{loss} \approx \frac{L \cdot \Delta I^2}{2K_C Q_{AF}^2} = \frac{L_{AF} \cdot \Delta I^2}{2Q_{AF}^2} \tag{2.32}$$

Since $Q_{AF} < 0.5$ always, the minimum loss in the SAF for the lossless case is $2L_{AF} \cdot \Delta I^2$. As Q gets larger, there is more ringing in the passive filter, $Q_{AF} \sim 1/(2Q)$, and the loss in the SAF increases quadratically.

The SAF may be placed close to the load, in which case L_{AF} is smaller at the cost of larger C_{AF} via (2.5), or well removed from the load, but then the loss increases. However, when the SAF is located close to the load, as is possible in IC applications, the space that is already allocated to bypass capacitors may be shared with the SAF capacitor. When the SAF is far removed from the load, which is typically a "hot-point", the extra power loss incurred may be easier to dissipate.

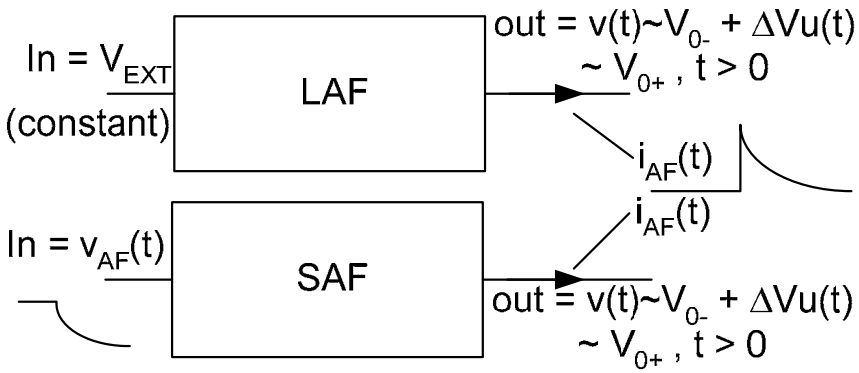

Figure 2.6. Comparison of LAF to SAF.

If the input power supply were used to power a linear regulator or a linear active filter, instead of the SAF (Figure 2.6) then, for comparison, the loss in the linear active filter or LAF, could be calculated in terms of SAF design values using the knowledge that, for perfect regulation, i_{AF} is the same regardless of AF implementation. So, using (2.3),

$$E_{lossLAF} \approx \int_0^{\infty} (V_{EXT} - V_{0+})\, i_{AF}\, dt = (V_{EXT} - V_{0+})\left[\mathbf{I_{AF}}\right]_{s=0} = C_{AF}\,(V_{AF} - V_{0+})(V_{EXT} - V_{0+}) \tag{2.33}$$

Not surprisingly, LAF loss increases linearly with V_{EXT}. However, given there must be some parasitic inductance, L_{AF}, the lowest possible V_{EXT} is bound to be close to the V_{AF} value set by (2.7), neglecting drop-out voltage. Therefore, for the worst case step-load, loss in the SAF is about half that of the best-designed LAF. Furthermore, since the design of the LAF for efficiency has not been essential in the past, supply voltages used in previously reported LAF's are considerably larger than the minimum value indicated by (2.7) and the resulting loss is also unnecessarily large.

However, energy saved by the SAF may easily be lost if the pre-charge circuit for the SAF is poorly designed. Resonant, switching converter techniques should be used in the design of the pre-charge circuit for optimum efficiency. These shall be discussed in Chapter 5.

Another indirect advantage of the energy-efficient design of the SAF is improved stability. Oscillations and instability are more likely to be fed by a stiff source with unlimited energy available. Since the SAF has stored only just enough energy to provide regulation of a single step, while it operates, it fades into a purely passive (an inherently stable) network.

2.5 Design Methodology

For convenient reference, the key design equations are summarized in Table 2.1. If it is possible to design the passive filter such that its equivalent single stage model satisfies the all-passive EAVP criteria given in Table 2.1, then no active filter is

necessary and loss may be kept to a bare minimum given by $\frac{1}{2}C\Delta V^2$ approximately. The all-passive EAVP filter design is discussed in Appendix A. However, more often than not, the values of the passive components required to ensure the EAVP criteria are met are not practical. The equivalent "DC" inductance, L, may be too large or the equivalent point-of-load capacitance, C, may be too small.

Table 2.1. Summary of design equations.

Definitions:

$$\Delta V = -\Delta I \cdot R \;; Q = \frac{\sqrt{L/C}}{2R};\; K_L = \frac{L_{AF}}{L};\; K_C = \frac{C_{AF}}{C};\; K_R = \frac{R_{AF}}{R};\; K_V = -\frac{V_{AF} - V_{0+}}{\Delta V}$$

	All-Passive EAVP	Damped Passive and SAF	Lossless Passive and SAF $Q_{AF} < 0.35$	Lossless Passive and SAF $0.35 < Q_{AF} < 0.5$
Absolute	$R_C = R_L = R$ $\frac{L}{R} = RC$ or $L = R^2C$	$R_C = R_L = R$ $Q_{AF} = \frac{2Q}{1+4Q^2}$ $L_{AF}C_{AF} = LC$ $R_{AF}C_{AF} = L/R + RC$ $V_{AF} = \Delta I \frac{L/R - RC}{C_{AF}} + V_{0+}$	$\Delta V_{\max} \approx 2QQ_{AF}\Delta V$ $Q_{AF} = 1/(2Q)$ $L_{AF}C_{AF} = LC$ $R_{AF} = \frac{\sqrt{L_{AF}/C_{AF}}}{Q_{AF}}$ $V_{AF} = R_{AF}\Delta I + V_{0+}$	$\Delta V_{\max} \approx 2Q\Delta V/e$ $L_{AF}C_{AF} = LC$ $R_{AF} = \frac{\sqrt{L_{AF}/C_{AF}}}{Q_{AF}}$ $V_{AF} = R_{AF}\Delta I + V_{0+}$
Normalized		$K_L = 1/K_C$ $K_R = (4Q^2+1)/K_C$ $K_V = (4Q^2-1)/K_C$	$\Delta V_{\max} = 4Q^2 \frac{K_L}{K_R}\Delta V$ $K_L = 1/K_C$ $K_R = 4Q^2/K_C$ $K_V = K_R$	$\Delta V_{\max} \approx 2Q\Delta V/e$ $K_L = 1/K_C$ $K_R = 4Q^2/K_C$ $K_V = K_R$
Loss	$\sim \frac{1}{2}C \cdot \Delta V^2$ $= \frac{1}{2}L \cdot \Delta I^2$	$\frac{L_{AF} \cdot \Delta I^2}{2} \cdot \frac{(4Q^2-1)^2}{4Q^2}$ (SAF loss only)	$\frac{L_{AF} \cdot \Delta I^2}{2Q_{AF}^2}$ (SAF loss only)	

Inability to achieve the desired L and C mandates the use of an active filter. If it is still possible to design precise damping resistors in the passive filter, then the "damped

passive and SAF" design equations may be used to obtain an SAF that would complement the passive filter to achieve a specified EAVP response. As noted already, there is one degree of freedom in the SAF design. Once an appropriate value of L_{AF} is chosen, the rest of the SAF design is generated by the equations. It is further possible to state the SAF design equations in normalized form as shown in the second-to-last row of Table 2.1.

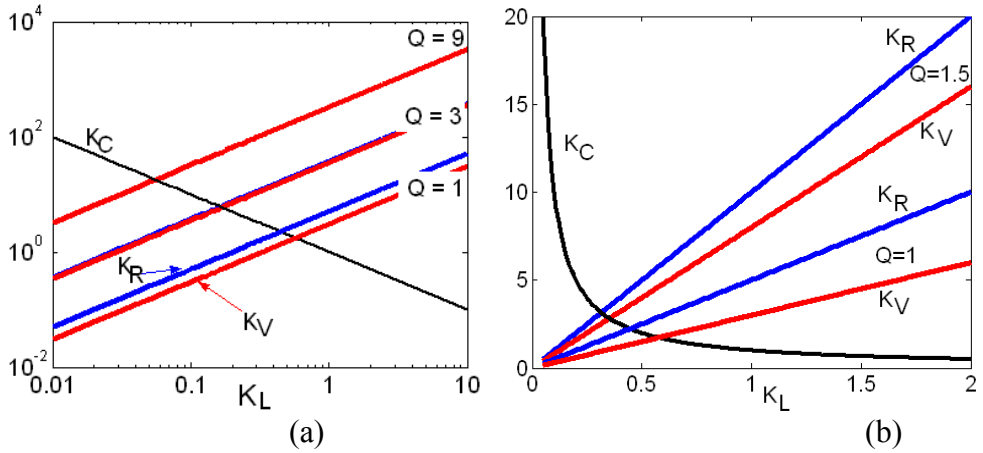

Figure 2.7. Normalized SAF parameters: (a) log-log plot and (b) linear plot.

This enables the SAF design parameters to be plotted against the relative size of L_{AF}, K_L, for various values of Q (Figure 2.7). The relative size of C_{AF}, K_C, is inversely proportional to K_L. All other SAF parameters vary linearly with K_L. For high Q, K_R and K_V coincide.

If precise damping in the passive filter is not possible, then the SAF may be designed assuming the passive filter is lossless since this is the case that exhibits the worst ringing. For variation of esr's between the lossless and perfect EAVP damped cases, only minor tweaking of the design would be necessary to achieve the desired regulation. The shape of the step-load response would of course vary depending on the

esr's but the output voltage transient can be maintained within the tolerance band using design equations for either the damped or lossless case, Table 2.1.

The choice of the appropriate set of design equations, whether damped or lossless, is determined by the actual esr values and tolerance specifications in the system. However, either set of equations would yield improvement in regulation performance. In fact, for high Q, when there is a lot of ringing in the passive filter by itself, both sets of design equations are equivalent and can be represented by plots as in Figure 2.7(a). This can be seen more clearly by comparing the columns of Table 2.1. For Q less than 1.5, Q_{AF} is sufficiently close to 0.5 that the SAF may automatically improve regulation beyond the design specifications. However, it should also be noted that for low Q, ringing in the passive filter by itself is also not that bad. In fact, for Q less than 1, the voltage excursion exceeds the tolerance band by so little that it may be possible to not include any SAF with the design, if reducing circuit complexity and losses is more important than achieving perfect load regulation.

The design methodology may now be summarized as follows:

- Find impedance at "injection node".
- Model each resonant peak with single stage equivalent.
- Design SAF to overcome resonance.
- Calculate Q.
    - Q < 0.5 SAF not needed for regulation.
    - (Q < 1, small tweaks in passive filter vs small loss with SAF).
    - 0.5 < Q < 1.5, use EAVP design equations if branch esr's ~ R.
    - Q > 1.5, EAVP and Non-EAVP design equations are equivalent.

- Choose L_{AF} based on layout/loss considerations.
- Design SAF for dominant peak (or an SAF for each peak).
- Test and tweak design if necessary
    - If esr's cause an additional droop of δv outside specification, then design for a new ΔV that is δv less than the original value used in the design.
    - If nonlinearities cause the output voltage to vary outside the limit, iteratively tweaking the design as for a slightly different value of L_{AF} without actually changing L_{AF} may improve the performance.

Usually however, tweaking is not necessary.

Loss incurred in the SAF is typically much larger than that incurred in the passive filter as it goes with the square of $2Q$, which is approximately the desired factor of improvement in regulation. The loss may be reduced by lowering L_{AF} but there is no evading the fact that there is a tradeoff between regulation performance and loss. Without the SAF and its associated loss, regulation specifications will not be met.

2.6 Design Methodology Illustrated by Examples

2.6.1 Damped Passive Filter Example

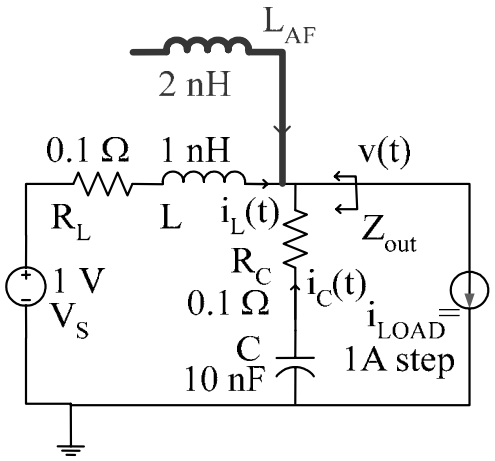

Figure 2.8. Sample design specification of an SAF and damped passive, hybrid filter.

Figure 2.8 shows an example of a design problem in which the numerical values are typical of low-current high-quality PCB applications and which is also, roughly, a scaled (~1/100x) high-frequency equivalent of a high-current microprocessor power delivery system. [Luo02]. Typically, the model of the passive filter, along with an estimate of the total inductance in the SAF path, will be known.

In the example of Figure 2.8, the DC bias from the equivalent voltage source is 1 V. Since initially load current is zero, $V_{0-} = 1$ V . The equivalent point-of-load capacitance is 10 nF and the DC path inductance is 1 nH. The esr in both branches of the passive filter is designed to be $0.1\,\Omega$ which meets the EAVP criterion (2.4) with $R = 0.1\ \Omega$. With the load set to be a 1 A step, $\Delta V = -0.1$ V from (2.1). L_{AF}, which is usually fixed by the layout, is 2 nH.

Using (2.5), $C_{AF} = LC/L_{AF} = \dfrac{1\text{ nH}\cdot 10\text{ nF}}{2\text{ nH}} = 5$ nF. Equation (2.6) yields

$$R_{AF} = \left(1/C_{AF}\right)\left(L/R + RC\right) = \frac{1\text{ nH}/0.1\ \Omega + 0.1\ \Omega\cdot 10\text{ nF}}{5\text{ nF}} = 2.2\ \Omega$$

while (2.7) gives $V_{AF} = \left(1/C_{AF}\right)\left(L/R - RC\right) + \left(V_{0-} + \Delta V\right) = 2.7$ V

The completed design is shown in Figure 2.9 and the waveforms showing the circuit response to a 1 A load step at t = 10 ns is given in Figure 2.10.

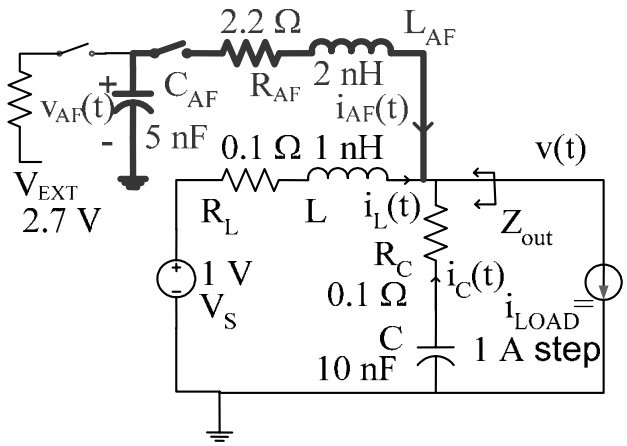

Figure 2.9. SAF and passive, hybrid filter design for the example of Figure 2.8.

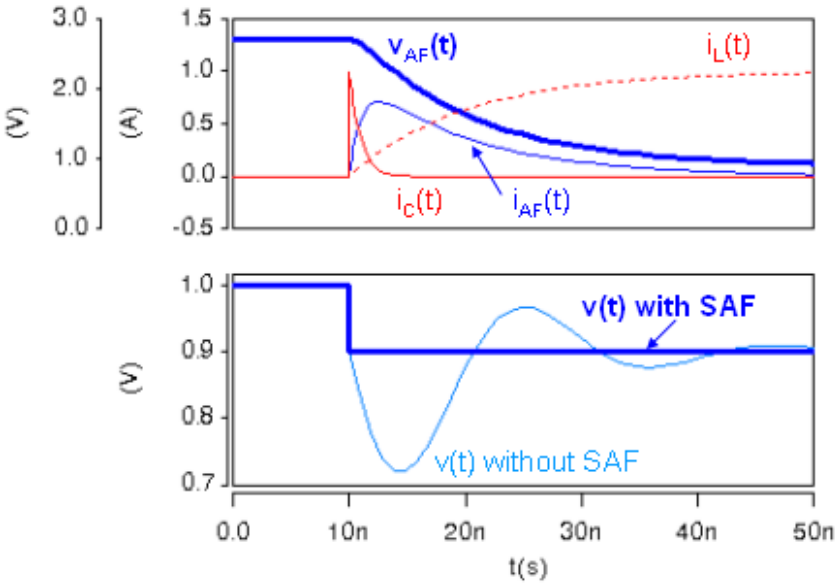

Figure 2.10. Hybrid filter response to 1 A step load.

Without the SAF, the output voltage (i.e. the voltage across the current source load) drops by almost 300 mV; more than twice the desired 100 mV deviation. With the hybrid design, however, perfect EAVP response (i.e. exactly 100 mV droop) is achieved. Initially all load current is drawn from the point-of-load capacitor C. However, the SAF quickly takes over thereby preventing the capacitor, C, from being discharged too much. While the current $i_L(t)$ builds up in the main DC path inductance, the SAF provides just

the right current, $i_{AF}(t)$, to maintain regulation. By the time that the full load current builds up in the DC path, the voltage across the SAF capacitor, $v_{AF}(t)$, decays to the steady state output voltage, 0.9 V in this example.

Thus far it has been assumed that the designer has precise control over all parameters in the circuit. However, usually parasitic elements cannot be tightly controlled. A discussion of a few of the possible variations from the ideal EAVP design follows.

2.6.1.1 Design for unequal esr's

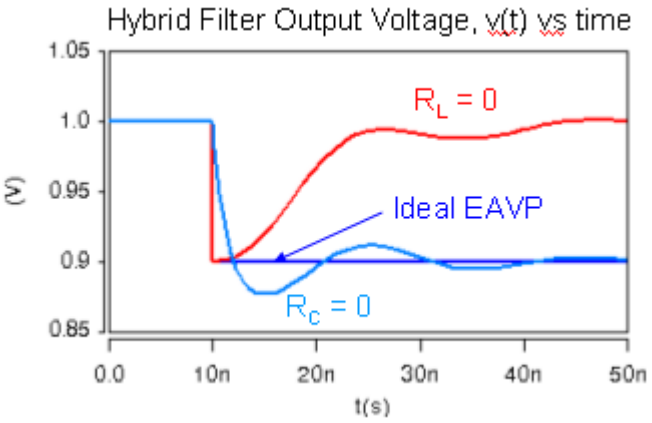

Figure 2.11. Hybrid filter response for unequal esr's.

Final steady state output voltage depends on R_L, the resistance in the DC power path. Furthermore, the maximum voltage deviation is set by the larger of R_C or R_L. Figure 2.11 shows the effect of zero resistance in L and then in C for the example design above. For zero R_L, the output voltage initially drops by $\Delta I \cdot R_C = 0.1$ V and then rises back to V_S without significant overshoot. The original SAF design, therefore, is sufficient to maintain regulation in this case. With R_C set to zero instead, there is not an instantaneous drop in output voltage and the initial di/dt provided by SAF is slightly

lower than ideal. This, combined with the fact that the actual Q of the passive filter is now larger than the original value, causes a bit more voltage deviation which causes the output voltage to drop out of specification. However, it should be noted that even in these two extreme cases of esr imbalance, the SAF still improves the regulation significantly ($\Delta V_{\max} = 120$ mV) over the passive filter alone ($\Delta V_{\max} > 240$ mV).

If low esr in the capacitor, R_C poses a problem, the extra deviation incurred by this imbalance can be corrected by simply designing the SAF for a lower ΔV. In this example, since the worse case output voltage deviation, 120 mV, was 20 mV more than the desired $\Delta V = 0.1$ V; redesigning the SAF for $\Delta V = 0.08$ V (Figure 2.12) corrected the problem. The output voltage transient can be seen in Figure 2.13.

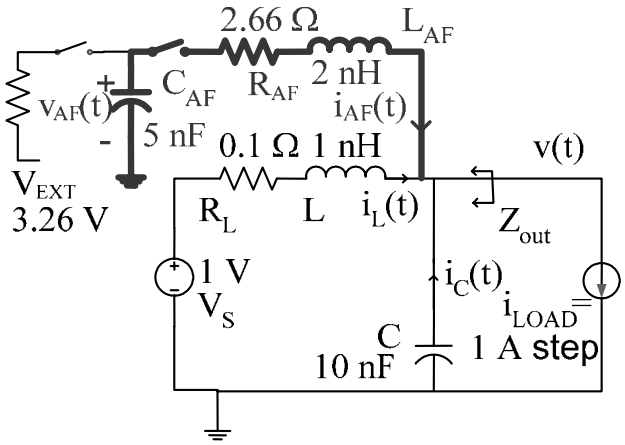

Figure 2.12. Corrected SAF design for $R_C = 0$ and $R_L = 0.1\ \Omega$.

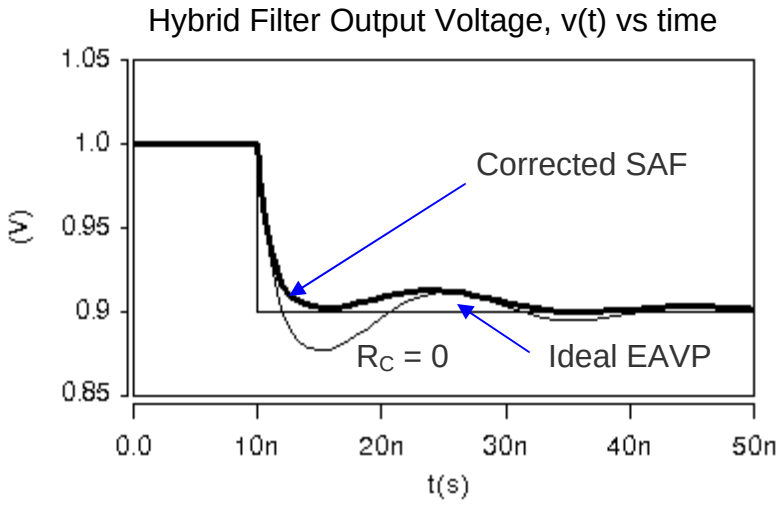

Figure 2.13. Corrected hybrid filter transient response.

With slightly larger SAF precharge voltage, V_{AF}, and total path resistance, R_{AF}, the output voltage response was improved by 20 mV to achieve the desired load regulation. The shape of the response is not perfectly sharp and flat as in the ideal design but, instead, exhibits a very small amount of ringing that quickly dies out. Using (2.30), the corrected SAF design incurs 14 nJ of loss per step load compared to just 8 nJ of loss in the original design. However, additional loss in the SAF is only natural since the passive filter is less damped with $R_C = 0\ \Omega$.

In summary, the iterative SAF design rule for unequal esr's can be stated as follows:

- Design the SAF for the greater of the esr's ($R = \max(R_L, R_C) = -\Delta V/\Delta I$) first and simulate the circuit performance.
- Then if necessary redesign the SAF for lower ΔV (or alternatively, lower R) to correct for excess voltage deviation.

  In most cases however redesign would not be necessary.

2.6.1.2 Gate drive delay, capacitor esl and load di/dt

Negative gate drive delay or turning on the SAF in advance of the load step must be avoided in order to prevent output voltage spikes outside the acceptable voltage range. For instance, if the SAF is turned on in advance of a pull-up transient it will tend to cause an overshoot equal to the anticipated voltage drop it is designed to compensate. Therefore, any delay fault that turns on the SAF when it is not needed could effectively double the original magnitude of the ripple. This however is usually not a problem in causal systems but timing precautions must be taken when the system is non-causal. If, for example, an additional 10% voltage deviation is allowed for the effects of delay, then the allowable delay can be approximated as $t_d = C\,\Delta V/(10\Delta I) = 0.1RC$ which is the approximate time it takes for the voltage across the capacitor C to drop by $0.1\,\Delta V$. Somewhat larger delay is tolerable if R_C is kept low and the SAF is designed for lower than required ΔV as previously described.

In practical applications, infinite di/dt does not occur; nor is the esl of the point-of-load capacitor, L_C, ever exactly zero. One may approximate a transition from one current level to another by an ideal step change if the transition time is much smaller than one quarter the resonant frequency of the equivalent passive filter (i.e. $t_{r,f} << 0.5\pi\sqrt{LC}$). However, even when this approximation is made, the relative effects of the esl-induced voltage spike, L_C di/dt, must be considered. In most cases $L_C\,di/dt < \Delta V/2$ is tolerable. The difficulties in dealing with the negative effects of delay, esl and transient edge di/dt are well recognized in general. For example, some microprocessor manufacturers find it necessary to allow for the power supply voltage to exceed the DC load regulation specification for a short burst of time just after a load transient edge [Int05b].

On the other hand, if the transient di/dt is so low that the step load approximation is no longer valid, then, instead of a switching function, the transistor in the SAF may operate as a controlled variable resistance. In addition the SAF pre-charge voltage may be adaptively varied in response to a slower transient. The issues of esl, delay and control are left as future topics to be discussed in Chapter 4.

2.6.1.3 Overall design sensitivity

Keeping all else fixed, individual hybrid filter design parameters were increased to obtain roughly 25% voltage deviation from ideal in order to investigate design sensitivity. As summarized in Figure 2.14, the circuit is most sensitive to variations in V_{AF}. Fortunately, this parameter can be controlled tightly. All other parameters may be varied by 20% to 50% without significant degradation in performance. Furthermore, the graphs show that increases in L_{AF}, R_{AF} or L cause additional voltage drops just after a transient step-up in load current. These effects may be compensated by increases in C, V_{AF} or C_{AF} respectively and are consistent with the design equations.

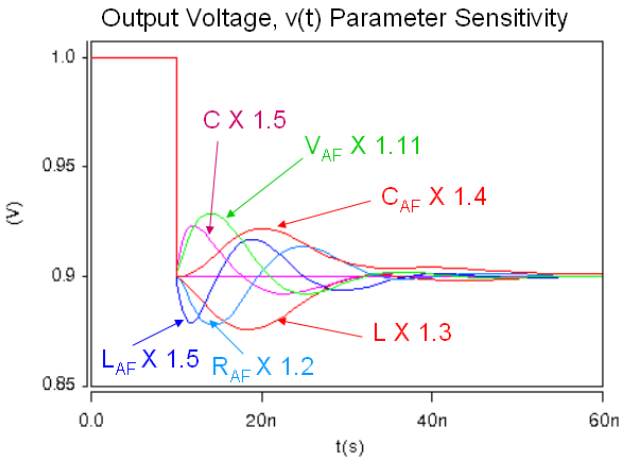

Figure 2.14. Output voltage parameter sensitivity of circuit in Figure 2.9.

Assuming 20% tolerance on all seven circuit elements in the hybrid filter, a Monte Carlo, worst case, 128 run simulation was performed. By comparing the results with the output voltage waveform without the SAF (Figure 2.15) it was observed that even with 20% variation on all circuit elements the hybrid filter performs better than the passive filter alone.

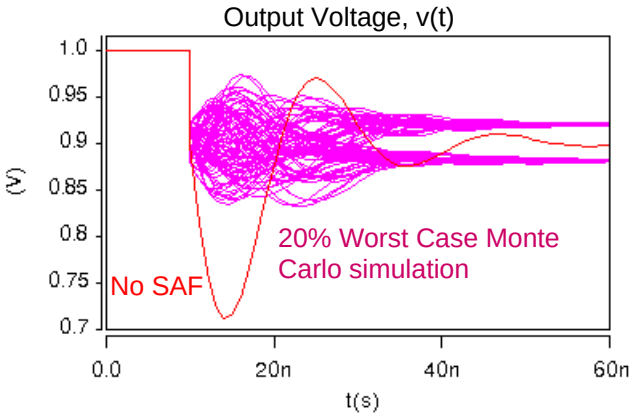

Figure 2.15. Worst case, 20%, 128 run Monte Carlo simulation of circuit in Figure 2.9.

2.6.2 Lossless Passive Filter Example

Given the example hybrid filter circuit shown in Figure 2.16, in which the passive filter is lossless (since there are no damping resistors), an SAF design to achieve $\Delta V = -0.1\,\text{V}$ for $\Delta I = 1\,\text{A}$ is desired. With $L_{AF} = 2\,\text{nH}$, the same as before, $C_{AF} = 5\,\text{nF}$ is also unchanged, via (2.5).

Nominally, $R = -\dfrac{\Delta V}{\Delta I} = -\dfrac{(-0.1\,\text{V})}{1\,\text{A}} = 0.1\,\Omega$ via (2.1). Hence, from the definitions of (2.2), $2Q = \dfrac{1}{R}\sqrt{\dfrac{L}{C}} = \dfrac{1}{0.1\,\Omega}\sqrt{\dfrac{1\,\text{nH}}{10\,\text{nF}}} = 3.16$ and the passive filter would ring with an amplitude $\Delta V_{pk} = 2Q\Delta V = 316\,\text{mV}$ or roughly three times the load regulation ripple

specification. In this particular example a few possible design solutions present themselves.

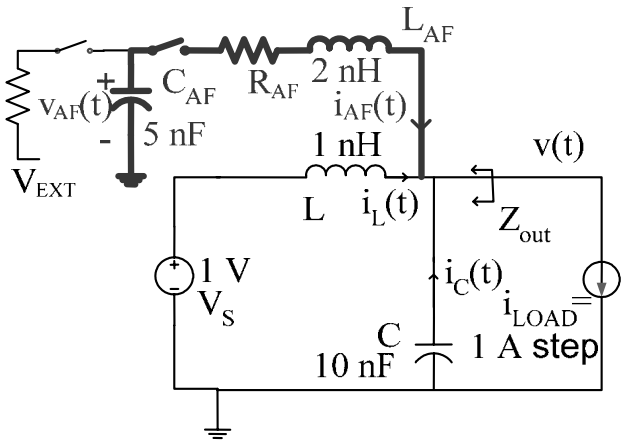

Figure 2.16. Example of SAF design for lossless passive filter.

Since the desired correction factor, $\frac{1}{2Q} = \frac{1}{3.16} \approx \frac{1}{3} \approx \frac{1}{e}$, one may design the SAF for $Q_{AF} = 0.5$ according to (2.27) if ΔV in the load regulation specification is somewhat loose. The advantage of this design is that it offers the lowest loss. For $Q_{AF} = 0.5$,

$R_{AF} = \frac{\sqrt{L_{AF}/C_{AF}}}{Q_{AF}} = \frac{\sqrt{2\text{ nH}/5\text{ nF}}}{0.5} = 1.265\ \Omega$, using (2.28), and

$V_{AF} = R_{AF}\Delta I + V_{0+} = (1.265\text{ V})(1\text{ A}) + 1\text{ V} = 2.265\text{ V}$, via (2.29).

Another option is to iteratively design for lower and lower values over the range $0.35 < Q_{AF} < 0.5$ until an exact value of Q_{AF} is found to meet the ΔV requirement. If a solution does exist over this range, then one may conservatively design for $Q_{AF} = 0.35$ and forego the iterative process.

Alternatively, one may simply find Q_{AF} using the conservative design equation (2.26) already derived. For this example (2.26) yields $Q_{AF} = \frac{1}{2Q} = \frac{1}{3.16} = 0.316$.

Assuming L_{AF} is given, then, after Q_{AF} is determined, C_{AF} may be calculated from (2.5); R_{AF} from (2.28) and V_{AF} from (2.29). SAF designs for the three values of Q_{AF} obtained are summarized in Table 2.2.

The design for $Q_{AF} = 0.5$ has the lowest loss but does not quite meet the ΔV specification.

Since $Q = 1.58 > 1.25$, the most conservative design is approximately the same as that derived for the damped passive filter of the same quality factor (Figure 2.9). Output voltage plots for the three designs are shown in Figure 2.17.

Table 2.2. Parameter and performance values for the circuit of Figure 2.16.

Q_{AF}	V_{AF} [V]	R_{AF} [Ω]	ΔV_{max} [mV]	E_{loss} [nJ]
0.5	2.265	1.265	111	4
0.35	2.8	1.8	87	8.2
0.316	3	2	81	10

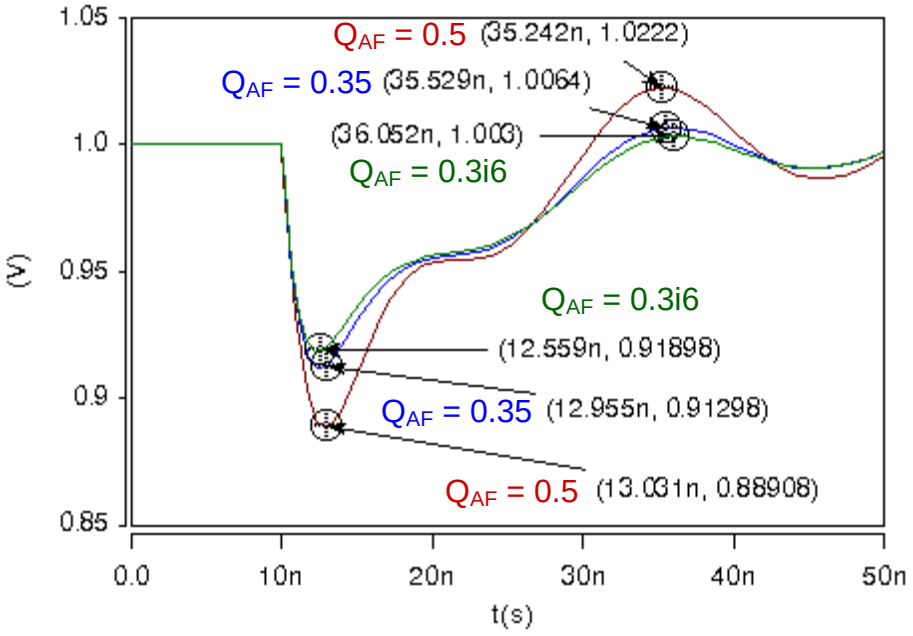

Figure 2.17. Output voltage response of the SAF designs for Figure 2.16.

For $Q_{AF} = 0.5$, the maximum voltage deviation is consistent with that predicted by (2.27). For all designs, the voltage waveshape roughly follows the corresponding analytic expressions derived from the form (2.16). However, there is a small exponentially decaying sinusoid superimposed on the ideal waveshape that is due to very slight numerical deviation from the exact design and is indicative of the fact that the SAF is trying to damp an infinite-Q network. Even so, the unwanted "ringing" is much less than ΔV (22 mV in the worst case), it decreases with lower Q_{AF} and it dies out within 1.5 to 3 cycles. It is insignificant when compared to the oscillation induced in the lossless passive filter alone by the step-load (Figure 2.18). A 32-run, 20%, worst case Monte Carlo simulation shows that even with 20% variation on all hybrid filter elements, the regulation is much better than the lossless passive filter alone and is comparable in performance to a lossless passive filter with 10 times the capacitance (Figure 2.19).

