

Filters with Active Tuning for Power Applications

by

Joshua W. Phinney

B.A., Wheaton College (1995)
B.S., University of Illinois at Chicago (1999)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2001


© Massachusetts Institute of Technology, MMI. All rights reserved.


Author _____
Department of Electrical Engineering and Computer Science
July 26, 2001

Certified by _____
David J. Perreault
Assistant Professor, Department of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by _____
Arthur C. Smith
Chairman, Departmental Committee on Graduate Students

BARKER





Filters with Active Tuning for Power Applications

by
Joshua W. Phinney

Submitted to the Department of Electrical Engineering and Computer Science
on July 31, 2001, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

EMI filters for switching power converters rely on low-pass networks — with corner frequencies well below the ripple fundamental — to attenuate switching harmonics over a range of frequencies. Tight ripple specifications imposed to meet conducted EMI specifications can result in bulky and expensive filters which are detrimental to the transient performance of a converter and often account for a substantial portion of its size and cost. This thesis focuses on two techniques for reducing the size of passive elements required to mitigate converter ripple: active tuning of resonant filters utilizing phase-sensing control, and hybrid reactive structures which develop low shunt impedances through a passive inductance cancellation.

Resonant networks provide extra attenuation at discrete frequencies, easing the filtering requirement of an accompanying low-pass section. By exchanging “brute-force” attenuation for selective attenuation, resonant filters can realize substantial volume and weight savings if they are aligned with switching harmonics. Manufacturing tolerances and operating conditions readily push narrow-band tuned circuits away from their design frequencies, and such filters are rarely employed in switching power converters. This thesis explores the design and application of a phase-lock control scheme which makes resonant filters practical by aligning a converter’s switching frequency with a filter immitance peak, or vice-versa. Such a resonant filter — an actively tuned filter — is distinct from an active filter because it does not directly drive waveforms within acceptable limits, and is not dissipation-limited. The applications and limitations of resonant filters are discussed, and experimental results from DC-DC converters are presented.

Integrated filter elements, hybrid capacitor/transformer structures, can cancel (to a large extent) the parasitic inductance of power capacitors. Parasitic inductance limits the effectiveness of shunt filter elements by increasing their impedance to ripple currents at high frequencies. Integrated filter elements can be constructed from wound foil and dielectric layers in the same manner as capacitors, and their magnetically coupled windings can be incorporated into capacitor packages. Experimental results from integrated elements are presented which demonstrate improvements in filtering due to shunt inductance cancellation and the accompanying introduction of series reactance.

Thesis Supervisor: David J. Perreault

Title: Assistant Professor, Department of Electrical Engineering and Computer Science

Acknowledgements

My thanks are due to the US Office of Naval Research and to the MIT/Industry Consortium on Advanced Automotive Electrical/Electronic Components and Systems for supporting my assistantship in the Laboratory for Electromagnetic and Electronic Systems (LEES); to Drs. Thomas Keim, Jeffrey Lang, and John Kassakian for their expert advise on technical matters; to various friends and colleagues — particularly Jamie Byrum, Tim Neugebauer, Chris Laughman, Ivan Celanovic, Amy Ng, Steve Shaw, and Dave Wentzloff — whose company made my years in LEES most enjoyable; but above to my advisor, David Perreault, without whose creativity and engineering judgement this thesis would not exist. I would also like to single out Vahe Caliskan, Tim Denison, Ernst Scholtz, and John Rodriguez for their patient explanations on numerous occassions. The responsibilty for any shortcomings remains entirely mine, but without the support of all these people this would have been a much poorer work.

Contents

1	Introduction	11
1.1	Filter topologies: Active Tuning	13
1.1.1	Filtering components: Integrated Filter Elements	15
1.2	Thesis Objectives and Contributions	17
1.3	Organization of the Thesis	17
2	Phaselock Basics	19
2.1	PLL Components	19
2.1.1	Phase detector (PD)	20
2.1.2	The four-quadrant multiplier as a phase detector	21
2.1.3	Voltage-controlled oscillator (VCO)	23
2.1.4	Loop filters	25
2.2	Linearized model for the PLL	27
2.3	PLL Operating ranges	29
2.3.1	Lock range $\Delta\omega_L$	30
2.3.2	Pull-in range $\Delta\omega_P$	31
2.3.3	Noise performance	33
2.4	PLL design	35
2.4.1	PLL Design with negligible noise power	36
2.4.2	PLL design when noise must be considered	37
2.4.3	Design example	38
3	Resonant-network design	43
3.1	Constraints on resonant network design	45
3.1.1	Impedance constraints: Quality factor Q and characteristic impedance Z_0	45
3.1.2	Harmonic constraints: Antiresonance	47
3.1.3	Duty-ratio constraints	50
3.1.4	Component-rating constraints	51
3.2	Resonator design	51
3.2.1	Parallel-tuned series resonator	53

Contents

3.2.2	Series-tuned shunt resonator	54
3.3	Magnetically coupled shunt resonator	56
3.4	Summary	61
3.5	Design examples	62
3.5.1	Design example: low ripple current	62
3.5.2	Design example: high ripple current	63
4	Phase-lock Tuning	67
4.1	Phase-lock tuning	67
4.1.1	Equivalence of phase and impedance tuning conditions	71
4.1.2	Tuning system dynamics	73
4.2	Application to a DC-DC converter	74
4.3	Alternative implementations	79
5	Integrated Filter Elements	83
5.1	Principle of Operation	84
5.1.1	Implementations of an Integrated Filter Element	86
5.2	Experimental results	88
5.2.1	Results from the switching converter	90
5.2.2	Series inductance	90
5.2.3	Shunt inductance cancellation	92
5.3	Manufacturing	93
5.4	Further work	95
6	Conclusions	97
6.1	Conclusions: actively tuned filters	97
6.2	Conclusions: integrated filter elements	99
6.3	Further work	100
A	MATLAB files	103
A.1	coredata.m	103
A.2	wiredata.m	109
A.3	inductance.m	110
A.4	coredesign.m	112
A.5	convergence.m	115
A.6	coreloss.m	116
A.7	acresistance.m	117

A.8 dcreistance.m	117
A.9 permBpk.m	118
A.10 permH.m	119
A.11 iripple.m	119
A.12 par.m	120
A.13 general.m	120
B Phase-lock tuning circuit	123

List of Figures

1.1	Dimensions of EHPS power converter	11
1.2	EHPS converter schematic	12
1.3	SAE J1113/41 Class 1 EMI specification	13
1.4	Resonator impedance and admittance	14
1.5	Resonant power-stage example	15
1.6	Two methods of realizing maximum resonant attenuation	16
1.7	Integrated-filter-element construction	17
1.8	Inductance cancellation	18
2.1	The basic components of a phase-lock loop	20
2.2	Phase-detector characteristic	21
2.3	Phase-error example	22
2.4	Multiplier phase-detector characteristic	23
2.5	VCO characteristic	24
2.6	Linearized model of the PLL with offsets	24
2.7	Four common loop filters	26
2.8	PLL root-locus diagrams	27
2.9	Linearized AC model of the PLL	27
2.10	PLL closed-loop phase transfer function	28
2.11	PLL static error transfer function	29
2.12	PLL operating ranges	29
2.13	Depiction of pull-in and lock-in processes	32
2.14	Phase jitter	33
2.15	PLL loop-noise bandwidth	35
2.16	Application of the AD633 multiplier	38
2.17	Application of the XR2206 function generator	38
2.18	Schematic of the PLL used in the prototype tuning system	41
3.1	Three resonant-network topologies	43
3.2	Applications of parallel- and series-tuned resonators	44
3.3	Q and characteristic impedance	45
3.4	Relationship of Q and tuning-point immitance	46
3.5	Antiresonance	47

3.6	Normalized performance surface for resonators	48
3.7	Loci of minimum reactive energy storage	49
3.8	Harmonic magnitudes vs. duty ratio	50
3.9	Normalized immitances used in generalized design	52
3.10	Parallel- and series-resonator designs using performance locus	54
3.11	Shunt-resonator design through resonance-shifting	56
3.12	Circuit models for the magnetically coupled shunt resonator	57
3.13	Transfer function from converter voltage to AC-winding voltage	58
3.14	Parameterized attenuation for the magnetically coupled shunt resonator . .	59
3.15	Filtering performance of resonant and “zero-ripple” designs	60
3.16	Resonant input filter design for a 300W buck converter	61
3.17	Design of a resonant power stage	63
3.18	Inductor loss model	64
4.1	Resonator impedance and admittance	68
4.2	Block diagram of phase-lock tuning system	69
4.3	Alternate resonant-excitation topologies	70
4.4	Tuning points of parallel resonator with asymmetric loss	71
4.5	Linearized model for the phase-sensing tuning system in lock.	72
4.6	Schematic of the tuning circuitry used in the prototype tuning system . . .	74
4.7	Tuning system applied to a buck converter	75
4.8	Size and performance comparison for power-stage example	77
4.9	Size and performance comparison for input-filter example	78
4.10	Application to other power converters	79
4.11	Hybrid inductive-capacitive elements that exhibit resonances	79
4.12	Structural diagram of a cross-field reactor	81
5.1	Circuit models of the integrated filter element	83
5.2	Magnetically coupled windings	84
5.3	T-model of the integrated filter element with negative branch impedance . .	85
5.4	Construction of an integrated filter element	86
5.5	Conducted-EMI test setup	88
5.6	LISN power spectra	89
5.7	Inductor-capacitor structure compared to integrated-filter element	90
5.8	Performance gains from series inductance	91
5.9	Experimental setup for inductance-cancellation measurement	92
5.10	Measured voltage gain demonstrating inductance cancellation	93
5.11	ESR and ESL histograms	94
5.12	Incorporation of the coupled windings into the structure of a power capacitor. . .	95

Introduction

PASSIVE filters for switched-mode power converters rely on low-pass networks — with corner frequencies well below the ripple fundamental — to attenuate switching harmonics over a range of frequencies. Ripple specifications imposed to observe conducted EMI limits (Fig. 1.3) or application constraints, however, can result in heavy, bulky filters which are detrimental to the transient performance of a power converter and contribute significantly to its cost.




Figure 1.1: Physical dimensions the power converter module and EMI-filter components for an Automotive EHPS (Electro-Hydraulic Power Steering) system. The converter module is mounted in the hydraulic fluid reservoir, so little extra volume is required for heatsinking. The volume of the 5 principle EMI filter elements (3 capacitors and 2 inductors) is 5.65 in³, compared to about 6 in³ converter volume (control and power devices) within the depicted enclosure.

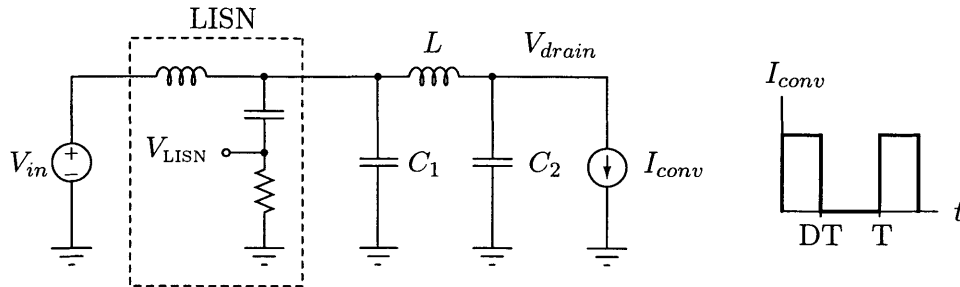


Figure 1.2: Conceptual schematic of the EHPS converter input filter (details of common-mode filtering removed). The converter power stage draws a large pulsed current I_{conv} that C_2 must pass in order to provide hold-up at V_{drain} . During testing, a line-impedance stabilization network (LISN) terminates the input filter in a known AC impedance, typically the 50Ω channel-input impedance of an oscilloscope or spectrum analyzer.

Consider the 1kW power converter of an automotive electro-hydraulic power-steering system (Fig. 1.1). AC impedance mismatches — low AC shunt impedance and high AC series impedance — are provided by a π filter (Fig. 1.2) to divert the ripple component of I_{drain} away from the input source V_{in} . V_{drain} is usually considered to be the converter input for control purposes, and must be held close to its average value as the power stage draws large pulsed currents. C_2 must therefore have low impedance at the converter switching frequency (and its first few harmonics), and a ripple-current rating high enough to accommodate the majority of the AC current drawn by the converter switching cell. Electrolytic capacitors are typical choices for C_1 , and may be placed in parallel to increase their current-handling capability. Such components are physically large, and with the accompanying series inductor account for the majority of a typical filter’s volume.

Stringent conducted EMI specifications (as low as $8\mu\text{A}$ at 3MHz, see Fig. 1.3) impose different constraints on the capacitors of subsequent stages (cf. C_1). Such capacitors need not handle much current, but must have low impedance at EMI frequencies above a few multiples of the switching frequency (hundreds of kilohertz to a few megahertz). Multilayer ceramic (MLC) or multi-layer polymer (MLP) capacitors are typical choices for this filter stage because of their large capacitance and small volume. Typically, such capacitors contribute negligibly to the volume a filter but significantly to its cost; an MLP or MLC capacitor can cost as much as a large electrolytic 50 to 100 times its size.

This automotive example highlights a general trend in converter design: passive components in a converter filter or power stage can dominate the volume of a system (see the volume comparison in caption, Fig. 1.1) and contribute significantly to its cost. This thesis explores topology- and component-level techniques for reducing the volume of passive

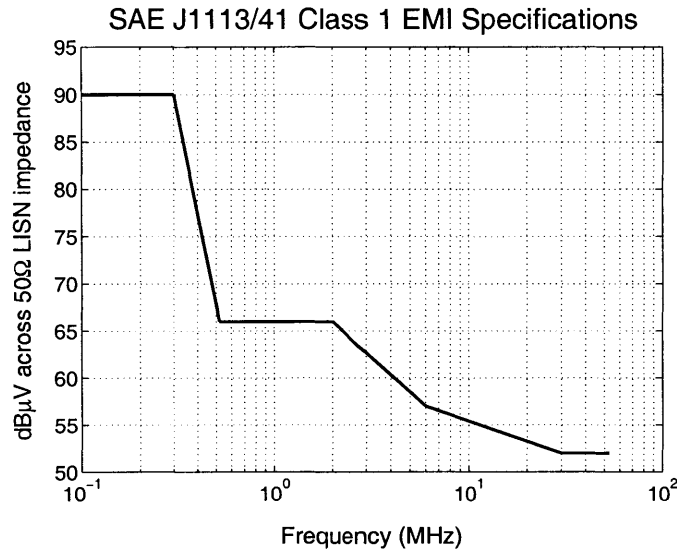


Figure 1.3: The SAE J1113/41 Class 1 specification for narrowband signals. The input of the EHPS power converter (Fig. 1.3) must meet this conducted EMI specification. 90 dBμV corresponds to just 31.6 mV across 50Ω, or 0.632 mA. The more stringent requirements at 3 MHz allow only 8 μA in conducted emissions.

filter elements required for a given level of performance. The first technique — a topological approach — employs resonant networks in conjunction with a phase-sensing tuning system. The second method investigates the incorporation of film-wound transformers in capacitor packages, reducing the high-frequency shunt impedance provided by typical high-ripple-current capacitors.

1.1 Filter topologies: Active Tuning

Resonant ripple filters offer attenuation comparable to low-pass networks — for less volume and weight — using the immitance peaking of parallel- and series-tuned circuits (Fig. 1.4) to introduce transmission nulls at discrete frequencies. Consider, for example, the buck converters of Figs. 1.5c and 1.5d. In Fig. 1.5c, the buck inductor and output capacitor C_1 form a low-pass filter which attenuates the ripple generated by the converter switching stage (Fig.1.5a). In the converter of Fig. 1.5d, a much smaller output capacitor is placed in parallel with a trap tuned to the converter switching frequency, resulting in a “notched” attenuation characteristic (Fig. 1.5a). While both converter structures yield the same peak-to-peak output voltage ripple (Fig. 1.5b) and require about the same magnetic energy storage, the converter of Fig. 1.5d needs only about one-fourth the capacitive energy storage.




Figure 1.4: Frequency response of second-order tuned circuits, normalized to the natural frequency $\omega_n = 1/\sqrt{LC}$. The high admittance (impedance) of the series-tuned (parallel-tuned) network can effectively divert AC currents away from a port of interest when placed in shunt (series) with its terminals.

Inasmuch as suitably low-loss resonator components are available in a small volume, active tuning can reduce the overall size and cost of the filter compared to a conventional low-pass design.

Because resonant networks must typically have high Q to attenuate target harmonics sufficiently,¹ they provide only narrow-band attenuation. Operating conditions and manufacturing variations can readily cause narrow-band resonators to miss their design frequencies[2] and fail to attenuate ripple; for this reason they are rarely employed in switching power converters. Filters with active tuning control achieve reliable resonant excitation by placing a resonator’s frequency response or a converter’s switching frequency under closed-loop control (Fig. 1.6). In the frequency-control form especially, resonant filters can process high power because the tuning circuitry operates at signal power levels. By modulating the switching frequency to realize the maximum attenuation from a passive network, actively tuned filters never directly drive the waveforms they condition, and are not — like active filters ([3]–[9]) — dissipation-limited.

Effective use of resonance allows the filter designer to exchange attenuation across a selected range of frequencies for physically smaller reactive components. I.e., by ensuring effective attenuation at the ripple fundamental or a ripple harmonic frequency, active tuning eases the filtering requirement — and so lowers the volume and energy storage — of an ac-

¹Some high-power applications use damped, low- Q resonators precisely for their broad attenuation characteristic and insensitivity to detuning, at the expense of attenuation performance.[1]




Figure 1.5: (a) Transfer functions from switching voltage to output voltage for the converters of (c) and (d). (b) Output voltage waveforms for the converters of (c) and (d) operating with 50% duty cycle. The circuits in (c) and (d) have different filter arrangements but an identical power stage and load. $V_{in} = 42\text{V}$, $f_{sw} = 125\text{kHz}$, $L = 20\mu\text{H}$, $R_L = 1\Omega$. (c) System with capacitive low-pass filter: $C_1 = 10\mu\text{F}$. (d) System with attenuated low-pass and notch filters: $C_2 = 1\mu\text{F}$, $C_f = 1.6\mu\text{F}$, $L_f = 1\mu\text{H}$, $R_f = 50\text{m}\Omega$.

companion network. A reduction in the volume of the passive elements required for a given level of ripple performance must not necessarily be realized as volume savings. A designer can, for instance, maintain passive-component volume at a lower switching frequency — reducing switching loss and improving efficiency — without sacrificing performance. One could also maintain the volume of a conventional filter while achieving better ripple performance at a constant switching frequency. Better performance can alternately ease the need for large impedance mismatches, allowing the replacement of a small but expensive MLP or MLC capacitor with a less expensive capacitor of lower value.

1.1.1 Filtering components: Integrated Filter Elements

Typical electrolytic capacitors (cf. C_2 in Fig. 1.2) have a frequency response well-approximated by a series-tuned resonance (like that in Fig. 1.4a) to frequencies in excess of 10MHz. Such power capacitors present a relatively high impedance at EMI frequencies because of the rise in their impedance magnitude above a low self-resonant frequency (tens of kilohertz, typically). While suitable for hold-up, i.e. passing large currents at the switching fundamental, large power capacitors alone cannot meet stringent conducted EMI specifications like those shown in Fig. 1.3.

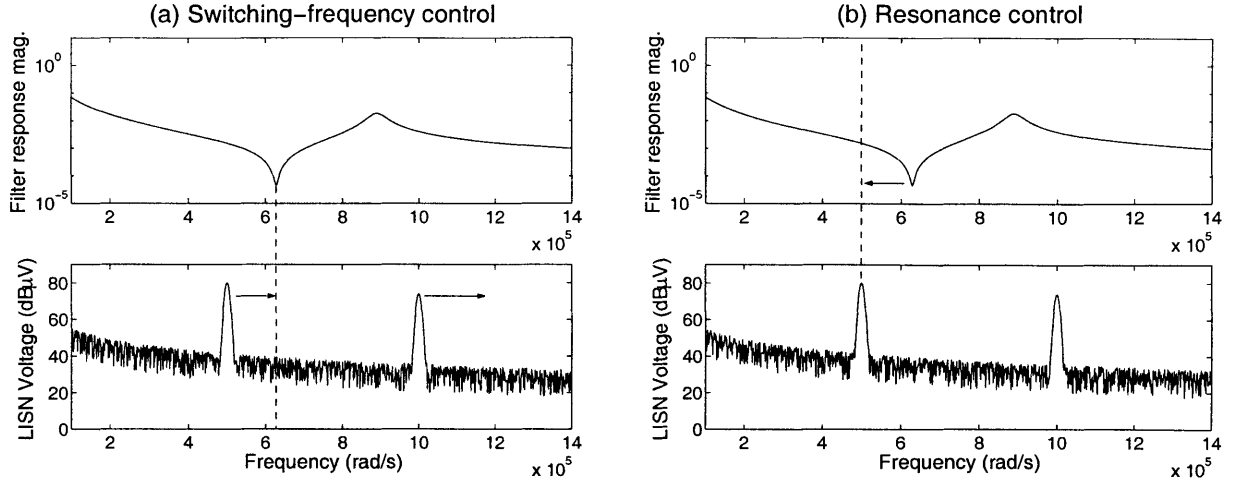


Figure 1.6: Two methods of tuning for maximum resonant attenuation. (a) Switching frequency control and (b) resonance control, in which a filter reactance is altered to adjust transmission nulls. Method (a) has the advantage of simplicity, while method (b) can independently tune multiple resonances to provide attenuation at several frequencies.

Integrated filter elements, proposed here, are transformer-capacitor structures that effectively cancel the equivalent series inductance (ESL) of a power capacitor, increasing the frequency of its impedance rise and making it useful at switching *and* EMI frequencies. The integrated element comprises a normal capacitor structure with a magnetically coupled film windings (Fig. 1.7). An equivalent T model for the coupled windings in an autotransformer configuration can be obtained from a Δ -Y transformation of the impedances measured at the three terminal pairs (Fig. 1.8b). The T model adds an internal node from which the inductance $L_{11} - L_M$ (self-inductance of the AC winding minus the mutual inductance) can be made negative by the proper choice of turns ratio $N_1 : N_2$.

When $L_{11} - L_M$ is chosen to be close to the capacitor's ESL (L_{ESL} in Fig. 1.8b), the shunt network reduces to a capacitor with excellent frequency response, i.e. lower ESL and lower impedance at high frequencies than the original C . Such an integrated element is not only useful for shunting high-frequency currents to ground: the addition of an inductive reactance (L_M or $L_{22} - L_M$) in series with either port results in a filter with higher-order roll-off than is possible with a simple capacitor. These two advantages — low shunt impedance at high frequencies and increased series impedance — can be obtained with almost no extra volume using inexpensive, repeatable manufacturing techniques.




Figure 1.7: Magnetically coupled foil strips— windings 1 and 2 — can be added over basic capacitor structure or made from extensions of capacitor foil. The integrated filter element is now a three-terminal device, with the extra 3rd winding brought out as Node A in Fig. 1.8a.

1.2 Thesis Objectives and Contributions

The first goal of this thesis is to elucidate the design of resonant filters with active tuning control. The discussion will include sufficient background and modelling information for the practicing engineer to design and evaluate resonant filters and phase-lock tuning controls. Active-tuning control is, in fact, a general technique for controlling the fundamental phase shift between periodic signals, and can be applied to many resonant excitation and detection problems.² The second goal of this thesis is to introduce integrated filter elements incorporating shunt inductance cancellation. The development here will focus on the feasibility and performance advantages of this hybrid reactive structure.

1.3 Organization of the Thesis

Chapter II presents the principles of the phase-lock circuitry utilized in active tuning control. Chapter III presents three resonant-network topologies considered for use in conjunction with the tuning system, and gives particular attention to the volume trends in each

²Resonant-beam chemical sensors offer an immediate mechanical analogy for the electrical systems discussed here. In such sensor problems, the absolute value of frequency command with phase-sensing feedback indicates the mass of adsorbed molecules. See [10] and [11].




Figure 1.8: (a) A schematic diagram of the integrated filter element, shown without parasitics. Capacitance C is the power capacitor whose ESL the transformer is intended to cancel. (b) The schematic redrawn, including important parasitics, but otherwise leaving terminal I-V relations unchanged. (c) When $L_{11} - L_M$ is chosen to be close to the capacitor’s ESL, the shunt network reduces to a capacitance with small ESL, i.e. $\Delta L = -L_M + L_{11} + L_{ESL} \approx 0$.

design. Chapter IV introduces the phase-sensing control system which aligns the switching frequency of a power converter with the resonant frequency of a filter. Experimental results from the power stage and input filter of a buck converter employing the phase-lock tuning approach highlight the weight and volume reduction achievable in converter magnetics. Chapter IV also considers additional applications and topologies of the phase-lock tuning system, including resonance-tuning methods (see Fig. 1.6b) implemented with cross-field reactors. Chapter V extends the discussion of magnetically coupled shunt resonators to integrated filter elements. Experimental results from prototype integrated structures are presented that demonstrate the great promise of this technology. Finally, Chapter VI summarizes the results of the thesis and suggest directions for continued work in this area.

Phaselock Basics

THE GOAL of this chapter is to provide an introduction to phase-lock loop (PLL) design for the practicing power-electronics engineer, along with its applications to the resonant-excitation problem. The discussion in the following pages follows a standard development of the subject found in Gardner [12], Best [13], and Wolaver [14], and should provide enough background for rapid design and troubleshooting of phase-sensing tuning systems. Section 2.4 on page 35 details a step-by-step design procedure for the PLL when noise power does not interfere with reliable lock-in. This design procedure is adapted from [13], and was used to design all PLLs used in the prototype controllers.

2.1 PLL Components

A PLL, fundamentally, is an oscillator whose output-signal frequency (ν_o in Fig. 2.1) is controlled to align with some frequency component of its input signal ν_i . Let θ_i be the phase¹ and ω_i the angular frequency of the component of interest within ν_i , with ν_o characterized similarly by θ_o and ω_o . The phase detector (PD) generates the detector voltage v_d , some frequency component of which is proportional to the phase difference $\theta_i - \theta_o$. Frequency tracking is achieved by driving the voltage-controlled oscillator (VCO) with a filtered version of this phase difference. In the limit of high loop-filter DC gain, a steady VCO command voltage v_c can be maintained with small steady-state phase error, resulting in “lock” between the input and output phase (and, hence, lock between ω_i and ω_o). The loop filter is usually chosen by the designer for a given VCO and PD, and adjusts, by altering control bandwidth, the range of input frequencies over which the PLL can reliably acquire lock.

In the following sections, detailed discussions of the PLL components will lead to the development of a linear model of the lock-in process. This model is the basis for predicting

¹Phase is the integral of frequency, the complete argument of a sinusoidal function. E.g., the phase of the signal $f(t) = \cos(\omega_o t + \phi)$ for constant ϕ is the linearly increasing function $\theta_o = \omega_o t + \phi$, which corresponds to a constant frequency ω_o .




Figure 2.1: The basic components of a phase-lock loop

lock range — the principle concern of the designer — at least when noise power is not “too high.” A step-by-step design guide for a PLL loop filter will follow a basic discussion of the frequency ranges which characterize the PLL, and under which circumstances the designer can alter these ranges by loop shaping. Finally, section 3.5.1 will present a PLL design example from the tuning system for a shunt-resonant filter.

2.1.1 Phase detector (PD)

A phase detector generates an output signal v_d which depends on the phase difference between its inputs. A plot qualitatively illustrating the relationship between detector voltage v_d and the *phase difference* θ_d (the difference between the input phase θ_i and the VCO phase θ_o) is shown in Fig. 2.2a. The curve is not, in general, linear, but 2π periodicity is typical for commonly used PDs for which a phase of ϕ is indistinguishable from any $\phi \pm 2n\pi$.² When no input signal v_i is applied to the PD, its output is the *detector offset voltage* V_{do} . Zero *phase error* θ_e is commonly referenced to the *phase offset* θ_{do} corresponding to $v_d = V_{do}$, as depicted in Fig. 2.2b and expressed below:

$$\theta_e = \theta_d - \theta_{do}$$

This shift in the point of zero phase error is usually carried over to the definition of input and output phase such that

$$\theta_e = \theta_i - \theta_o$$

E.g. even if v_i and v_o are phase-shifted sinusoids, the signals are “in phase” for analytic purposes — indeed we alter the description of their phases so that the shift is zero — when the magnitude of the phase difference is such that $v_d = V_{do}$. The curve of Fig. 2.2b is called the *PD characteristic* and has a slope K_d (the *PD gain*) at the point of zero tracking error $\theta_e = 0$. Even in cases where the PD characteristic is nonlinear, the PD output is

²This statement applies to two-state PDs and any memoryless PD. A generalized n -state PD with $n \geq 3$, is able to store enough information about cycle slips to maintain linear tracking to multiples of 2π radians.




Figure 2.2: (a) The PD characteristic with offset. (b) The same characteristic shifted such that the input signals to the PD, v_i and v_o are “in phase” for analytic purposes when their phase difference is such that v_d equals the detector offset V_{do} . (c) Linearized model of the PD with offsets, valid over the PD range ($\pm\pi/2$ for the depicted characteristic).

approximately

$$v_d = K_d \theta_e + V_{do}$$

(cf. the signal-flow graph in Fig. 2.2c). For linear PD characteristics with inflections or steps, or for any linearized PD characteristic, this model is sufficiently accurate over some PD range (e.g. $\pm\pi/2$ for the characteristic of Fig. 2.2b). Accurate prediction of PLL locking dynamics requires that the average phase error θ_{e0} be well within the PD range.

2.1.2 The four-quadrant multiplier as a phase detector

A four-quadrant multiplier acts as a phase detector by the trigonometric identity

$$\sin(\phi_1) \cdot \cos(\phi_2) = \frac{1}{2} \cdot \sin(\phi_1 - \phi_2) + \frac{1}{2} \cdot \sin(\phi_1 + \phi_2)$$




Figure 2.3: Phase error for the multiplication of the 8Hz and 7Hz sinusoids v_i and v_o . In the plot of $\tilde{v}_d = v_i v_o$, the 15 Hz sum frequency “rides” on top the 1 Hz difference frequency (the dashed line in the \tilde{v}_d graph), which is taken as the output v_d of the PD. The phase error θ_e is just the argument of the sinusoidal function necessary to produce v_d . In a feedback setting θ_e would (hopefully) never be allowed to increase beyond the PD range as shown.

Assume that the PLL has acquired lock to a purely sinusoidal input signal, and let the inputs to the multiplier be $v_i = V_i \sin(\omega_i t)$ and $v_o = V_o \cos(\omega_i t - \theta_e)$.³ The output of the multiplier is then $\tilde{v}_d = K_m v_i v_o$, where K_m is a constant associated with the multiplier as represented in the signal-flow graph of Fig. 2.4. The product expression becomes

$$\tilde{v}_d = \frac{1}{2} \cdot K_m V_i V_o \sin(\theta_e) + \frac{1}{2} \cdot K_m V_i V_o \sin(2\omega_i - \theta_e) \quad (2.1)$$

Fig. 2.3 plots \tilde{v}_d for θ_e increasing linearly with time. In most PLL applications, the sum-frequency term in the expression for \tilde{v}_d is at a high enough frequency ($2\omega_i$) that it is effectively removed by low-pass loop dynamics. The first term in Eqn. 2.1 is then considered to be the output v_d of the PD, and is just the average of the complete product waveform. This average is taken over a long enough period to eliminate the $2\omega_i$ term, but not so long as to affect the relationship $v_d = \frac{1}{2} \cdot K_m V_i V_o \sin(\theta_e)$ when θ_e is a function of time. For small values of θ_e , $\sin(\theta_e) \approx \theta_e$ and $v_d \approx K_d \theta_e$ where the PD gain, $K_d = \frac{1}{2} \cdot K_m V_i V_o$, depends on the amplitude of the input signals.

³Note that the phase error is defined with respect to quadrature phase, under which condition the average value of the PD output is zero.




Figure 2.4: (a) PD characteristic for the multiplier. (b) Full signal-flow model before linearization. K_m has units of volts⁻¹ and depends only on the multiplier.

2.1.3 Voltage-controlled oscillator (VCO)

A voltage-controlled oscillator generates a waveform whose frequency depends on a control voltage v_c . A schematic VCO characteristic is depicted in Fig. 2.5a. As with the PD, the curve need not be linear, though a linear relation is common in integrated VCOs and greatly simplifies an *a priori* prediction of the PLL lock range. In the locked state (i.e. when the average of ω_i equals the average of ω_o), the input to the VCO is a steady-state *control offset voltage* V_{co} . Unlike the detector offset voltage V_{do} , V_{co} is a function of the particular average input frequency ω_i .

The linearized dynamical treatment of the VCO parallels the PD analysis of Sec. 2.1.1. The *frequency deviation* $\Delta\omega_o$, a measure how far ω_o deviates from its average in lock, is given by

$$\Delta\omega_o = \omega_o - \omega_i$$

The frequency deviation characteristic Fig. 2.5b is just a shifted version of the VCO characteristic and is characterized by a slope — the *VCO gain* K_o — at the lock point. The frequency deviation can be modelled by the block diagram of Fig. 2.5c, where

$$\Delta\omega_o = K_o(v_c - V_{co})$$

Assume again that the PLL has acquired lock to a purely sinusoidal input signal, and let the inputs to the multiplier be $v_i = V_i \sin(\omega_i t + \theta_i)$ and $v_o = V_o \cos(\omega_i t + \theta_o)$.⁴

⁴The assumption of sinusoidal signals is actually not required by these expressions for v_i and v_o . Either signal, as written, can be made an arbitrary function of time by proper choice of $\theta_i(t)$ and $\theta_o(t)$. An assumption of purely sinusoidal v_o and v_i motivates the use of the phase notation $\omega t + \theta$ and relates the




Figure 2.5: (a) PLL output frequency vs. VCO command voltage v_c (b) Shifted VCO characteristic expressed in terms of frequency deviation $\Delta\omega_o$ (c) Signal-flow diagram for the VCO




Figure 2.6: Linearized model of the PLL with offsets

The average VCO output frequency in lock must be ω_i , but can also be expressed as the derivative of the output phase:

$$\omega_o = \frac{d}{dt}(\omega_0 t + \theta_o) = \omega_i + \frac{d\theta_o}{dt}$$

Rearranging terms, and applying the definition $\Delta\omega_o = \omega_o - \omega_i$, we arrive at

$$\Delta\omega_o = \frac{d\theta_o}{dt} \quad \text{or} \quad \theta_o(t) = \int \Delta\omega_o dt = K_0 \int (v_d - V_{co}) dt$$

expressions to previous equations. It is also worth stressing that ω_i and ω_o are the *average* frequencies of the input and output signals, where the *total* frequency is the derivative of phase, or $\omega + d\theta/dt$.