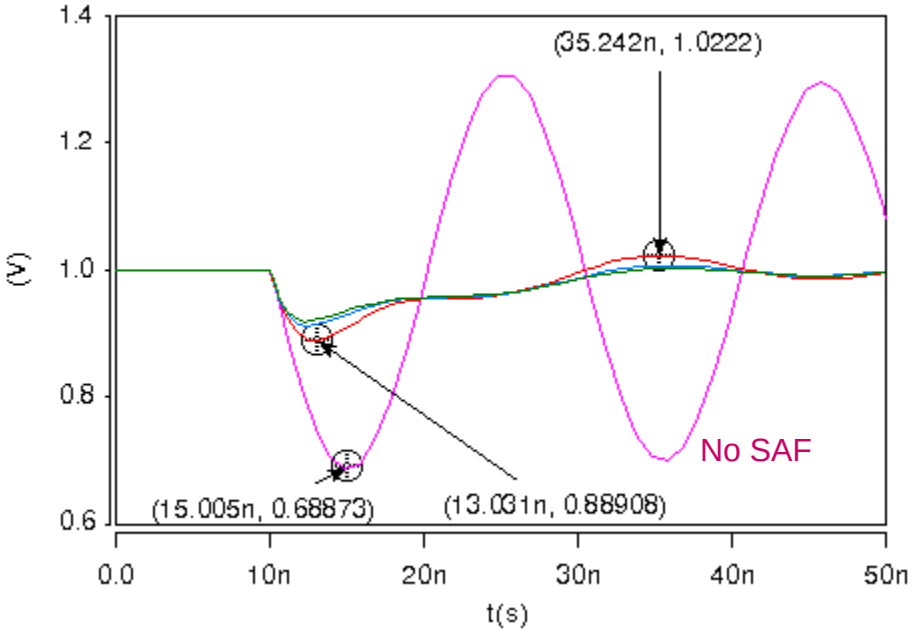

Figure 2.18. Relative improvement in regulation achieved by using SAF.

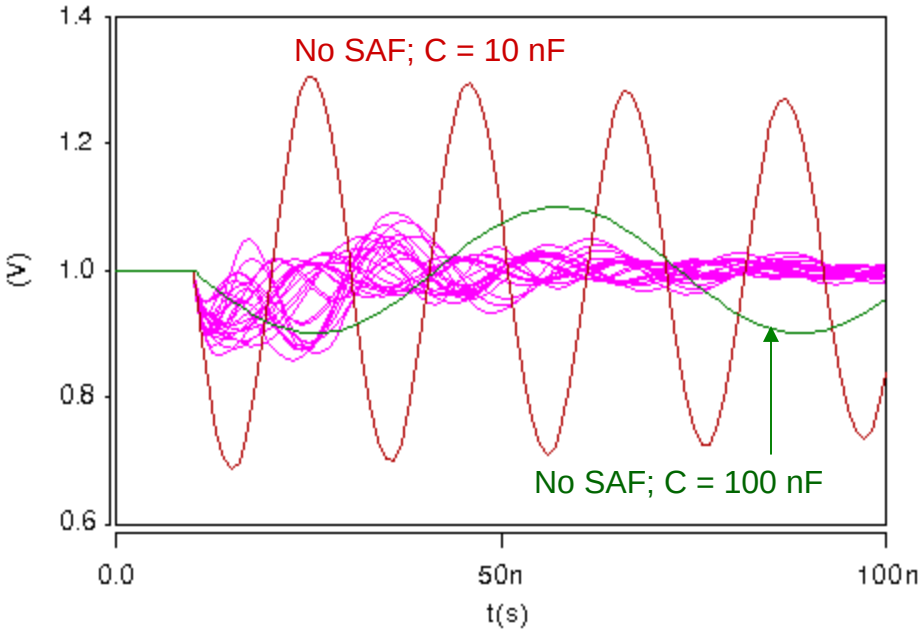

Figure 2.19. Monte Carlo simulations showing hybrid filter performs as well as lossless passive filter with 10 times the capacitance.

2.7 Design Synopsis

It is convenient to characterize SAF design by two key parameters- $2Q$ and K_L. The former is a measure of the desired improvement in ripple over the lossless passive filter: $\frac{\Delta V_{pk}}{\Delta V} = 2Q$, via (2.25); and governs the transient response. Figure 2.20 illustrates the use of the SAF in EAVP and non-EAVP designs to improve transient voltage response. The SAF designed for EAVP achieves a perfect step "ripple" in response to a step in load current. Without the SAF, the output voltage overshoots the allowed ΔV and may be almost as large as $2Q\Delta V$. The SAF designed for the loss-less (worst case) passive filter virtually eliminates all ringing without ever exceeding the ripple specification.

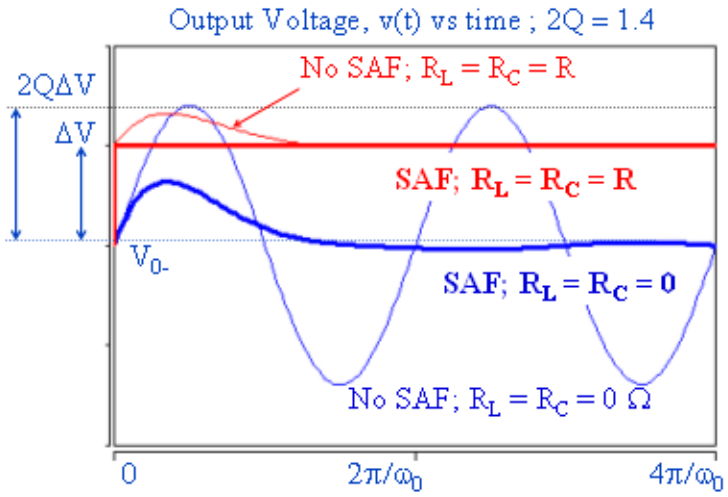

Figure 2.20. SAF response for 2Q = 1.4.

Power loss is dependent not only on Q but also on K_L. Although the transient response is somewhat insensitive to K_L, it is directly related to power loss via (2.31), (2.32) and (2.26).

2.8 Variations on Circuit Topology

The basic SAF circuit is depicted in Figure 2.21(a). It is a switched damped resonant circuit topology. As shown, the energy storage element, C_{AF}, is directly referenced to ground and the switch is mounted at the same location as C_{AF}. However, it is conceivable that if C_{AF} were allowed to float with respect to ground as in Figure 2.21(b), this will afford more flexibility in the design of the peripheral pre-charge and control circuitry. The switch will still be physically mounted in the same location as C_{AF} - away from the load. The fact that C_{AF} is mounted further from the load is a natural consequence of the basic problem- there is not sufficient space at the load for extra capacitors. However, it should be noted that it may be possible to share the available

space for package capacitors with C_{AF} in order to achieve much better performance than if no SAF were used and C_{AF} simply formed a fraction of C_{Pkg}.

The switch, which typically is smaller than C_{AF} may be positioned near the load as depicted in Figure 2.21(c). In fact, for microprocessor and similar IC applications, the switch may even be integrated onto the load IC.

Figure 2.21(d) shows a topology variation that consists of dedicated capacitors and SAF paths for pull-up and pull-down transients. Recall that the preset voltage depends on the expected transient. For operation in mid-range, either pull-up or pull-down transients may occur. When apriori transient information is unavailable (as typically is the case), one SAF capacitor will be precharged for worst case pull-up and the other for worst case pull-down. The feedback control circuit will be designed to activate the appropriate leg in response to any transient. This topology allows for quick response to either rising or falling edges at the cost of doubling the ancillary circuitry needed. It is envisioned that a simple manifestation of this topology may include two additional pre-charging switches as shown in Figure 2.21(e). While one half of the SAF is performing AF operation, the other half quickly pre-charges and vice versa. At DC steady state, both the SAF switches are off while the pre-charge switches are biased on.

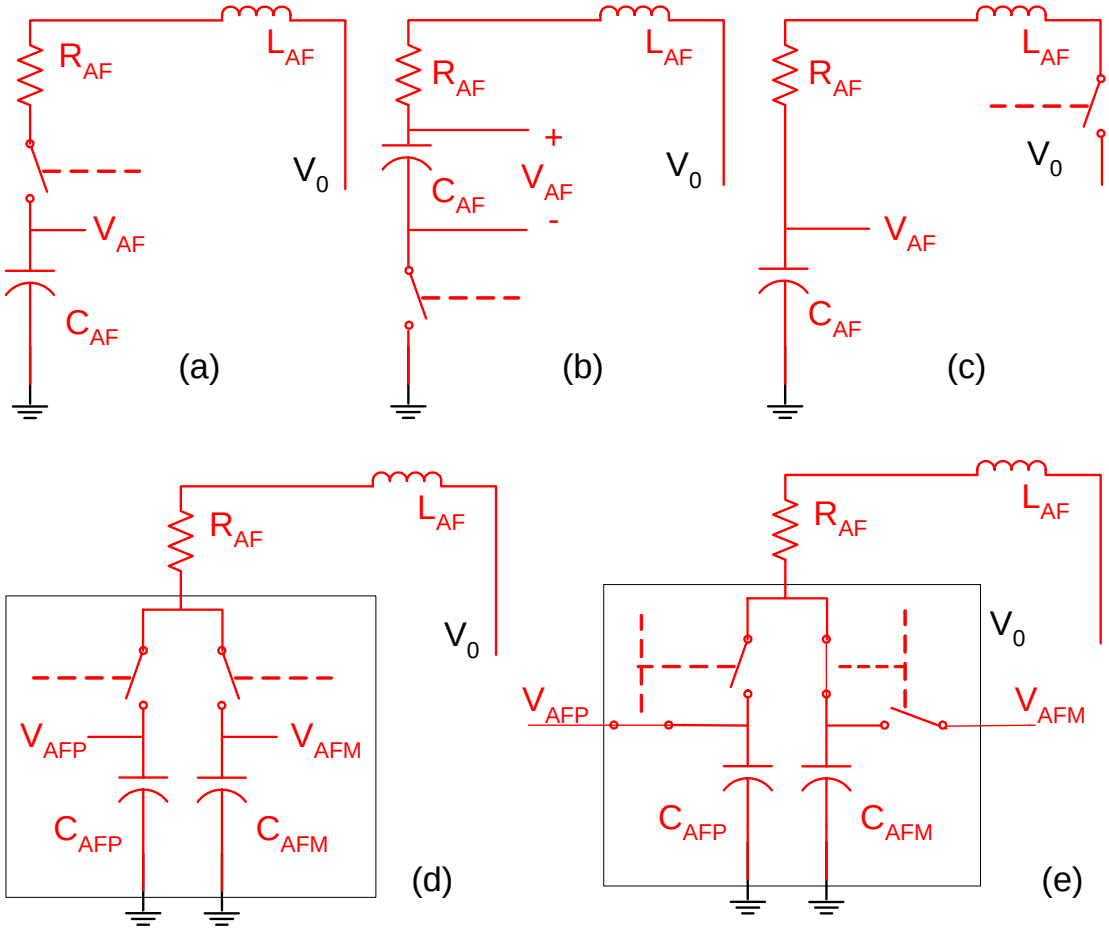

Figure 2.21. Variations of SAF: (a) original with grounded C_{AF}, switch near C_{AF}; (b) floating C_{AF}, switch near C_{AF}; (c) grounded C_{AF}, switch near load; (d) dual path SAF and (e) SCAF, switched capacitor active filter.

CHAPTER 3
EXPERIMENTAL RESULTS AND PRACTICAL IMPLEMENTATION ISSUES

This chapter describes the experimental verification of the hybrid filter design principles, measures improvement in performance over the passive filter alone, and demonstrates the issues that arise in practical implementation. In one practical scenario, the SAF injection point may be on the motherboard close to the BVR. Inductance in the DC path is relatively large while the esl in the capacitance at the injection node is negligible due to relatively low di/dt or slow speeds at that point. If, on the other hand, the SAF injection point is near a high-speed microprocessor load, for example, esl of the point-of-load capacitor should be considered even with the SAF assisting in regulation. With a wide range of applications in mind, low speed experiments, in which the load di/dt's and the esl of the filter capacitor can be ignored, are first described. These first experiments use a switched resistive load which differs from the ideal current source used to model the load on the pull-up edge. Thus it will be shown that the design methodology is applicable even when the load is not an ideal current source. Next, a 50-200 A/μs di/dt electronic load circuit which is used in the remaining experiments is presented. The use of this high di/dt load circuit to test the SAF in conjunction with a single stage passive filter shall demonstrate the negative effects of filter capacitor esl. Finally, the SAF is shown to operate independently of additional stages in typical multistage filter configurations.

The key list of symbols used in this work is given in Table 3.1.

Table 3.1. List of Symbols.

AF	Active filter
BVR	Bulk voltage regulator, such as a Voltage Regulator Module (VRM)
C	Equivalent POL capacitance in simplified model of passive filter; Figure 2.3
C_{AF}	Filter capacitor for the AF (LAF or SAF)
C_{die}	Capacitor associated with the microprocessor die
C_n	Capacitor associated with the n^{th} node of passive filter in Figure 2.1
C_{pkg}	Package capacitors, illustrated in Figure 1.8
C_{skt}	Socket capacitors, illustrated in Figure 1.8
EAVP	Extended Adaptive Voltage Positioning
E_{loss}	Energy loss in the SAF per step load
i_{AF}	Current from the AF (LAF or SAF)
i_C	Current through C
i_{Cdie}	Current through C_{die}
i_{Cpkg}	Current through C_{pkg}
i_L	Current through L
i_{Lpkg}	Current through L_{pkg}
i_{uP}	Current drawn by the microprocessor die
ΔI	Transient change in load current
L	Equivalent DC path inductance in simplified model of passive filter; Figure 2.3
LAF	Linear active filter
L_{AF}	Inductance between AF power source, V_{EXT}, and AF injection point
L_C	Equivalent series inductance of C
L_{Cpkg}	Equivalent series inductance of C_{pkg}
L_{Cskt}	Equivalent series inductance of C_{skt}
L_{mb}	Inductor between the BVR and C_{skt}
L_{pkg}	Inductor between Z_{Cpkg} and Z_{die}
L_{skt}	Inductor between Z_{Cskt} and Z_{Cpkg}
LVR	Linear Voltage Regulator
Q	Equivalent quality factor of damped passive filter
Q_{AF}	Quality factor of SAF network
R_{AF}	Resistance between AF power source, V_{EXT}, and AF injection point
R_C	Equivalent series resistance of C
R_{Cpkg}	Equivalent series resistance of C_{pkg}
R_{Cskt}	Equivalent series resistance of C_{skt}
R_{die}	Equivalent series resistance of C_{die}
R_L	Equivalent series resistance of L
SAF	Switched active filter
τ_{AF}	Settling time-constant of SAF
T_W	Width of esl-induced spike in output voltage
t_r	Gate drive rise-time
uP	Microprocessor
v_{AF}	Voltage across C_{AF}
V_{AF}	Initial voltage across C_{AF}
V_{EXT}	Power supply voltage for the AF
v_O	Output voltage or voltage across load
v_{pkg}	Voltage at the node where C_{pkg} is connected; see Figure 1.8
v_{skt}	Voltage at the uP socket
v_{uP}	Voltage across the microprocessor die (voltage across load)
ΔV	Allowable transient variation riding on V_O, the average of v_O
ΔV_{max}	Peak transient variation riding on V_O, the average of v_O
ΔV_{Pk}	Worst case ringing in v_O in lossless passive filter alone i.e. no SAF used
Z_{Caf}	Impedance of $C_{AF,}$ including equivalent series inductance & resistance
Z_{Cpkg}	Impedance of C_{pkg}, including equivalent series inductance & resistance
Z_{Cskt}	Impedance of C_{skt}, including equivalent series inductance & resistance
Z_{die}	Impedance of C_{die}, including equivalent series resistance
Z_{LAF}	Output impedance of the LAF
Z_{Lpkg}	Impedance of L_{pkg}
Z_{uP}	Impedance seen by i_{uP}

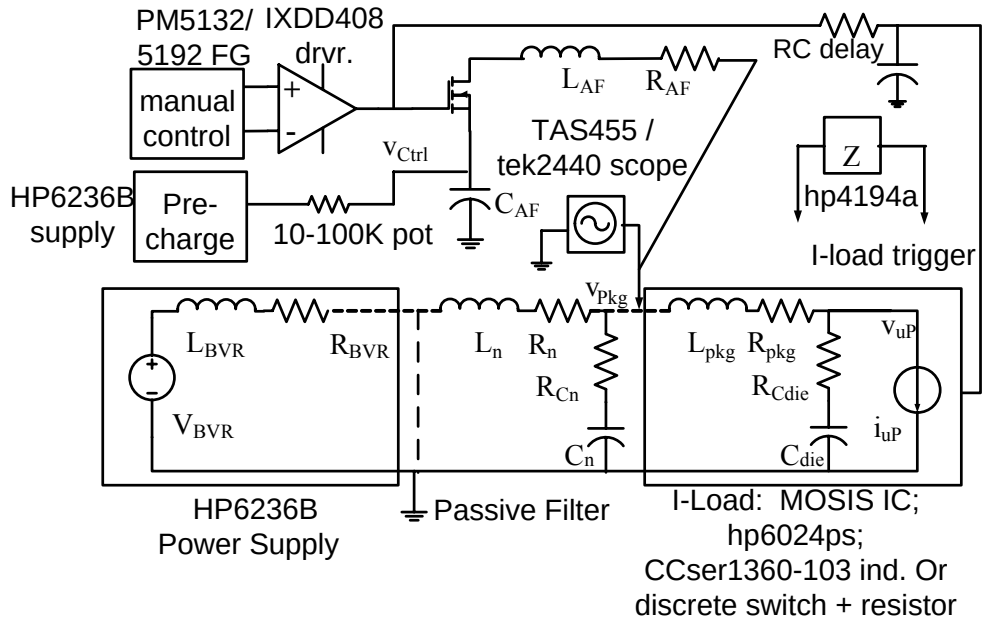

Figure 3.1. Experimental testbed.

Figure 3.1 depicts the experimental testbed. The passive filter is formed from discrete capacitors with L_n, R_n and R_{Cn} all parasitic elements associated with the components and interconnect layout. When necessary an HP6236 power supply is used as the BVR. When testing the SAF at high frequencies, the BVR is not necessary and may be replaced by a short circuit as shown by the bold dotted line. For low speed tests, the I-load circuit block was implemented by means of an STS9NF30L switch and 1205 chip resistors. The need to force high di/dt in the high-speed experiments was accomplished by means of a custom designed I-load emulator circuit in which a 10 mH inductor biased with high current was switched fast enough by means of a gate driver and mosfet IC to emulate a pulsed current source.

Various versions of the SAF were tested with open loop control. Function generator timing signals were sent to trigger gate drivers synchronously with load switching signals. Relative delays between load trigger and other gate signals were accomplished

by inserting RC delay networks at the appropriate signal line termination. Precharge was accomplished via a separate HP6236 power supply charging C_{AF} through a variable resistor. The resistance setting was such that the precharging time constant was much larger than the SAF time constant; thereby minimizing any effect of the precharging network on the actual SAF performance measurements.

In all experiments, inductances and capacitances were measured using the HP4194 impedance analyzer through a semi-rigid, twisted pair extension or a manufacturer-calibrated BNC extension cable where appropriate. Maximum measurement error was the larger of 4 nH or 20%. However, typically error was less than 10%. When the layout geometries were simple enough, inductance measurements were further cross-checked using analytic formulas. Capacitor measurements, lying in the range 1 nF to 0.1 mF and made between 100 KHz and 1 MHz, had less than 5% error.

Resistors were measured using hand-held multimeters. Small parasitic resistances were estimated from known current and voltage drops which were measured, along with time domain waveforms, on the Tek2440 (500 MS/s) or TAS455 (60 MHz) oscilloscopes. Data was recorded by hand and a camera was used to take snapshots of oscilloscope waveforms.

The bill of materials corresponding to the experimental testbed described above is now listed:

- PCB mill and soldering kit
- Dual power supply (HP6236B) Quantity: 2
- Power supply rated for over 3 A (HP6024A)
- 1-60 MHz Function Generator (PM5192/5132)
- 100 MHz Network/Impedance Analyzer (HP4194)

- 60/500 MHz Oscilloscopes (TAS455/TEK2440)
- Pulse-current probe (Isensortech 711-S)
- 10/30 mOhm, 10/30 nC Q_G switching gate charge NMOS/PMOS switches (IRF3415, STS9NF30L/STS5PF20V)
- 3 A, 10 ns risetime Gate driver (IXDD408SI)
- 1-10 uH smt coilcraft inductor kit
- 1-100 K potentiometer kit
- 1-2200 uF electrolytic capacitor kit
- 0.0001-10 uF ceramic (NPO/XRF) 0805 cap. kit
- 0.1-100 Ohm 1206, 0.5 W, resistor kit
- BNC-SMA adaptor kit
- SMA connectors and 50 Ohm cables and termination kit
- Standard DMM and workshop tools

Specific setup along with parts and quantities used are described for each experiment below.

3.1 Low Speed Experiments

An experimental PCB test circuit based on Figure 2.5 is given in Figure 3.2. The design parameters are given in the upper rows of Table 3.2. For $\Delta I = 1$ A, $\Delta V = 0.1$ V is chosen yielding $R = 100$ mΩ and $Q \approx 7$ via (2.1) and (2.2). However, it is understood that due to non-idealities the resulting worst-case $\Delta V_{\max}$ will be larger than ΔV. A plan view, a side-view and a labeled section of the prototype are depicted in Figure 3.3 and Figure 3.4. The load is formed by discrete design: an STS9NF30L n-channel MOSFET switch, two 1210, 1 Ω resistors and an 0805, 1 μF capacitor, which provides the "high frequency" power path. The "DC" power path starts from the power supply (not shown),

reaches SMA-2 through roughly 180 nH of twisted pair lead inductance, goes through another 16 nH of trace inductance between SMA-2 and SMA-1, and finally reaches the load through approximately 4 nH in the loop around the Isensortech 711-S current sensor. The parasitic resistances in the "DC" and "high frequency paths" were both determined to be less than 100 mΩ. Therefore, ideal EAVP is not to be expected. Separate SAF paths, each having about 40 nH of trace inductance, for pull-up and pull-down (shown without components) were laid out alongside the main power path for convenient testing. Since $Q \approx 7$, design values for an SAF can be calculated using the EAVP SAF design equations (2.4) through (2.7) even though it is understood the ideal EAVP response will not be obtained. Components that were already available in the lab were used to achieve the closest approximation to the calculated SAF design and simulations were then done based on the chosen component values. Table 3.2 lists values used in calculation, simulation and experiment.

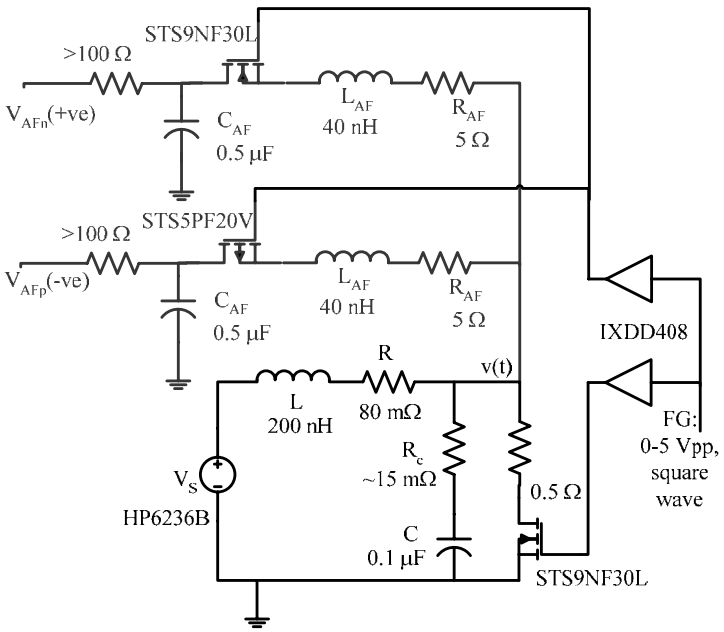

Figure 3.2. Hybrid filter circuit used in low-speed experiments.

Table 3.2. Hybrid filter parameters in low-speed experiments.

Symbol	Calculated	Simulated	Experimental
$\|\Delta V\|$	100 mV	100 mV	100 mV
$\|\Delta I\|$	1 A	1 A	0.8-1 A
t_r(switch)	0	20 ns	>20 ns
C	0.1 μF	0.1 μF	0.1 μF
R_C	0.1 Ω	0.05 Ω	<0.1 Ω
L	200 nH	200 nH	200 nH
R_L	0.1 Ω	80 mΩ	80 mΩ
L_{AF}	40 nH	40 nH	~40 nH
R_{AF}	4 Ω	5 Ω	5 Ω
C_{AF}	0.5 μF	0.5 μF	0.5 μF
V_{AF}	-3.4/4.5 V	-3.8/4.8 V	-3.8/4.8 V
ΔV_{max}	+/-100 mV	+200/-270 mV	+200/-450 mV
ΔV_{pk-pk}	2.8 V	2.5 V	2.5 V
E_{loss}	4 μJ	4.6 μJ	4.6 μJ

The plan of the full layout of the prototype is given in Figure 3.3(a). In order to highlight the layout of traces, the photographs were taken before soldering the pull-down SAF and the pre-charging circuit components. A side-view showing the relative positions of the Isensortech 711-S current sensor, which was used to measure the load transient, and the oscilloscope voltage probe is given in Figure 3.3(b). The pull-up SAF and load switches were implemented with STS9NF30L n-channel MOSFETS; the pull-down SAF switch was implemented with the STS5PF20V p-channel MOSFET. Switches were driven with IXDD408SI drivers. Using a single gate signal to trigger all the gate drivers, the relative delay between switching signal edges was negligible. In place of an ideal current source electronic load, a real switched $0.5\,\Omega$ load was used so that on pull-up the load switches from open to $0.5\,\Omega$. On pull-down, the load switch action very closely approximates an ideal current source since the switch suddenly opens and load current is forced to zero. Pre-charging was accomplished via a voltage source and resistor network

such that the charging time-constant was much greater than the SAF time-constant in order to minimize error due to the pre-charge network.

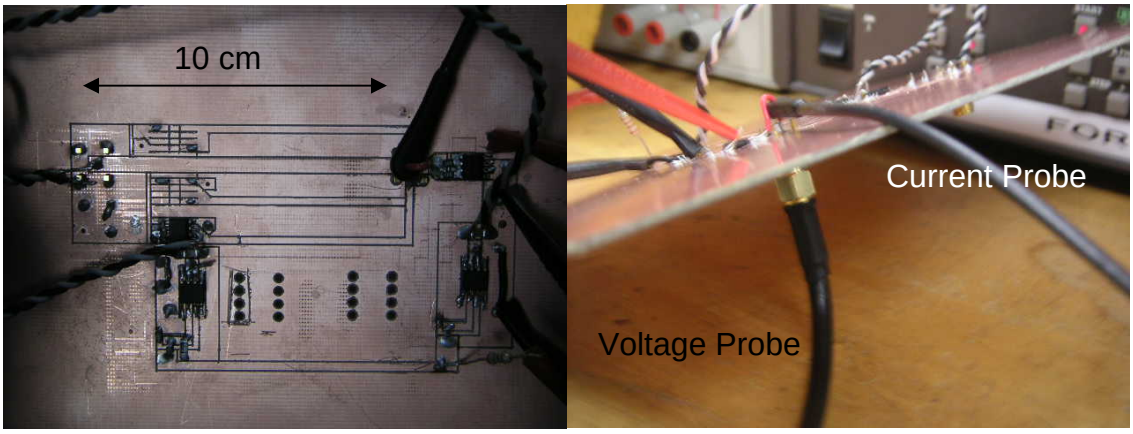

(a) (b)

Figure 3.3. Low speed prototype: (a) plan view and (b) side view.

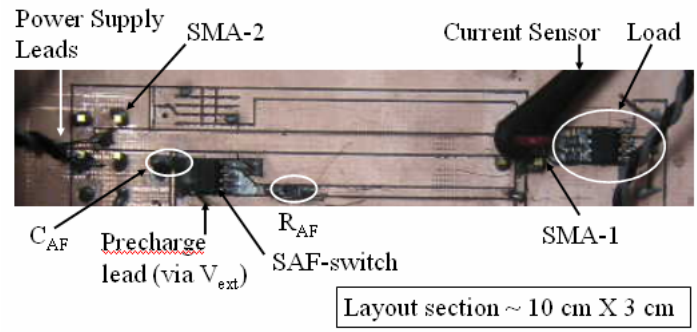

Figure 3.4. Labeled section of a hybrid filter layout.

The section view of Figure 3.4 shows the STS9NF30L MOSFET switch and parallel pair of resistors that formed the load. The wire loop around which the current sensor is mounted is perpendicular to the PCB plane and is only about 2 cm outstretched length. Underneath the current sensor, the point-of-load capacitor, C, is soldered between SMA-1 signal and one of the ground pins. Due to the proximity of SMA-1 and the load, the load voltage is practically given by the voltage probe measurement at SMA-1. Since the expected di/dt's are much less than 50 A/us, the esl of C, $L_C = 1$ nH, is not expected to induce greater than 50 mV spike. In fact, no noticeable spike was observed

which validated the assumption that L_C could be neglected. The SAF network is formed by C_{AF}, SAF-switch, and R_{AF} in that order from left to right. With R_{AF} about 4-5 Ohms as listed in Table 3.2, the switch on-resistance of about 25 mOhm (found in datasheets), esr of C_{AF} of about 15 mOhm (extracted from impedance curves in datasheets) and parasitic trace resistance of about 20 mOhm (measured by voltage-droop/forced-current) are altogether negligible. The parasitic SAF path inductance, L_{AF}, is dominated by pin and trace interconnect inductance which is approximately 40 nH. The esl of C_{AF} which can be extracted from the impedance curves in the component datasheet is about 0.8 nH and can be neglected.

As can be seen from Table 3.2, the values used in the experiment are not exactly equal to those generated by the design equations nor do the parasitic resistances in the passive filter satisfy (2.4). Therefore, the ideal EAVP response was not expected. However, although the R_{AF} value is larger than the calculated value, this is somewhat compensated by the larger V_{AF} magnitude. Saber simulation showed that even a non-ideal SAF design, (roughly 20% error on some parameters, and non-zero risetime, t_r), can reduce by over twenty times the point-of-load capacitance, C, required for a purely passive solution (Figure 3.5). Alternatively, the SAF assists the given passive filter by reducing ripple by more than five times that of the passive filter alone as shown in Figure 3.6(a). Load current (not shown) for the pull-down case is the same with and without the SAF and is simply a 1-0 A step down. Furthermore, recalling the SAF was derived using an ideal current source as the load since this is the worst case scenario that leads to the

most ringing, Figure 3.6(b) shows that the SAF works just as well with switched resistive loads, in pulling up the response to achieve the required step load.

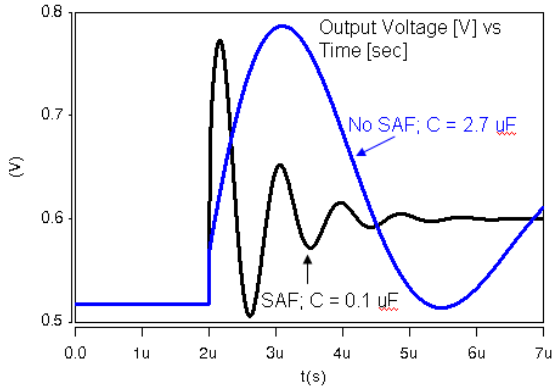

Figure 3.5. Hybrid filter simulated response to 1-0 A, 50 A/ μ s, step-down response showing usefulness in capacitor reduction for the same relative performance.

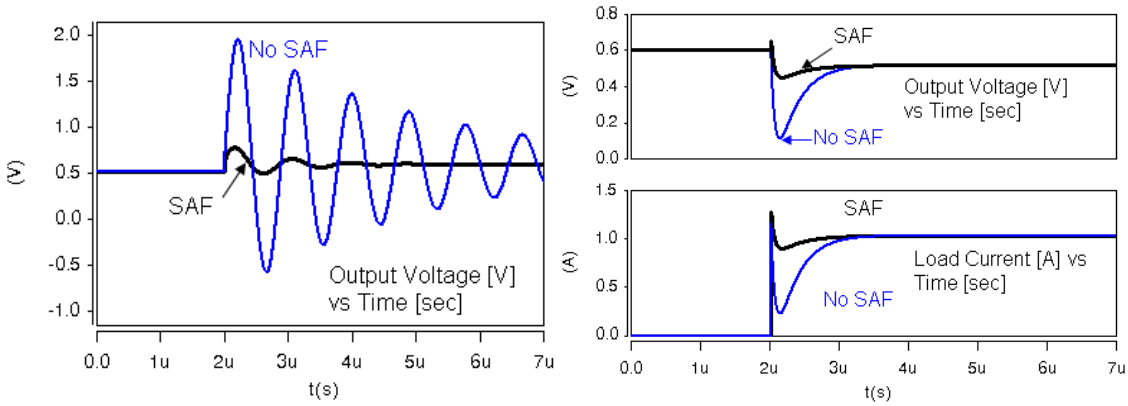

(a) (b)

Figure 3.6. Low-speed hybrid filter simulated response: (a) switched open, 1-0 A, 50 A/μs step-down, (b) switched 0.5 Ω, load pull-up.

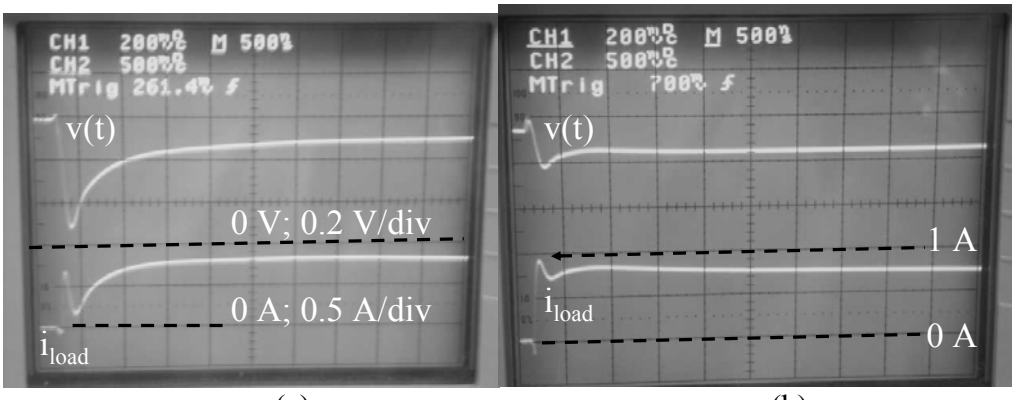

(a) (b)

Figure 3.7. Step-up output voltage response (a) without and (b) with the aid of the SAF.

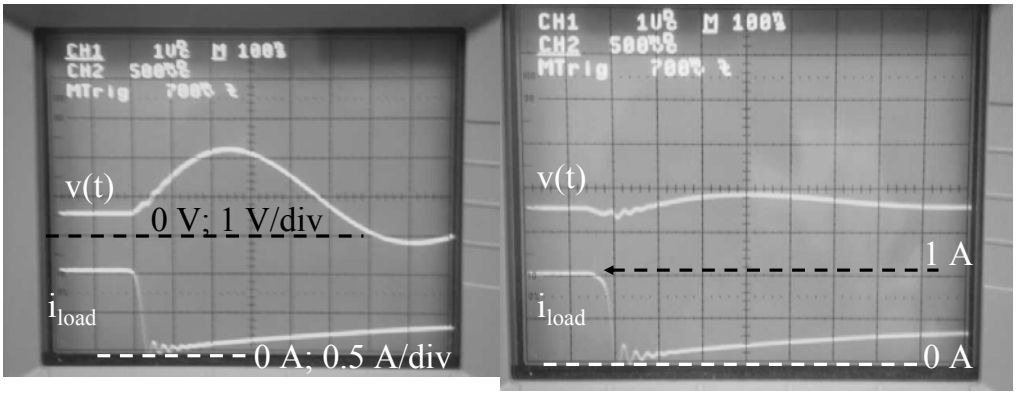

(a) (b)

Figure 3.8. Step-down output voltage response (a) without and (b) with the aid of the SAF.