2.1.4 Loop filters

The treatment of the PLL to this point has described every block in Fig. 2.6 except $F(s)$, the loop filter. The DC gain of this block decreases the steady-state phase error θ_{e0} needed to support a VCO command voltage v_c :

$$\theta_{e0} = -\frac{V_{d0}}{K_d} + \frac{V_{co}}{K_d \cdot F(s=0)}$$

The loop-filter should be low-pass in order to extract a moving average of the PD output, discarding — as much as is feasible — any high-frequency terms produced by the PD (e.g. the 15Hz signal in Fig. 2.3). More important to the designer, however, is the ability to use $F(s)$ to accommodate anticipated input signals. The loop filter is the designer's principle means of shaping the loop transmission and adjusting the PLL lock range $\Delta\omega_L$ (see Sec. 2.3.1).

Considering the loop filters presented in Fig. 2.7, all three can limit, with suitable choice of components, the bandwidth of the control signal v_c applied to the VCO. The active lag network of Fig. 2.7b has gain K_a at DC, and differs from the passive lag network (neglecting loading) only in the designer's freedom to change its magnitude response. All three loop filters in Fig. 2.7 have a zero at $1/\tau_2$ that inflects the low-frequency gain upward at 6 dB/octave. This rise in $|F(s)|$ is limited only by the open-loop gain of the op-amp in the proportional+integral (PI) case.

A high-frequency pole in each the active circuits of Fig. 2.7 — due to the finite unity-gain bandwidth of the op amps, if nothing else — will cause $F(s)$ to roll off at high frequencies. The designer may choose to introduce the high-frequency pole at some specified $\omega_3 = 1/\tau_3$ (cf. the circuit of Fig. 2.7d), to limit PLL phase jitter and improve locking performance (see Sec. 2.3.3). As seen in the root locus diagram of Fig. 2.8b for such a design, the low-frequency poles become ever more lightly damped as loop gain increases. If ω_3 is placed too close to the zero at $1/\tau_2$, the low-frequency singularities never enter far into the LHP, and are lightly damped at the natural frequency required for a usable lock-in range (see. Sec. 2.3). For this reason, ω_3 is at least four times the PLL crossover frequency, and this pole can be neglected in the control design but considered for noise analysis.




Figure 2.7: Schematic diagrams and transfer functions $F(s) = v_c(s)/v_d(s)$ for four commonly used loop filters. Filter (d) can be approximated by (c) for purposes of the control design.




Figure 2.8: Root locus for (a) distant high-frequency pole and (b) pole at $-1/\tau_3$.




Figure 2.9: Linearized AC model of the PLL

2.2 Linearized model for the PLL

The linearized AC model for the PLL, with averaged or DC quantities removed, is shown in Fig. 2.9. The loop transmission $G(s)$ for the loop filters presented in Fig. 2.7 is the product of the loop-filter transfer function $F(s)$, the VCO integrator transfer function $1/s$, and the gain product K . For the passive lag and active PI filters, $K = K_0K_d$. The designer must specify the DC gain of K_a for an active lag filter, so in this case $K = K_0K_dK_a$. K can be selected by the designer to choose the closed-loop pole locations (see the root-locus diagrams of Fig. 2.8).

Evaluating the small-signal transfer function $T(s)$ from input to output phase (the phase cosensitivity function) for the different loop transmission functions $G(s)$ and expressing the denominator in standard second-order form yields:




Figure 2.10: Closed-loop phase transfer function $T(s) = \Theta_o(s)/\Theta_i(s)$, plotted for various choices of damping ratio ζ .

$$\begin{aligned}
 \text{passive lag} \quad T(s) &= \frac{s\omega_n \cdot (2\zeta - \frac{\omega_n}{K}) + \omega_n^2}{s^2 + 2s\zeta\omega_n + \omega_n^2} & \text{where } \omega_n &= \sqrt{\frac{K}{\tau_1 + \tau_2}} & \text{and } \zeta &= \frac{\omega_n}{2} \cdot (\tau_2 + \frac{1}{K}) \\
 \text{active lag} \quad T(s) &= \frac{s\omega_n \cdot (2\zeta - \frac{\omega_n}{K}) + \omega_n^2}{s^2 + 2s\zeta\omega_n + \omega_n^2} & \text{where } \omega_n &= \sqrt{\frac{K}{\tau_1}} & \text{and } \zeta &= \frac{\omega_n}{2} \cdot (\tau_2 + \frac{1}{K}) \\
 \text{PI filter} \quad T(s) &= \frac{2s\zeta\omega_n + \omega_n^2}{s^2 + 2s\zeta\omega_n + \omega_n^2} & \text{where } \omega_n &= \sqrt{\frac{K}{\tau_1}} & \text{and } \zeta &= \frac{\omega_n\tau_2}{2}
 \end{aligned}
 \tag{2.2}$$

The $G(s)$ for practical PLLs satisfy the high-gain criterion $K \gg \omega_n$ so that the cosensitivity function for all loop filters can be rewritten

$$T(s) \approx \frac{2s\zeta\omega_n + \omega_n^2}{s^2 + 2s\zeta\omega_n + \omega_n^2}$$

A Bode plot of this second-order phase transfer function is shown in Fig. 2.10. The PLL is low-pass filter for input phase signals, acting like a “flywheel” responsive to modulation by signals with a frequency less than the PLL natural frequency ω_n . A Bode plot (Fig. 2.11) of the static error-transfer function

$$S(s) = 1 - T(s) \approx \frac{s^2}{s^2 + 2s\zeta\omega_n + \omega_n^2}$$




Figure 2.11: Static error transfer function $S(s) = 1 - T(s)$ plotted for various choices of damping ratio ζ .




Figure 2.12: PLL operating ranges

provides the same insight in command following: input modulation frequencies in excess of ω_n appear as phase error θ_e because the PLL loses static phase tracking at frequencies above ω_n . Note that ω_n should not be confused with the range of frequencies ($\Delta\omega_L$, introduced in Sec. 2.3) over which the PLL can acquire lock, though the two are related (see Eqn. 2.3).

2.3 PLL Operating ranges

As mentioned in the development of the linear VCO model in Sec. 2.1.3, the control offset voltage V_{co} is a function of a particular ω_i ($= \omega_o$) to which the PLL has acquired lock. So too, ω_o forms a shifting reference for the lock range $\Delta\omega_L$, pull-out range $\Delta\omega_{PO}$, and pull-in

range $\Delta\omega_P$ (Fig. 2.12), defined as follows:

- Lock range $\Delta\omega_L$** A PLL is normally designed to operate within its lock range. This is range of $\Delta\omega_o$ over which the PLL will acquire lock within one beat note between the VCO and input frequencies.
- Pull-in range $\Delta\omega_P$** A PLL will, in the absence of noise, always become locked for $\Delta\omega_o$ within the pull-in range, though perhaps after a slow “pull-in” process. See Sec. 2.3.2 for a more detailed description of this phenomenon, and reasons why the designer should avoid operation in the region outside $\pm\Delta\omega_L$
- Pull-out range $\Delta\omega_{PO}$** The pull-out range represents the dynamical limits to PLL stability, i.e. the frequency step which, when applied to the PLL input, causes lock-out. An exact expression for $\Delta\omega_{PO}$ has never been derived for the analog PLL, though simulations[15] yield good approximations, and verify that $\Delta\omega_P > \Delta\omega_{PO} > \Delta\omega_L$ in a typical design.
- Hold range $\Delta\omega_H$** The hold range indicates the static stability range of the PLL, and is determined by the absolute signal ranges of the PD or VCO. ω_o is considered to be in the middle of the VCO tuning range when computing $\Delta\omega_H$, as the limits of static frequency tracking — an absolute measure — does not depend on a frequency at which the PLL might previously have acquired lock.

2.3.1 Lock range $\Delta\omega_L$

The magnitude of the lock range $\Delta\omega_L$ can be computed accurately enough for design purposes using a few simple approximations. Consider for a moment that the PLL is not locked and that the PLL input is a sinusoid. Using the notation of Sec. 2.1.3, the input signal can be expressed as $\omega_i = \omega_o + \Delta\omega_o$. The detector voltage is then

$$v_d = K_d \sin(\Delta\omega_o t)$$

where the higher frequency terms of the linearization have been neglected due to the low-pass form of $G(s)$. The VCO command voltage is then approximately

$$v_c \approx |F(\Delta\omega_o)| \cdot K_d \sin(\Delta\omega_o t)$$

v_c is a time-varying signal which modulates the frequency of the VCO output, producing a peak frequency variation of $K_o K_d \cdot |F(\Delta\omega_o)|$.

Consider the case where $\Delta\omega_o$ is greater than the VCO's peak frequency deviation (Fig. 2.13a). The VCO command cannot support lock, at least not immediately, and so sweeps the VCO output at the beat note frequency $\Delta\omega_o$. If ω_i is brought closer to ω_o , so that $\Delta\omega_o$ just equals $K_o K_d \cdot |F(\Delta\omega_o)|$ (Fig. 2.13b), v_c is able to support the lock condition $\omega_i = \omega_o$ at the extreme edge of its modulation range. $\Delta\omega_L$ is therefore determined approximately by the nonlinear equation

$$\Delta\omega_L \approx K_o D_d \cdot |F(\Delta\omega_L)|$$

The solution to this equation can be found through some simplifying approximations for $|F(\Delta\omega_L)|$. First, the lock range of a practical PLL is always far greater than the pole or zero frequencies of $F(s)$, i.e. $\Delta\omega_L \gg 1/\tau_1$ or $1/(\tau_1 + \tau_2)$, and $\Delta\omega_L \gg 1/\tau_2$. The expressions for $|F(\Delta\omega_L)|$ then reduce to

$$\text{Passive lag filter} \quad |F(\Delta\omega_L)| \approx \frac{\tau_2}{\tau_1 + \tau_2}$$

$$\text{Active lag filter} \quad |F(\Delta\omega_L)| \approx \frac{K_a \tau_2}{\tau_1}$$

$$\text{Active PI filter} \quad |F(\Delta\omega_L)| \approx \frac{\tau_2}{\tau_1}$$

For many $F(s)$, τ_2 is much smaller than τ_1 , allowing the further simplification $|F(\Delta\omega_L)| \approx \tau_2/\tau_1$ for the passive lag filter. For each loop filter, then — assuming high gain product K — the lock range can be expressed as

$$\Delta\omega_L \approx 2\zeta\omega_n \tag{2.3}$$

2.3.2 Pull-in range $\Delta\omega_P$

A PLL is still able to acquire lock when ω_i lies outside of ω_o 's modulation range $\pm K_o D_d \cdot |F(\Delta\omega_o)|$. This process of acquisition — a pull-in process — takes place because the frequency devi-




Figure 2.13: Depiction of the (a) pull-in process and (b) the lock-in process.

ation $\Delta\omega_o$ is varied as the VCO output ω_o is cycled. Consider Fig. 2.13a. The duration in which ω_o is modulated in the positive direction, toward ω_i , is longer than the duration of modulation away from ω_i . I.e. the angular frequency deviation $\Delta\omega_o$ decreases as ω_o approaches ω_i , retarding the excursion of ω_o that decreases $\Delta\omega_o$ and pushing the mean of ω_o slightly above the average of its peak values. This peak-peak median is the mean of ω_o without a full consideration of frequency modulation, and is represented in Fig. 2.13a by the dotted horizontal line. Indeed, if ω_o were modulated by some function of *constant* $\Delta\omega_o$, the average value of ω_o would follow this dotted line, and a lock-in process would be the only means of decreasing the average value of $\Delta\omega_o$ for constant ω_i .

Note that the pull-in process, because it relies on a cycle-to-cycle decrease in $\Delta\omega_o$, is necessarily slower than lock-in. Expressions for PLL pull-in time and range can be found in [13], but because of approximations made in their derivation, viz. neglect of noise during pull-in, and because of the sluggishness of pull-in even under favorable conditions, the PLL should be designed to operate in its lock range exclusively, when possible. The tuning-frequency range required for resonant filters is sufficiently narrow to be covered by a PLL lock range, with proven lock reliability for VCOs modulated by $\pm 15\%$ from their center frequency. It should be noted that when the desired tuning range exceeds the $\pm\omega_L$, additional control circuitry can aid the pull-in process, enabling reliable acquisition of noisy signals [13].

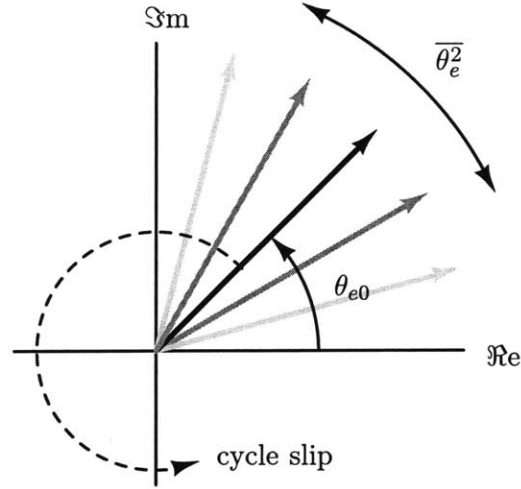


Figure 2.14: Phase jitter in PLL tracking error. An excessive probability of cycle slips will prevent the PLL from ever maintaining steady-state operation.

2.3.3 Noise performance

In steady-state operation and in the absence of noise, a PLL can maintain lock with a steady-state phase error θ_{e0} (Fig. 2.14). Phase noise in the PLL input signal θ_i introduces jitter in θ_e , but will not perturb θ_e far from its equilibrium point if the noise power is sufficiently low. Higher noise power can cause occasional cycle slips in θ_e , disturbances in which the PD output shifts a whole period (e.g. 2π in Fig. 2.2) but resumes operation around the equilibrium detector voltage. When cycle-slip disturbances become too frequent, no stable operating point can be achieved, and the PLL permanently locks out.

Locking performance is compromised whenever the PLL designer attempts to implement a lock range that is “too broad.” A quantitative bound on $\Delta\omega_L$ follows from a consideration of the PLL as a means of improving signal-to-noise ratio (SNR). Such noise reduction is regarded in some applications (viz. clock recovery) as the principle measure of merit for a PLL design. A PLL improves SNR by a ratio of noise bandwidths:

$$\text{SNR}_L = \frac{P_s}{P_n} \cdot \frac{B_i}{2B_L} = \text{SNR}_i \cdot \frac{B_i}{2B_L} \quad (2.4)$$

Where SNR_i is the SNR at the PLL input, the ratio of input-signal power P_s to noise power P_n

SNR_L is the closed-loop SNR at the PLL output

$B_i/2$ is the bandwidth of the input phase-noise signal. B_i is the bandwidth of the noise component in v_i , such that its spectral density W_i (assumed constant in where it is non-zero) is $W_i = P_n/B_i$ W/Hz.

B_L is the “loop noise bandwidth,” the equivalent noise bandwidth of the closed-loop phase transfer function $T(s)$.

The PLL improves the SNR of a phase signal as B_L decreases. B_L is the bandwidth of a fictitious low-pass filter with a constant magnitude of transmission equal to $T(0)$ (Fig. 2.15). B_L is selected such that the two filters — the rectangular filter and the PLL, which is a low-pass filter for phase signals — produce outputs with equal variance for white noise inputs of equal density. For a phase spectral density of Φ (rad)²/Hz, the output phase jitter (i.e. variance) is

$$\overline{\theta_{no}^2} = \Phi \int_0^{B_i/2} |T(s)|^2 ds \approx \Phi \int_0^{j\infty} |T(s)|^2 ds$$

From equal areas under the actual and equivalent squared frequency responses,

$$B_L = \frac{\Phi}{2\pi j} \int_0^{j\infty} \left| \frac{2\zeta\omega_n + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \right|^2 ds = \frac{\Phi}{2\pi} \int_0^\infty \frac{4\zeta^2\omega_n^2\omega^2 + \omega_n^4}{\omega_n^4 + 2\omega_n^2(2\zeta^2 - 1)\omega^2 + \omega^4} \cdot d\omega \quad (2.5)$$

$$B_L = \frac{\omega_n}{2} \left(\zeta + \frac{1}{4\zeta} \right) \text{ Hz} \quad (2.6)$$

Where $\overline{\theta_{no}^2} = \Phi B_L$.

Experiments performed with second order PLLs and reported by Gardner [12] reveal the useful limit on loop noise bandwidth. For SNR_L less than 4, lock-in may be possible, but is unreliable. SNR_L is a function and B_L and the noise characteristics of the signal source at the PLL input, so a design goal of $\text{SNR}_L > 4$ places an upper limit on the loop noise bandwidth in a given design setting. Because B_L is directly proportional — like the lock range $\Delta\omega_L$ — to ω_n , the designer may need to trade off acquisition and noise performance (see Sec. 2.4).

The best noise performance (the lowest B_L) is achieved at $\zeta = 0.5$ (Fig. 2.15). $\zeta = 0.707$, which is often selected for good control characteristics, increases B_L/ω_n a negligible

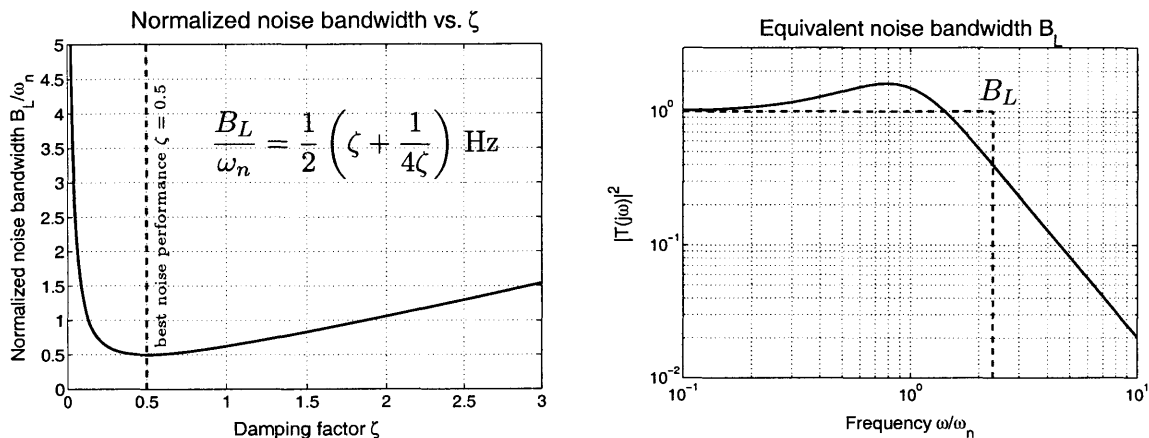


Figure 2.15: (a) Loop-noise bandwidth B_L normalized to ω_n , plotted for various ζ . (a) Equivalent noise bandwidth of $T(s)$. B_L describes a fictitious rectangular filter with the same variance in output phase as the PLL.

amount, viz. 0.53 compared to the 0.50 minimum. $\zeta = 0.7$ is a therefore suitable choice for any normal PLL application.

A high-frequency pole in $T(s)$ (e.g. $\omega_3 = 1/\tau_3$ in Fig. 2.7) decreases the loop noise bandwidth by decreasing the argument of the integral in Eqn. 2.5. This pole should be at the lowest possible frequency to decrease B_L as much as possible. For reasonable damping (see Sec. 2.1.4), ω_3 should only be placed four or more times higher than the crossover frequency ω_{co} . $4\omega_{co}$ is therefore a good choice for ω_3 .

2.4 PLL design

The loop noise bandwidth B_L and lock range $\Delta\omega_L$ cannot be specified independently. Both quantities are related, through the damping factor ζ , to the PLL's natural frequency ω_n . In the case of low noise power, B_L can be neglected and ω_n calculated using lock range alone. When noise power is larger, especially as the designer attempts to achieve a large lock range, an ω_n determined from $\Delta\omega_n$ alone can result in an excessive SNR_L — the closed-loop SNR at the PLL output — and prevent reliable lock-in. In such cases, ω_n must be kept within limits set by the noise design, and $\Delta\omega_L$ will be restricted.⁵

⁵If lower lock range is not acceptable, the designer has the option of implementing more complicated control strategies such as sweep acquisition and dynamic bandwidth limiting.[13]

2.4.1 PLL Design with negligible noise power

The following design procedure is useful whenever SNR_L is greater than about 4. Though the PLL input noise bandwidth B_i can be simply determined, it is often difficult to measure or approximate the input signal-to-noise ratio SNR_i . Because SNR_i is needed to check a lower bound on SNR_L , the practitioner may want to follow the procedure below without full justification. In a PLL design for resonant-filter tuning in a switched-mode power supply (Sec. 3.5.1), several designs were implemented without good prior knowledge of noise power levels. PLL designs with large lock ranges and inputs with large ringing noise voltages did indeed lock out, but *ad hoc* narrowing of the lock range solved the problem satisfactorily. These empirical design changes required the substitution of one capacitor and two resistors.

- Step 1** **Determine the center angular frequency ω_c of the PLL.** If this is not possible, specify some bounds ω_c^{\min} and ω_c^{\max} between which the PLL center frequency must lie.
- Step 2** **Choose the damping factor ζ .** Some amount of frequency-domain peaking is desirable to accommodate a fast rise time of the PLL phase-step response. $\zeta = 0.7$ is a good choice for most applications.
- Step 3** **Specify the lock range $\Delta\omega_L$.** This number should be larger than the largest shift in input frequency which the PLL should track using a lock-in process
- Step 4** **Determine the VCO range.** Representing the minimum and maximum VCO output frequencies as ω_0^{\min} and ω_0^{\max} , respectively, sensible choices might be $\omega_0^{\max} = \omega_c^{\max} + 1.5\Delta\omega_L$ and $\omega_0^{\min} = \omega_c^{\min} - 1.5\Delta\omega_L$.
- Step 5** **Determine the VCO gain K_o .** Typical integrated implementations (see Sec. 3.5.1) accommodate control voltages v_c between specified bounds v_c^{\min} and v_c^{\max} , where “max” and “min” refer to the values of v_c corresponding to the largest and smallest ω_o . Integrated VCOs also usually have a constant frequency/V slope as long as $v_c^{\min} \leq v_c \leq v_c^{\max}$. In such a case the VCO gain is

$$K_o = \frac{\omega_0^{\max} - \omega_0^{\min}}{v_c^{\max} - v_c^{\min}}$$

-
- Step 6** **Determine the PD gain K_d** as defined in Sections 2.1.1 and 2.1.2. K_d is often a function of input-signal magnitude, so a range of values may result. Some IC multipliers are packaged with gain networks (e.g. the AD633 in Sec. 3.5.1) that can contribute to K_d . Choose a large K_d , if possible, to improve tracking performance and support the assumption of high total gain K used in the analysis of Sections 2.2 and 2.3
- Step 7** **Determine the PLL natural frequency** from Eqn. 2.3. Check the high-gain and pole-zero-splitting assumptions of Sec. 2.3.1 if you plan on using a passive filter in Step 8.
- Step 8** **Select the loop filter topology** and, possibly, the DC gain K_a of the filter. Use Eqns. 2.2, the PLL component gains, ω_n , and ζ to solve for the time constants of the filter. The active PI filter has proven, from experience, to be a good choice with better performance than the lag designs.
- Step 9** **Choose R and C values for the loop filter** using the expressions in Fig. 2.7 as a guide. The component choice is under-constrained, and a common approach is to specify the capacitances first because of restrictions on readily available values. Excessively large or small resistances could result in driving or loading problems, and high-impedance nodes should always be avoided in power-electronics control circuits (viz. the op-amp summing nodes in Fig. 2.7.)

2.4.2 PLL design when noise must be considered

The $\Delta\omega_L$ specified in step 3, Sec. 2.4.1, may not be possible in a second-order loop. Choose instead the loop noise bandwidth B_L that yields a SNR_L greater than 4, as shown in Eqn. 2.4 in Sec. 2.3.3. Proceed to step 7. Calculate the natural frequency from Eqn. 2.6:

$$\omega_n = \frac{2B_L}{\zeta + \frac{1}{4\zeta}}$$

Eqn. 2.3 can now be used to estimate the lock range achievable with the allowable loop noise bandwidth.

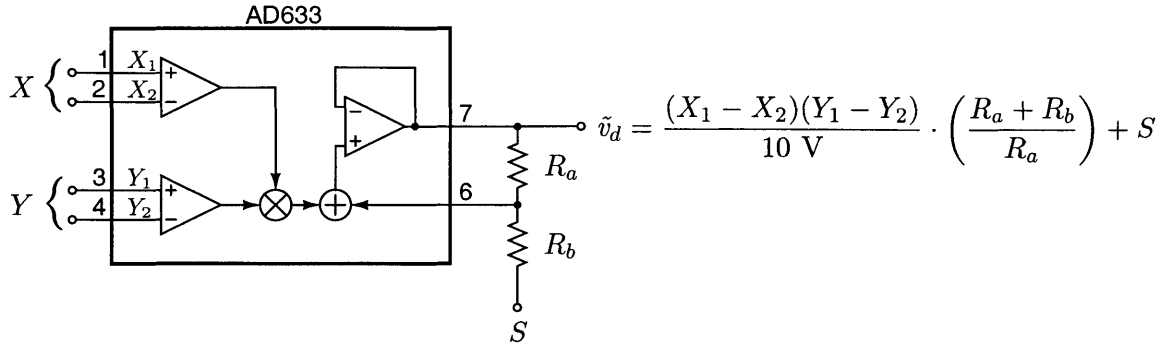


Figure 2.16: Use of the AD633JN with variable scale factor




Figure 2.17: The frequency-control network of the Exar XR2206 monolithic function generator. The XR2206 generates sinusoids with a frequency dependent on the effective resistance seen at its TR1 pin, $f_{out} = 1/(CR_{eff})$, where $R_{eff} = R \cdot (I_t - I_c)/I_t$.

2.4.3 Design example

In the following design example, a PLL must generate a replica — with as little steady-state phase error as possible — of a resonator-voltage fundamental frequency between 110 and 150 kHz. The tuning range represents the assumed variation in the resonant point of a series-tuned filter with a nominal $f_{res} = 130$ kHz. The VCO and PD devices to be used in the PLL are the Exar XR2206 function generator IC and the Analog Devices AD633, a low-cost analog multiplier with gain.

Step 1 Determine the center angular frequency ω_c of the PLL.

$$\omega_c = \frac{2\pi(160 \times 10^3 - 100 \times 10^3)}{2} = 816 \text{ krad/s}$$

Step 2 **Choose the damping factor ζ .** As explained in Sections 2.4 and 2.3.3, $\zeta = 0.7$ is a sensible choice.

Step 3 **Specify the lock range $\Delta\omega_L$.** At start-up, the output of the XR2206 is expected to sit in the center of its tuning range, which will be chosen to correspond to the PLL center frequency ω_c . The PLL will then only be expected to deal with 20 kHz frequency deviations $\Delta\omega_o = \omega_o - \omega_i$. More conservative choices for lock range, e.g. ± 40 kHz to cover a step between extremes in the tuning range, resulted in designs with poor lock-in. The tuning-point of the resonant filter, however, depends largely on manufacturing tolerances which vary unit to unit, but would never suffer a step change in a particular filter. The conservative ± 40 kHz design, then, is unrealistically cautious.

Step 4 **Determine the VCO range.** Following the notation of Sec. 2.4

$$\omega_c^{\min} = \omega_c^{\max} = 816 \text{ krad/s}$$

We can ensure VCO coverage of the tuning range by selecting

$$\omega_0^{\max} = \omega_c^{\max} + 1.5\Delta\omega_L = 816 + 1.5 \cdot 2\pi \cdot 20 \times 10^3 = 1007 \text{ krad/s}$$

$$\omega_0^{\min} = \omega_c^{\min} - 1.5\Delta\omega_L = 816 - 1.5 \cdot 2\pi \cdot 20 \times 10^3 = 627 \text{ krad/s}$$

Step 5 **Determine the VCO gain K_o .** The XR2206 can accept frequency commands into its base-frequency-adjust network anywhere between the power-supply rails.

$$v_c^{\min} = 15\text{V} \quad \text{and} \quad v_c^{\max} = -15\text{V}$$

$$K_o = \frac{\omega_0^{\max} - \omega_0^{\min}}{v_c^{\max} - v_c^{\min}} = \frac{(1007 - 627) \times 10^3 \text{ krad/s}}{-15 - 15 \text{ V}} = -12.6 \frac{\text{krad}}{\text{V}\cdot\text{s}}$$

See Fig. 2.17. This VCO range was implemented with $R_c = 75\text{k}\Omega$ and a fixed/variable combination R_d tunable around $4\text{k}\Omega$.

Step 6 **Determine the PD gain K_d .** $K_d = \frac{1}{2} \cdot K_m V_i V_o$ depends on the input-signal amplitudes and the gain K_m associated with the multiplier itself (see Sec.2.1.2). With signal amplitudes of 2.8 V and 7.2 V for v_i and v_o and the maximum recommended resistor ratio used with the AD633 (Fig. 2.16), $K_d \approx 25\text{V/rad}$. K_m was approximately ten with the choice of resistors $R_a = 1\text{k}\Omega$ and $R_b = 100\text{k}\Omega$.

Step 7 **Determine the PLL natural frequency.** For the high SNR case,

$$\omega_n = \frac{\Delta\omega_L}{\zeta} = 90 \text{ krad/s}$$

Step 8 **Select the loop filter topology.** An active PI filter with a high-frequency pole at ω_3 (Fig. 2.7d) was chosen because of superior noise and tracking performance. The time constants needed for the control design are

$$\tau_1 = \frac{K_o K_d}{\omega_n^2} \quad \text{and} \quad \tau_2 = \frac{2\zeta}{\omega_n}$$

Step 9 **Choose R and C values for the loop filter.** Neglecting the value of C_1 in comparison with C_2 ,

$$\tau_1 = R_1 C_2 = 28.9 \text{ ms} \quad \text{and} \quad \tau_2 = R_2 C_2 = 15.6 \text{ ms}$$

Choosing $C_2 = 0.22 \mu\text{F}$, the largest signal-level capacitor on hand, resulted in $R_1 = 177 \text{ k}\Omega$ and $R_2 = 70.9 \text{ k}\Omega$. 180 k Ω and 68 k Ω 5% resistors were selected for the final circuit (Fig. 2.18)

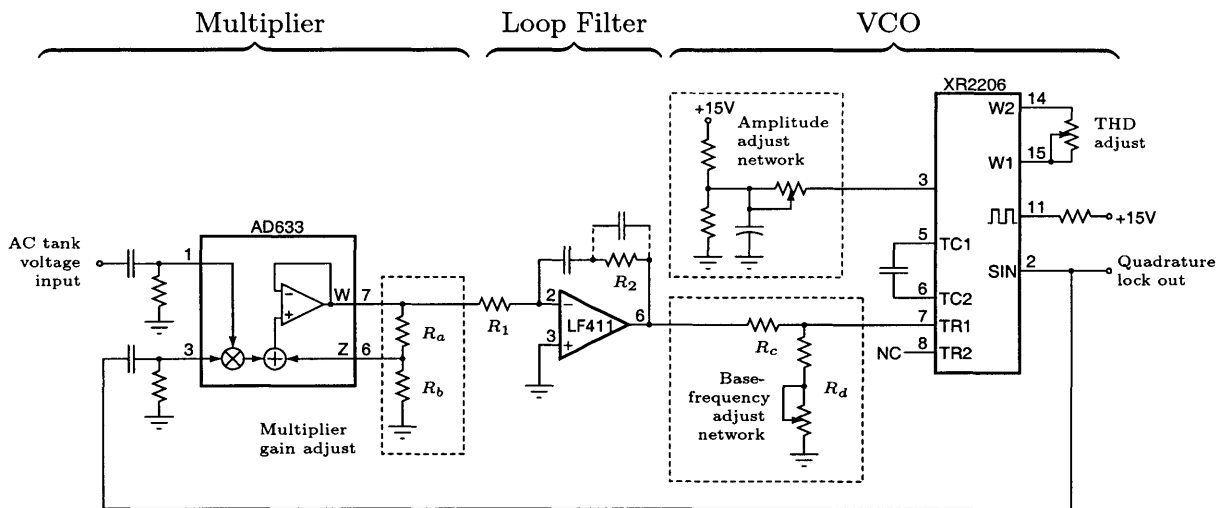


Figure 2.18: Schematic of the PLL used in the prototype tuning system. See the schematics and accompanying tables in Appendix B for component values.

Resonant-network design

“**R**ESONANT networks” or “resonant filters” include, in this discussion, any passive network which expressly incorporates resonant branch impedances. The self-resonant frequencies of passive elements must, as always, be taken into account when predicting the EMI performance of a resonant filter, but unless self-resonance is introduced to achieve filtering — rather than accepted as a parasitic — a filter is not resonant in the sense considered here. Three resonator topologies (Fig. 3.1) were considered for use with the phase-lock tuning system to accompany conventional low-pass networks. Note that the magnetically coupled resonator of Fig. 3.1c can be treated in terms of its equivalent T-model in Fig. 3.1d, so that each resonator design reduces to the series- or parallel-tuned case.

For power-stage or ripple-filtering applications in fixed-frequency power converters, resonant networks can provide small AC impedances shunted across (Fig. 3.2a), or large AC impedances in series with (Fig. 3.2b), a load or source. The resonant transmission null introduced by the resulting impedance mismatch reduces the attenuation requirement of an accompanying low-pass network, pushing its L-C corner ω_c to a higher frequency. Resonator design is complicated by this accompanying low-pass filter, which introduces trade-offs in attenuation performance, volume, reactance, and resonator frequency selectivity. Application constraints and component ratings, moreover, can bear critically on the choice of




Figure 3.1: (a) The series-tuned shunt resonator, (b) parallel-tuned series resonator, (c) magnetically coupled series resonator, with (d) its equivalent T model.




Figure 3.2: (a) A series-resonant input filter for a buck converter. The series-tuned leg provides a low-impedance current path (i.e. high attenuation) at a discrete frequency. (b) A parallel-resonant power stage presents a high impedance to ripple current at its tuning frequency. (c) Transfer function of switch drain current to input current or switch source voltage to output voltage.

component values, often favoring networks that make feasible the use of a particular type of element (viz. a small-valued, but reliable capacitor, see [16]). Furthermore, because resonators exhibit immittance peaking at discrete frequencies, the resonant-filter designer must usually account for performance changes over a range of anticipated switching waveforms.

The purpose of this chapter is to clarify the principle trade-offs of resonant network design. Section 3.1 will consider constraints common to resonator designs, regardless of topology, including the dependence of performance on conversion ratio, the introduction of anti-resonances, and the component limitations imposed by resonant circulating currents and ringing voltages. Sections 5.2.2–3.3 will present the optimization problems encountered in the design of particular resonant topologies, indicating the principle trade-offs of each design and suggesting outcomes of a search for an optimal network. The discussion of resonator optimization is followed by two parallel-tuned design examples in Sections 3.5.1 and 3.5.2, one for the low-ripple and one for the high-ripple case. For inductor-heavy resonators with high ripple currents, iterative network/magnetics co-design may be necessary to select filter components.