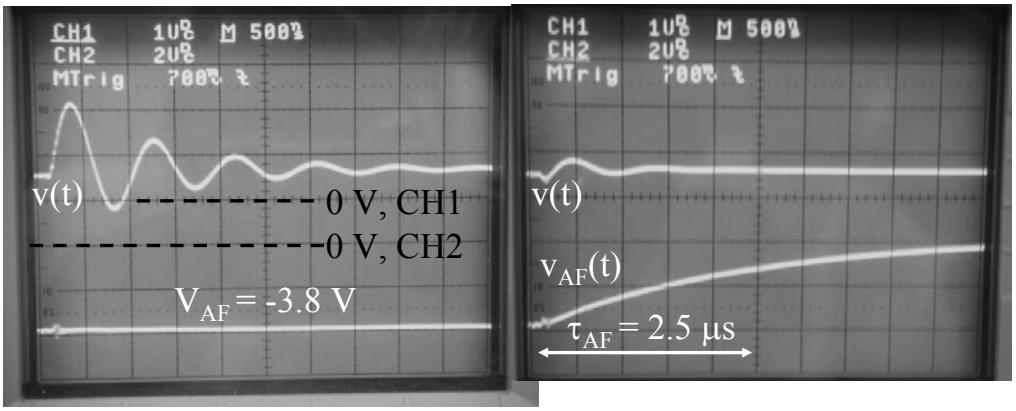

(a) (b)

Figure 3.9. Step-down output voltage and v_{AF} response (a) without and (b) with the aid of the SAF.

These simulation results were further corroborated in the actual experiments. The output voltage was measured using the Tektronix TAS455 60 MHz oscilloscope; first without the SAF connected and then with the SAF connected to the output node. Performance values are listed in the last rows of Table 3.2.

The step-up response in Figure 3.7 shows that the SAF was able pull-up the output voltage to within $\Delta V_{\max UP}$ = 200 mVpp of ripple (or 200% of the DC droop) thereby allowing for an almost square step in load current from 0 to 1 A. This result is the same as that predicted by simulations, as listed in Table 3.2. Note that the droop in the current waveform with time is merely an artifact of "DC settling" in the passive current sensor.

The slight peaking in output voltage and corresponding dip in current at the start of the transient are due to capacitive coupling of the gate-drive signal onto the output as well as timing error in turning on the SAF switch.

On step-down, the SAF reduced the ripple to within $\Delta V_{\max DN} = 450$ mVpp (or 450% of DC droop); an improvement of almost a factor of 5 compared to the case of the passive filter by itself (Figure 3.8). Although simulations showed 270 mVpp of ripple, this was largely due to waveshape deviation from the piecewise linear approximation and component errors, which, in this case, had a greater impact since the equivalent circuit quality factor is higher.

Without the SAF, the ringing frequency obtained from Figure 3.9(a) is about 1.2 MHz (6 peaks in 5 μs), and is consistent with the value predicted by model parameters listed in Table 3.2, $1/\left(2\pi\sqrt{LC}\right) = 1.12$ MHz. The SAF time-constant was verified to be about 2.5 μs as the v_{AF} waveform, given in Figure 3.9(b), covered 63% of maximum change in this time and settled to 0.6 V in about 12 μs.

As summarized in Table 3.2, the simulations and experiments agree fairly closely; and both support the underlying theory. With the exception of $\Delta V_{\max DN}$, the measured parameter values were within 20% of those calculated using the design equations. Since V_{AF} and R_{AF} were larger in the actual experiment the energy loss was also a bit larger than calculated: 4.6 μJ instead of 4 μJ. Even with significant variations from the ideal EAVP design, the hybrid filter reduced the ringing of a high-Q passive filter by more than a factor of five. The major sources of error were: $R_{AF} = 5\ \Omega$ instead of $4\,\Omega$; timing delay between SAF and load current switches; and R_C being much less than 100 mΩ.

3.2 Step-Current Electronic Load, I-load

Since the switched resistive load used in the previous section is incapable of forcing a step-up in load current, an IC was designed to be used in a switched inductor circuit to emulate a sharp current step-load for transient testing of the hybrid filter. Initially, the IC proposed was designed to achieve a step of 1-5 A. The on-chip slew rate specification was 1-5 A/ns on die and up to 1 A/ns off-chip through 0.5 nH equivalent package inductance. A sense resistor measurement path was built in for wafer probing but was never used since it was later decided that all the hybrid filter test circuits were to be mounted on the PCB level. Eventually, the I-load was verified for up to 2 A at 0.2 A/ns through an estimated 2 nH of package inductance. The I-load was still useful because, although lower than originally intended, the achievable di/dt was more than sufficient to reveal the effects of esl in the filter capacitor.

The complete circuit schematic is shown in Figure 3.10. The IC components are contained within the dashed box. The supplies, large (10 uH) inductance, L5, and current limiting resistor, R4, are external to the IC. The output filter capacitor, modeled as C2+R3+L8 is normally open-circuited. The IC consists of a two stage gate driver, a big MOSFET used as current switch, and another big FET and current-sense resistor which may be used to measure switching performance.

This circuit cannot be implemented on PCB due to inherently larger trace parasitic inductances but, instead, needs to be implemented on an IC technology with high frequency devices (f_t in the multi-GHz region). The IBM7WL process was therefore chosen.

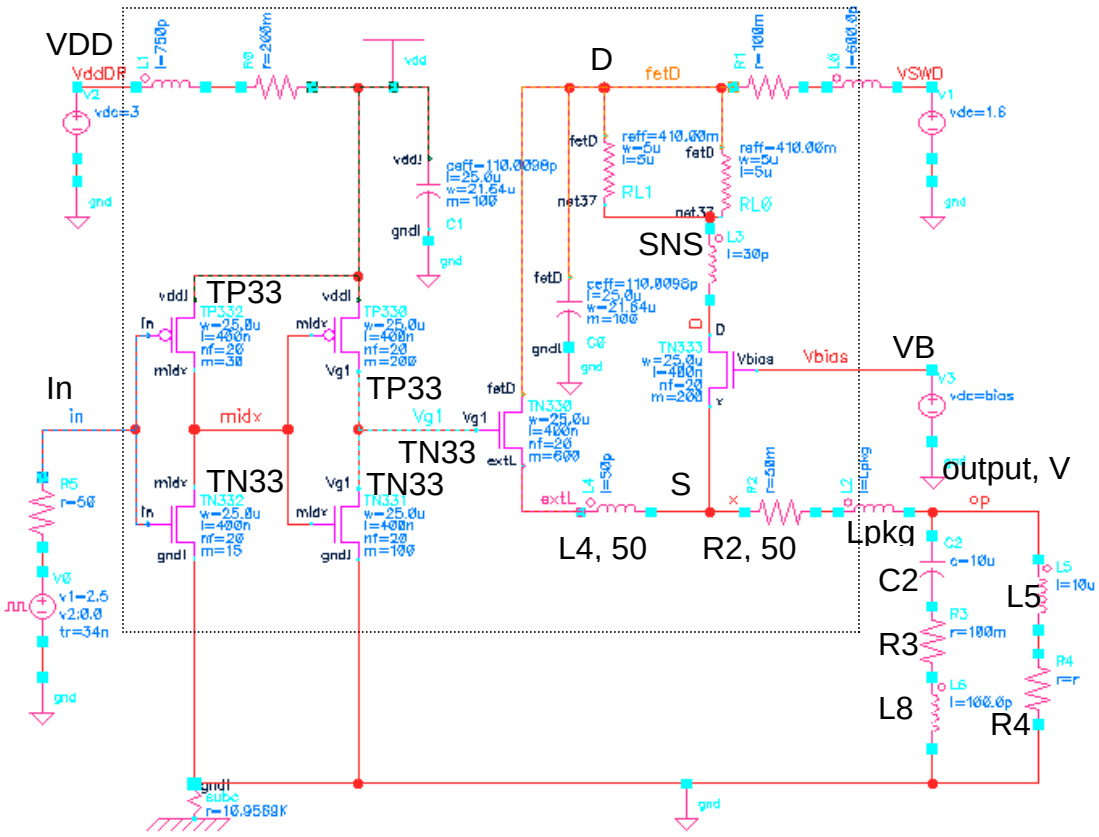

Figure 3.10. Circuit schematic of 5 A, 5 A/ns electronic load with built-in sensing.

The inductor is biased to carry up to 5 A which normally flows through the big FET switch, TN330, with the gate drive signal, “in”, normally high. The FET has been designed to dissipate a maximum of 250 mW. A standard 50 Ohm function generator with relatively slow (30-50 ns) rise-times will then cause TN330 to switch off very fast (in 1 ns), the external inductor acts like a current source as it keeps the current flow continuous, all the bias current must now switch to flow either through the C2+R3+L8 external load path OR the TN333 measurement path. If Vbias is set to zero and C2+R3+L8 connected then the external load path is selected. If Vbias is set to 2.5 V and C2+R3+L8 is disconnected then the measurement path is selected. Thus this circuit provides for switching in a very fast, positive current step to either an internal or an

external load. Of course the internal load is fixed for a given IC but the external load can be changed.

Figure 3.11 shows the transient performance obtained in Cadence simulation with C2, R3 and L8, 10 μF, 100 mΩ and 100 pH respectively. On chip, with VB set to 2.5 V, the preset current switched from the "DC" path to the sense path in 1 ns resulting in the corresponding voltage step shown. Externally, with VB set to 0 V, a resulting current step of 5 A at approximately 1 A/ns and the corresponding output voltage taken across the capacitive load are given. Since the esr of the capacitor was 100 mΩ with small esl, the transient voltage waveform closely resembles the current waveform on the initial ramp only it is inverted.

The actual design and layout of the I-load IC was done as a parallel array of over ten thousand smaller cells. The entire array was sectioned in a ratio of 1:3 in order to allow for isolated or independent use of each of the sections. The layout diagram is given in Figure 3.12. The smaller section can be seen to occupy the lower region of the layout.

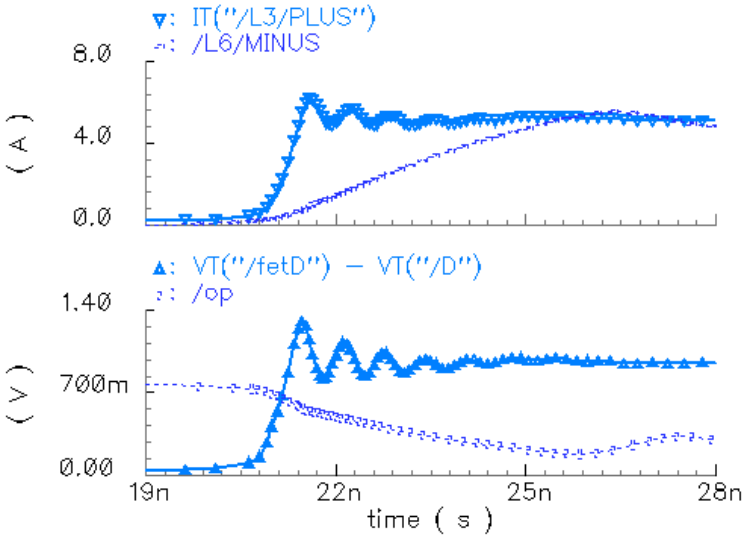

Figure 3.11. Internal and external I-load transient performance.

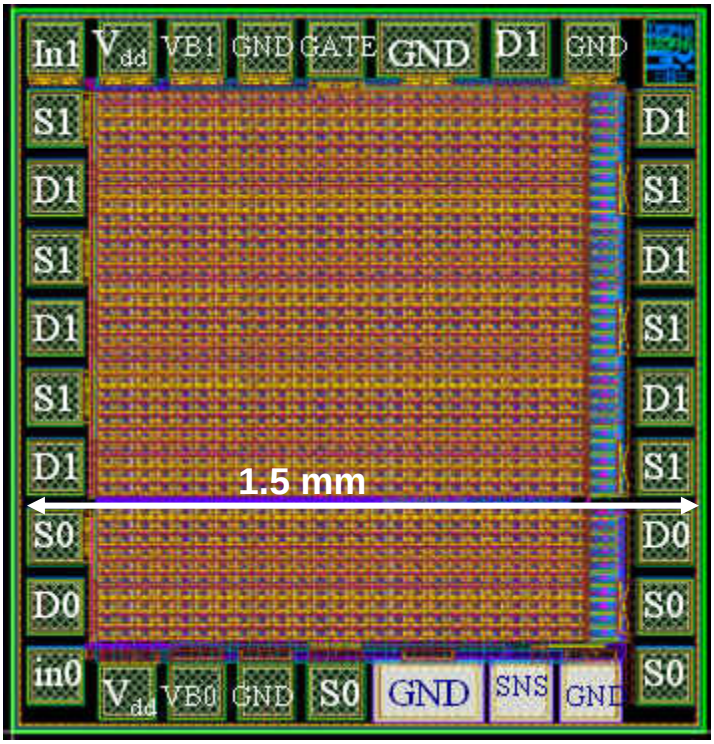

Figure 3.12. Layout of I-load IC.

The bonding diagram is shown in Figure 3.13. The IC is packaged in a 44 pin plastic quad flat-pack. Entering and exiting current paths were alternated in the pin-out to help lower the overall package inductance. The average inductance for the 1 mil diameter bondwire used in this package assembly is 1 nH/mm and the average bondwire length was 6 mm. Eight pins were allocated to carry the high current of up to 5 A. The complete I-load circuit, given in Figure 3.14, was modified from Figure 3.10 as follows:

- Bias resistor R4, $0.2\,\Omega$, was implemented with a 1 W, 1% current sense resistor.
- Load capacitor, C2, was $2\,\mu$F with lower esr and higher esl relative to the $10\,\mu$F capacitor with $100\,\text{m}\Omega$ esr and 0.1 nH.
- Lpkg was closer to 3 nH with PCB interconnect to L5 included and was much greater than the desired 0.5 nH.

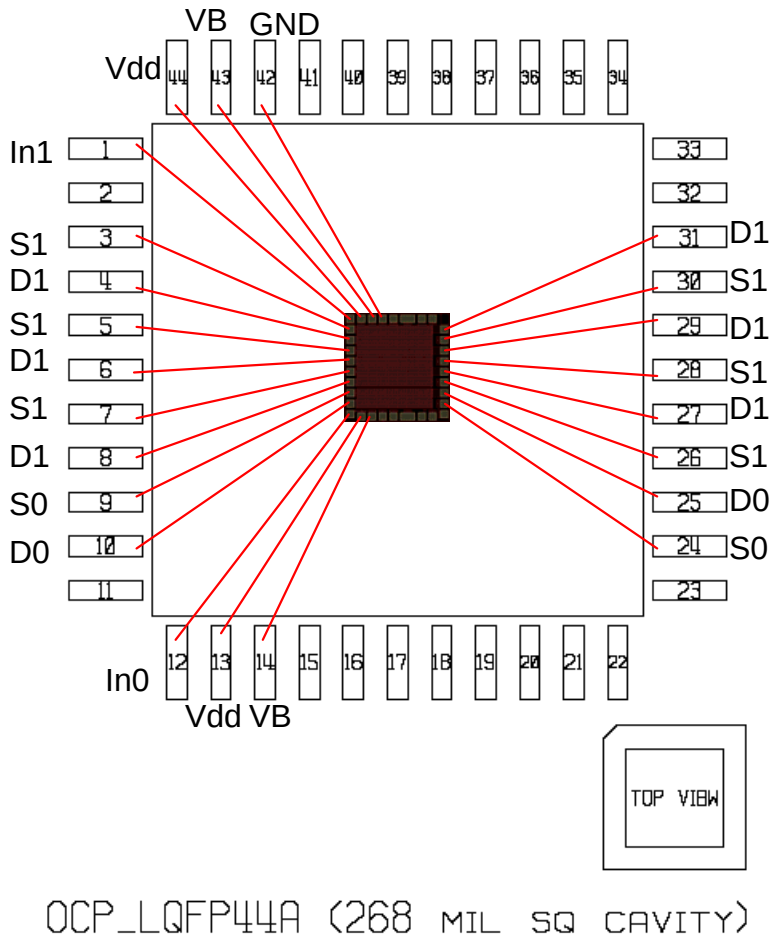

Figure 3.13. Bonding diagram for I-load IC.

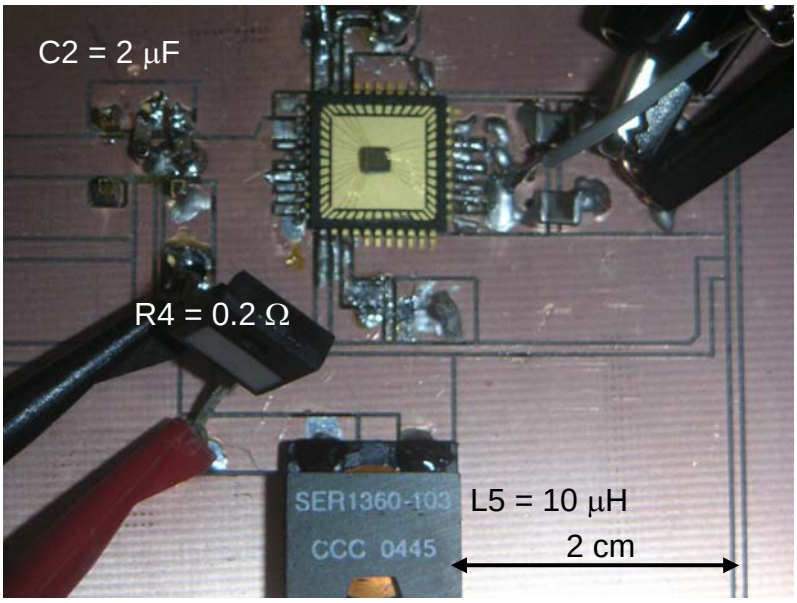

Figure 3.14. Completed I-load circuit.

To prevent breakdown without impacting initial transient response, a 2200 μF capacitor with 30 nH of interconnect inductance was connected across the 2 μF load capacitor.

With the above modifications, VDD = 3 V and VB tied to GND, the I-load circuit was given a 0-2.5 V square wave input and the voltage across the load capacitor, C2, was measured with the Tek2440 oscilloscope via an SMA connector and appropriately terminated 50 Ω coaxial cable. Note the 50 Ω termination on the coaxial probe cable is much greater than the impedance seen looking into the output node "V" and therefore the measurement apparatus does not impact circuit performance. Waveforms obtained for steps 1 A and 2 A are given in Figure 3.15 and Figure 3.16 respectively. The magnitude and the di/dt of the load current step were determined from the point-of-load capacitor voltage waveform, $v(t)$.

Since in general, the current through a capacitor, $i_C = C\frac{dv_C}{dt}$, where v_C is the voltage across the capacitor, then, referring to the waveforms of Figure 3.15 and Figure 3.16, the magnitude of the load current step, starting from zero and jumping up to ΔI, is given by

$$\Delta I = -\frac{dv}{dt}C2 \tag{3.1}$$

where $\frac{dv}{dt}$ is the slope just after the transient step and is relatively constant for some time. The minus sign in (3.1) indicates the load current is pulled out of the point-of-load capacitor. In principle the value thus obtained should be equal to the DC bias current in the inductor, L5.

As the current ramps up in the point-of-load capacitor, a downward, esl-induced voltage spike occurs. This spike peaks when the di/dt reaches a maximum and goes away when the current step transition is complete, i.e. when the di/dt returns to zero. Although the esl of the point-of-load capacitor, L6, may be unknown, a fairly accurate estimate of the di/dt can be obtained just from knowing the width of the spike, T_W, and the magnitude of the current transition:

$$di/dt \approx \Delta I/T_W \tag{3.2}$$

Really, this value is slightly lower than the maximum di/dt but as such it is still a useful measure. It can in turn yield an estimate of the esl of the point-of-load capacitor:

$$\text{esl}, L6 = \frac{V_L}{di/dt} \tag{3.3}$$

where V_L is the magnitude of the inductive spike.

Table 3.3. I-load measured data and calculated performance parameters.

Nominal DC current	V(R4) (R4 = 0.2 Ω)	dv/dt [V/ μ s]	T_W [ns]	V_L [mV]	ΔI [A] (C2 = 2 μ F)	di/dt [A/ns]	esl, L6 [nH]
1 A	201 mV	0.5	8	105	1	0.125	0.84
2 A	409 mV	1.0	10	160	2	0.2	0.8

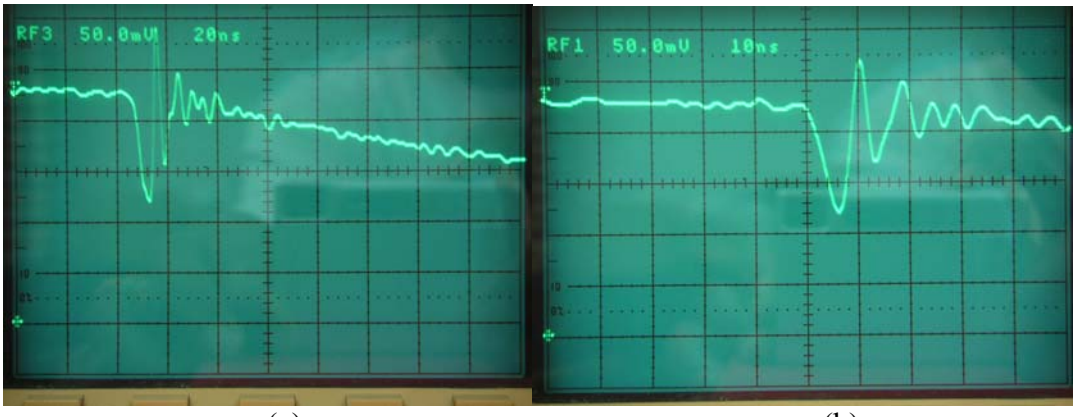

(a) (b)

Figure 3.15. Load voltage waveform for 1 A step (a) at 20 ns/div and (b) 10 ns/div.

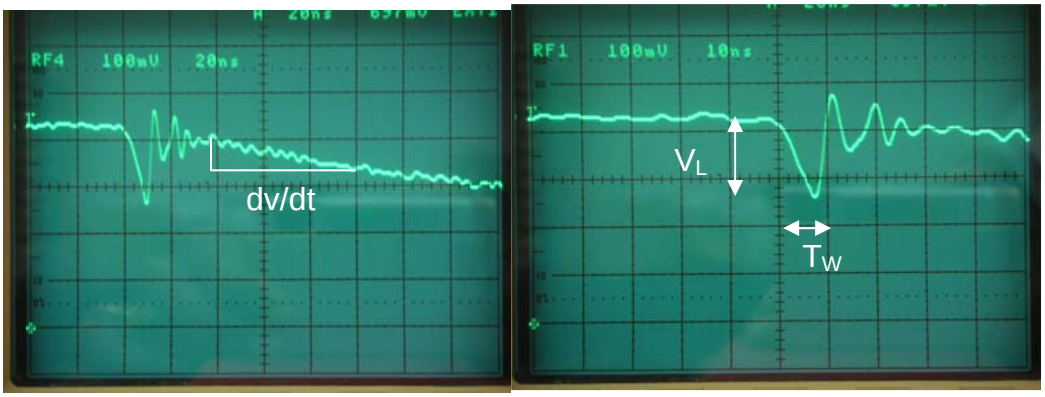

(a) (b)

Figure 3.16. Load voltage waveform for 2 A step (a) at 20 ns/div and (b) 10 ns/div.

Data taken from Figure 3.15 and Figure 3.16, together with other relevant DC measurements were used to calculate the magnitude and di/dt of the I-load circuit. The results are summarized in Table 3.3. The values of ΔI obtained from the oscilloscope measurements were consistent with the DC measurements. It was observed that the I-load switched in 8-10 ns which was more than twice as fast as the switched resistive load circuit. The maximum di/dt obtained was 0.2 A/ns and the esl values derived from both sets of measurements were consistent.

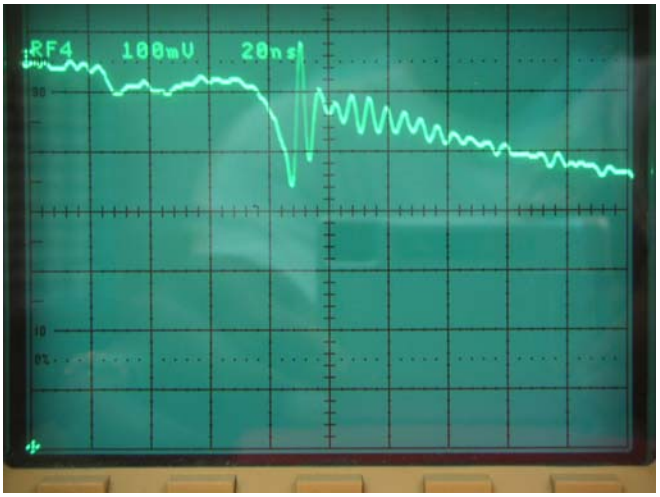

Figure 3.17. Load voltage waveform for 3 A.

At greater than 2.5 A however, the start of the inductive spike became less well-defined (Figure 3.17). This meant the switching of current flow from the IC to the point-

of-load capacitor started to be less sharp and indicates that the equivalent parasitic inductance (estimated to be about 3 nH) between the source of TN333 and the top of L5 was pulling the source of TN333 sufficiently below ground to turn on TN333, draw current through the substrate diode connection or both.

The decaying ringing that appears after the inductive spike is due to parasitics associated with the measuring probe. This ringing gets noticeably worse if a longer cable or a lower quality connector is used. It gets intolerably bad if the $50\,\Omega$ termination at the oscilloscope is removed.

Finally, it was observed that reducing VDD resulted in longer switching times and hence lower di/dt's.

3.3 Single Stage Hybrid Filter Experiment

The I-load circuit of the previous section was used to test a hybrid filter comprising a single stage passive filter and the complimentary SAF. The experimental setup is shown in Figure 3.18. The simplified circuit topology, along with design values, is given in Figure 3.19. The values for the passive filter in this experiment were chosen to match those related to the microprocessor power delivery model as listed in Table 1.3. The inductance in the SAF path was laid out to be twice that in the passive filter. The SAF was designed using the non-EAVP design equations (2.5), (2.26), (2.28) and (2.29). The design values are summarized in Table 3.4. It should be noted that once again the SAF design using the EAVP equations yield approximately the same values.

At the time these experiments were performed, the smaller I-load section had been disabled. Therefore as an added precaution against breakdown and bondwire fusing, for this experiment, the I-load was set to a switching time of about 20 ns by adjusting VDD (VDD = 2.4 V) and current was limited to a maximum of 1 A.

The 0.2 μF point-of-load capacitor was implemented by a parallel combination of two radial 0.1 μF capacitors. Even with the leads of these capacitors cut to less than 3 mm the equivalent esl of the soldered combination was about 3 nH, which dominates over the esr for the time window of concern. Hence the point-of-load capacitor esr could be ignored. Furthermore, this large value esl was beneficial in retaining the inductive spike at lower di/dt's thereby allowing for easy measurement of the load di/dt. In real application, however, the leaded capacitors would be replaced by low esl, surface mount parts.

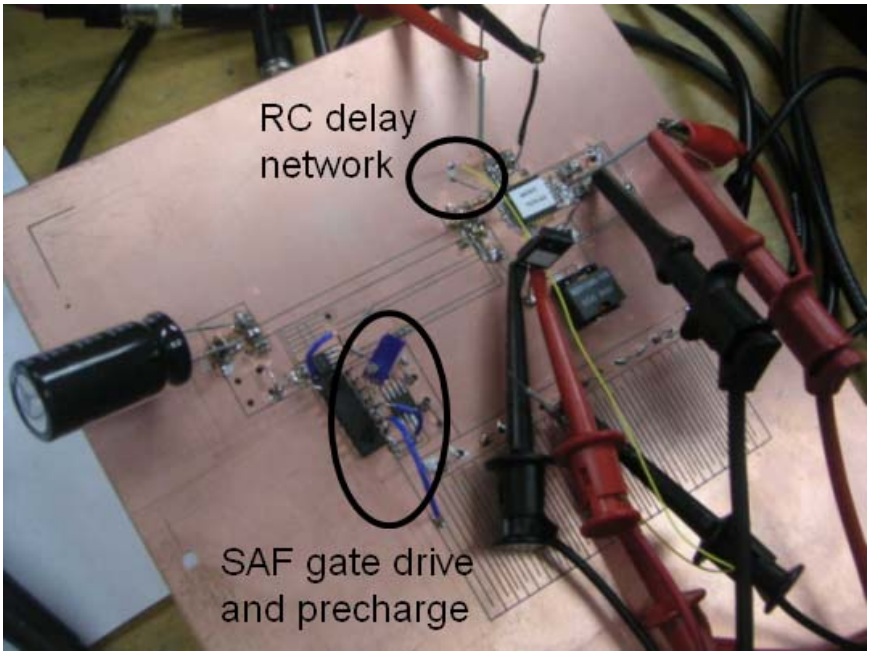

Figure 3.18. I-load circuit together with SAF and single stage passive filter.

The SAF switch was driven from the same input that went to trigger the I-load. However, this signal was first inverted using an HCT 7414 inverter and sent to an IXDD408 gate driver connected to the STS9NF30L switch in the SAF. (Use of the complementary gate driver is also an option but parts were not available at the time.) The 25 ns relative delay incurred by the signal in going through the extra circuitry was compensated by adding an RC delay on the input of the I-load circuit using 1 KΩ and 67

pF in a lowpass configuration i.e. the input signal entered through the resistor and the output, taken across the capacitor, went to trigger the I-load.

Note the 2200 μF electrolytic capacitor was not used as shown in Figure 3.18. It was removed and the leads were cut short to minimize esl before it was re-soldered. The esl and esr of the 2200 μF capacitor- about 4 nH and 60 mΩ - were lumped into L and R_L respectively. These values were obtained from measurement with the HP4194 impedance analyzer.

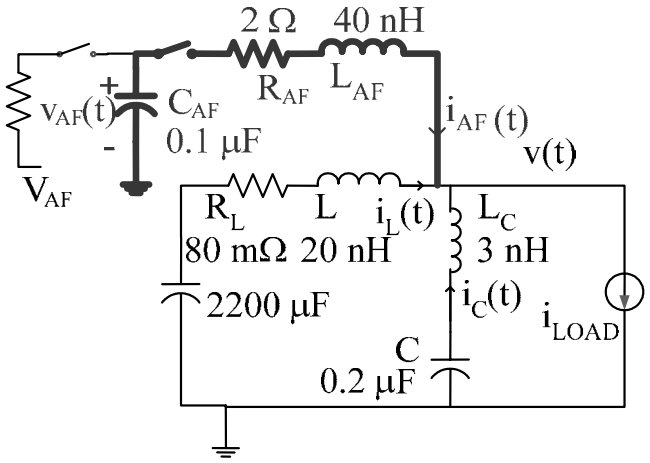

Figure 3.19. Simplified circuit topology for SAF and single stage passive filter test.

The SAF was never disconnected during testing as a safeguard against accidentally destroying the circuit in the process of unsoldering and soldering. In order to show the improvement in performance data was taken with the SAF precharge voltage, V_{AF}, set to zero. The response measured under this condition differs only slightly from the passive filter response by itself and therefore it serves as a good reference by which to judge improvement. The value of V_{AF} was adjusted to obtain good regulation and the performance parameters noted. The results obtained for tests at 0.5 A and 1 A are shown in Figure 3.20 and Figure 3.21 respectively. The key parameters are shown in Table 3.4.

Fortunately, components were readily available such that the calculated and actual designs values were equal within tolerance. Ignoring the first inductive spike caused by point-of load capacitor esl, it was possible to regulate the transient response to within 100 mV, the desired specification. For 0.5 A, $|\Delta V_{\max}| = 45$ mV and for 1 A it was 60 mV. However, the values of V_{AF} measured in the experiment were slightly higher than those calculated. These higher V_{AF} values in turn caused larger losses than calculated in the experiment- 44 nJ instead of 41 nJ at 0.5 A and 283 nJ instead of 162 nJ. The losses were obtained by first measuring C_{AF}, then, the initial and final steady state values of $v_{AF}(t)$, and substituting the measured values into equation (2.30).

Although the increase in loss seems high- more than 50% in the worse case- the regulation is significantly better, 60 mV instead of 100 mV. Also it may be partially the result of the parameter measurements being somewhat in error. For example the L_{AF} inductance measurements did not include the exact package inductance of the FET switch but only an estimate. Also the esl of the 2200 μF capacitor may have been measured with the leads close together and may have effectively been increased when the leads were bent further apart in soldering the component.

Table 3.4. SAF and single stage passive, hybrid filter design and performance.

$-\Delta V/\Delta I$	L	C	L_{AF}	C_{AF}	R_{AF}
0.1 Ω	20 nH	0.2 μF	40 nH	0.1 μF	2 Ω
t_r	R_L	R_C	L_C	SAF-load delay	$\|\Delta V_{pk}\|$
20 ns	80 mΩ	~10 mΩ	3 nH	+/-5 ns	2.2 V
ΔI	$\|\Delta V_{\max}\|$	Calc. V_{AF}	Exp. V_{AF}	Calc. E_{loss}	Exp. E_{loss}
0.5 A	45 mV	0.96 V	1.0 V	41 nJ	44 nJ
1.0 A	60 mV	2.2 V	2.5 V	162 nJ	283 nJ

The 20 MHz ringing that can be seen on the $v_{AF}(t)$ waveforms in Figure 3.20(b) and Figure 3.21(b) is caused by the measurement probe which had clip-leads and 1 MΩ scope termination.

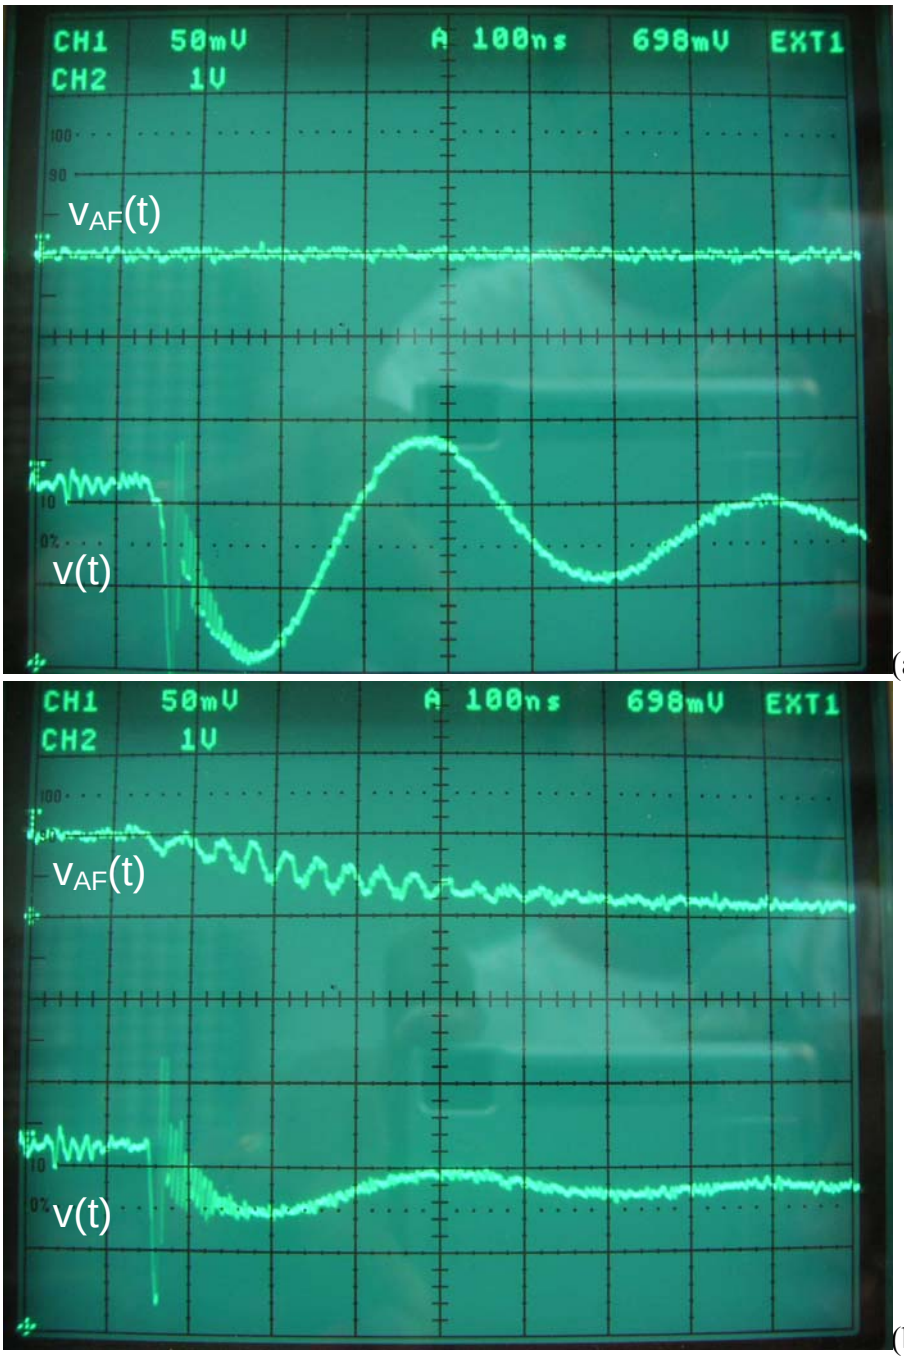

Figure 3.20. SAF and single stage passive, hybrid filter tested using I-load at 0.5 A with (a) V_{AF} = 0V and (b) V_{AF} = 1.2 V.

(a)

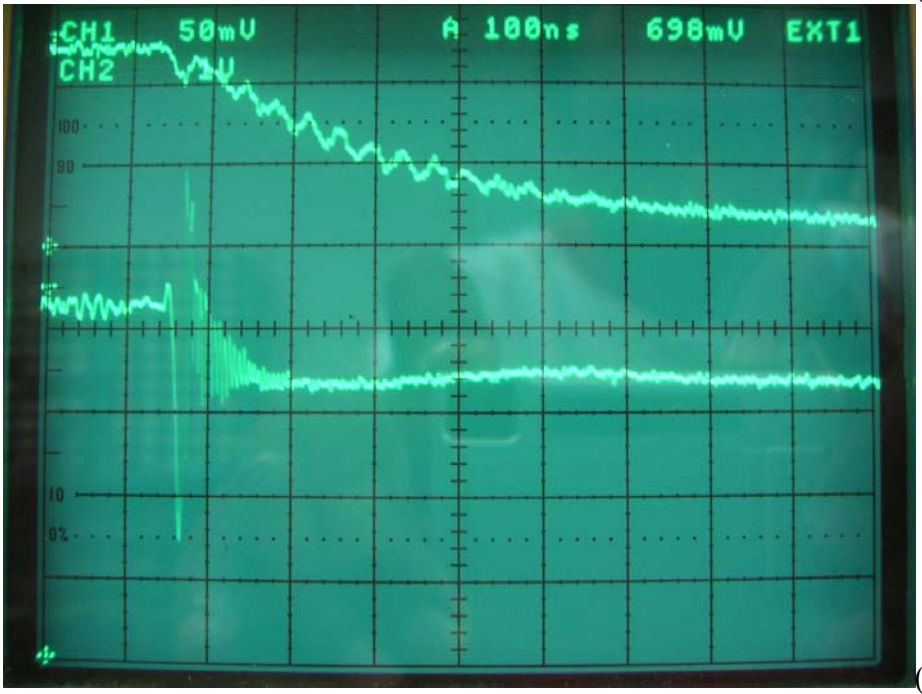

(b)

Figure 3.21. SAF and single stage passive filter tested using I-load set at 1 A with (a) $V_{AF} = 0$ V and (b) $V_{AF} = 2.3$ V.

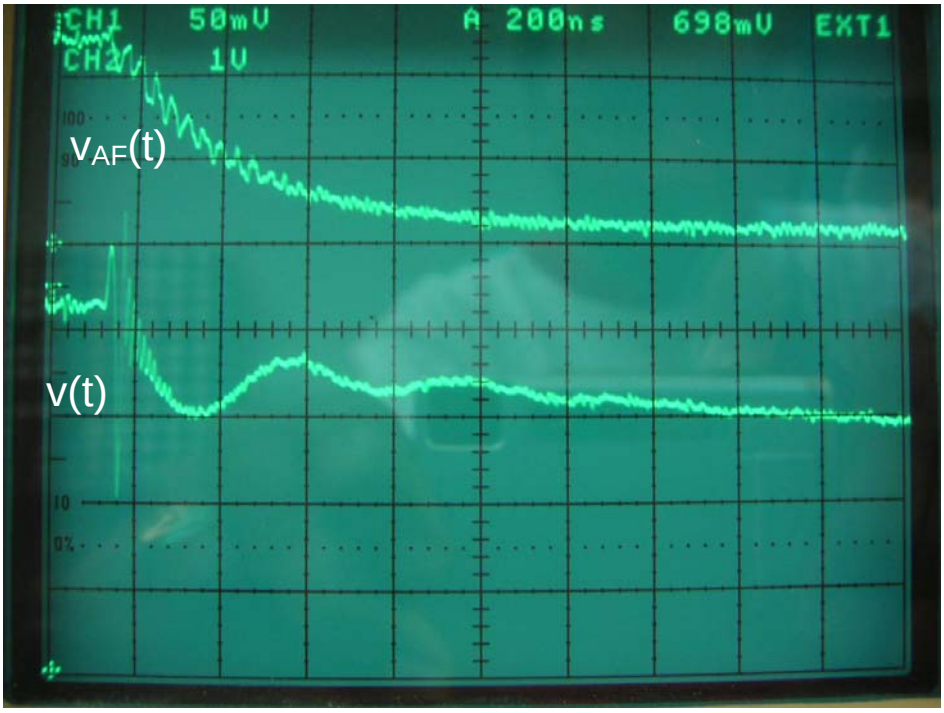

Figure 3.22. Effect of delay error in hybrid filter response.

In order to demonstrate the importance of proper timing in switching in the SAF, the relative delay was adjusted such that the SAF was switched on 20 ns in advance of the current step load. This resulted in a 40 mV overshoot spike which was immediately followed by the 140 mV undershoot spike caused by the esl of the point-of-load capacitor and then increased ringing in the transient for a few hundred nanoseconds thereafter. Note that switching in the SAF before the load transient edge did little to reduce the magnitude of the inductive first spike- it merely shifted its relative position up by 40 mV.

Finally Figure 3.23 has been included to show that the switching time of the I-load in these experiments was 18 ns and hence the maximum di/dt was about 50 A/μs for a step of 1 A.

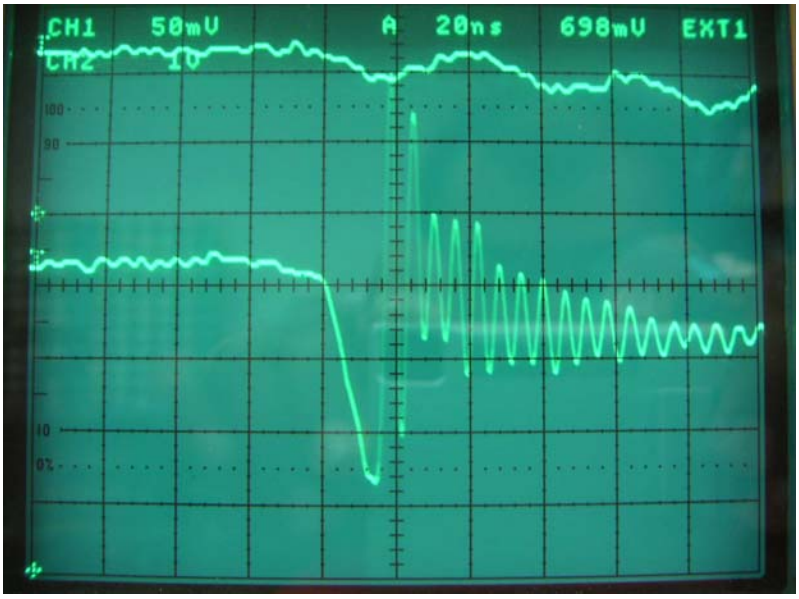

Figure 3.23. Spike induced by esl of point-of-load capacitor.

3.4 Operation of the SAF with Multi-stage Passive Filters

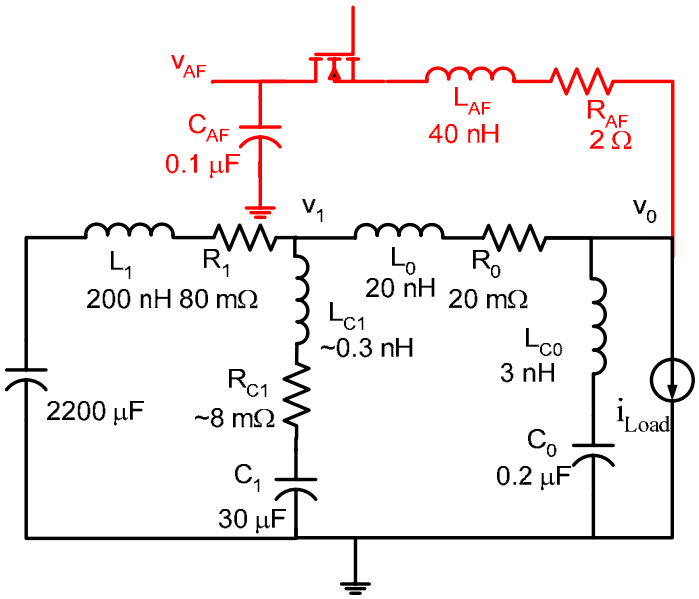

Figure 3.24. Circuit used to test SAF with multistage passive filter.

By extending the 2.2 mF electrolytic capacitor out an appropriate length of semi-rigid, twisted wire and soldering 30 μF in its place, the circuit of Figure 3.24 was implemented and used to test the SAF of the previous section with a two stage passive filter based on the scaled values of Table 1.3. Once it is shown that the SAF (designed for

an equivalent single stage) operates independently of the adjoining stage, it is a straightforward inference that lower stages have even less impact. Two stages, therefore, are sufficient to investigate SAF operation with multi-stage passive filters.

Note that designing the experimental two-stage circuit to exact specifications of a perfectly scaled model (i.e. any given column of Table 1.3) was not practical since it would be too time-consuming to design for exact capacitor esl. However the two-stage circuit was still developed using the design methodology in Chapter 2, with the impedance-scaled parameters of Table 1.3 as a guide, assuming SAF injection at the socket:

- Assume $L_0 = 20$ nH, $R_0 = 20$ mΩ, $L_1 = 200$ nH and $R_0 = 80$ mΩ are given. This is a justifiable assumption since, in full-scale application, these values are mainly set by the choice of socket and location of the BVR respectively. Little can be done to arbitrarily change these values. The 20 nH is consistent with (4 nH + 14 nH) of Table 1.3. The value, 200 nH is larger than the corresponding (10 nH + 70 nH) in Table 1.3 because it is assumed that fewer motherboard capacitors would be used in order to make room for the SAF.

- At the output node v_0, the DC path resistance, $R = 20 + 80 = 100$ mΩ. Designing for EAVP:
  $R = R_{C0} = R_{L0} = R_0 + R_{C1} \Rightarrow R_{C1} = 100\text{ m}\Omega - 20\text{ m}\Omega = 80\text{ m}\Omega$, $R_{C0} = 100$ mΩ;
  and $C_0 = \dfrac{L_0 + L_{C1}}{R_{L0}^2} \approx \dfrac{20\text{ nH}}{(100\text{ m}\Omega)^2} = 2\ \mu\text{F}$.

- At the output node v_1, the DC path resistance to the source (the 2.2 mF capacitor), $R = 80$ mΩ. Designing for EAVP: $R = R_{C1} = R_{L1} = R_1 = 80$ mΩ;
  and $C_1 = \dfrac{L_1}{R_1^2} \approx \dfrac{200\text{ nH}}{(80\text{ m}\Omega)^2} = 31\ \mu\text{F}$.

- Based on Table 1.3, C_0 is limited to 0.2 μF (the impedance-scaled value for C_{pkg}) roughly. Furthermore, practical implementation however, esr's and esl's cannot be controlled. In the experiment, significant esl in the point-of-load (output) capacitor was present, $L_{C0} = 3$ nH; R_{C0} was less than 15 mΩ; and R_{C1} was about 8 mΩ.

- The dominant peak in the impedance of this imperfect EAVP passive filter is caused mainly by the small value of C_0. Therefore, injecting the SAF at v_0, the key parameters in the equivalent single stage model of the passive filter were found: $L \approx L_0 = 20$ nH and $C \approx C_0 = 0.2\ \mu$F.

- Hence, the SAF parameters were determined, given that $L_{AF} = 40$ nH (a value typical in PCB level interconnects): $C_{AF} = LC/L_{AF} = (20\text{ nH})(0.2\ \mu\text{F})/(40\text{ nH}) = 0.1\ \mu$F,
$R_{AF} = (L/R + RC)/C_{AF} = (20\text{ nH}/0.1\ \Omega + 0.1\ \Omega \cdot 0.2\ \mu\text{F})/(0.1\ \mu\text{F}) \approx 2\ \Omega$; and
$V_{AF} = \Delta I/C_{AF}\,(L/R - RC) + V_{0-} + \Delta V$
$\approx (0.5\text{ A}/0.1\ \mu\text{F})(20\text{ nH}/0.1\ \Omega) + 0.1\text{ V} - 0.1\text{ V} = 1\text{ V}$; for $\Delta I = 0.5$ A
and
$\approx (1\text{ A}/0.1\ \mu\text{F})(20\text{ nH}/0.1\ \Omega) + 0.2\text{ V} - 0.1\text{ V} = 2.1\text{ V}$; for $\Delta I = 1$ A

- Note that the I-load circuit of Figure 3.10 was used not only to obtain a step current but also to set the bias voltage, $V_{0-} = \Delta I \cdot$ R4.

Table 3.5. Comparison of the passive filter in the experiment with the ideal EAVP filter.

	L_1	R_1	C_1	R_{C1}	L_0 [nH]	R_0	C_0	R_{C0}
Ideal	200 nH	80 mΩ	31 μF	80 mΩ	19.7	20 mΩ	2 μF	100 mΩ
Expt.	200 nH	80 mΩ	30 μF	8 mΩ	20	20 mΩ	0.2 μF	<15 mΩ

Table 3.5 allows for a comparison of the passive filter in the experiment with the ideal EAVP passive filter derived above using the method described in Appendix A. The lower stage of the passive filter is almost the same as the ideal EAVP design except that the esr, R_{C1}, is much lower than the desired 80 mΩ. This of course will lead to a bit of ripple at approximately the 65 KHz resonant frequency (or alternatively a 15 μs period) in the transient response. This is indeed observed in Figure 3.25. The oscillation causes an additional undershoot of 30%. However, the oscillation dies out in one cycle or 15 μs.

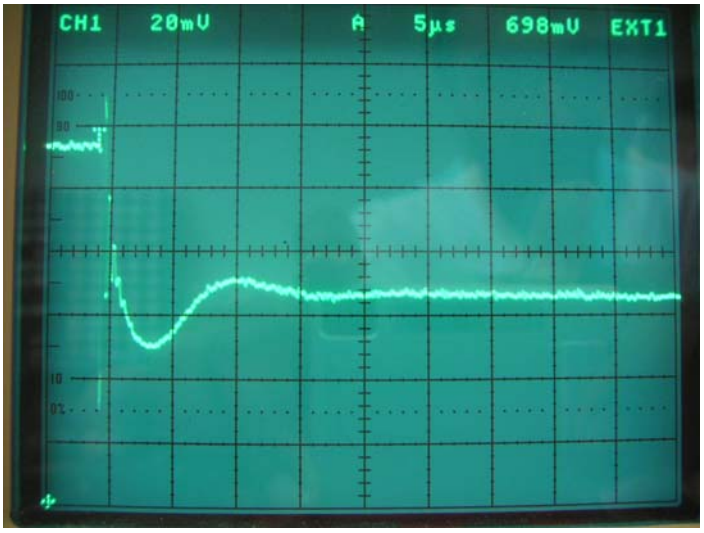

Figure 3.25. Output voltage, $v_0(t)$, on a long time scale

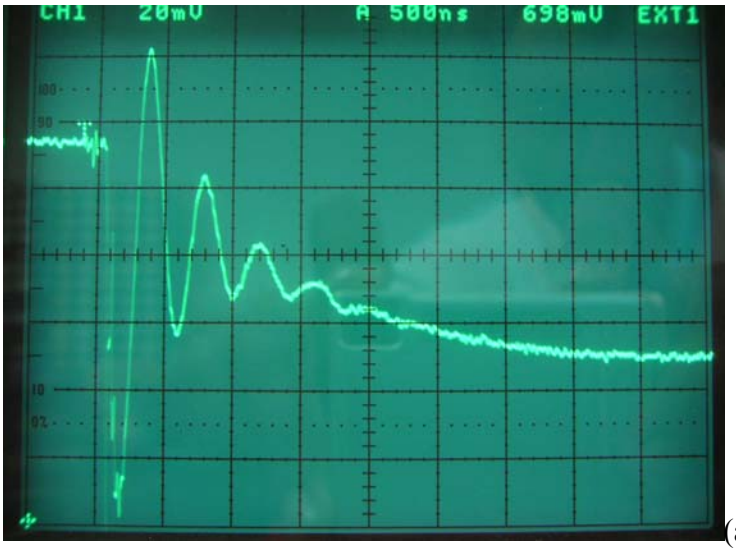

(a)

Figure 3.26. Waveforms of $v_0(t)$ showing the SAF working with multi-stage passive filter for (a) V_{AF} = 0V and (b) V_{AF} = 1 V at 0.5 A step load.

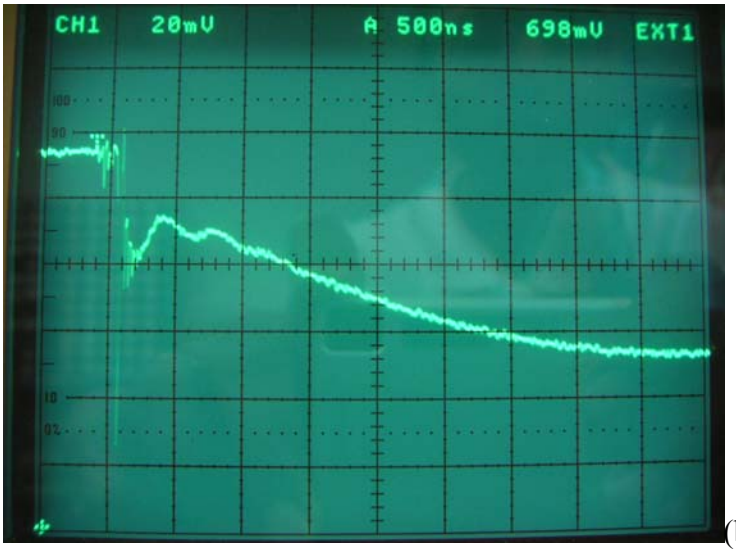

(b)

Figure 3.26. Continued

Riding on this low frequency decay is another decaying oscillation, seen in Figure 3.26(a), caused by the point-of-load capacitor, C_0, which is much lower than desired 0.2 μF instead of 2 μF. This oscillation is made worse by the fact that the esr R_{C0} is also much lower than ideal- less than 15 mΩ instead of 100 mΩ.

However, Figure 3.26(b) shows that the SAF designed for the equivalent single stage passive filter operates to regulate the output voltage to within 40 mV for the first couple of microseconds after the transient edge (ignoring the very first inductive spike caused by the 3 nH esl of the leaded point-of-load capacitors). After the first couple microseconds, the response follows almost exactly the trace of Figure 3.26(a). It droops just a little bit less because the SAF effectively adds C_{AF} in parallel with the lower frequency filter capacitors. The results show that the SAF operation is unaffected by lower frequency filter stages or by adjoining stages with interconnect inductance much greater than the esl of filter capacitor of that stage. Conversely, the SAF does not greatly affect the performance of the adjoining filter stages.

Similar results were also obtained at the 1 A step load setting (Figure 3.27). With $V_{AF} = 2.3$ V (as in the single stage experiment), the SAF regulated the output voltage to within 80 mV before falling into the lower frequency filter response. These findings indicate that multiple SAF's may be used with a multistage passive filter- a future topic that is left for Chapter 4.

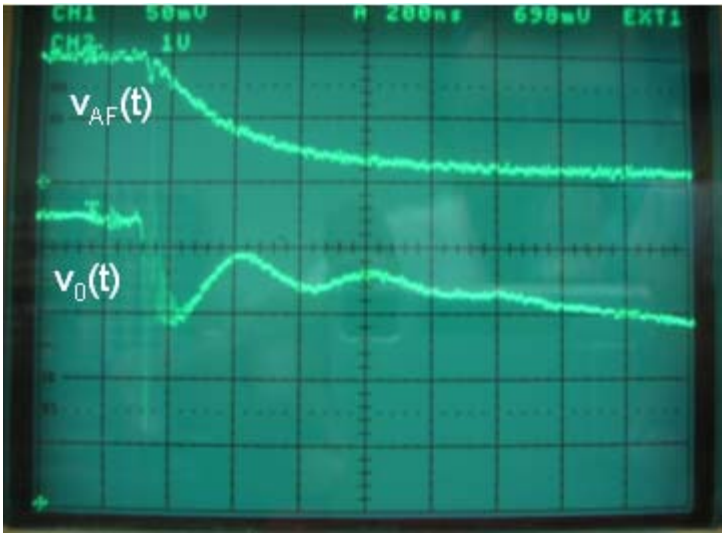

Figure 3.27. SAF operation with the multistage passive filter at 1 A step load, and with $V_{AF} = 2.3$ V.

3.5 Feedback Control Feasibility

Feedback control may be used to achieve continuous regulation when the load varies randomly provided the SAF is designed for the worst case step load that may occur in the system. The envisioned configuration is given schematically in Figure 3.28. The topology is similar to a linear active filter except for the extra switches that control the power supply feed to the circuit.

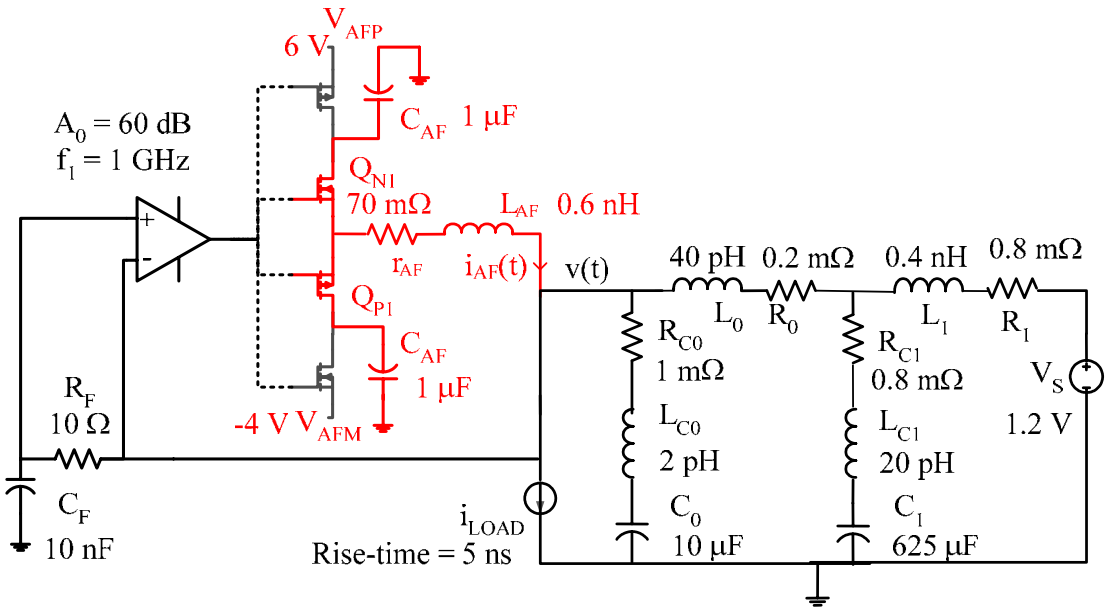

Figure 3.28. Hybrid filter with feedback control on the SAF.

The network comprising R_F and C_F may be a more sophisticated compensation circuit but in the simplest case, it extracts the desired reference voltage which in turn is maintained by the opamp in a negative feedback loop. The associated $R_F C_F$ time constant is chosen to be comparable with the risetime of the worst case step load. Bias circuitry between opamp output and FET gates are not shown. In simulation, voltage sources were used as level shifters to set the bias voltages such that both the SAF switches are just off while the pre-charge switches are biased on at DC. For a rise in load current, the V_{AFP} supply is disconnected leaving the upper C_{AF}, pre-charged to V_{AFP}, to provide the necessary energy through Q_{N1} in a controlled fashion (that is approximately linear) by actively varying the channel-resistance, while V_{AFM} pre-charges the lower C_{AF} to be ready for a fall in load current. In the event of a step load, the opamp saturates to the appropriate rail and the circuit effectively functions as a switched active filter once again.

The passive filter design is loosely based on the 2004 microprocessor model. The package capacitor, C_0, is 10 μF instead of 23 μF as it is assumed some additional space would be required for the pre-charge and control sections. The bulk capacitor has been removed and the equivalent resistance in the socket R_0, esr R_{C1} of the capacitor in the socket cavity, and equivalent resistance in BVR model R_1, have been tweaked to ensure 1 mΩ total resistance in the DC and high frequency paths. The esl of the socket cavity capacitor L_{C1} is 20 pH, in anticipation of better component packaging in the future.

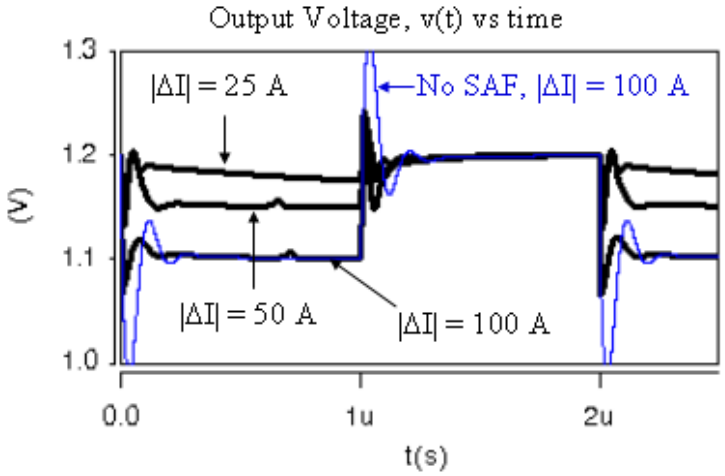

Figure 3.29. Output voltage of hybrid filter for various step loads.

Simulation waveforms show for the use of an SAF at the output node of this two-stage passive filter (the first stage of the passive filter is an EAVP design except for esl, L_{C1}, which is larger than desired) the ripple can be maintained under 150 mV for a range of step load magnitudes, from 25-100 A, and di/dt's, from 5-20 A/ns (Figure 3.29). Without the SAF, the voltage ripple exceeds 200 mV.

3.6 Applying the SAF to LAF Topologies

Figure 3.30 shows a hybrid filter comprising a single-stage passive filter and a linear active filter topology. Normally, when this circuit is drawn, the capacitors and inductors on the active filter supply are not shown- the supply is assumed to be ideal.

However, in reality, there is usually need for some filter capacitance near the circuit as there is physically some inductance in between the power source and the circuit. These components are shown explicitly in Figure 3.30.

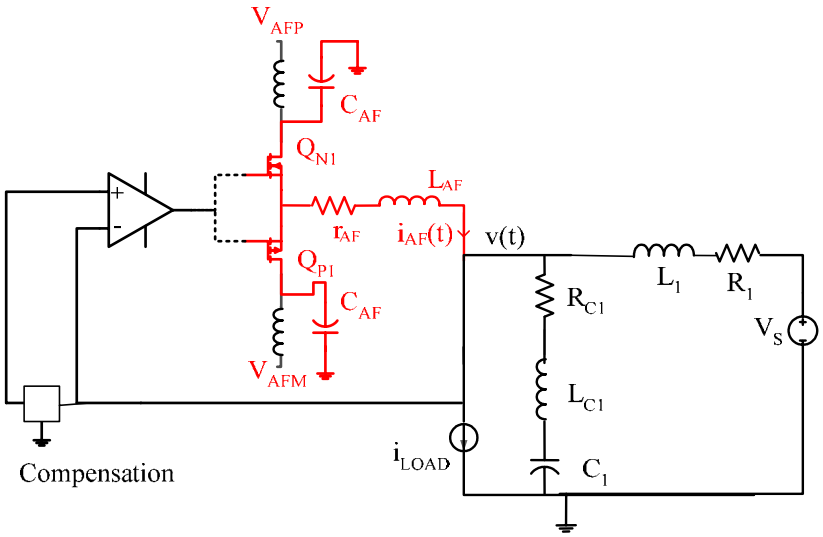

Figure 3.30. Linear active filter (left of iLOAD) operating with single-stage passive filter.

For a short time immediately after a sudden load transient, the parasitic inductances on the AF supply act like open circuits, the opamp saturates since the feedback loop is generally too slow to respond and the active half of the AF topology in Figure 3.30 becomes topologically the same as an SAF (Figure 3.28).

If the LAF design parameters are chosen to match those obtained using the SAF design equations, the output voltage would naturally be regulated over the time it takes the feedback control loop to respond or "catch up". Applying the SAF design values to an LAF may reduce the strain placed on the LAF control circuit design and operation.

Conversely, feedback control methods developed for LAF topologies may be adapted to SAF and hybrid filter control.

Basing experiments on a microprocessor power delivery architecture, and downscaling to inductances on the order of 10 nH, currents on the order of 1 A and

di/dt's on the order of 0.05 A/ns; the hybrid filter theory and design methodology was validated. Although the theoretically expected improvement was not achieved in every case, up to a factor of five improvement in regulation performance was obtained in both single stage and multistage passive filters. Deviations from ideal were consistent with the underlying theory as they were explained by component tolerance and timing errors. Remaining practical implementation issues have been investigated to some degree through simulation.

3.7 Pre-charging Mechanism

In order to take advantage of the energy savings in the SAF, one must find a way to efficiently pre-charge C_{AF}. One method is illustrated in Figure 3.31. By switching between charging and free-wheeling positions, C_{AF} (also called C_{PS} in the analysis) may be charged up from arbitrary initial voltage, V_{PSi}, to an arbitrary final voltage, V_{psAF} (which is also called V_{AF} in the larger hybrid filter circuit) that is less than $2\left(V_{EXT}-V_{PSi}\right)$. The rest state is shown in Figure 3.31(a). The source voltage, V_{EXT}, across a very large C_{ext} is roughly constant. Zero current is in inductor L_{PS}. When switched to the circuit of Figure 3.31(b), the current i_{PS2} builds up in L_{PS} while building up v_{PS2} across C_{PS}. Solving for v_{PS2} and i_{PS2},

$$v_{PS2}=\left(V_{EXT}-V_{PSi}\right)\left(1-\cos\omega_r t\right)+V_{PSi}=V_{EXT}-\left(V_{EXT}-V_{PSi}\right)\cos\omega_r t \tag{3.4}$$

$$i_{PS2}=\frac{V_{EXT}-V_{PSi}}{\sqrt{L_{PS}/C_{PS}}}\sin\omega_r t \tag{3.5}$$

are obtained in terms of the known variables in the system. Energy in L_{PS} is a function of the current through it, i_{PS2} whereas energy in C_{PS} is a function of the voltage across it,

v_{PS2}. When the stored energy in the inductor L_{PS} plus the stored energy in the capacitor C_{PS} is equal to the energy that must be stored in C_{PS} when the voltage across it is equal to V_{psAF},

$$\frac{1}{2}L_{PS}i_{PS2}^2 + \frac{1}{2}C_{PS}v_{PS2}^2 = \frac{1}{2}C_{PS}V_{psAF}^2 \tag{3.6}$$

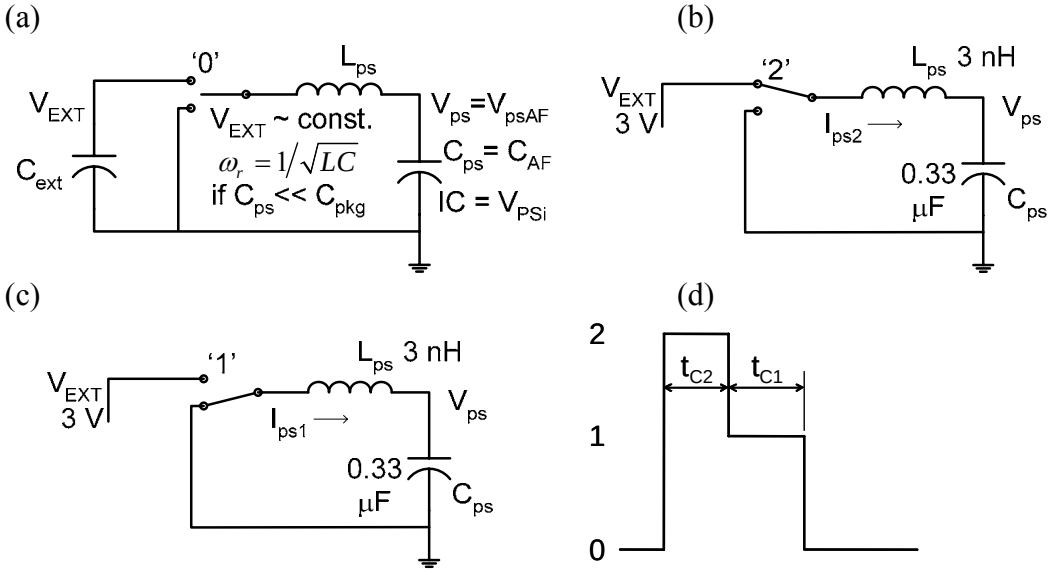

Figure 3.31. Pre-charging mechanism: Resonant converter (a) open, rest position '0', (b) charging position '2', (c) free-wheeling position '1', and (d) sample timing diagram.