Figure 3.3: (a) Q shown as the impedance peaking of the parallel tuned resonator (b). Z_0 is the characteristic impedance, the reactance $X(j\omega_{res}) = \sqrt{L/C}$ of the capacitor or inductor at resonance.

3.1 Constraints on resonant network design

The three resonator topologies of Fig. 3.1 share several design constraints, two of which address the resonator's incorporation in a low-pass network (Secs. 3.1.1 and 3.1.2), and two which pertain to the resonator itself (Secs. 3.1.3 and 3.1.4).

3.1.1 Impedance constraints: Quality factor Q and characteristic impedance Z_0

The many different, but equivalent, definitions of Q can be summarized by the expression

$$Q = \omega \cdot \frac{\text{energy stored}}{\text{average power dissipated}}$$

i.e., Q is proportional to the ratio of energy stored to energy lost, per unit time. In a low-loss electrical resonator, the peak energy stored in either the inductor or capacitor equals the total energy stored in the network at any given time. For a parallel-tuned network (Fig. 3.3) driven at resonance by a current $i_{in} = I_{pk} \cos(\omega_{res}t)$, the network impedance is purely resistive at resonance and the peak tank voltage is $I_{pk}R$. The total resonator energy is then just the capacitor energy when the voltage is maximum,

$$E_{tot} = \frac{1}{2}C(I_{pk}R)^2$$




Figure 3.4: (a) Q raised by an increase in $X_L(j\omega_{res}) = X_C(j\omega_{res})$ outstripping an increase in ESR. (b) locus of inductive impedances (for a fixed core geometry) corresponding to the situation in (a)

The average power dissipated at resonance can likewise be expressed simply in terms of I_{pk} :

$$P_{avg} = \overline{i_{in}^2} R = I_{pk} R \cdot \overline{\cos^2 \omega_{res} t} = \frac{1}{2} I_{pk}^2 R$$

The Q of the parallel-tuned network at resonance is the ratio

$$Q = \omega_0 \frac{E_{tot}}{P_{avg}} = \frac{1}{\sqrt{LC}} \frac{\frac{1}{2} C (I_{pk} R)^2}{\frac{1}{2} I_{pk}^2 R} = \frac{R}{\sqrt{L/C}}$$

The quantity $\sqrt{L/C}$ has the dimensions of resistance, and is the resonator's *characteristic impedance* Z_0 . Z_0 is the magnitude of the inductive and capacitive reactances at resonance, as can be verified by evaluating $|Z_C|$ and $|Z_L|$ at $\omega_{res} = 1/\sqrt{LC}$:

$$|Z_C| = \left| \frac{\sqrt{LC}}{jC} \right| = \sqrt{\frac{L}{C}} \quad |Z_L| = \left| \frac{jL}{\sqrt{LC}} \right| = \sqrt{\frac{L}{C}}$$

As shown in Fig. 3.3, Z_0 is the point from which resonant immittance peaking (also equal to Q) is referred. The Q of a series-tuned network (Fig. 3.1a) is the reciprocal of that derived for the parallel-tuned case, $Q = Z_0/R$, where the peak in series-tuned admittance is now bounded from above by the series resistance R .

For resonators constructed from practical components, high frequency-selectivity — high Q — is not necessarily the criterion of good performance. A resonator with infinite Q *might* develop an unbounded impedance mismatch and provide perfect filtering at its tuning




Figure 3.5: The introduction of an antiresonance at the frequency where the rising, predominately inductive impedance Z_{res} reacts with the falling capacitive impedance Z_{cap} .

frequency, *but only if its tuning-point impedance supports the mismatch*. Consider the series-tuned resonator impedance magnitudes of Fig. 3.4a. Q is a ratio which indicates nothing about absolute impedance magnitudes: the resonator with $Q = 150$ has greater frequency selectivity, but only because its characteristic impedance, compared to the $Q = 10$ design, has increased more than its ESR. The external network impedances seen at the resonator terminals relative to the resonator’s tuning-point impedance (a function of Z_0 and Q), determines the filtering effectiveness of the resonator, which is *poorer* for the higher- Q resonator in Fig. 3.4.

3.1.2 Harmonic constraints: Antiresonance

The similarity in the ripple-attenuation transfer functions of Fig. 3.2c — the fact that different resonant filters can have identical transfer functions — suggests that resonant-filter attenuation can be parameterized in terms of the relative depth of a transmission null and the low-pass filter’s high-frequency roll-off. In fact, the number significant parameters is higher because of the introduction of *antiresonance*.¹ Consider the series-tuned shunt

¹Note that in network theory many authors refer to any maximum-impedance resonance as an antiresonance. In this discussion, resonance refers to any series- or parallel-resonant tuning point which diverts ripple waveforms away from a port of interest. Antiresonance denotes any reactance-cancellation condition which worsens attenuation.




Figure 3.6: Achievable increase in corner frequency ω_c of a low-pass filter accompanying a resonant network, plotted versus resonator loss and Q. Peak-peak ripple performance at $D = 0.3$ limits how high the corner frequency can be increased in each case. The locus of largest reactive energy storage gain is indicated by the solid line.

resonator of Fig. 3.5b in parallel with the capacitor of a reduced low-pass network. The total shunt impedance magnitude (Fig. 3.5a) exhibits a series-tuned resonance (the minimum of $|Z_{total}|$) and an antiresonance (the maximum of $|Z_{total}|$), where the falling capacitive impedance reacts with the predominately inductive resonator impedance above ω_{res} .

The antiresonance can be close to the second harmonic of a converter's switching frequency, and may be pronounced enough to reverse, for a peak-peak ripple measure, performance gains from resonant attenuation of the fundamental. Figs. 3.6 and 3.7 detail, for normalized cases, the loci of resonator loss and Q which provide the greatest decrease in reactive energy storage over a conventional low-pass network, with no decrease in ripple performance and with the effects of antiresonance included. The surface of Fig. 3.6 is the result of phasor analysis of many resonant/low-pass networks parameterized in terms of resonator tuning-point impedance and Q. The accompanying low-pass shunt impedance or series admittance required to match the performance of a particular low-pass network at




Figure 3.7: Loci of maximum reactive-energy-storage improvement (cf. the solid line of Fig. 3.6) for duty ratios $D = 0.2, 0.3, 0.4,$ and 0.5 .

a specific duty ratio is expressed as a shift in ω_c , the low-pass corner frequency. The low-pass network is the basis for immittance normalization (Fig. 3.9): its shunt impedance or series admittance at the resonator tuning point (for the series- and parallel-tuned cases, respectively) is defined as 1Ω or 1 S . The resonator tuning point, likewise, normalizes the frequency axis. All immittances used in the normalized design, though they refer to either shunt or series elements, bear a one-to-one correspondence to the notch depth and high-frequency roll-off of the complete filter-network transfer function.

The surface of Fig. 3.6 shows many predictable features. For low Q and low loss — the large elevation at the furthest corner of the surface of Fig. 3.6 — the resonator will have a low characteristic immittance and a broad attenuation characteristic. The low-pass corner can only be increased slightly before bringing antiresonance close to the second harmonic. For cases with high resonator loss and moderate Q (> 8 , as shown in Fig. 3.6), the antiresonance will not hurt performance because the accompanying low-pass network dominates the filter immittance around the second harmonic. But even for the case of high Q (the near corner in Fig. 3.6), the designer cannot realize much gain in ω_c because the resonator provides little extra attenuation at its tuning point. The ideal design locus is indicated by the bold line in Fig. 3.6, representing designs which achieve the best balance between the breadth and depth of the transmission null.

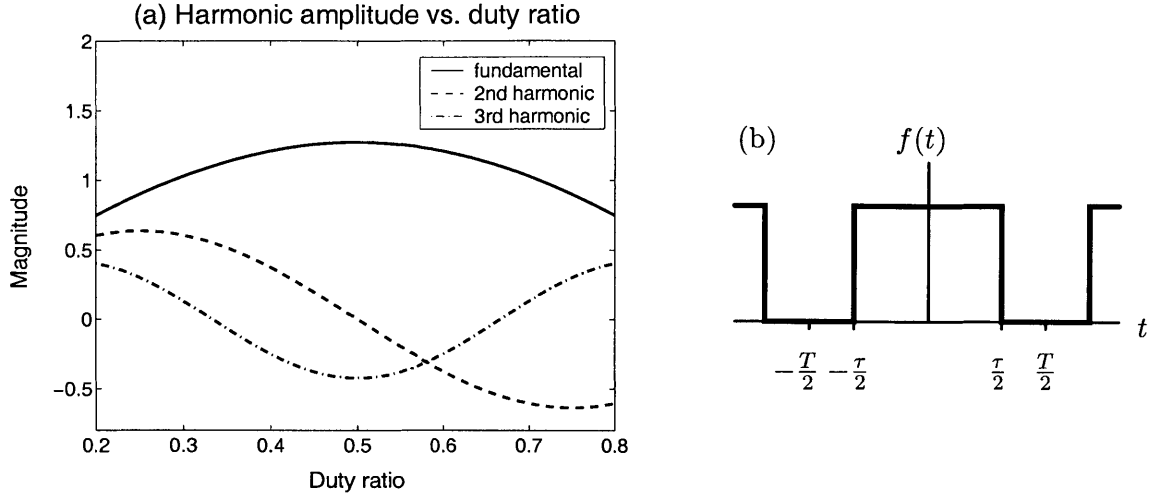


Figure 3.8: (a) Normalized harmonic amplitudes for the variable-width pulse train of (b).

3.1.3 Duty-ratio constraints

Resonant attenuation is naturally more effective for waveforms with a concentration of power at a resonator's imittance peak, and so depends on the harmonic content of current and voltage waveforms produced by the converter switching cell. For fixed-frequency converters under duty-ratio control, switch drain currents and source voltages (cf. Fig. 3.2) are periodic gate functions like that plotted in Fig. 3.8b . Given a converter switching frequency $\omega_{sw} = 2\pi/T$ and duty ratio $D = \tau/T$, the harmonic content of the switch waveforms (Fig. 3.8a) can be expressed as an exponential or cosine Fourier series:

$$\begin{aligned}
 f(t) &= \sum_n F_n e^{jn\omega_{sw}t} & F_{n \neq 0} &= \frac{\tau}{T} \cdot \frac{\sin(\frac{n\pi\tau}{T})}{\frac{n\pi\tau}{T}} = D \cdot \text{sinc}(n\pi D) \\
 &\Updownarrow \\
 f(t) &= \sum_{n=0}^{\infty} c_n \cos(n\omega_{sw}t + \phi_n) & c_{n \neq 0} &= 2|F_n| \quad \phi_n = \tan^{-1} \frac{\Im\{F_n\}}{\Re\{F_n\}} = 0
 \end{aligned}$$

For the case in which a converter's switching (viz. fundamental) frequency is aligned with a filter resonance, the maxima of Fig. 3.8a indicate that resonant attenuation is most effective for duty ratios centered around $D = 0.5$.

3.1.4 Component-rating constraints

The AC branch currents and internal node voltages of the parallel- and series-tuned resonators, respectively, differ substantially from the ripple seen at the resonator terminals. For the parallel-tuned case driven at resonance, as already mentioned, the AC voltage across the network is $Ri_{in} = RI_{pk} \cos(\omega_{sw}t)$. $|Z_C|$ and $|Z_L|$ are equal at ω_{res} , so their branch currents must be equal in magnitude:

$$|i_L| = |i_C| = |i_{in}| \frac{R}{Z_0} = Q|i_{in}|$$

I.e. the tank circulating currents are Q times higher than the AC currents at the resonator terminals. A similar result holds for the internal node voltage of the series resonator, where the common current through equal tuning-point-reactance magnitudes can develop large AC voltages:

$$|v_L| = |v_C| = |v_{in}| \frac{Z_0}{R} = Q|v_{in}|$$

In resonant networks developed for a 300 W DC-DC down converter with an 18–60 V input, Q values of 20 have been achieved. Ripple currents of 1 A peak-peak and 2 V peak AC voltages were typical in this application when conventional passive networks were employed. The 20 A circulating currents and 40 V internal node swings in such a case may either determine resonator component ratings, or force the designer to consider lower-ripple networks (e.g., deep continuous-conduction mode) to maintain high resonator Q .

3.2 Resonator design

The selection of resonator components, considered rigorously, is a volume optimization problem subject to the constraints of Sec. 3.1. The discussion here will detail resonator design in a normalized setting, treating parallel- and series-tuned networks in a common framework. As mentioned in Sec. 3.1.2 and depicted in Fig. 3.2, series- and parallel-tuned designs can have identical ripple-attenuation transfer functions. Generalized resonator design amounts to a parameterization of this filtering action in terms of the depth and breadth of the resonant transmission null. In such generalized coordinates, the resonators which decrease the volume of the accompanying low-pass network the most (i.e., those designs which, allowing for the antiresonance of Sec. 3.1.2, permit the greatest increase in ω_c for constant ripple performance at some minimum duty ratio) are described by the lines shown in Fig. 3.7. By considering the intersection of these loci with a locus of practical resonators, the designer




Figure 3.9: Normalized immitances for the generalized resonator/low-pass design. Resonator Q and normalized resonator loss bear a simple correspondence to the depth and breadth of the transmission null of Fig. 3.2c. Note that the “depth” shown here is not the resonator Q , but a measure of the resonator’s immitance excursion beyond the normalizing low-pass network’s shunt impedance or series admittance.

can quickly narrow the search for a suitable network. See Sec. 3.5.1 for an example of normalized design in the parallel-tuned case.

The “locus of practical resonators” mentioned above describes, for some limit on materials and total volume, a “ragged half-plane” in coordinates of Q and tuning point loss. That is, for some assortment of components, the highest achievable resonator Q is a *single-valued* — but perhaps not well-known — function of loss. Lower- Q resonators at constant tuning-point immitance (those that lie away from the border of the half-plane) are realizable by exchanging inductance for capacitance.² An “optimal” resonant-filter design resides where the maximum- Q locus (i.e., the set of resonators for which $Q = Q_{max}$ as a function of loss) intersects the minimum-volume locus of Fig. 3.6 in normalized coordinates.

²For practical purposes — in the absence of application restrictions — such an exchange can be made with no increase in loss, and in general will decrease tuning-point immitance. To maintain constant loss in such a case, of course, resistance can be introduced into the network.

³. The designer is free, of course, to shift the function Q_{max} by altering the constraints in its computation, e.g., by adjusting the volume or type of the resonator components.

3.2.1 Parallel-tuned series resonator

For the parallel-tuned resonator in series with the ripple source, ripple attenuation increases with Q because the resonator's tuning-point impedance increases with selection of a larger inductance ($Q = \omega L/R$). Larger resonant inductances can develop greater impedance mismatches since the resonator characteristic impedance Z_0 increases with inductance ($Z_0 = \sqrt{L/C}$), and larger inductances are themselves generally realizable with larger Q .⁴ Designing for the highest possible Q results in a narrowband resonator that will generally *require* tuning to excite reliably. Because the attenuation performance of a resonator network can always be improved by increasing inductance, the task of design — in the absence of application constraints — reduces to a trade-off between volume and performance.

Sections 3.5.1 and 3.5.2 detail the design of two parallel-tuned series resonators, one for use in an input filter (low ripple-current case) and one in a switching stage (high ripple-current). An *a priori* limit on total inductance or core volume is used in these examples to balance attenuating performance and filter size. In both cases, filter volume and mass savings of approximately three times are realized.

Fig. 3.10a shows, schematically, the parallel-tuned design in normalized coordinates. Given a selection of resonator components with “small” total volume (about $6\times$ less than the passive elements of the low-pass network), the achievable combinations of resonator Q and loss are depicted by the shaded half plane. The best design — the network which achieves high fundamental attenuation with its large tuning-point impedance, avoids accentuation of the second harmonic, *and* maximally reduces the size of its accompanying low-pass network — resides on the performance locus (solid line) at the high- Q edge of the region of realizable inductors.

³Note from Fig. 3.6, that increases in ω_c always accompany increases in Q .

⁴For a given core geometry, losses increase roughly as turns and inductance as turns squared, so larger inductances are achievable with (relatively) lower loss. See Fig. 3.4b. A quick perusal of the Q curves in [17] will convince the reader that in the more general case of minimal magnetic dimensions, large inductances are necessary to achieve high Q , particularly at switching frequencies.

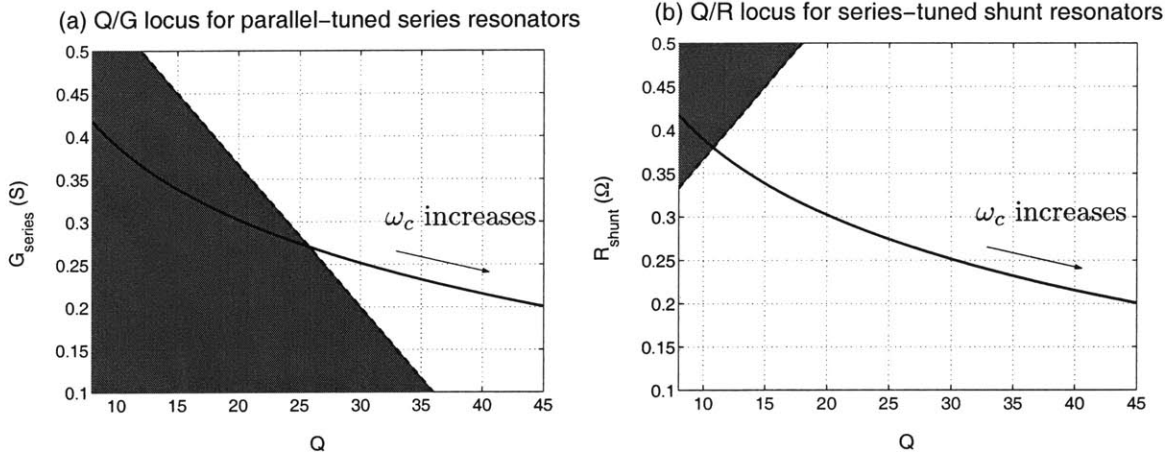


Figure 3.10: (a) An “ideal” series-resonator design resides at the intersection of the solid performance curve with the dashed line, schematically indicating the locus of maximum achievable resonator Q as a function of normalized loss. (b) For appreciable volume savings in the shunt-resonator case, low Q and little increase in ω_c were achievable.