Substituting (3.4) and (3.5) in (3.6), it becomes possible to solve for the time, t_{C2}, at which the energy in the inductor is exactly the additional amount required to build up the voltage across C_{PS} to V_{psAF}. The resulting equation (3.11) gives t_{C2} explicitly. Recall that $\omega_r = 1/\sqrt{L_{PS}C_{PS}}$.

After time t_{C2} in the forcing mode of Figure 3.31(b), the circuit is switched to the freewheeling mode given in Figure 3.31(c)

Hot-switching to the circuit of Figure 3.31(c),

$$i_{PS1} = I\cos\left(\omega_r t + \theta\right) \tag{3.7}$$

where

$$I = \sqrt{i_{PS2}\left(t_{C2}\right)^2 + v_{PS2}\left(t_{C2}\right)^2 / L_{PS}C_{PS}} \tag{3.8}$$

and

$$\theta = \tan^{-1}\frac{v_{PS2}\left(t_{C2}\right)}{i_{PS2}\left(t_{C2}\right)\sqrt{L_{PS}/C_{PS}}} \tag{3.9}$$

From (3.7), the first time i_{PS1} is zero is when

$$\omega_r t + \theta = \pi/2 \tag{3.10}$$

Solving (3.10), the time to zero-current switch-off is given by (3.12).

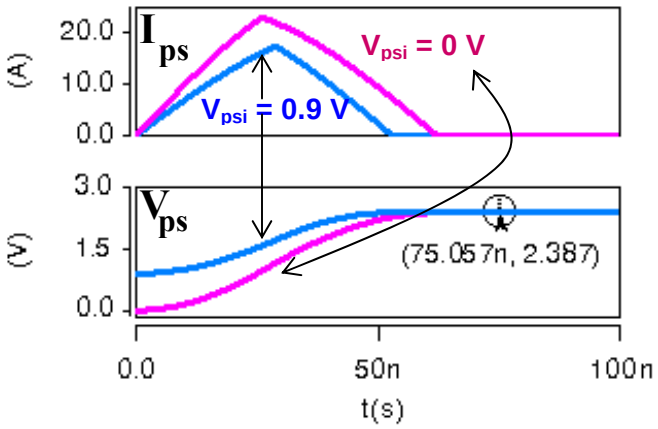

Figure 3.32. Pre-charging waveforms.

The timing diagram is given in Figure 3.31(d) and switching times for charging are given by (3.11) and (3.12). Sample waveforms when charging to 2.4 V are shown in Figure 3.32 for initial voltage equal to 0 V and then 0.9 V. Thus, analytic equations are available to pre-charge a capacitor from an arbitrary initial voltage (the microprocessor power supply voltage for example) to a desired value (V_{AF} for example) using a given supply (3.3 V perhaps). Similar equations may be derived for discharging but these are

not essential. Since the process is resonant, there can be negligible energy loss. The key design equations used in implementing the precharge mechanism of Figure 3.31 are listed here

$$t_{C2} = \frac{1}{\omega_r} \cos^{-1} \frac{V_{EXT}{}^2 + \left(V_{EXT} - V_{PSi}\right)^2 - V_{PS}{}^2}{2V_{EXT}\left(V_{EXT} - V_{PSi}\right)} \tag{3.11}$$

$$t_{C1} = \frac{1}{\omega_r}\left(\frac{\pi}{2} - \tan^{-1}\left[\frac{1}{\sqrt{L_{PS}/C_{PS}}} \frac{v_{PS2}(t_{C2})}{i_{PS2}(t_{C2})}\right]\right) \tag{3.12}$$

CHAPTER 4
APPLICATIONS OF THE HYBRID FILTER

Step loads occur in almost every electronic system, usually the result of turn-on or turn-off of all or part of the system, and often force high power surges and sags. Efficiently regulating step loads is a challenge in the design of power delivery systems [Zho00]. Furthermore, as mentioned in Chapter 1, linear active filters tend to be lossy while previously reported switched methods are not flexible enough to find widespread application. Thus the SAF may readily be applied in a variety of systems, from microprocessor and laser power supplies to electromechanical machines, to assist in regulation, energy-saving and high slew-rate power delivery under step load conditions. Furthermore, if feedback control is developed for the SAF, circuits can be designed to perform these functions continuously for arbitrary load transients.

Since the SAF theory and hybrid filter design methodology described herein arose out of research into power delivery systems for future generation microprocessors, the majority of this chapter will discuss applications to microprocessor power delivery systems. However, this work may find application in any electrical system where there is need for efficient load regulation but where there is insufficient space for point-of-load filter capacitors. Other applications- retrofit filter circuits that circumvent tight spots to provide regulation, active damping of switched coils and complementary regulation circuits for EMI filters- will be briefly discussed last.

The architecture and the corresponding lumped element model for the microprocessor power delivery path are first described in order to identify the “passive

filter" that exists in a real microprocessor power delivery system. Possible locations of the SAF are then discussed with specific regard to the validity of the lumped element model at the injection points. Factors impacting the choice of injection point are also discussed. The SAF is designed to work with the passive filter at each injection point for which the lumped element model is valid and Saber simulation results are given for each design.

4.1 System Architecture and Lumped Element Model

Presently computer motherboards share a general power supply chain-architecture similar to Figure 4.1 [Ren04], where the power path is kept as short and as heavily "bypassed" as physically possible. The corresponding lumped element model is given in Figure 4.2.The bulk voltage regulator or BVR, usually a switching regulator, is located furthest from the microprocessor and can be represented by an equivalent linear model [Vor89] comprising V_{BVR}, L_{BVR} and R_{BVR}. Filter capacitors are judiciously placed on the motherboard around the BVR and socket (C_{mb}); on the motherboard within the socket cavity (C_{skt}); and on the back surface of the microprocessor package (C_{pkg}). The processor core die also has some on-chip filter capacitance (C_{die}). The layout of these capacitors, power and ground return connections are carefully optimized to keep the parasitic inductances low [Li00].

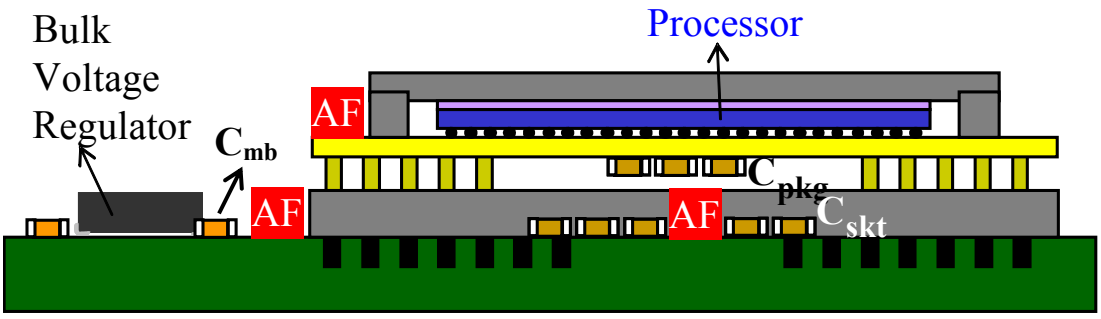

Figure 4.1. The 2004 Power delivery architecture system diagram.

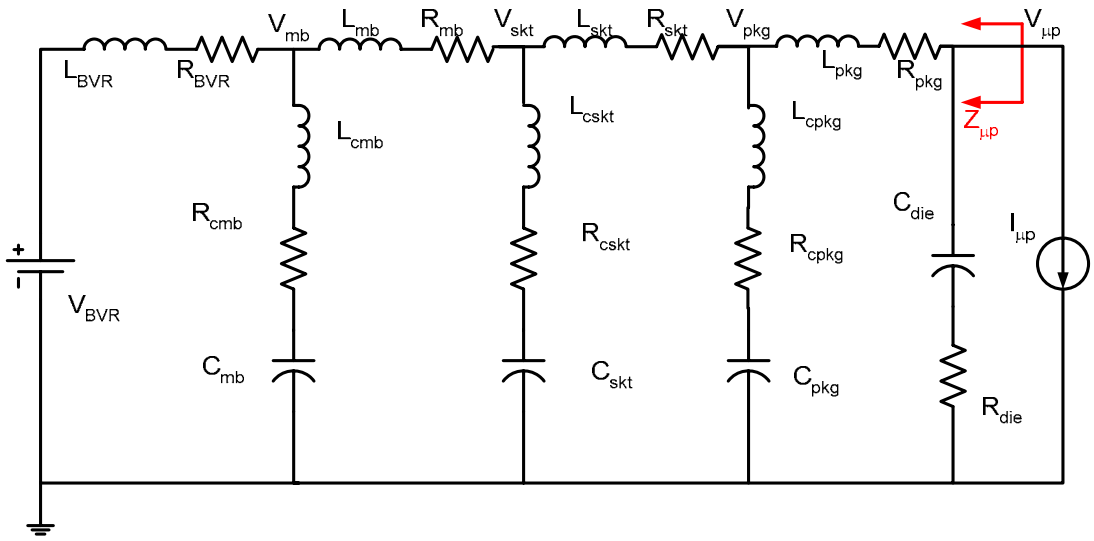

Figure 4.2. Lumped element model corresponding to Figure 4.1.

Model parameters vary with motherboard, VRM, package and microprocessor. Table 4.1 lists the parameters for three of the most recent power delivery path models. The 1998 model refers to the Pentium 4 in the 478 pin Pin Grid Array or PGA package [Int03, Luo02]. The 2004 model was obtained through private communications with the Center for Power Electronics Systems, CPES and are consistent with the datasheets and design guidelines for the Pentium 4 (Prescott) microprocessor in the 775 Land Grid Array or LGA package [Int05a, Int05b].

In reality, the model parameters L_{die}, L_{pkg}, and L_{cpkg} have little physical relevance and are kept merely to be consistent with the linear scaling associated with the lumped element approach. These sections can only truly be modeled by a distributed circuit model approach. One method discussed in [Rah04] subdivides the die into a large array of smaller cells each of which can be modeled with a lumped element network. The entire array then models the distributed circuit effects. Adopting this approach, the SAF injecting at the die would also need to be subdivided into an array of smaller cells, each with an SAF and injection point dedicated to a particular blocked area of the power grid

on the die. The effects of interaction between adjacent blocks on the power become negligible if current loops for each block are well isolated in the layout via star connections for example and the designed ΔV for the SAF is sufficiently lower than the actual ripple specification. Although tractable, the optimization problem for such a design has not been solved. Therefore direct injection into the die is not considered in this work. A similar distributed effect may occur at the package capacitor, C_{pkg}, but it is not expected to be as significant. On that plane, the di/dt is lower. Also additional voltage variation here may still be tolerable as it is followed by another filter stage on the die.

Table 4.1. Power Delivery Path Model Parameters

Model Element	Model Parameters	1998 Model [Luo02] 478 pin PGA	2004 Model 775 pin LGA
BVR (linear model)	V_{BVR}	1 V (estim.)	1 V (estim.)
	L_{BVR}	Data varies with design (0.4 nH, 1 mΩ used in simulations) (1 nH, 1 mΩ typical)	
	R_{BVR}		
C_{mb}	C_{mb} [uF]	5 600	6 700 (12 x 560)
	L_{cmb} [pH]	800	700 (50 + 8000/12)
	R_{cmb} [mΩ]	1	0.6 (0.02 + 7/12)
L_{mb}	L_{mb} [pH]	21	100
	R_{lmb} [mΩ]	0.1	0.2
C_{skt}	C_{skt} [uF]	240	850 (18 x 47)
	L_{cskt} [pH]	340	140 (26 + 2000/18)
	R_{cskt} [mΩ]	0.17	0.2 (0.03 + 3/18)
L_{skt}	L_{skt} [pH]	120	40
	R_{lskt} [mΩ]	1.1	0.6
C_{pkg}	C_{pkg} [uF]	26	23
	L_{cpkg} [pH]	4.6	1 (25/23)
	R_{cpkg} [mΩ]	0.54	Unavailable (~0.02)
L_{pkg}	L_{pkg} [pH]	6	Unavail. (~3)
	R_{lpkg} [mΩ]	.03	Unavail. (~0.015)
Micro-processor	C_{die} [uF]	0.53	Unavail. (~1)
	R_{die} [mΩ]	0.1	Unavail. (~0.7)
	L_{die}	Assumed to be negligible	
	I_{up}	120 A, $\pm$150 A/ns	

For the rest of this chapter the 2004 model will be used.

4.2 SAF Design for Various Locations

The need to consider different points of SAF injection arises for both theoretical and practical reasons. Theoretically, the best injection node is at the specific filter stage which is responsible for the dominant resonant peak in the output impedance. As a practical matter, however, injection at certain nodes may be difficult or impossible. In the case where injection at the preferred node is not possible, injection at adjacent nodes may also improve performance; although by not as much as would be obtained if injecting at the problem node.

In the fastest microprocessors, the main cause of peaking in the output impedance, Zup, is the relatively small effective die capacitance, C_{die}. Therefore, injection at the die would be preferred. However, presently, microprocessor layout and packaging are not designed to allow this. The possible locations for the SAF, or any AF, are shown in Figure 4.1. If injecting into C_{pkg} the best location for the SAF is on the microprocessor package as the inductance may be lower. Injection at the socket is also possible. In fact, it is a more feasible option. Since the di/dt at this node is lower, the design of a feedback control circuit would in principle be easier as the control bandwidth would be less than that necessary to control the package-mounted SAF injecting into the package capacitor. Furthermore the lumped element model has been shown to be valid at this node [Ren04]. For injection at the socket, the SAF is best placed within the socket cavity although locating it outside on the motherboard is also possible. However, when mounted outside the cavity on the motherboard, the SAF path inductance and hence the loss are expected to be larger.

As described in Chapter 2, after simplifying the multistage passive filter (Figure 4.2) with respect to the injection node to a single stage model, the passive filter could be

augmented by an SAF when a purely passive EAVP design is not realizable. Design examples for two cases are given below. A circuit of a 4-phase buck converter, obtained through collaboration with CPES at Virginia Tech, was used as the BVR in the simulations. The extraction of the linear model of the BVR is given in Appendix B.

4.3 SAF Mounted On Package and Injecting into the Package Capacitor

The design process is summarized in Table 4.2. The first step is to find the equivalent impedance the SAF would meet when injecting current. In the present example, with the SAF injecting into the package, the injection node is labeled 'V_{pkg}', as shown in Figure 4.2. The equivalent impedance is found by injecting a test current source at V_{pkg} without breaking the circuit, i.e. the model from source to load is kept intact.

Table 4.2. SAF design for injection at the package.

Step	Description	Equations	Key Parameters	Result
1	Find impedance at injection node		$Z_{out\ n}$ i.e. Z_{pkg}	See Figure 4.3
2	Model fit to dominant peak		L, C	150 pH, 21 μF
3	Calculate R, Q	2.2	R, Q	1 mΩ, 1.34
4	Approx. as lossless or damped		Q, R_L, R_C	lossless
5	Choose L_{AF}			600 pH
6	Derive SAF using appropriate set of design equations	2.5; 2.26; 2.28; 2.29	$C_{AF} = 5.3\ \mu$F; $Q_{AF} = 0.374$; $R_{AF} = 33$ mΩ; $V_{AF} = 3.43$ V	

Saber simulation generated Z_{pkg} (Figure 4.3) which exhibits a dominant resonant peak at 2.8 MHz, with a 10 MHz "bandwidth" in excess of the desired 1 mΩ impedance magnitude. The SAF will be designed to combat the effects of this peak. Note that beyond 100 MHz the impedance curve rises at 20 dBΩ/decade due to the 3 pH, L_{pkg}, inductance. This is not a resonant peak and the SAF cannot be designed to overcome it. In fact this rise in impedance further indicates that the SAF injecting at the package

cannot overcome any higher frequency resonance associated with a filter stage beyond L_{pkg} and nearer the load.

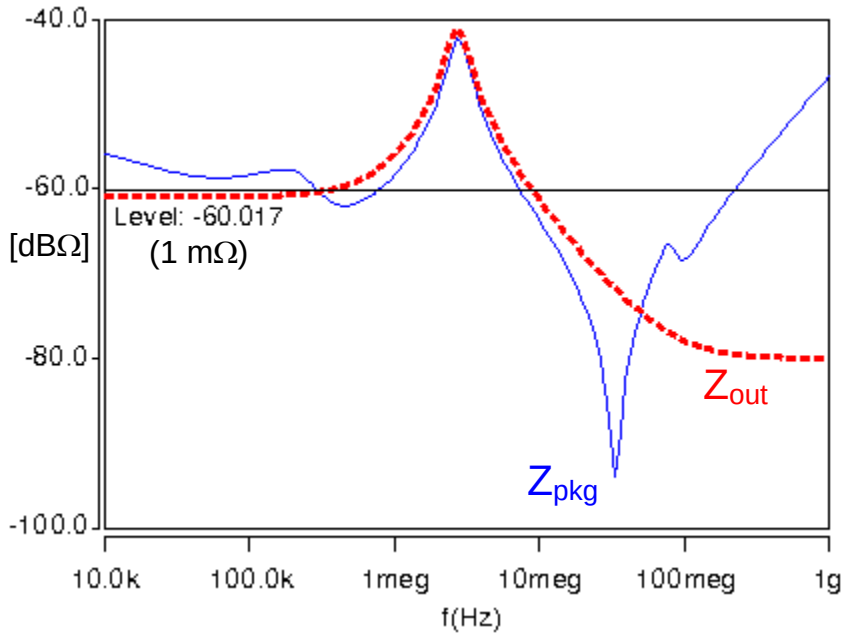

Figure 4.3. Modeling the impedance seen by SAF injecting into package node.

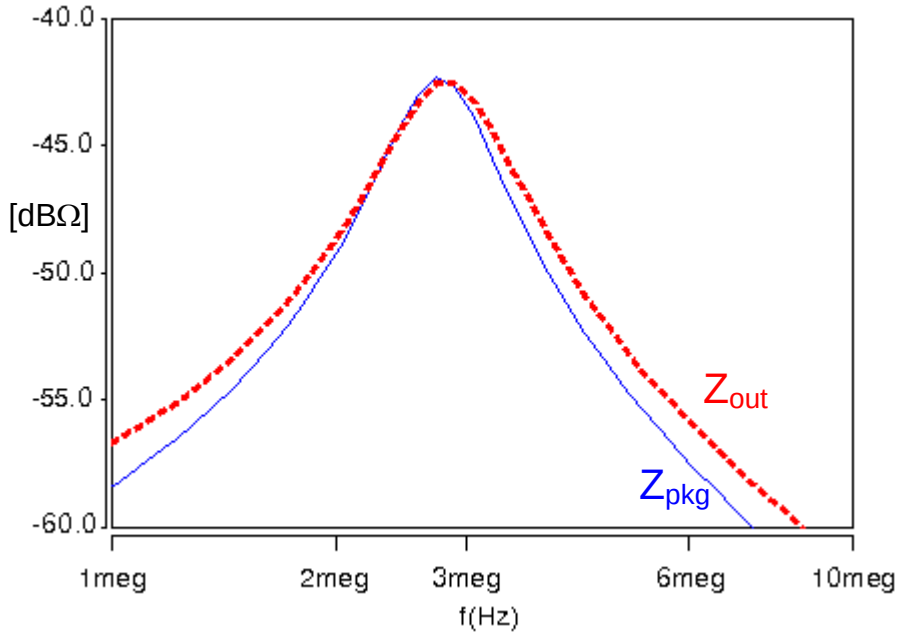

Figure 4.4. Magnified view of Figure 4.3 showing resonance peak approximation.

Focusing on the resonant peak (Figure 4.4) a second order model fit was obtained after a few iterations in simulation. This is step 2 in Table 4.2. The corresponding single-

stage circuit model is given in Figure 4.5. The value of $C = 21\ \mu\text{F}$ is slightly lower than what one would obtain by inspection of Figure 4.2 and Table 4.1 due to the effects of the esl, L_C. In this sense, the SAF design methodology is capable, to an extent, of dealing with esl in the filter capacitors. Similarly, the value of $L = 150$ pH is slightly lower than the 180 pH one would crudely extract by inspection due to the effects of the other filter stages nearer the source.

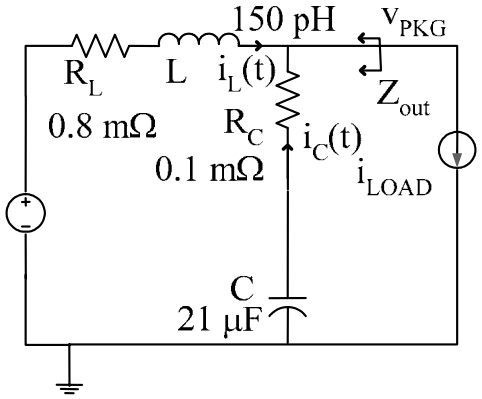

Figure 4.5. Equivalent passive filter with respect to Vpkg injection point.

By calculating $Q = 1.34$ via (2.2) as listed in step 3 of Table 4.2, it can be determined whether it is better to use the design equations for the damped (EAVP) or the lossless (non-EAVP) case.

Since $Q = 1.34$ and $R_C << R$ (0.1 mΩ << 1 mΩ), it is better, from the standpoint of achieving the best regulation, to approximate the passive filter as lossless (see step 4 of Table 4.2) and therefore choose to design the SAF using non-EAVP set of design equations.

The next step- choosing L_{AF} - allows the designer some flexibility. Using $L_{AF} = 0.6\,\text{nH}$ (since this value is consistent with the other parasitic inductances in the model), $\Delta I = 100\,\text{A}$ and $\Delta V = 100\,\text{mV}$, the corresponding SAF was designed in the final

step of Table 4.2 using the non-EAVP SAF equations and the simplified circuit topology is given in Figure 4.6. All SAF parameters appear to be reasonable- they can be implemented practically. Switches with far less than 33 $\mathrm{m}\Omega$ exist. A few microfarads of capacitance can occupy a very small footprint and a voltage of about 3-4 V can easily be taken from the motherboard.

So for $L_{AF} = 0.6\,\mathrm{nH}$, $V_{AF} = 3.23\ \mathrm{V}$, a value consistent with available voltages in the system. If V_{AF} were much larger than the available supply voltages, the designer would be forced to reduce L_{AF} by adjusting the circuit's location and layout. The switch on-resistance and SAF path parasitic resistance must total $33\ \mathrm{m}\Omega$, the R_{AF} value. Given that $C_{AF} = 5.3\ \mu\mathrm{F}$, it is a small fraction of C_{pkg} which may occupy an even smaller fraction of surface area since narrower terminal devices with larger esl may be used.

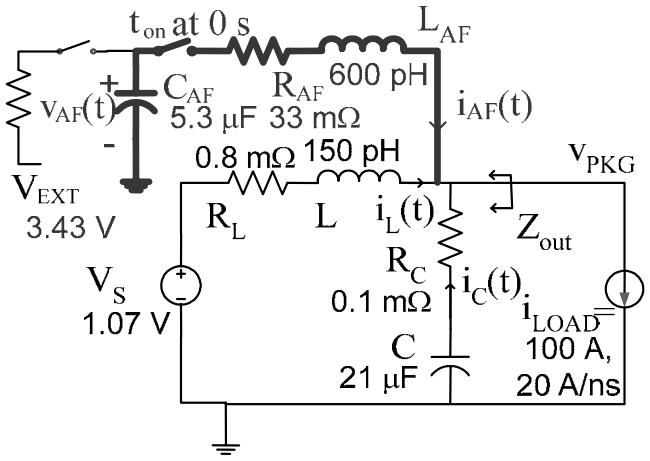

Figure 4.6. SAF design using the single stage model of Figure 4.5.

Simulating the SAF with the full linear 2004 power path model of Figure 4.2, the improvement in load regulation performance was investigated. With the SAF designed for injection at the package, load regulation to within 115 mV, or 95% of the steady state droop was observed at the package (V_{pkg}) for 100 A, 20 A/ns step load at the die. Notice

that the steady state droop is more than 100 mV because in the linear model, the DC resistance between source and load is greater than 1 $m\Omega$. Feedback control in a switched regulator will correct this phenomenon. Without the SAF, there is much more ringing- about 250 mV- which takes almost 2 μs to settle. At the actual microprocessor die load, the transient response, V_{up}, is essentially the same as V_{pkg} except for a short burst of high frequency ringing just after the step load.

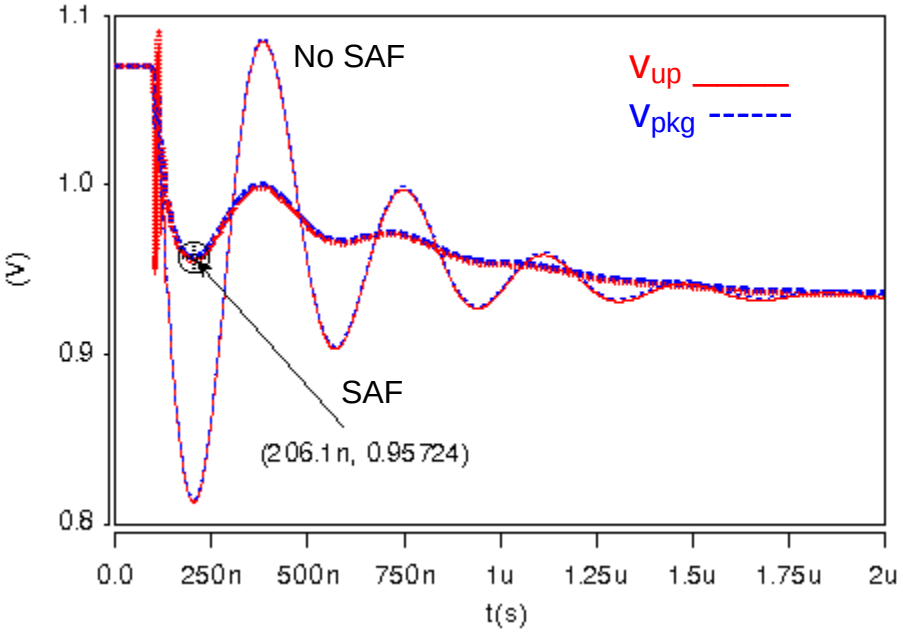

Figure 4.7. Performance improvement with SAF; simulated with linear power path model.

Zooming in Figure 4.8 to observe this high frequency (about 100 MHz) effect at the die, it is clear that the SAF injecting at the package does little to reducing this 100 MHz ringing. This observation is consistent with the fact that from 100 MHz and beyond, the impedance of L_{pkg}, effectively isolates the SAF from the load as indicated in Figure 4.3. However, such a narrow spike (less than 10 ns wide) may not be as severe a performance and reliability concern.

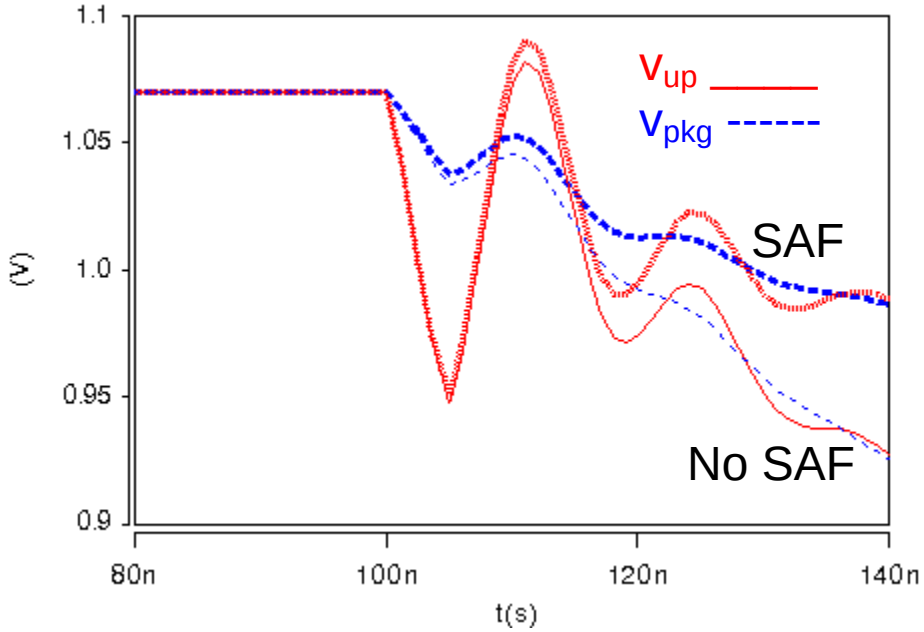

Figure 4.8. Performance improvement with SAF; simulated with linear power path model; magnified view.

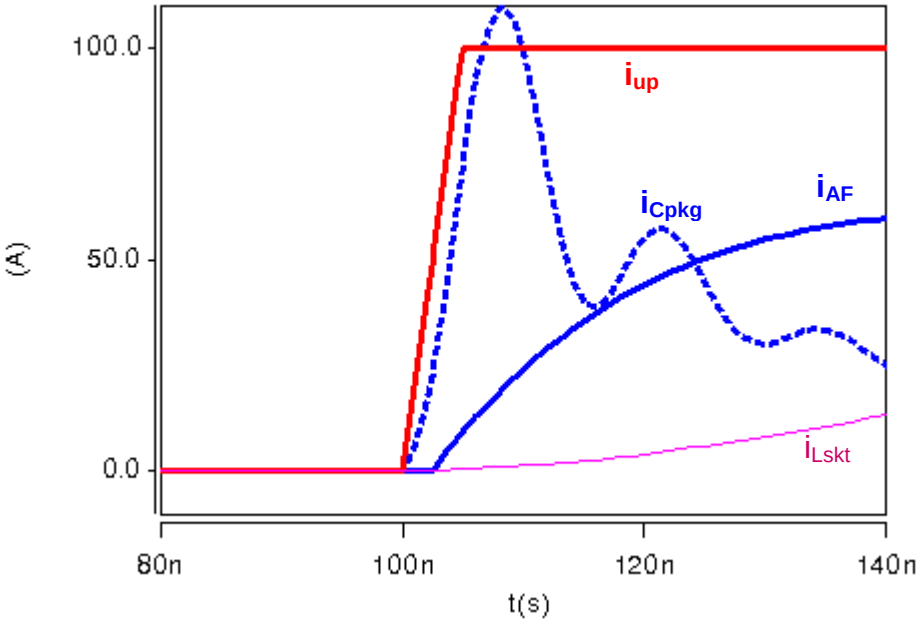

Figure 4.9. Load current relative to currents associated with branches of the package node with SAF operating.

Figure 4.9 highlights the currents at the package node relative to the load current. It can be seen that after a very short delay the package capacitor, C_{pkg}, provides the required di/dt to the load. It overshoots a little due to the package inductance, L_{pkg}. Even though

the di/dt provided by the SAF through i_{AF} is relatively low, it does not matter to the load current which has reached its steady state level by the time the SAF has turned on. The SAF really assists in regulation not by forcing high di/dt but by taking over from C_{pkg} in supplying the current to the load before the voltage across C_{pkg} falls out of specification.

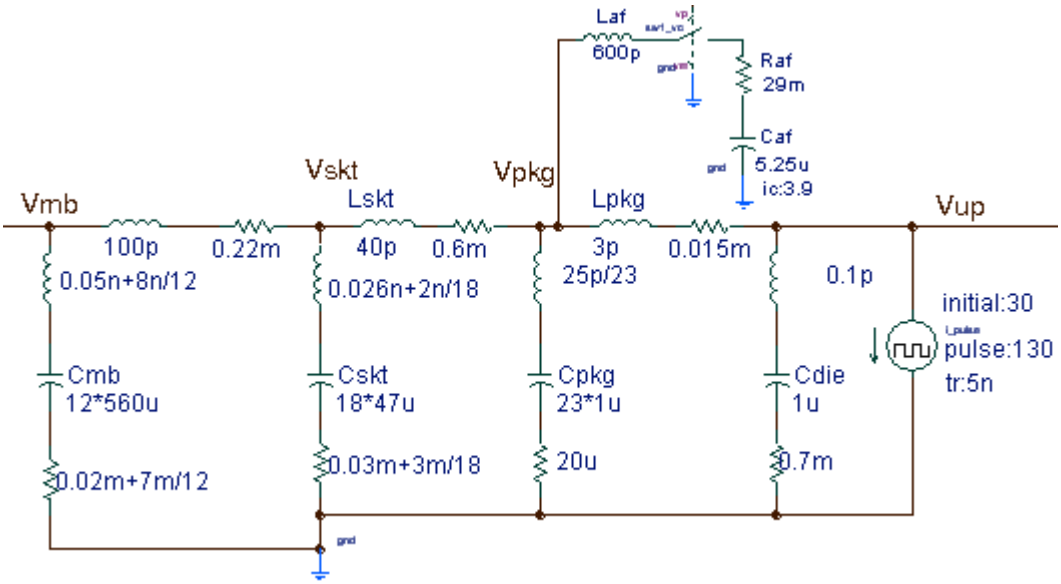

Figure 4.10. Schematic section showing the SAF injecting into package. Vmb is connected to BVR output and Vup is connected to BVR sense.

Replacing the linear model for the BVR with an actual 4 phase switching regulator, obtained through collaboration with CPES at Virginia Tech, the overall hybrid filter performance was simulated in a more realistic setting. Figure 4.10 shows the hybrid filter section of the power delivery system. The line extending out from V_{mb} connects to the output of the switching regulator whereas the line extending from V_{up} goes back as a sense signal to the motherboard. The output voltage (v_{pkg}) response to a 100 A, 20 A/ns load transient with and without the SAF are given in Figure 4.11. Without the SAF, the present-day 2004 power delivery system is inadequate. There is significant ringing on the output voltage which causes almost 100 mV of overshoot and more than 270 mV (270% ideal design specification) of initial undershoot. More than 2 μs settling time is apparent.

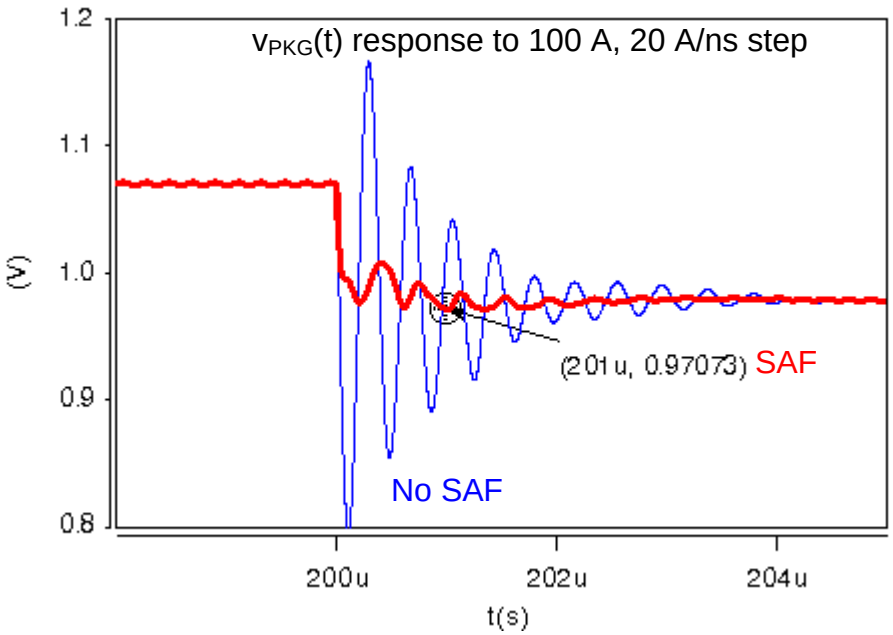

Figure 4.11. Improvement in v_{pkg} using the SAF; simulated with switching regulator.