3.2.2 Series-tuned shunt resonator

Shunt networks divert ripple current by presenting low AC impedance at the switching frequency and its harmonics. The impedance magnitude of a series-tuned network at its resonant point, not its Q , is therefore the metric of resonator performance. Considering the locus of practical inductive impedances (Fig. 3.4b), note that increased inductive reactance is always accompanied — for constant magnetic geometry — by an increase in ESR. The lowest-loss resonator is the network with the smallest inductance, given some limit on the total resonator volume. A low tuning-point impedance is achieved by spoiling Q , not in the sense that loss is increased, but that by designing for low ESR, the designer is driven to low characteristic impedance Z_0 and hence low $Q = Z_0/R$.

From these considerations, the best shunt resonator is the one with the lowest possible characteristic impedance: a capacitor that is self-resonant at the ripple frequency of interest. Where a discrete inductor must be added to the resonator design, this performance trend is fundamentally opposed to the need for tuning, as the design with the lowest tuning-point impedance will have a low characteristic impedance and a low series-tuned Q . Shunt-resonator design therefore amounts to a complicated trade-off of volume, performance, and Q . By retreating from a large resonant capacitor and high impedance mismatch, the designer must search for low-volume resonators with acceptable attenuation *and* a Q large enough to make phase-lock tuning attractive.

In a search for shunt-resonant components for a 300 W, 125 kHz converter, many combinations of low-pass and resonator capacitors were considered. Given the losses and volume of these elements, the resonant-inductor value, volume, and loss combinations that would make a shunt-resonator design attractive were computed. No suitable designs were found, because large resonator Q ($Q > 6$, large enough perhaps to make locking desirable) required large volume at the 300 W power level. The capacitors considered for use in the series-tuned network were self-resonant in the 500 kHz–1 MHz range, and therefore could serve as effective shunt resonators for converters operating in this frequency range.

A normalized, schematic depiction of the series-tuned problem (at 300 W and 125 kHz) is shown in Fig. 3.10b. The Q /loss locus for resonators with volume savings comparable to that of the parallel-tuned case ($\approx 3\times$) intersects the volume-savings curve at low Q and low normalized ω_c . The search for an attractive design amounts to the selection for resonator components which can support higher Q and lower loss for as little additional volume as possible (retreating from the initial aim of $3\times$). In the parallel-tuned case, significant volume savings were achieved using readily available components, and possible improvements were never pursued: the border of the resonator locus was never determined and no attempt was made to find the best possible resonator. In the series-tuned case, “casual” designs met with failure repeatedly, and the search for lowest possible tuning-point loss, conjoined with substantial Q and low volume, required a wide-ranging search over component types and values.

The search for a high- Q shunt network with low tuning-point impedance and low loss can also be reduced to the problem of shifting the self-resonance of a capacitor downward in frequency (Fig. 3.11). A suitable design must preserve Q by minimizing added loss, and provide enough inductance L_{shift} to move the series-tuned resonance to a useful frequency. No suitable shunt network was found for the 300 W, 125 kHz converter using capacitors naturally resonant around 1 MHz: one must go to extreme lengths to design a low-value inductance (several tens of nH) that limits the increase in tuning-point impedance to a decade for each decade shift in resonance (note the dashed ESR line in Fig. 3.11a). In several experiments, large-volume air-core Litz inductors were required to implement the inductance L_{shift} with sufficiently low loss. These inductors were impractically large, and conventional capacitive shunt elements could always provide comparable ripple performance in the same volume.

Note that if application constraints require the designer to use small shunt capacitances, series-tuned resonators can be the best shunt-element alternative. High-power inverters, for instance, make exclusive use of series-tuned harmonic traps for ripple filter-

ing. In such cases, if the designer is willing to expend volume to achieve low tuning-point loss (hence high Q, for small capacitance), phase-lock tuning may be attractive.

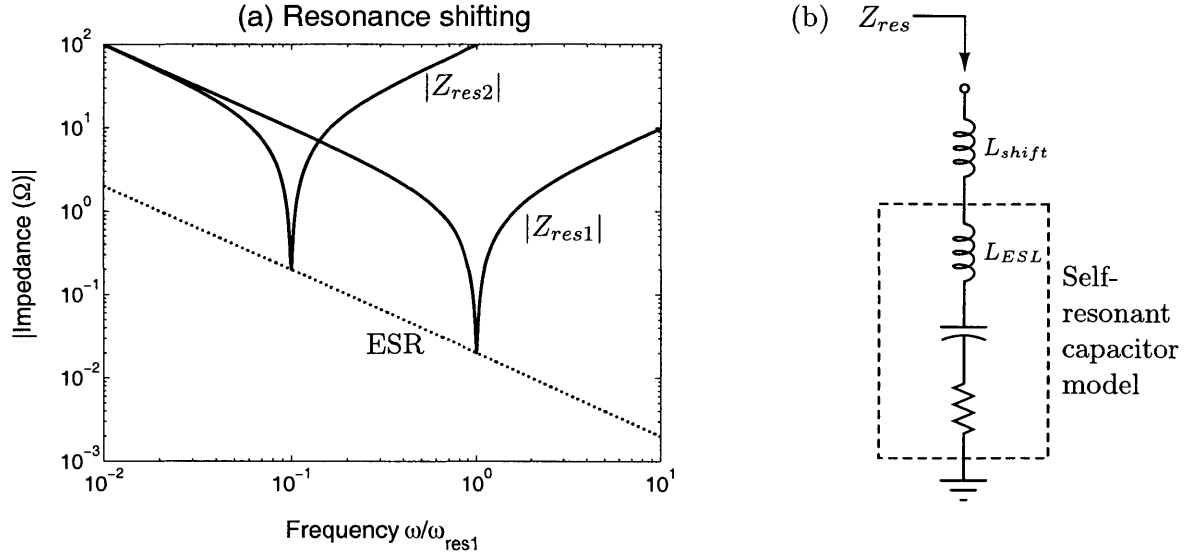


Figure 3.11: (a) Impedance magnitudes for a capacitor before ($|Z_{res1}|$) and after ($|Z_{res2}|$) a shift in its resonant frequency, controlled by adding a small inductance L_{shift} in series with the capacitor.

3.3 Magnetically coupled shunt resonator

Magnetically coupled series resonators (Fig. 3.12a) were originally considered for use with phase-lock tuning because of their decreased sensitivity to resonator loss compared to shunt designs. Consider the equivalent T-model of Fig. 3.12c. The overall ripple performance is less sensitive to the resonant-element loss in the series-tuned leg because of the additional series AC impedance seen from either port. I.e., large impedance mismatches are easier to achieve with relaxed requirements on the resonant loss, compared to the simple series-tuned shunt case.

Consider the output filter for a converter shown in Fig. 3.12b. The converter is represented by a voltage source v_c with output resistance r_c . The resonant “AC leg” of the transformer comprises C_r , the winding with self-inductance L_{ac} , and r , which models the total loss of C_r and L_{ac} . The transfer function between converter output voltage v_c the ac-winding voltage v_{ac} is

$$\frac{V_{ac}(s)}{V_c(s)} = \frac{s^2 L_{ac} C_r}{s^2 L_{ac} C_r + s r C_r + 1} \quad (3.1)$$




Figure 3.12: (a) Self- and mutual-inductances of the magnetically coupled resonator. (b) Use of the resonant structure in the output filter for a converter. (c) Equivalent T-model for the transformer.

the bode plot of for which is shown in Fig. 3.13. For frequencies considerably higher than the low-pass corner $\omega_0 = 1/\sqrt{L_p C_r}$, θ is close to 0° when r is negligible compared to the principally inductive resonator reactance $X_{ac} \approx \omega L_{ac}$. I.e. in the limit of high Q, the majority of the AC voltage drop applied at the transformer common node drops *in phase* across the AC winding. Since the phase shift between the primary and secondary windings of the transformer is ideally zero, all three phasor voltages V_c , V_{ac} , and V_{dc} are in phase at for $\omega \gg \omega_0$. The load voltage phasor V_L can then be expressed as the algebraic difference between V_c and V_{dc} such that

$$\frac{V_L}{V_c} = 1 - \frac{V_{dc}}{V_c}$$

Denoting the turns ratio $n = N_{ac}/N_{dc} = V_{ac}/V_{dc}$,

$$\frac{V_L}{V_c} = 1 - \frac{V_{ac}}{nV_c}$$

V_{ac}/V_c has a magnitude of n at some frequency — call it ω_m , for maximum attenuation — where the load ripple transfer function V_L/V_c is minimized. Assuming low AC-leg loss and equating Eqn. 3.1 (evaluated at ω_n) to n yields

$$\omega_n = \frac{1}{\sqrt{LC}} \sqrt{\frac{n}{n-1}} = \omega_0 \sqrt{\frac{n}{n-1}}$$

$|V_L/V_c|$ is plotted in Fig. 3.14, with the principle attenuation corner ω_0 and attenuation

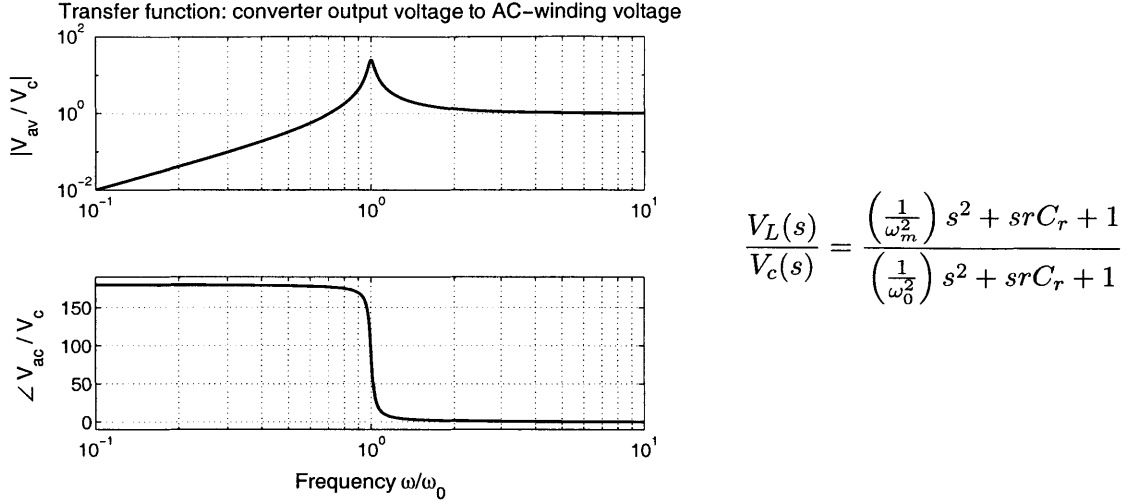


Figure 3.13: Voltage transfer function from the converter voltage v_c to the voltage v_{ac} across the transformer AC winding, in the limit of low resonator loss ($r \rightarrow 0$)

maximum ω_m indicated. The ripple attenuation function can be written

$$\frac{V_L(s)}{V_c(s)} = \frac{\left(\frac{1}{\omega_m^2}\right) s^2 + srC_r + 1}{\left(\frac{1}{\omega_0^2}\right) s^2 + srC_r + 1}$$

or, in terms of the reactive components of the T model of Fig. 3.12c,

$$\frac{V_L(s)}{V_c(s)} = \frac{s^2(L_{ac} - M)C_r + srC_r + 1}{s^2L_{ac}C_r + srC_r + 1}$$

from which $\omega_m = 1/\sqrt{(L_{ac} - M)C_r}$ is clearly seen to be the series-resonant frequency of the shunt-path resonator in Fig. 3.12c. In the absence of parasitics, the high-frequency ripple attenuation for $\omega \gg \omega_m$ is a constant:

$$\lim_{\omega \rightarrow \infty} \left| \frac{V_L(s)}{V_c(s)} \right| = \frac{\omega_0^2}{\omega_m^2} = \frac{n-1}{n} \quad (3.2)$$

The maximum attenuation is given by

$$\left| \frac{V_L(j\omega_m)}{V_c(j\omega_m)} \right| = \frac{\frac{1}{Q}}{\sqrt{\left(1 + \frac{\omega_m}{\omega_0}\right)^2 + \left(\frac{1}{Q}\right)^2}}$$

This expression summarizes the “promise” of the magnetically coupled shunt resonator, i.e. that attenuation increases with Q, unlike the simple shunt case. Here, it would seem, is




Figure 3.14: Parameterized transfer function V_L/V_c , showing voltage attenuation of the magnetically coupled shunt resonator. ω_ℓ and ω_h are derived in [16] as a means of placing restrictions on the principle attenuation corner ω_0 given some attenuation requirement h over the frequency range $\omega_\ell < \omega < \omega_h$:

$$\omega_\ell = \omega_0 \sqrt{\frac{(1+h)n}{(1+h)n-1}} \quad \omega_h = \omega_0 \sqrt{\frac{(1-h)n}{(1-h)n-1}}$$

a series-tuned resonator that requires tuning to excite reliably when designed for the best attenuating performance. Because of practical relationships between Q and ω_m , however, this network performs best for *low* Q .

The transformer arrangement of Fig. 3.12a is, in fact, the so-called “zero-ripple” filter.⁵ A transformer turns ratio n just large enough to make up for imperfect coupling injects the inverse of voltages which drive ripple current toward the load. (E.g., In the case of perfect coupling, $n = 1$ nullifies the ripple at the output port.) The transformer network is an effective ripple filter when its AC leg presents a low impedance to as many frequencies as possible (i.e., for a spoiled resonant-leg Q). This performance trend is fundamentally opposed to the need for tuning, as with the pure shunt case. Because better attenuating performance requires larger volume — ever larger AC-leg capacitance — a complicated trade-off results. Retreating from the best-performing ripple-steering design, the designer can search for networks with low volume and a resonant-leg Q sufficiently high to make locking attractive.

⁵The “zero-ripple,” or ripple-current-steering topology has, interestingly, been rediscovered under various names on about four occasions [18] dating back to 1928 [19]. See also [20]–[23].




Figure 3.15: Voltage attenuation for two magnetically coupled resonators employing the same magnetic elements. The “zero-ripple” design has a turns ratio n close to one, but requires a larger resonant capacitor C_r to bring its maximum attenuation frequency ω_m down to the switching frequency ($f_{sw} \approx 100$ kHz in this example).

In a search conducted for a 300 W buck, 125 kHz converter, no attractive magnetically coupled shunt-resonator designs were found. See Fig. 3.15. With constant magnetic dimensions, better ripple performance could always be achieved by bringing n closer to one, increasing high-frequency attenuation and decreasing the shunt-path inductance $L_{ac} - M$ (see the the T-model in Fig. 3.12). The resulting shunt-leg Q was always spoiled in such a case (note the shallow attenuation of the “zero-ripple” curve in Fig. 3.15). Though a larger capacitance C_r was required to resonate with $L_{ac} - M$ near the switching frequency, the extra capacitor volume necessary was small, given the size of the magnetics necessary to carry the filter’s DC current. As with the shunt-resonator case, however, high AC-leg Q and phase-lock tuning may be natural design choices when application constraints encourage the use of small capacitance.




Figure 3.16: Resonant input filter design for a 300W buck converter. In the original low-pass design, $L'_3 = 70 \mu\text{H}$. In the resonant design, $L'_1 = 15.4 \mu\text{H}$, $L'_2 = 5.4 \mu\text{H}$, and $C_r = 10\mu\text{F}$.

3.4 Summary

Parallel-tuned

Placed in series with a source or load, parallel-tuned networks can typically reduce the volume of magnetics by about a factor of three. Resonator Q and ripple attenuation increase together, so resonator design amounts to a trade-off of performance and volume.

Series-tuned

Low-Q, capacitor-heavy resonators provide the best ripple attenuation, a trend opposed to the need for tuning. A difficult optimization of volume, Q, and attenuation performance will likely be necessary to arrive at a suitable network. Applications restricted to relatively small capacitances, or those which can take advantage of the natural inductance of available capacitors, could benefit from tuned, resonant filters.

Magnetically coupled

As with the series-tuned shunt resonator, low Q and large capacitance provide the best ripple filtering, and a multi-dimensional search may be necessary to find suitable networks. A limit on capacitor values, again, can favor high-Q magnetically coupled shunt resonators.

3.5 Design examples

3.5.1 Design example: low ripple current

Consider the component selection for a resonant input filter, Fig. 3.16, in which the inductors L'_1 and L'_2 replace the π -filter inductance L'_3 . The filter is intended for a converter which must deliver up to 300 W to a 12 V load, and meet a 90 dB μ V (0.632 mA) input ripple specification across 50 Ω for 24 V < V_{in} < 40V.