With the SAF, there is significant improvement. The voltage dips only 100 mV with minimal ringing and less than 1 μs settling time. Voltage waveforms at the die is virtually the same as at the package (Figure 4.12); except, at the die, the additional high frequency ringing is present (Figure 4.13), just as it was in Figure 4.7.

Even when the switch action was delayed by 10 ns and then 20 ns, the degradation in response was relatively small (Figure 4.14). Up to 20 ns of delay could be tolerated for 150 mVpp of voltage deviation.

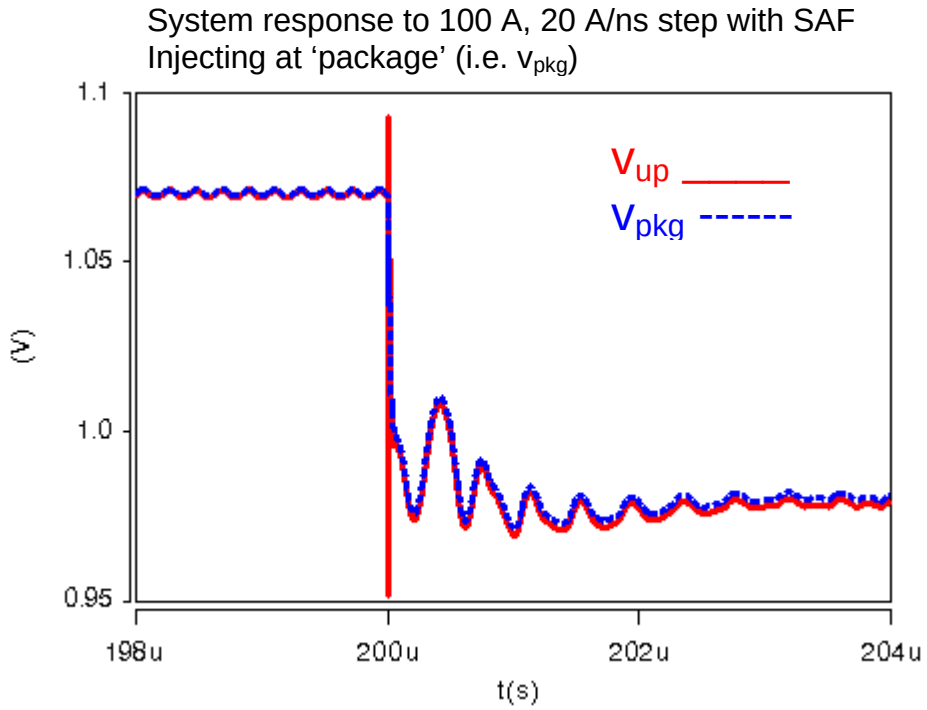

Figure 4.12. Comparison of waveforms at package and die with SAF injecting at the package.

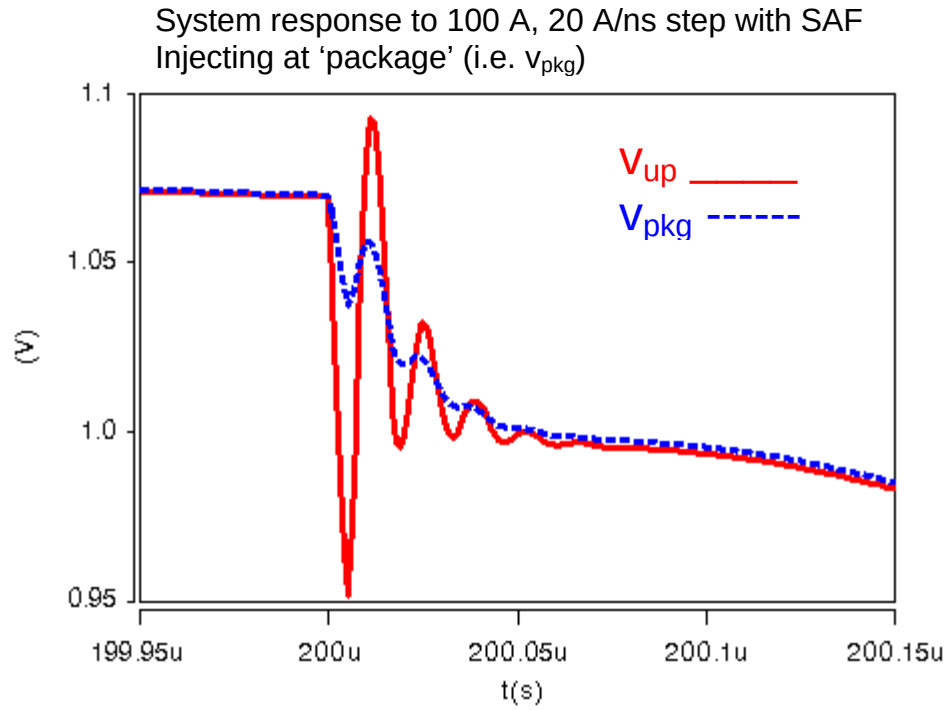

Figure 4.13. Comparison of waveforms at package and die with SAF injecting at the package; magnified view.

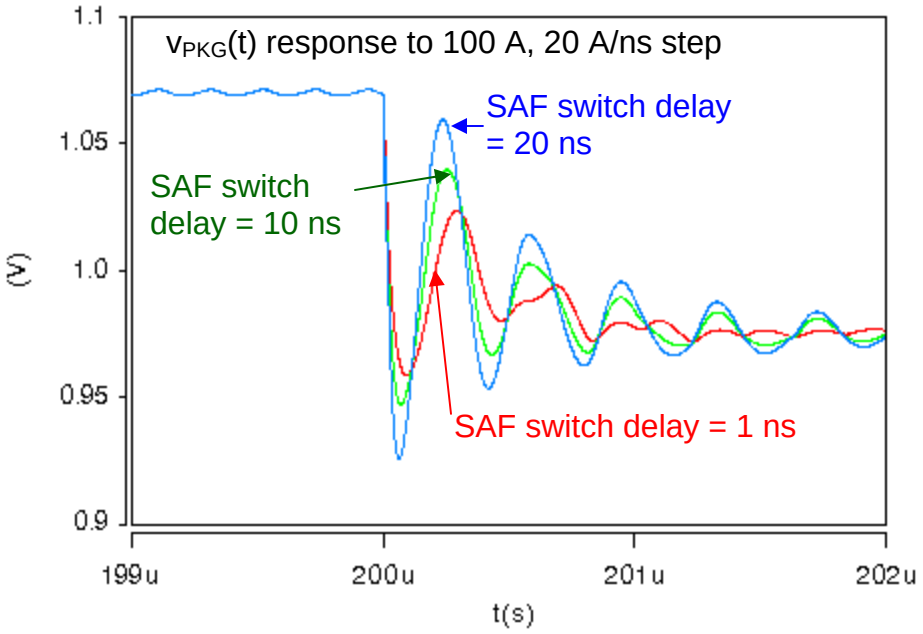

Figure 4.14. Effects of delay in the circuit of Figure 4.10.

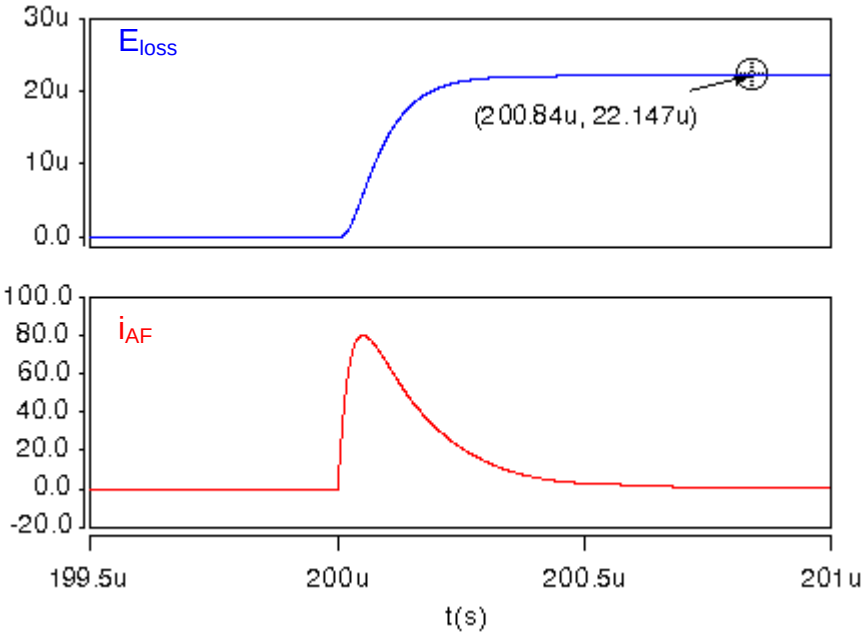

Figure 4.15. SAF current and energy loss in the circuit of Figure 4.10.

The SAF current waveform in Figure 4.15 indicates a peak current of 80 A and an SAF operation time of about 1 μs (~5 $R_{AF}C_{AF}$) but note that most of the energy transfer occurs within 2 $R_{AF}C_{AF}$. The shortest transient cycle, which includes turn-on and turn-off, is therefore about 2 μs. The maximum transient event frequency, assuming there is no

feedback control, is 0.5 MHz. The energy loss per edge is 22 μJ and maximum power loss in the SAF resistor will be about 22 W or roughly 0.4 W/MHzA. If thermal considerations require it, the maximum transient event frequency will need to be reduced or else, the SAF must be redesigned to have lower path inductance, L_{AF} .

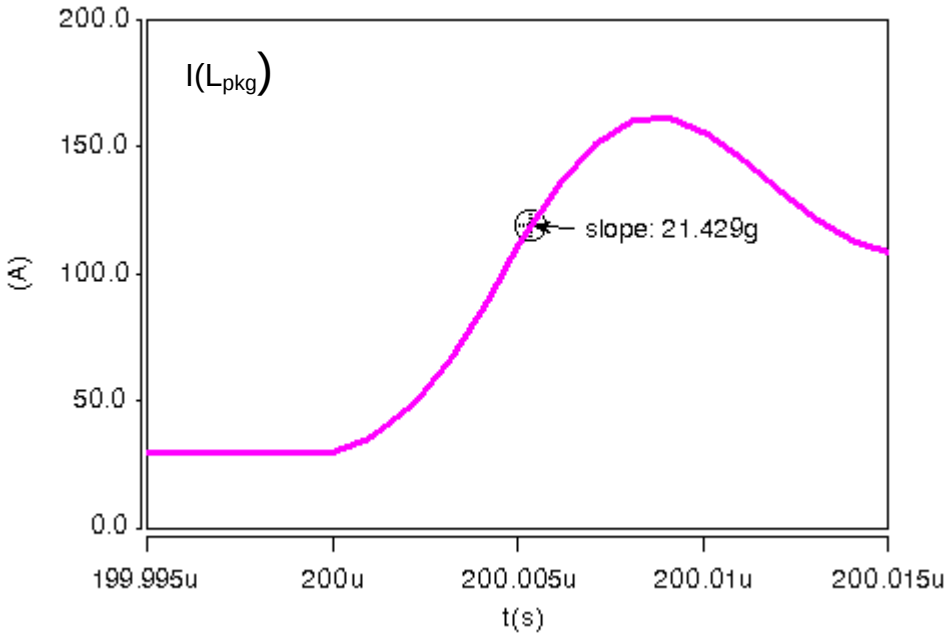

Figure 4.16. Graph of current in L_{pkg}, i_{Lpkg}, for the circuit of Figure 4.10 with 20 A/ns di/dt at the die (load).

Figure 4.16 reveals an interesting feature in the profile of the current drawn into the microprocessor package. With 100 A, 20 A/ns load current (i_{up}), the current (i_{Lpkg}) through the equivalent package inductance may experience an instantaneous di/dt of slightly greater that 20 A/ns when the SAF is injecting into the package. The higher di/dt and current overshoot is consistent with the ringing seen at the output. However, it also implies that, in this example at least, the di/dt at the package can be as high as that at the load. Therefore, without loss of generality, an ideal step current source at the package node, V_{pkg}, may serve as an equivalent load at that node. Once V_{pkg} is found under this condition, V_{up} may be taken as approximately equal to V_{pkg} but with the additional high

frequency ringing seen in Figure 4.13. Thus, only V_{pkg} waveforms are used in the remainder of the chapter.

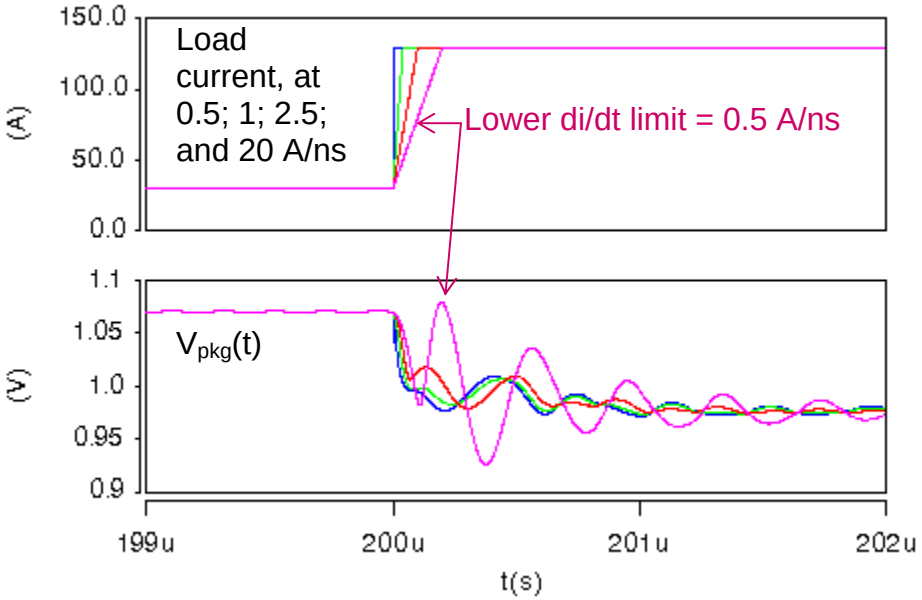

Figure 4.17. Allowable range of di/dt when injecting at package.

Since the SAF is designed for ideal step loads, it works well to maintain regulation for high di/dt's; provided the switch turn-on is coincident with the middle of the load transient edge. Figure 4.17 shows that, as long as there is no unwanted delay and the point-of-load capacitor voltage does not drop out of specification in the time it takes the SAF to turn on, the SAF can handle di/dt as low as 0.5 A/ns (or roughly a risetime equal to one quarter of the ringing period) without requiring any design adjustment. For lower di/dt's, regulation can be maintained by either increasing R_{AF} or decreasing V_{AF} in magnitude via a programmable scheduler or continuous feedback control. As an illustration, in Figure 4.18, with a low di/dt of 0.5 A/ns, the SAF designed for an ideal step load causes overshoot and ringing in the V_{pkg} response. However, V_{pkg} may be brought back within specification by changing R_{AF} from 33 mΩ to 50 mΩ. This is achievable by adjusting the gate voltage on the FET that forms the SAF switch.

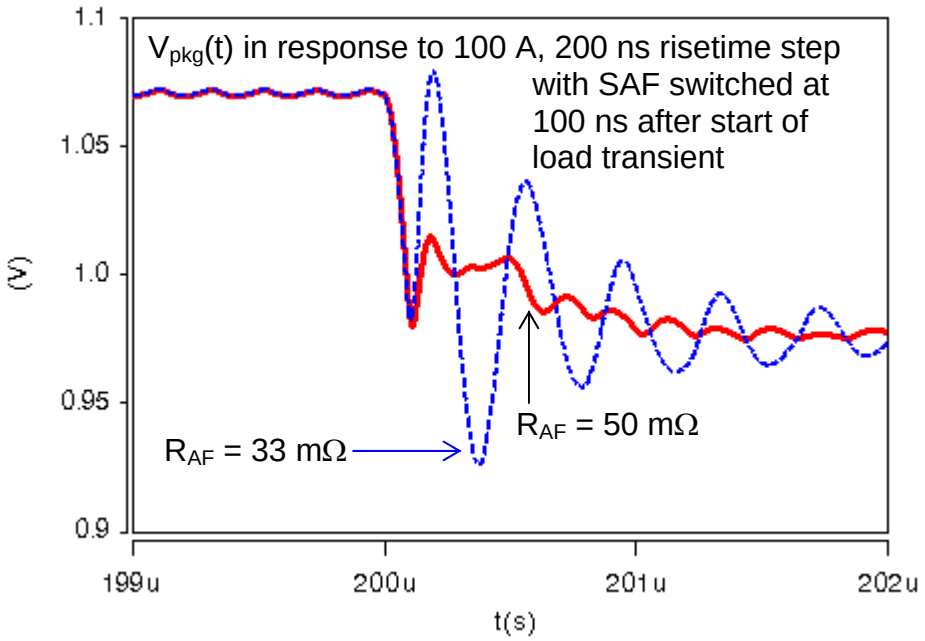

Figure 4.18. Adjustment of R_{AF} to correct for low di/dt.

4.4 SAF Injecting into the Socket

Presently it is inconvenient to mount the SAF on the microprocessor package as it requires extensive cooperation between motherboard designers and the microprocessor manufacturer. As a result, hardware for testing the SAF in microprocessor power supply applications, will most likely be at the motherboard level. In this section, it is shown that the SAF mounted on the motherboard is helpful in transient load regulation- although the improvement is not as much as when the SAF injects directly into the package. Furthermore, it is also shown that the SAF may be used to reduce the amount of filter capacitors in the socket cavity, thereby opening more space for signal routing.

Reducing the number of capacitors that make up C_{pkg} implies the equivalent series inductance or esl in C_{pkg} (L_{Cpkg}) is bound to go up. Interaction between filter stages is now more likely- there is further deviation from the EAVP design guidelines and it is not advisable to attempt obtaining an estimate of the single stage model by inspection using

the all-passive EAVP assumptions mentioned in Appendix A. Instead simulation or numerical curve-fitting techniques must be employed in obtaining the simplified model.

The circuit of Figure 4.19 was used to obtain the impedance of the passive filter when injecting at V_{skt}. The number of 47 μF capacitors in the socket cavity was reduced to 8 instead of 18. Less than 6 capacitors in the cavity resulted in switching regulator instability that appeared as a constant oscillation in output voltage. It is expected that redesigning the regulator compensation network to the reduced capacitor values will improve performance and allow for even further reduction of capacitors in the socket cavity. However, this task has been left as part of the coordinated control problem associated with designing the SAF for this particular application.

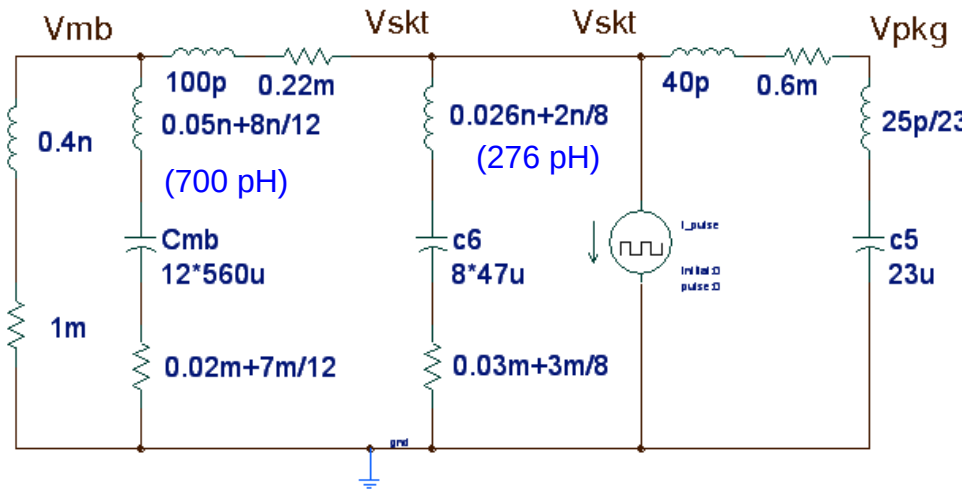

Figure 4.19. Circuit used to find equivalent passive filter with respect to V_{skt}.

Table 4.3. SAF design for injection at the socket.

Step	Description	Equations	Key Parameters	Result
1	Find impedance at injection node		$Z_{out\ n}$ i.e. Z_{skt}	See Figure 4.20
2	Model fit to dominant peak		L, C	180 pH, 23 μF
3	Calculate R, Q	2.2	R, Q	1 mΩ, 1.46
4	Approx. as lossless or damped		Q, R_L, R_C	lossless
5	Choose L_{AF}		L_{AF}	600 pH
6	Derive SAF using appropriate set of design equations	2.5; 2.26; 2.28; 2.29	$C_{AF} = 7\ \mu\text{F}; Q_{AF} = 0.356;$ $R_{AF} = 26\ \text{m}\Omega; V_{AF} = 3.7\ \text{V}$	

The design procedure is concisely summarized in Table 4.3. Using the single-stage model, parameter values were chosen to fit the multistage impedance curves given in Figure 4.20. This was done manually in Saber. The resultant model is shown in Figure 4.21. The value of L based on the curve-fit done with Saber simulation is 180 pH instead of 150 pH estimated by inspection (assuming the inductances to dominate over the frequencies of concern).

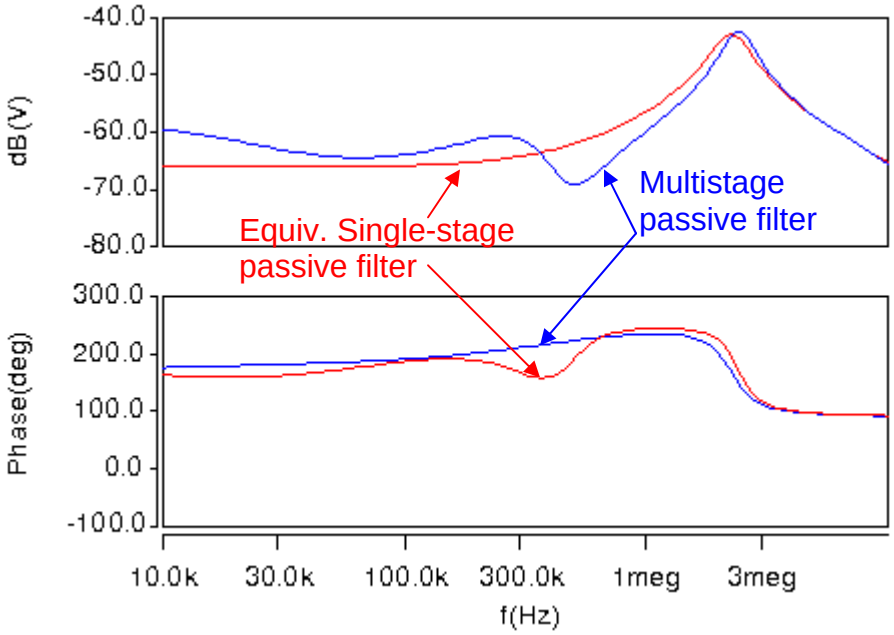

Figure 4.20. Impedance with respect to V_{skt} of the multistage passive filter and its equivalent single stage model.

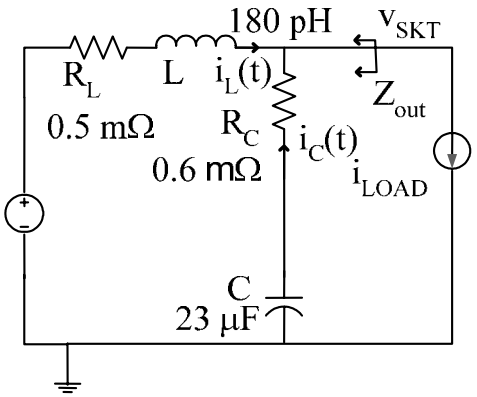

Figure 4.21. Equivalent single stage model for the circuit of Figure 4.19.

Unlike the ideal EAVP filter, the inductances in the series and parallel branches are comparable and function as one equivalent inductance near resonance. As a result, the package capacitor, C_{pkg}, must still be included in the model. In fact, the values of C, R_C and L_C are formed by C_{pkg} and associated interconnect parasitics. The 2.5 MHz resonance is associated with the 180 pH and 23 μF components. Essentially, the 40 pH socket inductance is treated like a short in the design. However, as will be seen, it does degrade performance.

Since $R_L < R_C < R$, the model was approximated as lossless although, with $Q = 1.46$, both sets of design equations may be equally valid.

Keeping $L_{AF} = 600\,\text{pH}$, the rest of the SAF design was derived and is listed in Table 4.3.

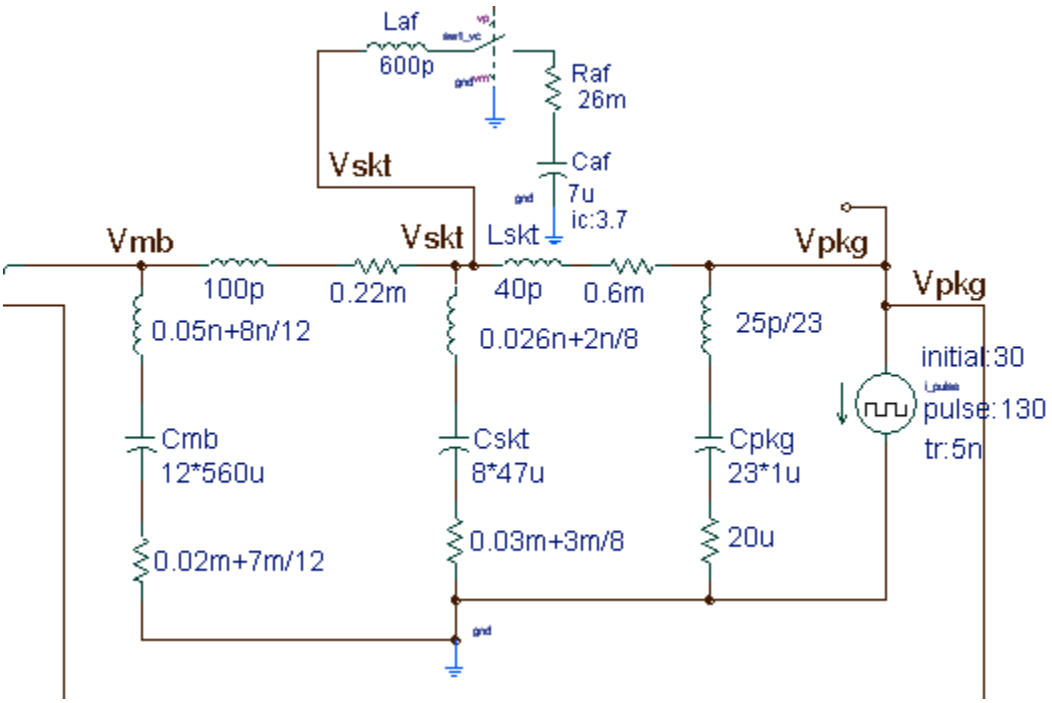

Figure 4.22. Schematic section showing an SAF used for injection at the socket.

The values of the SAF injecting into the socket given in Figure 4.22, are similar to those in the SAF injecting into the package.

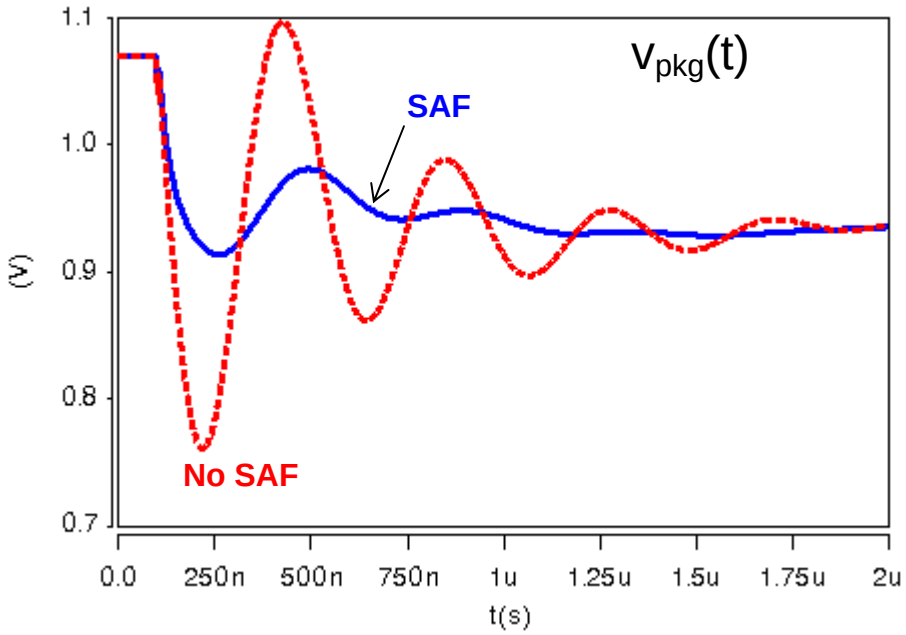

Figure 4.23. Performance of the SAF used for injection at the socket; simulated with 2004 linear model.

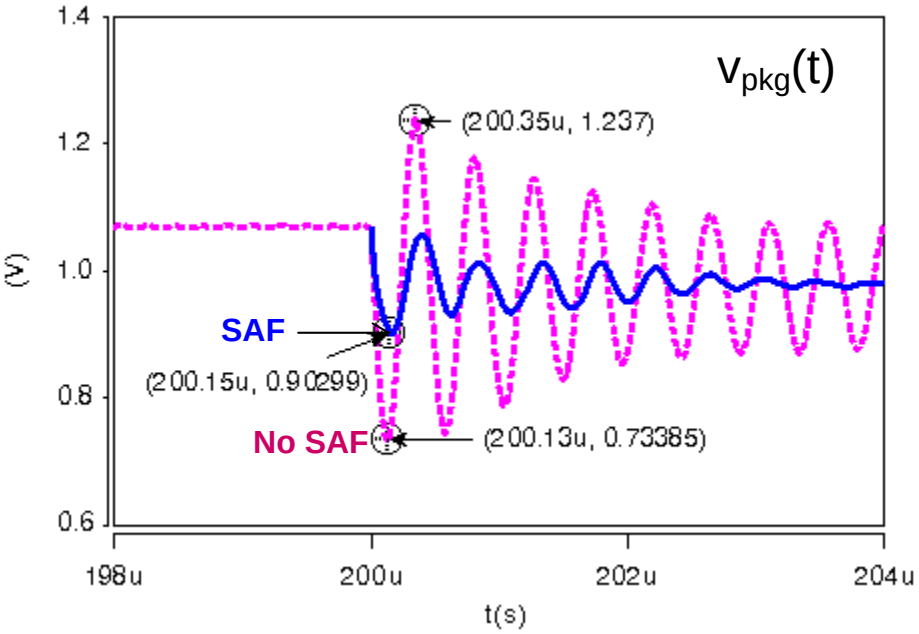

Figure 4.24. Performance of the SAF used for injection at the socket; simulated with switching regulator.

Simulating the hybrid filter using the lumped element, linear model of the power delivery path resulted in almost a factor of two reduction in output voltage deviation when compared to the case with no SAF (Figure 4.23).

When simulations were carried out using a realistic switching regulator as the BVR, there was noticeably more ringing. With the SAF, the output voltage, again measured at the package, shows 170 mVpp ripple and 4 μ s settling time (Figure 4.24). This reflected almost a factor of three reduction in the peak to peak voltage deviation. Note the switching noise is slightly more as well due to the reduction in socket capacitance. The fact that the decay in the ringing appears uneven indicates the compensation in the switching regulator should be adjusted if possible.

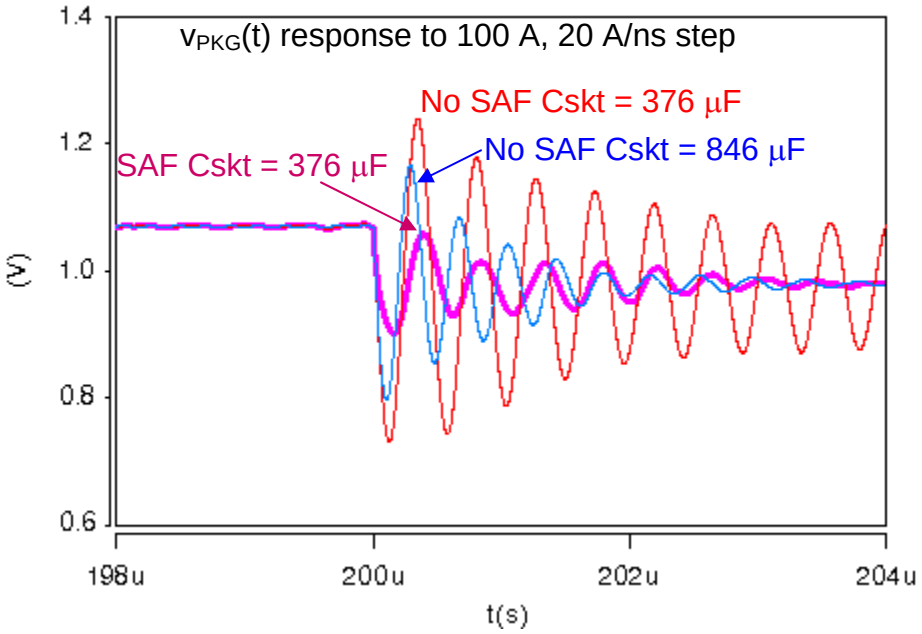

Figure 4.25 Output voltage, v_{PKG}, with SAF used for injection at the socket.

Figure 4.25 shows the hybrid filter performance in context with the passive filter alone. The peak to peak ripple values are listed in Table 4.4 for the three cases depicted in Figure 4.25. The SAF, as part of the hybrid filter, achieves significantly better regulation than the passive filter alone- 170 mVpp instead of 500 mVpp. Even if an

attempt is made to place the maximum amount of capacitance that would fit within the socket area, without regard for keep-away zones, the resulting passive filter by itself would still have too much ringing. For the LGA775 package, the maximum value of Cskt given present-day component technology is about 846 μF. Without the SAF, the passive filter with the maximum value of 846 μF in the socket would still oscillate up to 270 mVpp in response to the step load (specifically, 100 A, 20 A/ns).

Table 4.4. Voltage deviation for the conditions given in Figure 4.25.