Designing L'_3 establishes a basis for filter-volume comparison. For adequate hold-up at the MOSFET drain — a typical control requirement — C must be large enough to pass the maximum pulsed drain current with a ripple voltage of, say, 20% of the DC drain voltage at 50% duty cycle.

$$v_{drain} \approx \frac{1}{C} \int_0^{T/2} i_{dc} dt = \frac{I_{out} T}{2C}$$

A 30 μ F film capacitor for C will limit the drain-voltage ripple to 2.08 V p-p. If C' is omitted, L'_3 must limit the ripple in I_{in} to 0.632 mA. Approximating the current through L'_3 as

$$i_{L'_3} = \frac{v_{drain}}{2\pi f_{sw} L'_3}$$

would require $L'_3 = 5.2$ mH, too large to be practical. The addition of $C' = 10$ μ F shunts almost 99.7% of $i_{L'_3}$ away from the 50 Ω source impedance, and increases the allowable ripple current through L'_3 to 0.2 A. L'_3 can now be 70 μ H, a reasonable value given the size of other components in the filter. In this case, L'_3 is the dominant physical element in the system.

The damping leg of Fig. 3.16 comprises C_d and R_d , and prevents ringing at the low-pass corner $\omega_c = 1/\sqrt{L'_2 C}$. C_d is a large-valued electrolytic capacitor with no extraordinary ESL requirements: its impedance magnitude can begin to rise above ω_c , a decade or more below the switching frequency.

The resonant/low-pass inductors L'_1 and L'_2 can be chosen from normalized-design considerations. The fundamental-frequency reactance of L'_3 under full DC bias was predicted from spectral measurements to be about 38 Ω . To match the low-pass performance at $D = 0.3$ with an assumed resonator Q of 30 requires, as seen from the appropriate curve in Fig. 3.7, a normalized tuning-point resonator loss of 0.25 S. This loss corresponds to a




Figure 3.17: Design of the resonant/low-pass (L_1 and L_2) and simple low-pass (L_3) inductors for the power stage of a prototype 300W buck converter

normalized resonant-inductor reactance at 100 kHz of

$$\frac{X_{L'_2}}{X_{L'_3}} = \frac{1}{QG} = \frac{1}{0.25 \cdot 30} = 0.133$$

corresponding to $X_{L'_2} = 5.09 \Omega$, or $L'_2 = 8.10 \mu\text{H}$. L_2 was chosen near this value, $5.4 \mu\text{H}$ to resonate just under 100 kHz with a $0.47 \mu\text{F}$ capacitor. Such changes in inductance are frequently necessary to accommodate available capacitance values. L'_1 was selected by iterative phasor analysis to meet, along with the resonator, the filtering performance of L'_3 at $D = 0.3$. See Sec. 4.2 for details of the cores chosen for L'_1 – L'_3 and measurements of their filtering performance.

3.5.2 Design example: high ripple current

The choice of parallel-tuned resonator components for a power stage (Fig. 3.17) can be more complicated if ripple currents are high. Consider again a design for a 300 W, 100 kHz buck converter with 12 V output. The inductors must support 25 A DC currents, and their total inductance should not exceed $200 \mu\text{H}$. This *a priori* limit on inductance was chosen as a means of limiting the size of the complete low-pass/resonant network.

In the case of high ripple currents ($> 1\%$ of the corresponding DC values), manufacturer's sizing charts for DC-biased inductors can lead to the selection of undersized, lossy cores. Core selection from manufacturer's permeability and loss charts is complicated by design interdependence (in the high-ripple case, at least) whenever multiple cores (e.g., L_1 and L_2) must be selected. I.e., losses depend strongly on AC currents which themselves depend critically on Q. The designer needs impedances to find currents, and currents to




Figure 3.18: A set of achievable inductor Q's used to derive a loss model for volume-minimized inductors at 25 A DC bias with 100 kHz ripple currents. The surface $Q = 4.4025e \cot I_{rpp}^{-2.378}$ is a least-squares fit of the log-transformed set. All inductor designs were based on Micrometals iron-powder toroidal cores.

find impedance, even if inductances are approximately known. The mutual dependence, moreover, is steep in the design of resonant filters: transfer functions with lightly damped poles relate impedance and current, and core loss (hence inductor impedance and Q) is a strong function of peak to peak AC current. The latter dependence is especially strong for minimal magnetic geometry, when core loss becomes especially significant relative to winding loss.

Iterative magnetics design and network simulation is a possible, but cumbersome, means of predicting the performance of resonant networks in the search for a low-volume choice of cores. Rather than designing magnetics at each iteration, however, the designer can construct a loss model for a large set of inductors for the resonant network. The strategy, essentially, is “break the design loop” at AC current, and select volume-minimized cores for many inductors given their DC current, ripple-current frequency, and ripple magnitude. The inductance must be specified too, or at least approximated, to obtain initial guesses the field quantities used in the iterative magnetics design. A loss model can be derived from a plot the computed Q versus ripple magnitude and inductance (close to the initial guess for inductance, hopefully), *assuming* that voltages are applied by the external circuit which

drive the AC currents used to compute loss. A surface results (Fig. 3.18) that indicates the achievable Q of volume-minimized inductors. Should this surface be well-approximated by a simple function, the task of minimizing the dimensions of a multi-inductor network (with the DC bias and ripple frequency assumed in the loss model's derivation) is simplified considerably.

The loss model derived for powdered-iron toroids with 25 A DC bias and 100 kHz ripple is shown in Fig. 3.18. The results show that Q depends much more strongly on peak-peak AC current than on a particular value of inductance. A simple exponential loss model, $Q = 4.4025e^{I_{rpp}^{-2.378}}$, is a least-squares fit of the log-transformed data set which summarizes the dependence of Q on AC current for volume-minimized inductors. The loss model was applied to the evaluation of low-pass/resonant networks which divided 200 μH in various proportions between L_1 and L_2 . Phasor analysis showed that the best ripple attenuation, with realistic loss, was achieved with $L_1 = 120 \mu\text{H}$ and $L_2 = 80 \mu\text{H}$. $L_3 = 600 \mu\text{H}$, computed with the loss model and verified with core manufacturer's software, was designed to match the resonant network's performance at $D = 0.3$. See Sec. 4.2 for details of the cores chosen for L_1 - L_3 and measurements of their filtering performance.

For high ripple-current resonator designs, the selection of cores can be viewed as a trade-off between resonator Q and component rating. In a parallel-tuned power stage, the designer can achieve continuous conduction for less total inductor volume inasmuch as Q is large. A Q that is too high eventually increases total inductor volume, however, because a resonant inductor must increase in size (for a given core material) to carry circulating currents with low loss.

*Phase-lock Tuning***4.1 Phase-lock tuning**

RESONANT excitation is equivalent to maintaining a resistive phase relationship (0°) between resonator voltage and current (note the impedance angles in Fig. 4.1). Because the phase response of a series- or parallel-tuned circuit monotonically increases or decreases around the 0° tuning point, it can be used as an error signal to control for excitation at the point of maximum immitance. The phase-lock tuning system presented here employs this method precisely, feeding back the phase difference between resonator voltage and current to drive a voltage-controlled oscillator (VCO) toward the resonator's tuning frequency.

A control topology to excite a series resonance at its minimum-impedance point is shown in Fig. 4.2a. The dual of this tuning system is shown in Fig. 4.2b, which drives a parallel resonance at its maximum-impedance point (its resistive-impedance point). The control circuitry in either case generates the frequency command shown at the right. This command specifies the fundamental — but not the harmonic content or DC level — of the sources on the left of the block diagrams. In a power converter, the frequency command would represent an adjustable PWM frequency.

To excite the parallel- and series-tuned resonators, the controller must, in either case, adjust the fundamental drive frequency such that the resonator current and voltage *fundamentals* are in phase. The inner-loop PLLs in Figs. 4.2a and b serve two functions in this regard. They provide, first of all, a 90° phase shift in lock, which allows a subsequent phase detector (multiplier 1) to develop zero average output for a 0° V-I resonant condition in the resonators. By itself, this phase shift is poor motivation for introducing the complexity of a PLL, as the designer could employ a phase detector with 0° offset.¹ The more significant

¹In fairness to an inner-loop PLL, phase detectors with 0° phase offset (state-machine detectors, typically) can be confounded by the edge timing of PWM waveforms. Signal conditioning of some sort (a filter or PLL) will probably be necessary to develop a signal with zero-crossings in phase with the fundamental component of such waveforms.

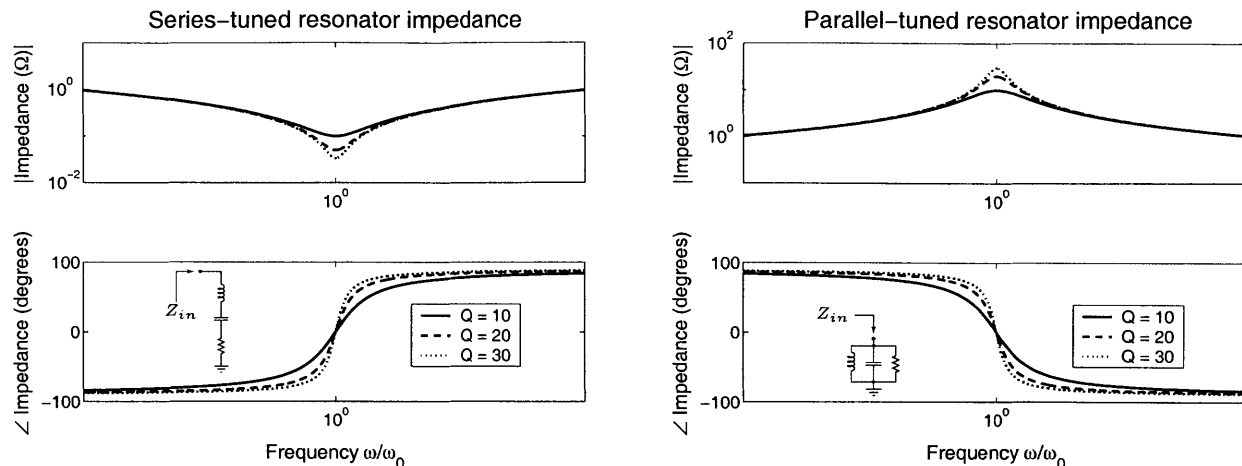


Figure 4.1: Frequency response of second-order tuned circuits, normalized to the natural frequency $\omega_n = 1/\sqrt{LC}$. The impedance magnitude at a single frequency can indicate proximity to resonance (with calibration) but not whether resonance lies above or below the stimulus frequency. The impedance phase, however, increases or decreases monotonically, and its difference from 0° is an error signal indicating the distance and direction to resonance.

function of the inner-loop PLLs, then, is to reject harmonics by locking on the *fundamental* component of an input waveform.

Consider, for example, the parallel-resonator tuning system of Fig. 4.2b. The differential amplifier measures the AC voltage across parallel-tuned tank, a signal with, presumably, a large fundamental component. The current through the parallel-tuned circuit, however, is dominated by its harmonic content, since the resonator suppresses the fundamental inasmuch as its Q is large. The PLL effectively filters this harmonic content, extracting a signal proportional to the fundamental current only. The AC tank currents can have arbitrary harmonic content as long as the phase-lock loop employs a VCO with sinusoidal output. I.e., as long as one input of multiplier 2 is sinusoidal, the product waveform is a useful phase-detector signal when the AC tank current is *any periodic waveform* with roughly the same fundamental frequency. The low-pass PLL dynamics ensure that the multiplier develops an average detector voltage proportional to the phase error between the fundamental frequencies of its inputs: all other sum or difference frequencies are effectively attenuated. With proper selection of gains, the PLL will apply negative feedback to drive its phase error to zero, producing a *quadrature replica* of the fundamental AC resonator current even when this current is dominated by harmonics.

Multiplier 1 accepts at its inputs the phase-lock replica of the fundamental AC tank-current waveform (shifted by 90° from the original), and a measurement of the AC tank




Figure 4.2: Block diagram of the phase-lock tuning system, demonstrating two possible methods of sensing resonator AC voltages and currents. Equivalent tuning controls can be implemented by switching the sensing connections, e.g., by phase-locking to the tank voltage rather than the tank current in the upper diagram. Such an approach, however, does not take advantage of phase-lock loop's ability to cleanly extract the fundamental component from the signal most dominated by harmonics,

voltage. Again, only the fundamental components of the the multiplier inputs produce an average output, a product in this case proportional to the phase difference between the fundamentals of resonator voltage and current. Multiplier 1 has zero average output (zero error) for a 90° phase shift between its inputs, or zero error for a 0° V-I phase relationship at the resonator. Because of the resonator's monotonic phase slope, loop gains with the proper sign always push the outer-loop VCO (and hence the controlled AC source) towards the resonator's tuning frequency.

To illustrate the need for harmonic rejection, consider the resonant-excitation system of Fig. 4.3a. The power amplifier (PA) drives a series-tuned resonator at the fundamental frequency commanded by the VCO, but with harmonics not present in the VCO output (e.g., as in a class D amplifier or a switching converter). The capacitive divider presents a high

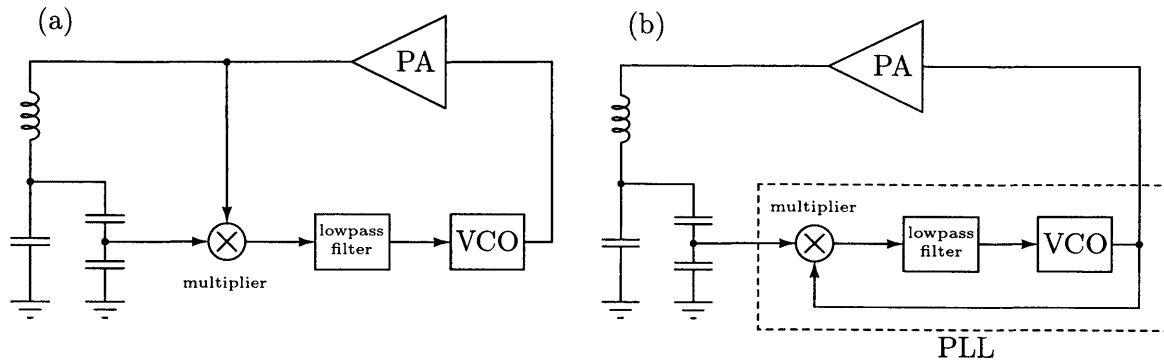


Figure 4.3: Alternate resonant-excitation topologies. The PA blocks are power transimpedance amplifiers with current-drive outputs. The capacitive dividers sense the current driving the series-tuned resonators (far left in each subfigure) with a 90° phase shift.

impedance to resonator currents relative to the resonator capacitance, and provides a scaled version of the resonator’s internal node voltage which, at the series-resonant frequency, is a measure of the PA driving current shifted 90° . Were the PA drive purely sinusoidal, the system of Fig. 4.3a provides everything need to acquire lock: voltage and current measurements with a phase relationship at resonance corresponding to zero phase-detector error. Because the PA output voltage is harmonic-rich in most switching converters, however, the system does not provide a perfect measure of a phase shift between voltage and current *fundamentals*. I.e., some harmonic-current signals will appear at the capacitive-divider output, which — when multiplied by the corresponding voltage harmonics from the PA output — will produce many low-frequency product terms within the tuning system’s control bandwidth. Though a detailed treatment would require knowledge of the PA output impedance and resonator Q, these harmonic product terms might be significant compared to the fundamental term if, for instance, the PA were loaded so heavily at the resonator’s tuning frequency that its fundamental-voltage output became small. With aggressive filtering in the signal paths leading to the multiplier, though, such a tuning system might be feasible. Tuning with no special filtering may nevertheless be feasible in systems with relatively small power-amplifier harmonic content and low-Q resonators [24].

The topology of Fig. 4.3a accentuates harmonic-term FM of the VCO; the scheme of Fig. 4.3b avoids this problem by multiplying the sensed current with the sinusoidal output of the VCO, eliminating all low-frequency product terms except that contributed by the V-I fundamentals. Such a design, however, is only useful when the commanded and driven phase (the phase at the input and output of the PA) are equal. For a switching power amplifier under duty-ratio control, the PA introduces duty-ratio-dependent phase shift for




Figure 4.4: Tuning points of a parallel-resonant circuit with low unloaded Q. **Note:** for inductor Q values above 10, these resonant points all converge to within 1% of frequency.

any $D \neq 0.5$. So, while useful for square-wave excitation of a resonance [25], the depicted scheme does not have the flexibility of the inner-PLL topology of Fig. 4.2.

4.1.1 Equivalence of phase and impedance tuning conditions

In the tank of a practical parallel-resonant filter, as suggested in Fig. 4.4(a), the inductor is the chief source of loss. Such an “almost parallel” tuned circuit can have a low unloaded Q (< 20) in power applications, and may exhibit the multiple resonant conditions shown in Fig. 4.4(b):

Equal-reactance resonance

$X_C = X_L$ is the tuning condition for series resonators. In the parallel case, the impedance of the inductive leg is composed of X_L and R , an impedance which is greater than — and not 180° out of phase with — X_C . The total current i_T is greater than its minimum value and not in phase with the voltage.




Figure 4.5: Linearized model for the phase-sensing tuning system in lock.

Anti-resonance

(*viz.* maximum impedance resonance). By altering the value of the inductor slightly (and holding its Q constant), a new frequency is found where i_T is minimized and the total parallel reactance is maximized. Again, i_T is out of phase with the voltage.

Unity-power-factor resonance

The 0° V-I resonant point, found by adjusting the inductance at constant Q so that $X_L + R$ just cancels the capacitive reactance. The value of a parallel-equivalent inductor for this condition is always smaller than the L shown in Fig.4.4(a), resulting in a resonant frequency different than the other two cases.