Condition	SAF C_{skt} = 376 μF	No SAF C_{skt} = 376 μF	No SAF C_{skt} = 846 μF
Peak-peak ripple	170 mVpp	500 mVpp	270 mVpp

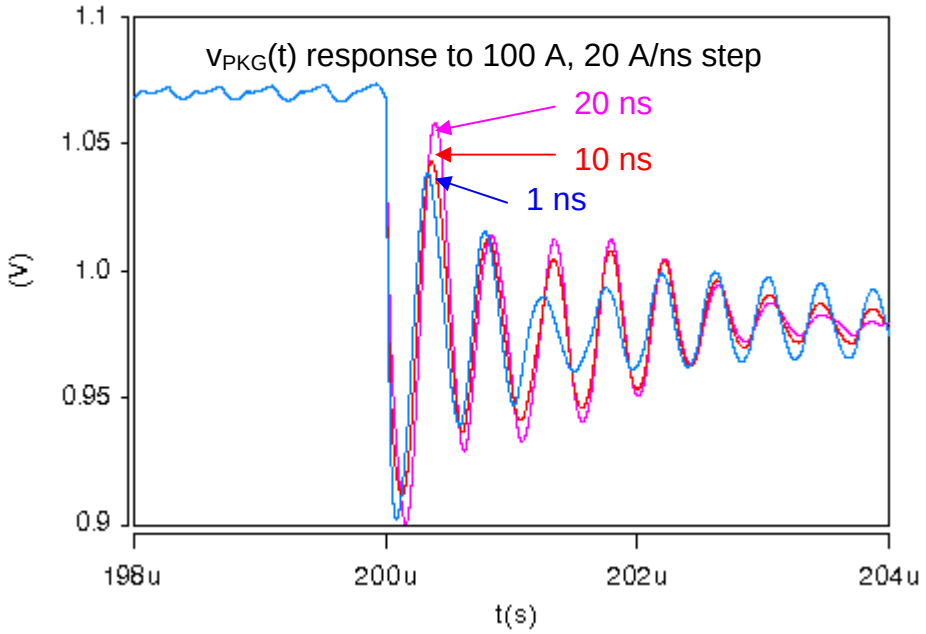

Figure 4.26. Output voltage, v_{PKG}, with 1ns, 10 ns and 20 ns SAF switch delays.

As seen in Figure 4.26, delays of up to 20 ns can be tolerated at 170 mVpp ripple. There was no trend to the rate of ripple decay with respect to delay. The operation time of the SAF was about 1 μs and the corresponding energy loss was 25 μJ (Figure 4.27).

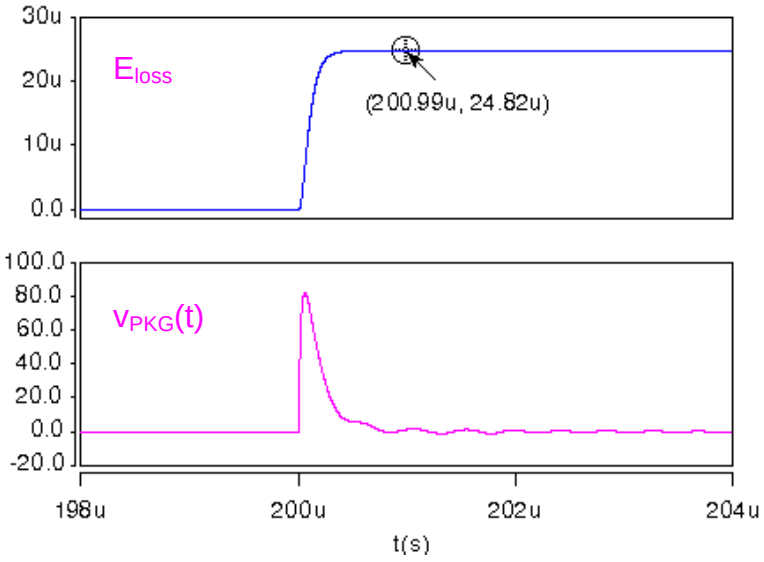

Figure 4.27. SAF current waveform and energy loss for the circuit in Figure 4.22.

4.5 Applying the SAF to Multiple Stages

A single SAF is designed to correct for a single resonant peak in impedance. If there are multiple peaks in the output impedance curve of a multistage passive filter, more SAF's may be applied- one to handle each resonant peak- to improve load regulation.

A two-stage passive filter based on the 2004 power path model was used to test the application of multiple SAF circuits injecting at successive nodes of the passive filter. The circuit diagram is given in Figure 4.28. Once again, the BVR linear model resistance, R2, was reduced to keep the total DC path resistance at 1 $\mathrm{m}\Omega$; and the bulk capacitor was removed. To obtain an EAVP response, the capacitance needed across V_2 to ground is $L_2/R_2^2 = 2.5\,\mathrm{mF}$, which would require much more space than is available in the socket cavity. The choice of 376 μF stemmed from the fact that the corresponding esr R_{C2} = 0.4 $\mathrm{m}\Omega$, was consistent with the EAVP requirement. Plus, this value was obtained using just

8 components; which is in keeping with the desire to reduce the number of capacitors in the socket cavity.

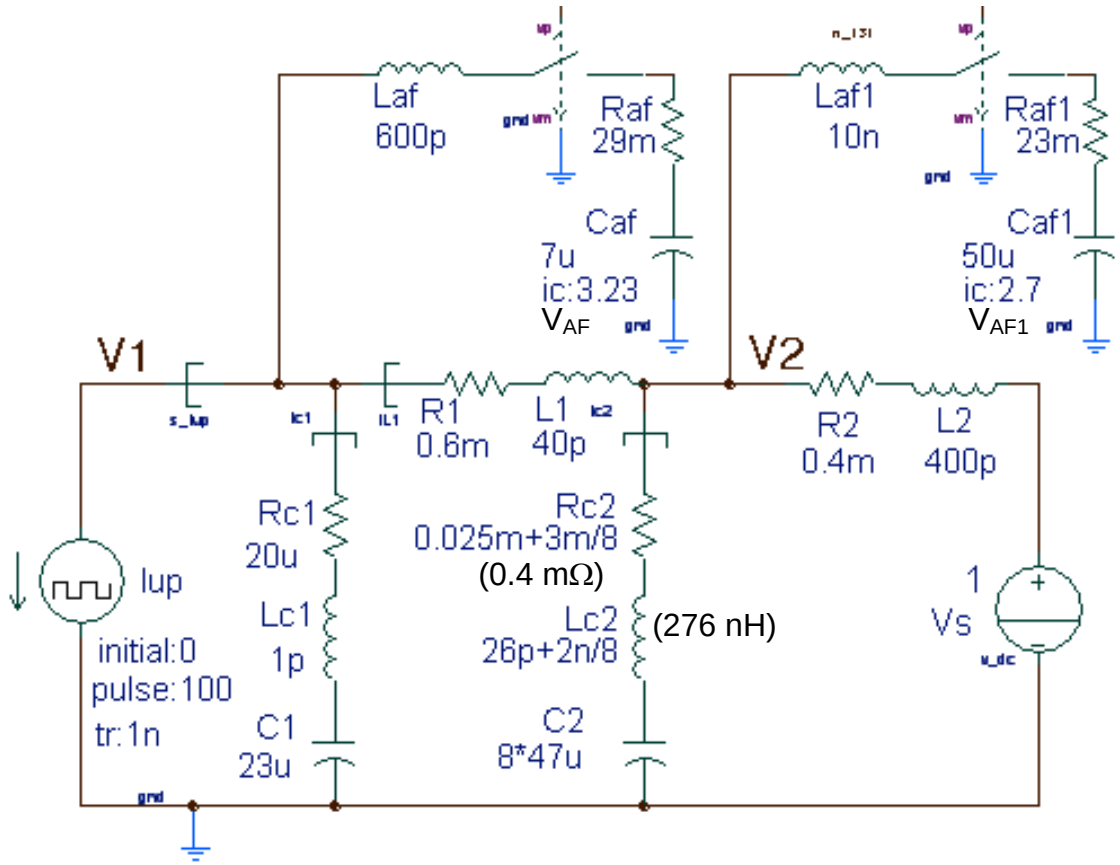

Figure 4.28. Two-stage hybrid filter with an SAF on each stage.

Table 4.5. Multistage SAF design parameters for Figure 4.28

	L_{AF} [nH]	C_{AF} [μF]	R_{AF} [mΩ]	V_{AF} [V]
SAF injecting at V1	0.6	7	29	3.23
SAF injecting at V2	10	50	23	2.7

At high frequency, looking into V1 towards the source, impedance appears as that of a 200 nH inductor (L_1+L_2//L_{C2}) which implies $L = 200$ nH. This is about the same value used in the SAF design above, in section 4.3. Using $L_{AF} = 600$ nH as above, the SAF designed for injection at V1 was almost the same as the design derived above. The SAF designed for injection at V2, used $L_{AF} = 10$ nH, a practical value that enabled the $C_{AF} = 50$ μF, a relatively low value. The two SAFs' parameters are given in Table 4.5.

The values for R_{AF1} and V_{AF1} were comparable to those used in the SAF designed for injection at V1.

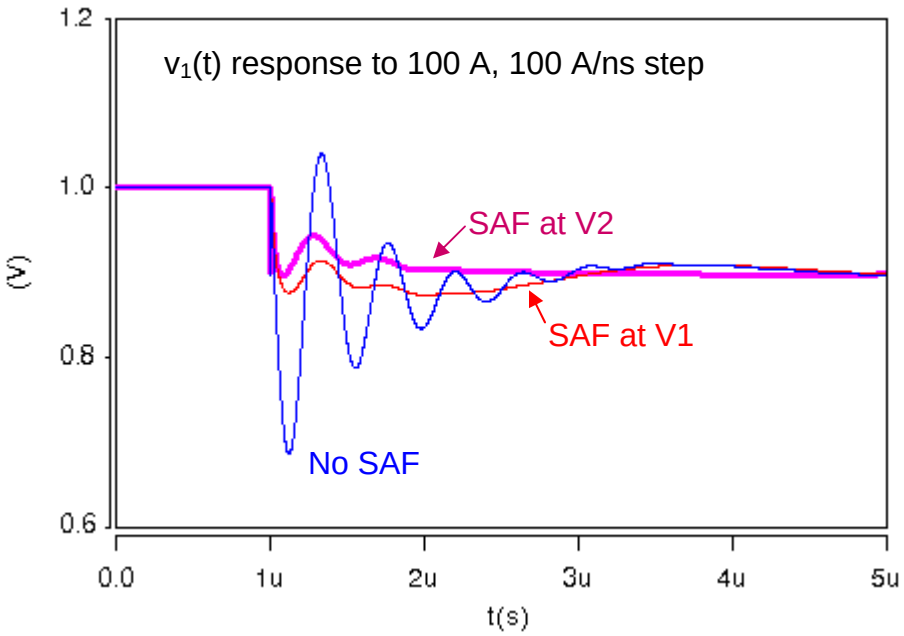

Figure 4.29. Performance improvement using two-stage hybrid filter.

Simulations showed that both SAF's switched on at the same time (t = 1000.5 ns) operated independently to correct the high and the low frequency ringing associated with the first and second stages respectively (Figure 4.29). The SAF injecting at V1, when switched alone, corrected the high frequency ringing, but the low frequency ripple was still seen.

4.6 Alternative Injection Points and Relative di/dt's

The previous section illustrated the use of multiple SAF's at multiple injection nodes, with the high-speed SAF injecting right at the point-of-load. The question naturally arises as to whether the high-speed SAF may be injected at another node further from the load while still assisting in load regulation. Following the same design methodology, but injecting both SAF's at the node V2, an alternative hybrid filter design was obtained (Figure 4.30).

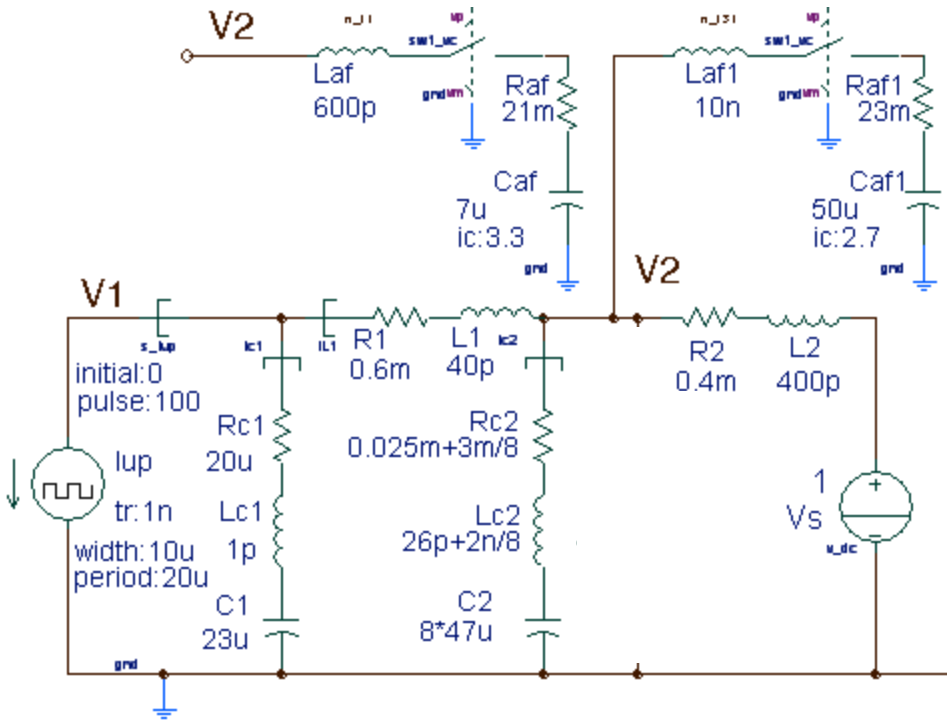

Figure 4.30. Alternative hybrid filter design to that in Figure 4.28.

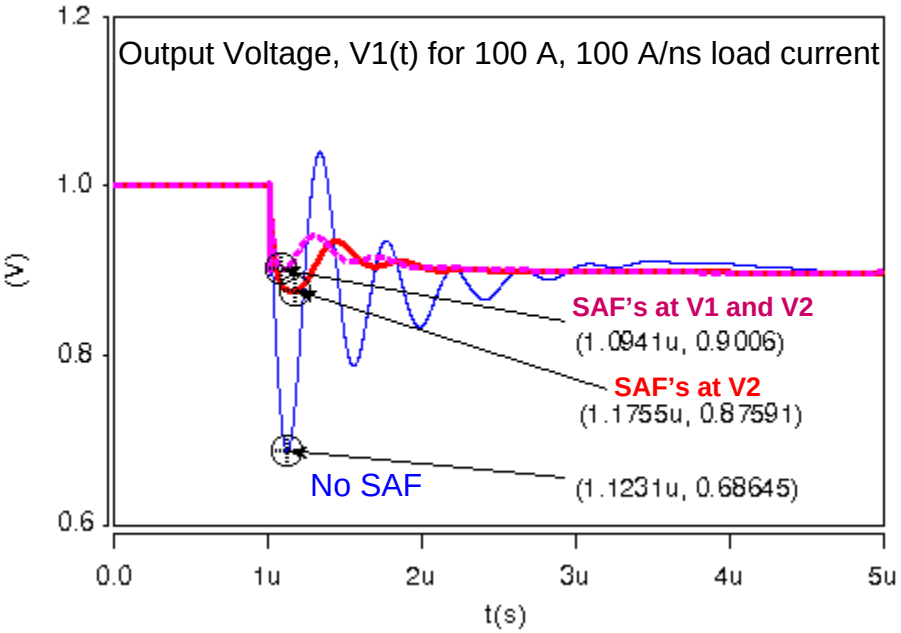

Figure 4.31. Relative performance of hybrid filter design given in Figure 4.30.

In shifting the high-speed SAF injection point from V1 to V2, the value of R_{AF} decreases from 29 mΩ to 21 mΩ while V_{AF} increases slightly from 3.2 V to 3.3 V. The performance of this hybrid filter in response to a 100 A, 100 A/ns step in load current

is given in Figure 4.31. Instead of 100 mVpp deviation, there is 125 mVpp deviation at the output,V1 when injecting solely at V2. Nonetheless, this performance is still significantly better than the passive filter alone.

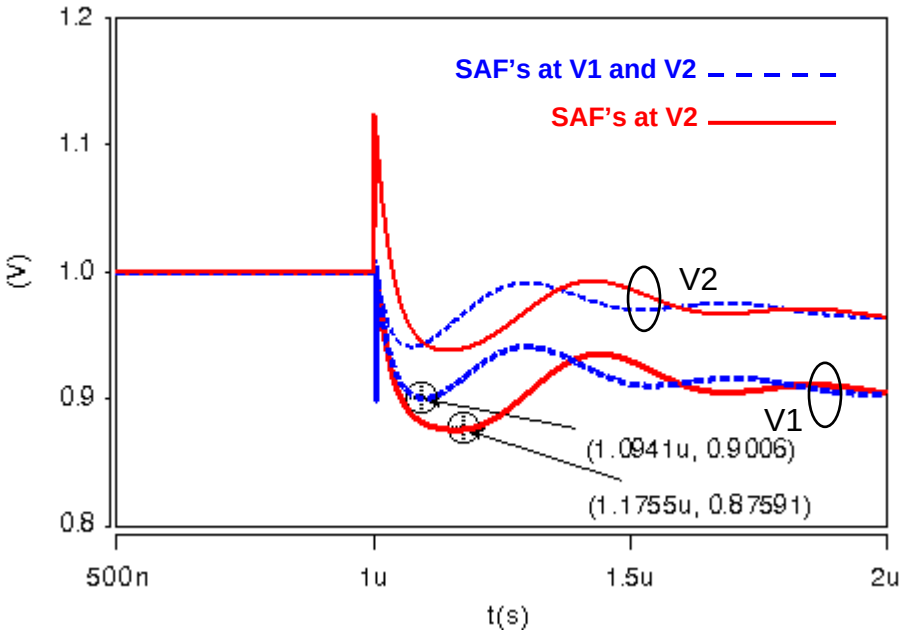

Figure 4.32. Comparison of node voltages.

Since the high-speed SAF injecting into V2 must overcome a relatively large L1 in order to assist C1 in maintaining the load voltage, there is necessarily an overshoot spike in V2 in order to maintain regulation at V1 (Figure 4.32). This phenomenon is absent if using the design of Figure 4.28 where there is an SAF injecting at V1 and V2 each.

Figure 4.33 shows the relative di/dt's of the key currents in the hybrid filter of Figure 4.30 while it operates to maintain regulation in response to a 100 A, 100A/ns step load. Initially, this full load di/dt is supplied by the point-of-load capacitor, C1. Injecting at V2, the voltage V2(t) spikes up to allow current to build up in L1 at an initial di/dt of 3 A/ns to quickly allow the SAF to take over from C1 and maintain regulation at V1. As the SAF current, $i_{AF}(t)$ dies down, the current in L1 is maintained via the lower frequency capacitors and the DC source. Thus, it can be seen that the SAF may be injected at a

lower frequency node, forcing a lower di/dt, while still aiding in load regulation at a high di/dt node closer to the load. Further analysis of the potential range (i.e. the design space) of such SAF designs is an interesting topic for future work.

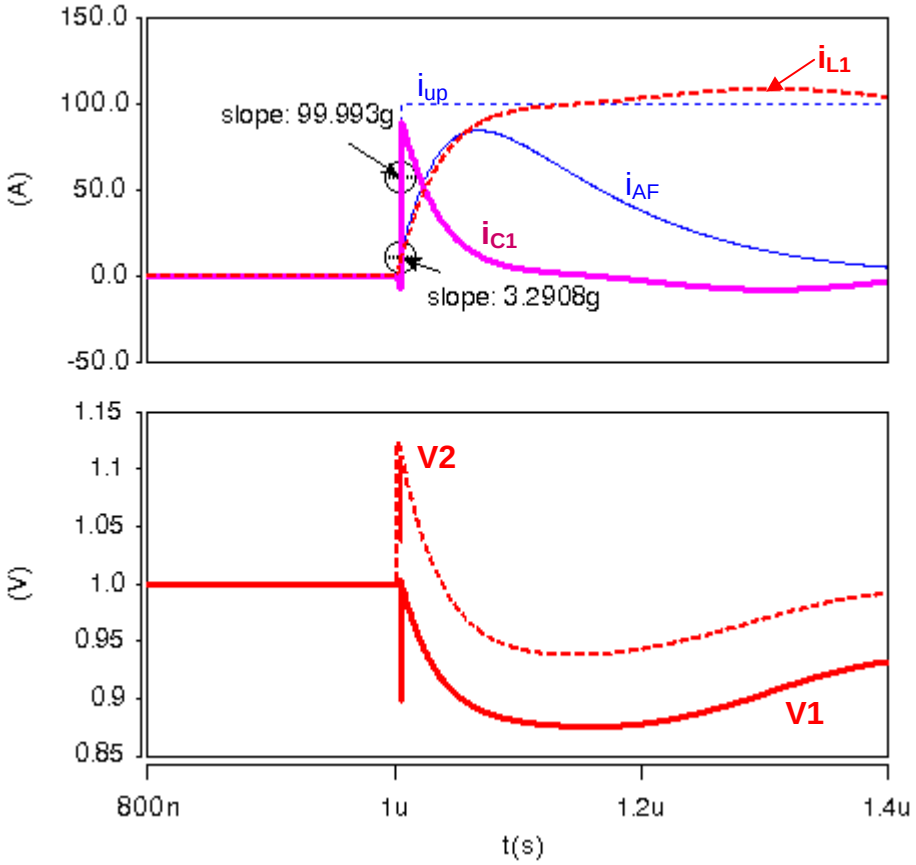

Figure 4.33. Current and voltage waveforms showing relative di/dt's with respect to load regulation performance.

4.7 Summary of SAF Application to Microprocessor Power Delivery Systems

The results described in this chapter do indicate the potential for the SAF to work in a switched regulator system to reduce ringing on sharp load transients. There is flexibility in the location and injection points for the SAF, although injection nearer the load (where the filter stages are easier to separate and analyze) is preferred. Even without strict control of the parasitic esl and esr values, the SAF is useful. In deriving the simplified model when stages interact, some tweaking or numerical optimization may

enhance performance. Still, this may not be worthwhile if the parasitics cannot be tightly controlled in the physical design. Simply using the key model parameters, L and C, together with ΔV, chosen in advance to allow some additional "head-room" or margin for error, yields useful designs. When esr's cannot be controlled, the design obtained using the EAVP SAF equations tends to reduce overshoot while that obtained from the non-EAVP equations tend to reduce the magnitude of undershoot, on a pull-up transient. Coordinated design of the control and compensation schemes is necessary.

4.8 Extension of EAVP to a Transmission Line Solution

As actual load current slew rates or di/dt's go up, esl of the capacitor, no matter how small numerically, becomes dominant. Figure 4.34 depicts one possible method of slowing down the signal at a given node so that the capacitor esl is again made insignificant thereby allowing the SAF to be used at that node.

By designing the interconnect between the node of concern and the load to be a transmission-line that has ultra-low characteristic impedance equal to the desired output resistance with acceptable losses, even an ideal step at the load end appears with a delayed dispersed edge at the SAF node. In this way, extremely wide bandwidth power delivery systems may be designed.

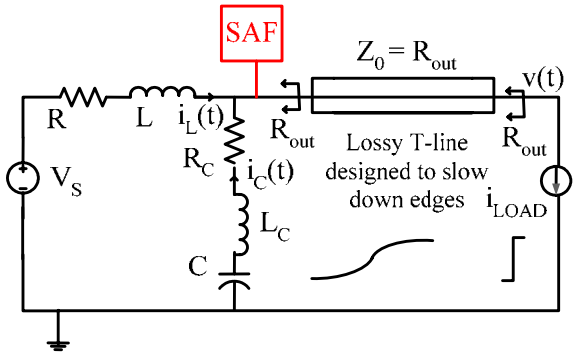

Figure 4.34. Extension of EAVP to transmission line.

4.9 Further Applications of the SAF

4.9.1 Switched Coil Damping

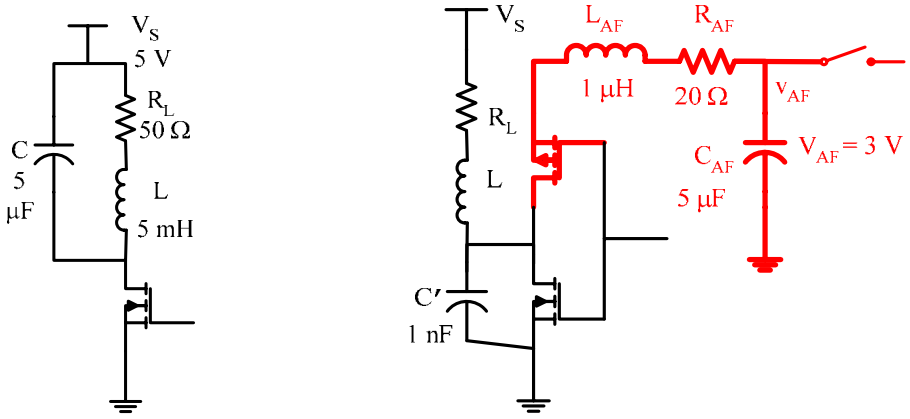

(a) (b)

Figure 4.35. Damping a switched coil: (a) using single capacitor and (b) an SAF.

The circuit in Figure 4.35(a) is a good model for typical coil-driving circuits for relay and (neglecting back-EMF) motor applications. The coil is modeled by inductance L and winding resistance R_L. The capacitance C is used to absorb the coil energy and prevent switch voltage over-shoot when the switch is turned off. Sometimes a "free-wheeling" diode is used in place of, or in parallel to C. However, in certain cases, it may be difficult to find a fast-switching diode. Regardless, passive methods of damping or clamping are lossy. Only the case shown in Figure 4.35(a) will be considered here. The loss associated with the charging and discharging of C is given by $E_{Closs} = CV_S^2$ which works out to be 125 μ J for the numerical example shown in Figure 4.35(a).

In principle, C may be switched into the circuit just when the driver switch is turned off, allowed to absorb energy until the coil current is exactly zero and then turned off. This "resonant" method affords the lowest loss but would require at least another switch and current sensing circuits, in addition to the capacitor, near the coil. The SAF may be

used as a mid-range option that allows the capacitor and switch to be placed away from the coil and timing circuits to be simplified while still affording reasonably low-loss.

Removing the original damping capacitance, leaving only parasitic capacitance C', a design of an SAF was obtained Figure 4.35(b). For this SAF design, the switch may be a PMOS device and easily driven by the coil drive signal. No additional timing circuit is needed to turn off the SAF as it naturally settles without overshoot. The PMOS switch, L_{AF} , R_{AF}, and C_{AF} may be placed far away from the coil itself as L_{AF} is large and easily formed from any combination of trace inductance and miniature ferrite bead inductors. The SAF performance is shown in Figure 4.36.

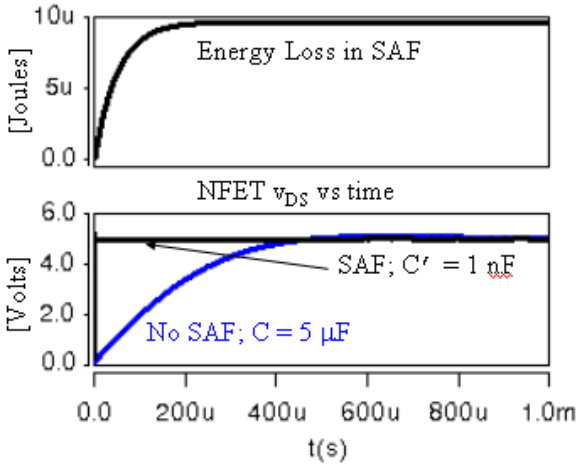

Figure 4.36. Effects of SAF on voltage across switch and circuit losses.

For the same capacitance and the same overshoot, the SAF incurs only 10 μJ of loss compared to 125 μJ lost in the case of purely passive damping, while 40 μJ is available to be recycled using resonant converter techniques. Although most applications today do not need the SAF, future designers may take advantage of the combined relaxed spatial constraints and energy savings of the SAF to work with larger coils or a large number of coils.

4.9.2 Retrofit Filters

Silicon IC technology achieves relatively low capacitance per unit area (about 1 μF/cm^2 [Bla04]). Dense discrete PCB designs may have "keep-away" zones that limit placement of capacitors at the load-point. In cases such as these, the SAF may be used to compensate for insufficient filter capacitance at the point-of-load, provided there is an access path to the load.

Standard manufacturability guidelines prevent any other components sharing area with through-hole parts such as the SMA connectors used in the experiments. Furthermore, depending on mounting fixtures and other mechanical considerations, large sections of PCB surface area are often forced to be free of electrical components. Fortunately, it is also a common manufacturability guideline to provide at least one accessible test point to every key node in the circuit. Therefore an SAF may be safely mounted some distance from the load but connected to the load test-point to provide the necessary aid in regulation.

4.9.3 Auxiliary Power Regulation to Complement EMI Filters

For circuit protection and switching noise reduction, EMI filters which may consist of large inductors are often necessary. As a simple example, ferrite bead inductors are often placed in series with the power supply path. Unlike the examples using the microprocessor power delivery path, these inductances are necessary. The SAF may be used as a method for providing regulation for sharp load transients which demand current at a rate that is prohibited by the EMI filter.

CHAPTER 5
CONCLUSION

The main body of this work detailed the design oriented analysis, experimental verification and potential application of a switched active and passive hybrid filter that provides transient voltage regulation.

Taking a design-oriented approach, the switched-active filter (SAF) design equations were derived based on a single stage, 2^{nd} order model of the passive filter. One set of design equations were derived for EAVP assuming resistances in the circuit could be set precisely if necessary- a best-case scenario. Another set of design equations were derived assuming absolutely no resistive damping in the passive filter model- a worst case scenario. It was then found that for almost every case where the SAF would be used (i.e. cases with high-Q passive filters where there is a lot of ringing), the design equations gave roughly the same solutions. It was therefore concluded, without having to do extremely tedious analysis, that the SAF design equations would also be useful in assisting a wide range of passive filter designs over a wide range of component tolerance.

This realization was borne out extensively by experiments and simulations for both single-stage and multi-stage filters. Experiments verified hybrid filter performance where the SAF section was designed based on passive filters that were high-Q (>3) and moderate-Q (0.8 to 1.6). Step-up as well as step-down performance was verified. Even with large tolerance on components, and deviations from ideal, SAF designs based on the simplifying assumptions proved to be useful. For more than 20% variations on multiple components, hybrid filter performance was shown to be close to ideal- less than 20%

degradation in performance. In fact the “ideal” design equations themselves rendered qualitative insight into adjusting the design to overcome some of the non-ideal effects, such as imbalance in esr’s and error in component values. In the experiments, it was never necessary to attempt tweaking the esr of the capacitors to obtain good hybrid filter performance. It was concluded therefore, that unlike design for an all-passive EAVP, which calls for precise control of resistances (and other components), the use of the SAF alleviates the demands placed on component tolerance. In one sense, this can be considered an effect of “over-designing”. Even so, the fact that the experiments show that the SAF used in the hybrid filter can be over-designed is itself a desirable freedom. All-passive EAVP filters are difficult to implement at nominal specifications. Having the flexibility on the SAF portion of the design, allows the hybrid filter design methodology to be useful. Once the best efforts at doing an all-passive EAVP design have been spent, the designer can expect an easier task of doing a practical design for an SAF to go along with the imperfect EAVP passive filter.

Simulations supported the application of the SAF to microprocessor power supplies. Over the sections that can be adequately modeled as a lumped element passive filter, the SAF was shown to augment performance in a variety of ways. The SAF, connected to the passive filter near the load was able to form a good hybrid filter. With the SAF injecting at a stage further away from the load, the equivalent hybrid filter still performed much better than the passive filter alone, despite additional deviations from the ideal single stage model due to parasitics. Multiple SAF circuits were able to operate at the same time to correct resonance associated with multiple stages. Simulations showed that multiple SAF’s could be placed at a single node as well as at multiple nodes.

Furthermore, it was shown that an SAF injecting a relatively low di/dt (3 A/ns) at a node far away (40 pH separation) from a high di/dt (100 A/ns) load could perform almost as well as when injecting right at the point-of-load (125 mV vs 100 mV). Even when the passive filter stages, far from the ideal EAVP design assumptions, interacted with each other, an SAF based on an approximate single-stage equivalent was able to improve load regulation. In no test case where there was proper timing of the control signals did the hybrid filter perform worse than the passive alone.

Indeed, the design, simulations and experiments revealed the importance of proper timing or control. While shown to be feasible in simulation, feedback control of the hybrid filter poses a few challenges. The most critical issue in designing the feedback control is one that plagues all active filtering methods- component delays. In the experiments it was shown that the delays were on the order of tens of nanoseconds. A step of 100 A drawn unchecked from 1 μF over 10 ns leads to 100 mV voltage drop. Advances in high-speed, high-current component technology are needed. Compounded with the knowledge that esl in the point-of-load capacitor cannot be completely overcome by active filter methods, in the end, a return to an all-passive solution via an extension of EAVP design to transmission lines may be required.

For lower speed specifications, a detailed study of compensation schemes for the hybrid filter is required before a control design methodology can be developed. Fortunately, much of the research into controlling Linear Active Filters may be relevant to the SAF and hybrid filter discussed in this document due to the similarities in topology and function.

The issue of power loss is related to pre-charging which in turn places demands on the control. The SAF is inherently less lossy than linear regulators and linear active filters due to the capacitor voltage variation on discharge. However, these savings go away if simple resistive charging networks are used (as they were in this research). Resonant converter design equations have been presented to facilitate the eventual design of efficient pre-charging circuits. The precise timing involved in these converter techniques however does place stringent demands on the associated control. For the foreseeable future, designers may prefer to use resistive pre-charging while taking advantage of the fact that the additional losses do not need to be dissipated near the load or any other hot-point on the circuit.

Finally, further applications of the SAF have been suggested that do not share the design concerns specific to microprocessor power supplies. Most notably, the SAF may find direct application to switched coil damping in a wide range of applications including electric machines, and relays.

Given the elegance of the underlying theory, the striking improvement in step-load regulation observed in simulations and experiments and the potential for a wide range of application, the SAF (or altogether, the hybrid filter) is believed to be a valuable contribution to the available methods for load regulation and resonant circuit damping.

APPENDIX A
ACHIEVING EAVP WITHOUT AN ACTIVE FILTER

A.1 AVP and EAVP

Adaptive Voltage Positioning or AVP is the common term given to a special type of dynamic load regulation that allows for a small static output voltage shift proportional to the static load current [Red98]. AVP takes full advantage of the voltage tolerance window to reduce the effects of inductive spikes as depicted in Figure A.1. It also provides size and cost advantages as it uses fewer capacitors with larger esr than would be required if the DC output were rigidly fixed [Yao04],[Red98].

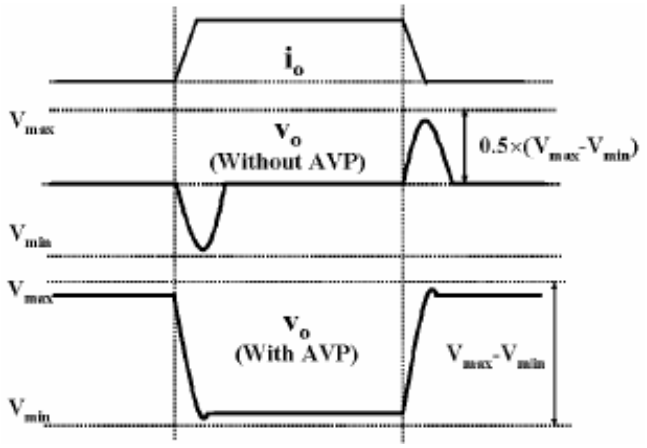

Figure A.1. Transient response with and without AVP Source: [Yao04].

Designing the BVR for AVP, has been addressed [Red98], [Yao03] and is fundamentally a matter of designing the control feedback to get the BVR output impedance to be equal to R_{droop} from DC to about 10-100 kHz, depending on the loop gain bandwidth. Beyond the high corner frequency of the BVR, the passive bypass filter

takes over. The process of designing successive passive bypass stages to achieve AVP overall is termed Extended AVP or EAVP [Wai01].

The condition for AVP can be expressed as

$$Z_0 = R \tag{A.1}$$

where Z_0 is the output impedance and R is a small constant resistance.

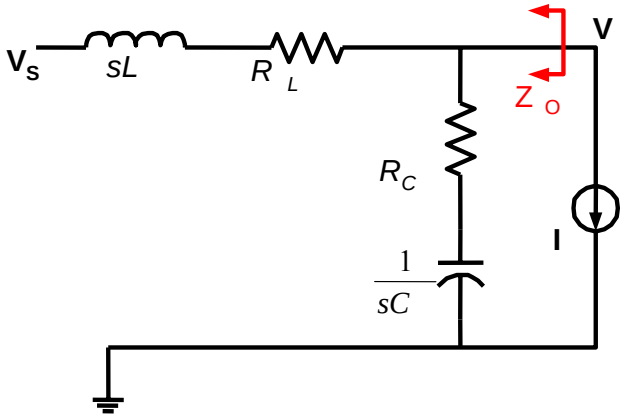

Figure A.2. Single stage model.

Given the LC model shown in Figure A.2, one may obtain

$$Z_0 = \frac{(R_L + sL)(R_C + 1/sC)}{R_L + sL + R_C + 1/sC} = R_L \frac{1 + s(R_C C + L/R_L) + s^2 LC R_C/R_L}{1 + s(R_L + R_C)C + s^2 LC} = R \tag{A.2}$$

which holds true only when

$$R_L = R_C = R \tag{A.3}$$

and

$$L = R^2 C \text{ or } C = L/R^2 \tag{A.4}$$

These are the key design criteria stated in [Wai01].

Before even attempting to design the SAF, it must first be determined if an SAF is required at all. All-passive, filter networks can be designed to achieve EAVP. The next

section summarizes the EAVP passive filter design methodology that first appeared in [Wai01], was further developed in [Man04] and was applied in [Par05].

A.2 Passive Filter Design Equations

The equations for designing a single bypass filter or decoupling stage, together with the component relations that allow for multiple stages to be easily cascaded are given in Figure A.3. For flat impedance the equivalent series path resistance, R_{Tn} of each stage should equal the esr, R_{Cn}, of the bypass capacitor, C_n. The time constant associated with the series interconnect path should also equal the time constant of the parallel bypass path. For successive stages the corner frequency set by the esl of the bypass capacitor must be progressively higher closer to the load.

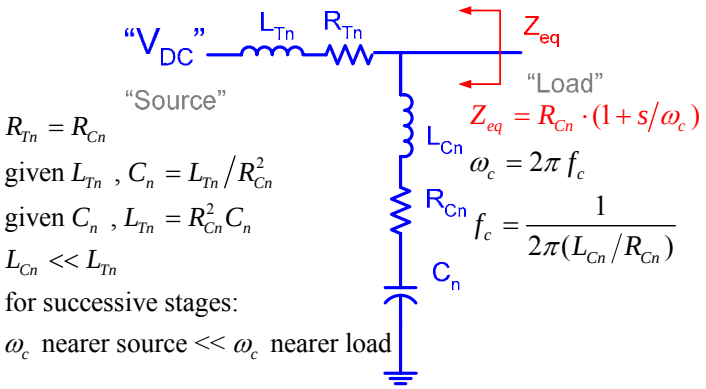

Figure A.3. Design equations for single filter stage.

In general, there is some flexibility in designing the filter stages even when the BVR equivalent model at the source end and the C_0-R_{C0} at the load end are given. However, depending on the number of decoupling stages and the actual design values, no practical solution may exist. One procedure to design the passive filter is as follows:

1. Given two parameters in equation: $\Delta V_0 = \Delta I_0 \cdot R_{C0}$, find the unknown.

2. Form Thevenin equivalent network for every stage using the approximation that the bypass capacitor of the stage closer to the source behaves as a DC voltage source with esr and esl- all else is effectively open due to large interconnect inductance.

3. Budget series interconnect resistance, R_n , over all stages keeping $\sum R_n = R_{C0}$

4. Calculate esr of successive bypass capacitors using $R_{Tn} = R_{Cn} = R_{Cn+1} + R_n$

5. Estimate physically realizable interconnect inductance, L_n over all stages.

6. Where applicable, calculate desired bypass capacitor values, $C_n = L_{Tn} / R_{Cn}^2$; design C_n to have lowest possible esl, L_{Cn} ; find $L_{Tn-1} = L_{Cn} + L_{n-1}$ and use to progressively determine remaining capacitances OR …

7. Alternatively, calculate desired Thevenin Equivalent stage interconnect inductance, $L_{Tn} = R_{Cn}^2 C_n$; find $L_{Cn+1} = L_{Tn} - L_n$; obtain largest possible C_{n+1} with esl, L_{Cn+1}, and use this value of C_{n+1} to progressively design the stages.

8. Verify values satisfy all constraints including,
$$Q_{Cn} = \frac{\sqrt{L_{Cn}/C_n}}{R_{Cn}} < \frac{1}{2} \Rightarrow L_{Cn} < \frac{1}{2} R_{Cn}^2 C_n \Rightarrow L_{Cn} < L_{Tn}/2$$

9. If design is not practical, modify path resistance or inductance budget and repeat steps.

Using the 2004 Power Delivery Network, which is relabeled in Figure A.4 for generality, a detailed example of a passive bypass filter design will now be worked out.

The topology, symbol names and load regulation specifications are shown in Figure A.4.

The BVR source model is given.

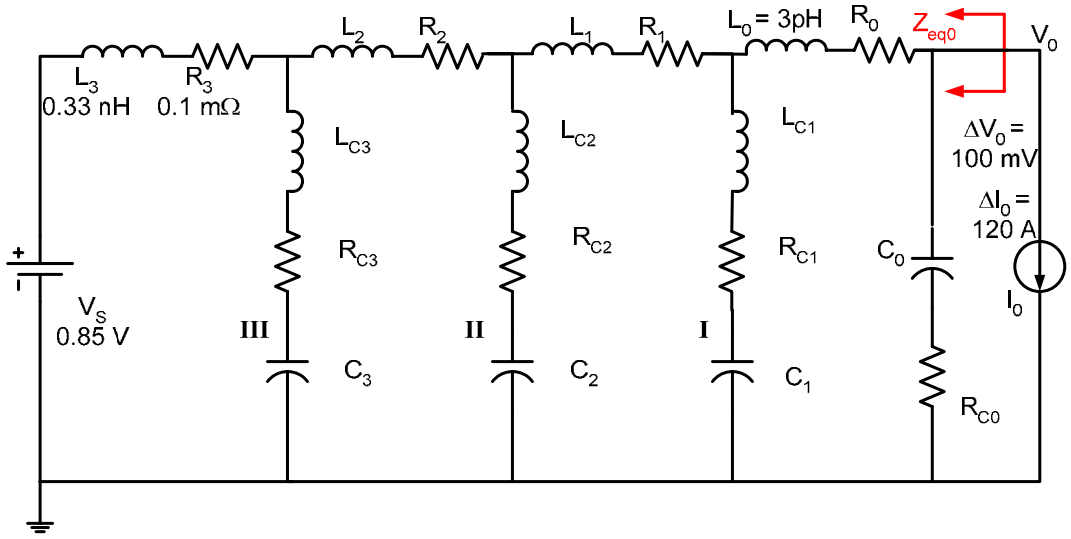

Figure A.4. Example of a 4-stage Power Delivery Network.

1. $\Delta V_0 = \Delta I_0 \cdot R_{C0} \Rightarrow R_{C0} = \Delta V_0 / \Delta I_0 = 100\text{mV}/120\text{A} \approx 0.8\ \text{m}\Omega$

2. The Thevenin equivalent circuits for the four stages are given in Figure A.5

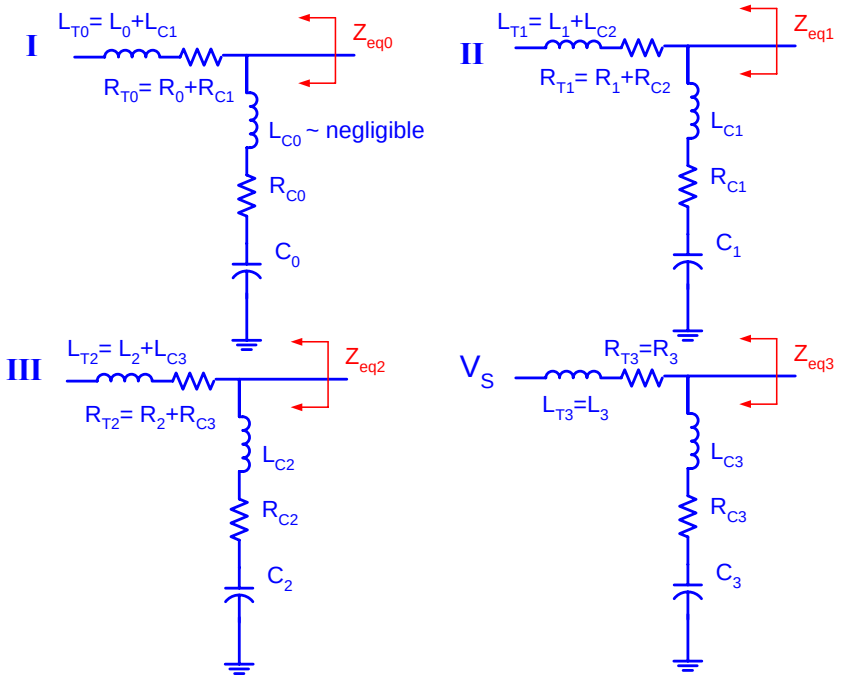

Figure A.5. Equivalent circuits of decoupling stages starting from load to source.

3. Based on the 2004 Power Delivery Path model, it is reasonable to divide the required 0.8 mΩ total series resistance among the interconnect resistances as follows: $R_0 = 0.1\ \text{m}\Omega$; $R_1 = 0.4\ \text{m}\Omega$; $R_2 = 0.2\ \text{m}\Omega$; $R_3 = 0.1\ \text{m}\Omega$

4. Starting with the predetermined value of R_{C0}, the other capacitor esr's are found:

$$R_{Tn} = R_{Cn} = R_{Cn+1} + R_n \Rightarrow R_{Cn+1} = R_{Cn} - R_n$$
$$R_{C0} = 0.8\ \text{m}\Omega;\ R_0 = 0.1\ \text{m}\Omega \Rightarrow R_{C1} = 0.7\ \text{m}\Omega$$
$$R_{C1} = 0.7\ \text{m}\Omega;\ R_1 = 0.4\ \text{m}\Omega \Rightarrow R_{C2} = 0.3\ \text{m}\Omega$$
$$R_{C2} = 0.3\ \text{m}\Omega;\ R_2 = 0.2\ \text{m}\Omega \Rightarrow R_{C3} = 0.1\ \text{m}\Omega$$

It is fortuitous that R_{C3} is 0.1 mΩ, the same value as R_3. A lower value may mean R_3 would have to change. In a real power delivery system this would force a redesign of the BVR to have lower output resistance. Similar issues may arise with interconnect resistances and, more likely, interconnect inductances.

5. Again, based on 2004 Power Delivery Architectures, an estimated set of inductance values may be: $L_0 = 3\ \text{pH}$; $L_1 = 40\ \text{pH}$; $L_2 = 50\ \text{pH}$; $L_3 = 330\ \text{pH}$

6. $$C_3 = L_{T3}/R_{C3}^2 = L_3/R_{C3}^2 = (330\ \text{pH})/(0.1\ \text{m}\Omega)^2 = 33\ \text{mF}$$

Assuming high-density capacitors of 5 nH esl at 1 mF, which is consistent with the 2004 model,

$$L_{C3} = 5\ \text{nH}/33 = 150\ \text{pH}$$
$$\Rightarrow L_{T2} = L_{C3} + L_2 = 150\ \text{pH} + 50\ \text{pH} = 200\ \text{pH}$$
$$\Rightarrow C_2 = L_{T2}/R_{C2}^2 = (200\ \text{pH})/(0.3\ \text{m}\Omega)^2 = 2.2\ \text{mF}$$

Assuming low-esl capacitors of 100 pH esl at 1 mF, which is also consistent with the 2004 model,

$$L_{C2} = 90\ \text{pH}/2.2 \approx 40\ \text{pH}$$
$$\Rightarrow L_{T1} = L_{C2} + L_1 = 40\ \text{pH} + 40\ \text{pH} = 80\ \text{pH}$$
$$\Rightarrow C_1 = L_{T1}/R_{C1}^2 = (80\ \text{pH})/(0.7\ \text{m}\Omega)^2 = 160\ \mu\text{F}$$

Assuming ultra-low-esl capacitors of 25 pH at 1 uF, also based on the 2004 model:

$$L_{C1} = 25\ \text{pH}/160 \approx 0.16\ \text{pH}$$
$$\Rightarrow L_{T0} = L_{C1} + L_0 = 0.16\ \text{pH} + 3\ \text{pH} = 3.16\ \text{pH}$$
$$\Rightarrow \text{desired } C_1 = L_{T0}/R_{C0}^2 = (3.16\ \text{pH})/(0.8\ \text{m}\Omega)^2 = 5\ \mu\text{F}$$

7. As the design was constructed from the source model specifications, it isn't necessary to carry out this step.

8. Reviewing Table A.1, the inductances satisfy the "resonant-free" constraint, $L_{Cn} < L_{Tn}/2$. However note that at stage three (n = 3), the condition is just barely met. The overall design is shown in Figure A.6. Simulation results are given in The output impedance is maintained at 0.8 mΩ except near 80 KHz where it dips as low as 0.65 mΩ which is acceptable as the transient response does meet the load regulation specifications- it never droops more than 100 mV.

Table A.1. Design Values.

n =	0	1	2	3
L_n [pH]	3	40	50	330
L_{Cn} [pH]	0	0.16	40	150
L_{Tn} [pH]	3.16	80	200	330
R_n [mΩ]	0.1	0.4	0.2	0.1
R_{Cn} [mΩ]	0.8	0.7	0.3	0.1
C_n [uF]	5	160	22 000	33 000

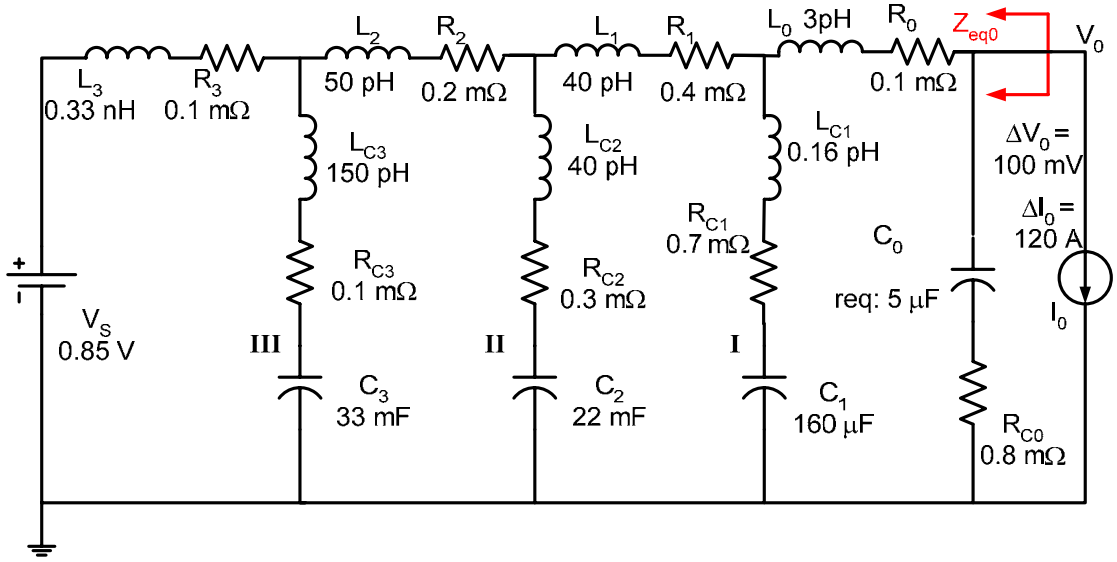

Figure A.6. Resultant passive filter design.

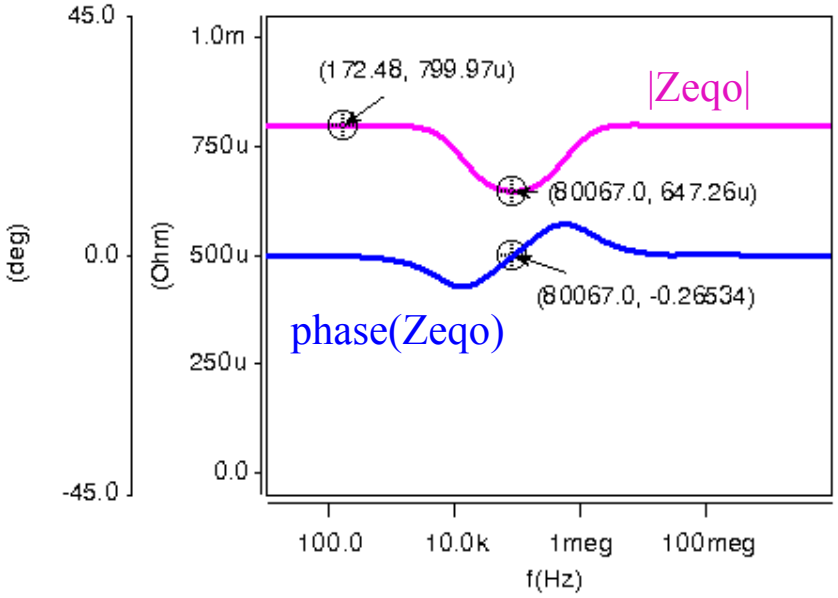

Figure A.7. Output impedance of circuit in Figure A.6.

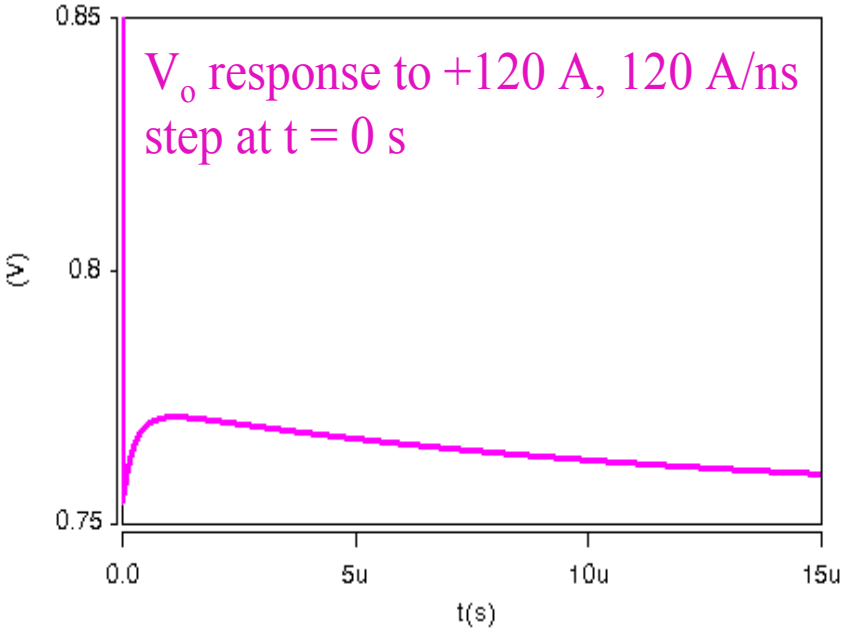

Figure A.8. Pull-up transient response of circuit in Figure A.6.

9. There may be many other design constraints not considered in this example. In fact a solution may not exist. As a practical exercise, an additional constraint of $C_0 = 1\ \mu$F is now imposed. It is now to be expected that at the very least L_0 must

also change. Keeping all else the same, and working from step 7:

$L_{T0} = R_{C0}^2 C_0 = \left(0.8\ \text{m}\Omega\right)^2 \left(1\mu F\right) = 0.64$ pH

$\Rightarrow$ desired $L_0 = L_{T0} - L_{C1} = 0.64 - 0.16 = 0.48$ pH

or at the very least, design MUST have $L_0 < 0.64$ pH

So here again it is obvious that $L_0 = 3$ pH is prohibitive

APPENDIX B
LINEAR MODEL EXTRACTION OF BULK VOLTAGE REGULATOR

The output voltage response of the BVR to a 100 A step in load current, Figure B.1, together with the data for the output filter capacitor, was given. The open circuit voltage is V_{BVR}. The steady state voltage droop yielded an initial guess for the output resistance of the BVR, R_{BVR} . R_{BVR} = (voltage droop) ÷ (current step) = 100 mV ÷ 100 A = 1 mΩ . The time constant of the response, which is equal to L_{BVR}/R_{BVR} , was then used, together with the value obtained for R_{BVR} to get an initial guess for L_{BVR}:

$L_{BVR} = R_{BVR}$ (time constant) = 1 mΩ x 0.4 μs = 0.4 nH.

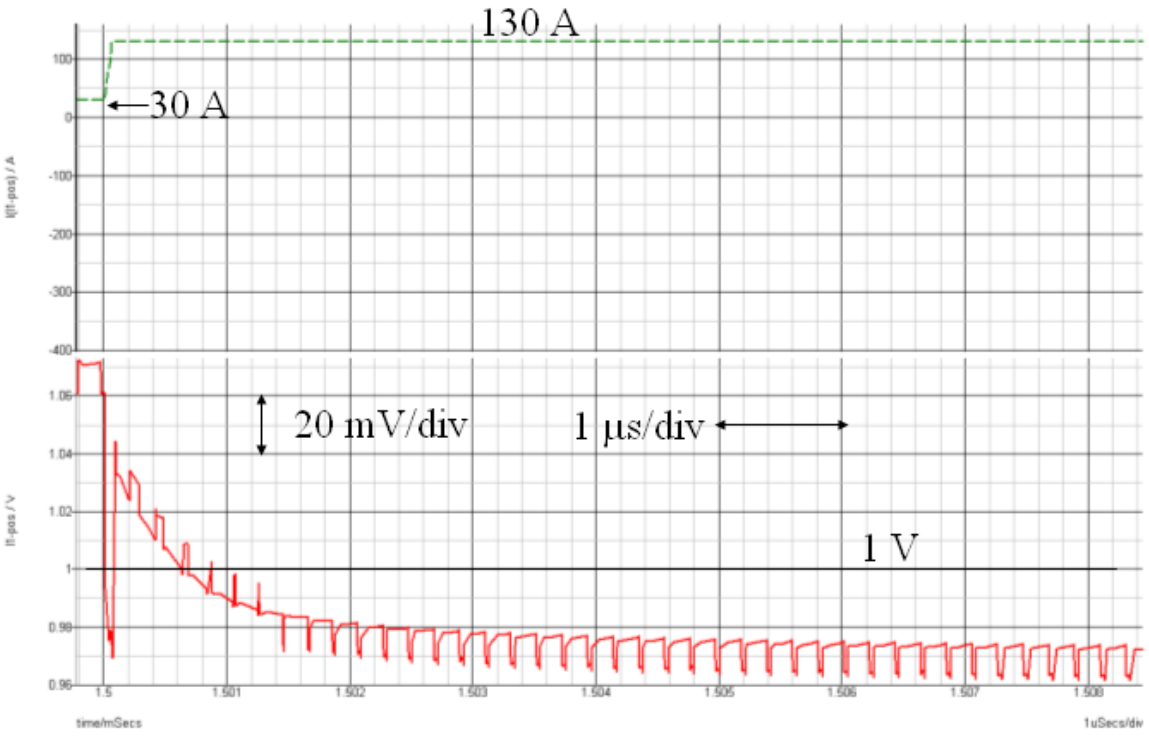

Figure B.1. BVR step-load response with known output filter capacitor.

These values were entered into the basic linear model topology, Figure B.2, and simulated. The response shown in Figure B.3 was compared to Figure B.1 to verify that the initial guesses were correct and that no additional tweaking to the model was necessary.

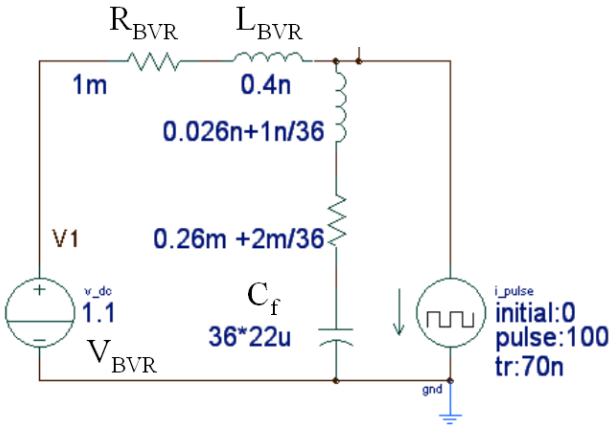

Figure B.2. Equivalent linear model of the BVR.

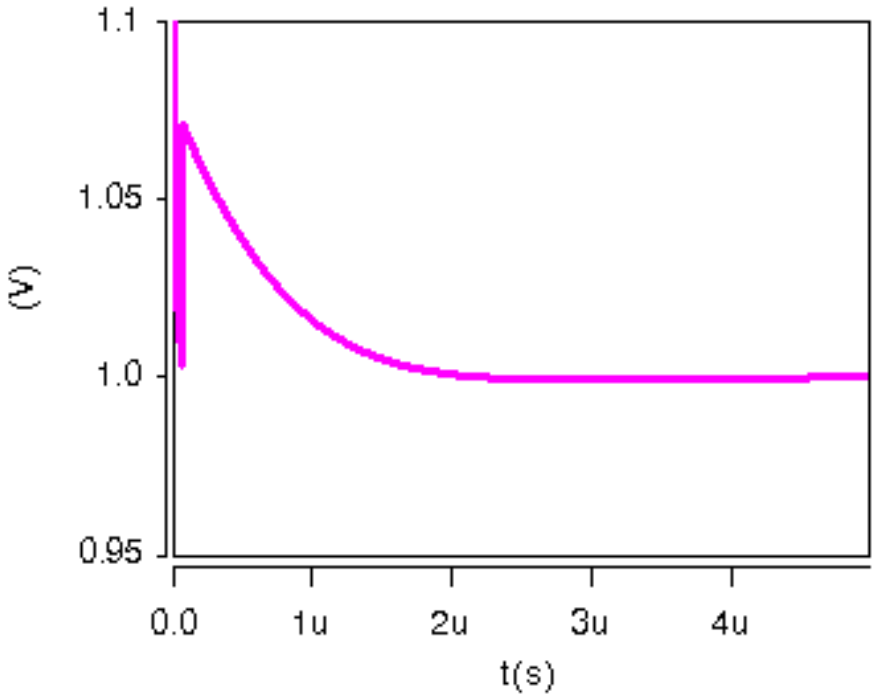

Figure B.3. Linear model step-load response.

By simply inserting this model as extracted into the lumped element power delivery path of the microprocessor, including the interconnect parasitics, the resulting model can be used to design and test the hybrid in simulations. The advantage of using the linear model for the BVR in simulations is that it greatly reduces simulation time. However, switching noise and other BVR control non-linearity cannot be studied with the simple linear model.

LIST OF REFERENCES

[Amo99] L. Amoroso, M. Donati, M.X. Zhou, and F. C. Lee, "Single shot transient suppressor (SSTS) for high current high slew rate microprocessor," in Proceedings of IEEE Applied Power Electronics Conference and Exposition, Dallas, TX, 1999, pp. 284-288.

[Ang00] M. Ang, R. Salem, and A. Taylor, "An on-chip voltage regulator using switched decoupling capacitors," IEEE International Solid State Circuits Conference, Castine, ME, 2000, pp. 438-439.

[Bar05] A. Barrado, A. Lazaro, R. Vasquez, V. Salas, and E. Olias, "The fast response double buck DC-DC converter (FRDB): operation and output filter influence," IEEE Transactions on Power Electronics, vol. 20, pp. 1261-1270, November 2005.

[Bla04] C. Black, K. Guarini, Z. Ying, K. Hyungjun, J. Benedict, E. Sikorski, I. Babich, K. Milkove, "High-capacitor, self-assembled metal-oxide semiconductor decoupling capacitors," IEEE Electron Device Letters, vol. 25, pp. 622-624, September 2004.

[Int02] VRM 9.0 DC-DC Converter Design Guidelines, Intel Corporation, Santa Clara, CA, 2002.

[Int03] Intel Pentium 4 Processor in the 478-pin Package/Intel 850 Chipset Family Platform Design Guide, Intel Corporation, Santa Clara, CA, 2003.

[Int04] Voltage Regulator-Down (VRD) 10.0: for Desktop Socket 478 Design Guide, Intel Corporation, Santa Clara, CA, 2004.

[Int05a] Intel® Pentium® 4 Processor 660, 650, 640, and 630Δ and Intel® Pentium® 4 Processor Extreme Edition Datasheet, Intel Corporation, Santa Clara, CA, 2005.

[Int05b] Voltage Regulator-Down (VRD) 10.1 Design Guide, Intel Corporation, Santa Clara, CA, 2005.

[Li00] Y. Li, T. Yew, C. Chung, and D. Figueroa, "Design and performance evaluation of microprocessor packaging capacitors using integrated capacitor-via-plane model," IEEE Transactions on Advanced Packaging, vol. 23, pp. 361-367, August 2000.

[Lim03] F. Lima, A. Geraldes, T. Marques, J. N. Ramalho, and P. Casimiro, "Embedded CMOS distributed voltage regulator for large core loads," Conference on European Solid-State Circuits, Estoril, Portugal, 2003, pp. 521-524.

[Luo02] J. Luo, I. Batarseh, X. Gao, and T. Wu, "Transient current compensation for low-voltage high-current voltage regulator modules," in Proceedings of IEEE Applied Power Electronics Conference and Exposition, Dallas, TX, 2002, pp. 223-228.

[Man04] O. Mandhana, "Modeling, analysis and design of resonant free power distribution network for modern microprocessor systems," IEEE Transactions on Advanced Packaging, vol. 27, pp. 107-120, February 2004.

[Par05] H. Park, H. Kim, D. G. Kam, and J. Kim, "Co-modeling and co-simulation of package and on-chip decoupling capacitor for resonant free power/ground network design," in Proceedings of the Electronic Components and Technology Conference, Orlando, Fl, 2005, pp. 727-731.

[Poo99] F. N. K. Poon; C. K. Tse, and J. C. P. R. Liu, "Very fast transient voltage regulators based on load correction," in Records of 30th Power Electronics Specialists Conference, Charleston, SC, 1999, pp. 66-71.

[Rah04] T. Rahal-Arabi, J. Gang, M. Ma, A. Muhtaroglu, and G. Taylor, "Development and validation of an electromagnetic distributed power grid model for the 90nm Pentium/spl reg/ 4 processor," Digest of Technical Papers, VLSI Circuits Symposium, Honolulu, HI, 2004, pp. 110-113.

[Red98] R. Redl, B.P. Erisman, and Z. Zansky, "Optimizing the load transient response of the buck converter," in Proceedings of IEEE Applied Power Electronics Conference and Exposition, Anaheim, CA, 1998, pp. 170-176.

[Ren04] Y. Ren, K. Yao, M. Xu, and F.C. Lee, "Analysis of the power delivery path from the 12 V VR to the microprocessor," in Proceedings of IEEE Applied Power Electronics Conference and Exposition, Anaheim, CA, 2004, pp. 285-291.

[Sta04] E. Stanford, "Microprocessor voltage regulators and power supply trends and device requirements," Proceedings, International Symposium on Power Semiconductor Devices and ICs, Kitakyushu, Japan, 2004, pp. 47-50.

[Vor89] V. Vorperian, R. Tymerski, F.C. Lee, "Equivalent circuit models for resonant and PWM switches," IEEE Transactions on Power Electronics, vol. 4, pp. 205-214, April 1989.

[Wai01] A. Waizman and C. Chung, "Resonant free power network design using extended adaptive voltage positioning (EAVP) methodology," IEEE Transactions on Advanced Packaging, vol. 24, pp.236-244, August 2001.

[Wu01] Albert M. Wu and Seth R. Sanders, “An active clamp circuit for voltage regulation module (VRM) applications,” IEEE Transactions on Power Electronics, vol. 16, pp. 623-634, September 2001.

[Yao03] K. Yao, K. Lee, M. Xu, and F.C. Lee, “Optimal design of the active droop control method for the transient response,” in Proceedings of IEEE Applied Power Electronics Conference and Exposition, Miami Beach, Fl, 2003, pp.718-723.

[Yao04] K. Yao, Y. Ren, J. Sun, K. Lee, M. Xu, J. Zhou, and F.C. Lee, “Adaptive voltage position design for voltage regulators,” in Proceedings of IEEE Applied Power Electronics Conference and Exposition, Anaheim, CA, 2004, pp.272-278.

[Zho00] X. Zhou, P. Wong, P. Xu, F. Lee, and A. Huang, "Investigation of candidate VRM topologies for future microprocessors," IEEE Transactions on Power Electronics, vol. 15, pp. 1172-1182, November 2000.

BIOGRAPHICAL SKETCH

Attma Sharma was born in Trinidad and Tobago in 1974. He migrated to the United States in 1989. He is currently a PhD candidate, and a student of Professor Khai D.T. Ngo, in electrical and computer engineering at the University of Florida. His dissertation research into transient voltage regulation circuits for integrated power supplies led to new methods of active filtering for power supply regulation applications. Some of his work has been published at regional power electronics conferences, as well as in the *IEEE Power Electronics Transactions*. As part of a joint project led by Prof. Ngo and Prof. Toshikazu Nishida, he did research in micropower converters for piezo-electric transducers. Working in the labs of Prof. Toshikazu Nishida, he gained experience in microfabrication techniques. He has also done some research in passive networks for RF applications under the guidance of Professors Khai D.T. Ngo and William R. Eisenstadt. He received the M.S. degree in electrical and computer engineering from the University of Florida in 2003 and the B.S. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, in 1995. He has worked as a full-time lab instructor in the electrical engineering laboratories of the University of the West Indies, St. Augustine Trinidad and Tobago and as an analog design engineer with the semiconductor test division of Teradyne Inc., Agoura Hills, CA. His research interests are in analog electronics, RF and power IC's, and microelectronic devices.