The relative frequencies of these three tuning points are not in the order shown for all cases, but above a Q of ten, they converge to within a percent of frequency. Tuning for zero phase difference between the current entering the resonator and the voltage at its terminals results in operation not appreciably different from tuning for maximum impedance: the proposed control scheme can effectively maintain operation where the resonator provides maximum ripple attenuation.

4.1.2 Tuning system dynamics

Fig. 4.5 shows a linear model for the tuning system dynamics in lock. The ω_{res} reference is the 0° V-I tuning frequency of the resonant filter. This reference can vary from resonator to resonator due to component tolerances, and can experience abrupt changes during load-step inductance swings. The K_ϕ block represents the fundamental-frequency phase-sensing action of the inner-loop PLL. When the inner-loop PLL dynamics are fast compared to the overall tuning dynamics, phase-shift sensing — i.e. the generation of a quadrature replica of resonator voltage or current, and the operation of the outer-loop phase detector — can be represented algebraically. K_ϕ is then just the small-signal phase gain of the resonator, the incremental change in phase-shift for an excursion $\Delta\omega$ of the switching frequency away from ω_{res} :

$$K_\phi = \left. \frac{d\phi}{d\omega} \right|_{\omega=\omega_{res}} = \left. \frac{d}{d\omega} \left[\tan^{-1} \left(\frac{\frac{\omega R}{L}}{\frac{1}{LC} - \omega^2} \right) \right] \right|_{\omega=\omega_{res}} = \left. \frac{-\frac{R}{L} \left(\frac{1}{LC} + \omega^2 \right)}{\left(\frac{1}{LC} - \omega^2 \right)^2 + \left(\frac{\omega R}{L} \right)^2} \right|_{\omega=\omega_{res}}$$

$$K_\phi = -\frac{2L}{R} = -\frac{2}{\omega_{res}^2 RC} = -\frac{2Q}{\omega_{res}}$$

The three remaining blocks, K_d , $H(s)$, and K_0 are identical to their counterparts in the linearized model of the basic PLL (Sec. 2.2). $H(s)$ is a loop-shaping filter which, in the prototype systems, contains an integrator to support the VCO command with zero steady-state frequency error, and a low-frequency pole to further limit the bandwidth of phase-error signals from the outer-loop multiplier (see Sec. 2.1.2). Note that no integrator is required in the “nested-PLL” model because the tuning system does not operate on phase signals, but aligns frequencies. The output of the K_0 block, the converter frequency command, can sustain arbitrary phase shift (e.g., from converter dynamics) without affecting the operation of the tuning system. A more detailed view of the frequency-tuning circuitry used in the prototype systems is shown in Fig. 4.6 and in Appendix B.

Though the block diagram of Fig. 4.5 has proven satisfactory for developing tuning controls, it does not indicate all details relevant to locking and hold-in performance. In particular, the K_ϕ block summarizes the result of multiplications involving two sensed signals, each of which is a potential source of destabilizing disturbances. Though it is difficult to determine whether the performance of the system depicted in Fig. 4.5 is limited by the inner loop, the tuning controller exhibits the locking and hold-in trade-off typical of a simple PLL (Sec. 2.3.3). I.e., choice of a narrow outer-loop bandwidth can have been observed to *prevent* (not just slow) the tuning process, as might be expected for PLL pull-in acquisition of a noisy signal. The converter ripple waveforms sensed by the tuning controller can be

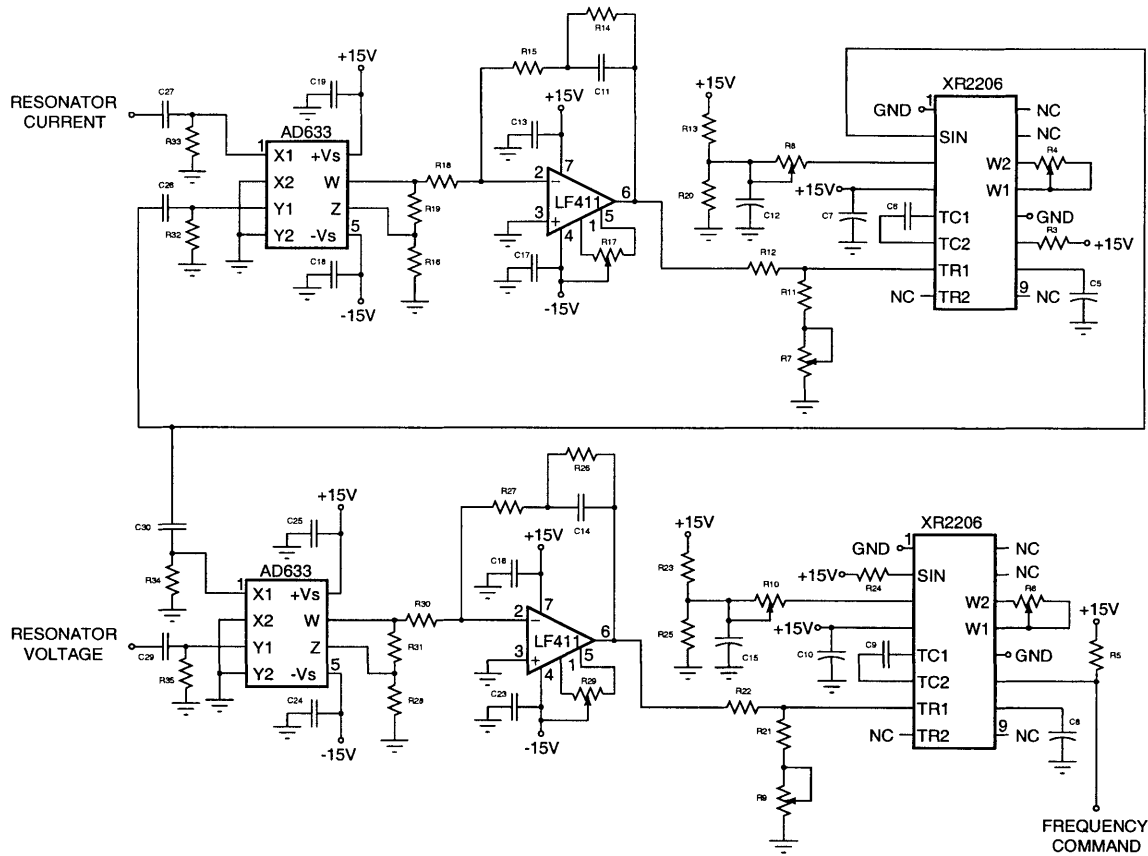


Figure 4.6: Schematic of the tuning circuitry used in the prototype tuning system. The resonator voltage and current inputs refer to the topology of Fig. 4.2b. See Appendix B for component values.

small signals ($< 10\text{mV}$) with large ringing voltages around switch transitions. Though controllers that reliably achieved resonance-lock in less than a second have been tested, the noise theory of the PLL (a cumbersome topic) has not been extended to explain these results or guide the selection of outer-loop bandwidth.

4.2 Application to a DC-DC converter

Fig. 4.7 depicts the phase-lock tuning system applied to a buck converter with a resonant power-stage filter network. By sensing voltage and current in the parallel-tuned network, the tuning controls align the switching frequency with the transmission null of the resonant power stage. The control circuitry is amenable to integration because it processes only signal power levels. The AC tank voltage can be measured with a single control connection

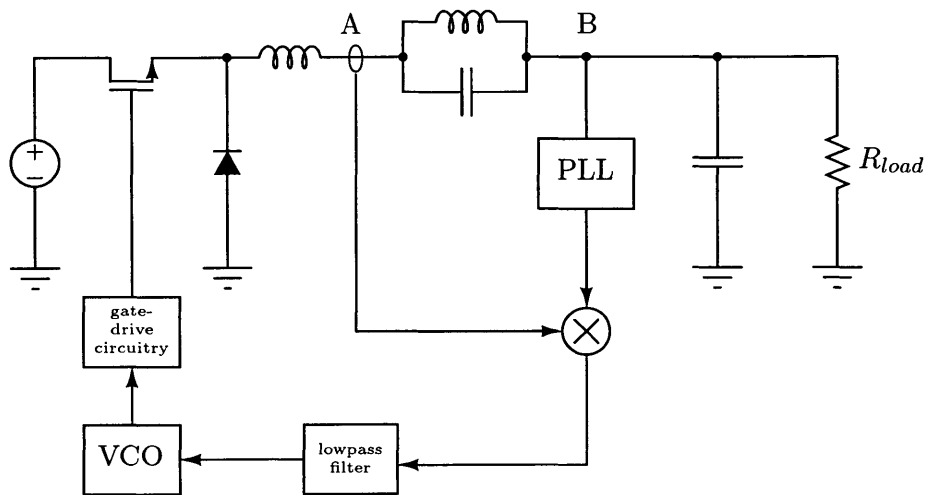


Figure 4.7: Block diagram of phase-lock tuning system used to align a buck converter's switching frequency to the maximum-impedance resonance of its output filter. A ground-referenced voltage measurement is made at point A, from which the AC voltage across the tank is determined. The resonator current is measured at point B.

(assuming the load voltage is nearly constant). Inductor current is often sensed and used by commercial controllers. The VCO and gate-drive circuitry likewise represent no additional complexity in the converter control circuitry: the VCO, in particular, can be implemented with simple modifications to a conventional controller's free-running oscillator.²

To demonstrate the benefits of controlled resonant excitation, resonators were incorporated in a 300 W buck converter with 12 V output and an input voltage ranging from 24–40 V. Consider first the power-stage resonant network designed in Sec. 3.5.2 (see Fig. 4.7). As seen from the measured current ripple in Fig. 4.8, both the resonant and low-pass networks meet a flat 120 mA maximum-ripple-current specification for all duty ratios greater than 0.38. The ripple-current fundamental has the largest magnitude at $D = 0.5$, around which point the resonant filter obtains the greatest benefit from its parallel-tuned network and outperforms the single inductor at each duty ratio in the range $0.38 < D < 0.62$. The actively tuned filter achieves this performance for 3.7 times less total filter volume (151 cm^3 , compared to 557 cm^3) and 3.6 times less total filter mass (0.691 kg, compared to 2.50 kg) than the conventional single inductor. If smaller core volume and weight improvements are acceptable, the resonant system could match the performance of L_3 over a wider range of

²Note that the prototype tuning controller (See Appendix B) is unnecessarily complex because it was implemented with discrete control elements. E.g., the frequency command is communicated to the PWM controller through modulation of the controller's RC oscillator. A fully integrated controller could employ a single oscillator, avoiding altogether the problem of synchronizing the VCO and PWM-controller frequencies.

duty ratios. L_1 has a full single-layer winding and is about 50% saturated at full load current, so a further decrease in its ripple current (an increase in its inductance) would require a shift to an even larger core or a more expensive core material. These results demonstrate that the proposed active-tuning approach can provide dramatic improvements in performance for a given limit on passive-element size. The conventional filter mass (2.5 kg) necessary to match the deep continuous-conduction performance of the resonant filter is impractically large. The resonant filter, at 691 g, is, perhaps, of acceptable size, and enables operation at lower ripple ratios than might otherwise be practical.

Deep continuous conduction — or, equivalently, low ripple ratios — favor resonant filters. Circulating tank currents are Q times larger than ripple currents at the resonator terminals, and can produce large peak AC flux density (hence high loss) in a resonant inductor. Should the designer choose to operate with higher power-stage ripple, a resonant inductor using the same core material may have to be made larger to support circulating currents with acceptable Q . Nevertheless, resonant filter designs with typical ripple ratios of 5–20% can still be dramatically smaller than conventional low-pass designs of comparable performance.

Furthermore, it should be noted that very low ripple ratios are common in input and output filters for switching converters: in such applications, resonant networks offer clear volume and mass savings. Consider the resonant input filter designed in Sec. 3.5.1 (see Fig. 4.9). The AC currents measured through L'_1 or L'_3 , respectively, are plotted in Fig. 4.9 as a function of duty ratio. The resonant and low-pass networks have essentially identical performance and, when incorporated in a π -section filter with a capacitance C' across a $50\ \Omega$ source impedance, meet a flat $-90\ \text{dB}\mu\text{V}$ ripple specification. The resonant filter matches the performance of the conventional design with 92 g and $19\ \text{cm}^3$, 3.0 times less total filter volume and 2.8 times less mass than with L_3 alone.

The converters in Figs. 4.8 and 4.9 were never tested with simultaneous output-voltage and switching-frequency regulation. Though switching-frequency adjustments necessary to track abrupt resonance shifts might perturb output voltage slightly, duty-ratio control will be able to reject such disturbances in the same manner as load steps. No complicated control interactions are anticipated because duty-ratio variation cannot alter the base switching frequency or filter resonant frequency. A time-varying duty ratio capable of producing ripple subharmonics will, furthermore, not affect the phase-lock circuitry for well-designed VCO ranges. Subharmonic frequencies can always be excluded from the inner PLL's hold-in range, so that all components of the phase-detector voltage not produced by the base switching frequency will lie outside the tuning system's control bandwidth.

4.2 Application to a DC-DC converter





Figure 4.8: Comparison of the core sizes and peak-peak inductor current for the resonant/low-pass (L_1 and L_2) and simple low-pass (L_3) magnetics for the power stage of a prototype 300W buck converter. Micrometals core data: L_1 OD 3", Ht. 1", Vol. 7.069 in³/115.8cm³, Wt. 0.533 kg; L_2 OD 2.25", Ht. 0.55", Vol. 2.187 in³/35.82 cm³, Wt. 0.158 kg; L_3 OD 5.2", Ht. 1.6", Vol. 33.98 in³/556.6 cm³, Wt. 2.50 kg

Phase-lock Tuning


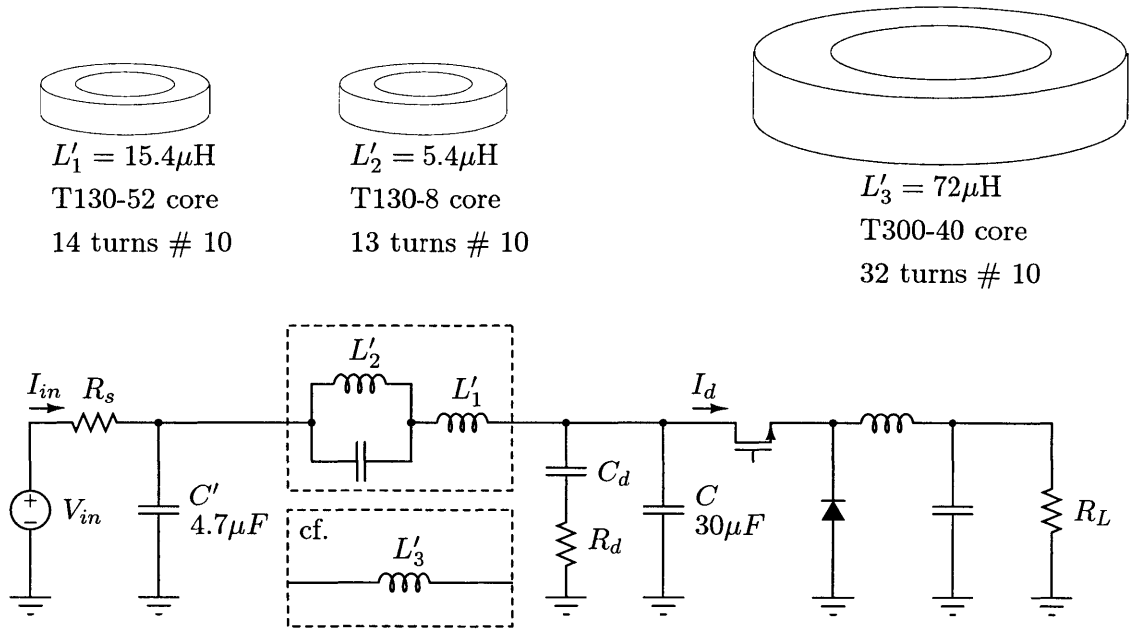


Figure 4.9: Comparison of the core sizes and peak-peak inductor-network current of the single-inductor and resonant filters. Micrometals core data: L'_1 and L'_2 OD 1.3", Ht. 0.437", Vol. 1.16 in³ total, Wt. 0.092 kg total; L'_3 OD 3", Ht.=0.5", Vol. 3.53 in³, Wt. 0.259 kg. Filter capacitors were small ITW Paktron Capstick parts, and did not contribute greatly to filter volume.




Figure 4.10: Parallel-tuned resonant filters can be applied to the power stages of most major PWM switching-converter topologies.

4.3 Alternative implementations

Resonators constructed from discrete components can be incorporated in most major converter topologies, direct or indirect, isolated or non-isolated (Fig. 4.10). More elegant, however, is the use of tuning for structures which expressly incorporate resonances in a hybrid reactive element. Phase-lock tuning can realize the benefits of such resonant structures that, like the lumped resonator in the example filter, are otherwise limited by component tolerances. Single-resonant[26] and multi-resonant[27] inductors, for instance, use magnetically




Figure 4.11: Hybrid inductive-capacitive elements that exhibit resonances. (a) and (c), self-resonant inductors with large inter-turn or inter-coil capacitance. (b) schematic depiction resonant foil capacitor, similar to a typical film capacitor without extended-foil contacts.

coupled tuned circuits to produce high impedances at discrete frequencies. Such structures exploit mutual inductance and inter-winding capacitance (M and C in Fig. 4.11a) between coils to introduce impedance peaks at one or more frequencies. Additional coils that do not carry DC current can be magnetically coupled to the principle winding, and two-terminal devices with as many as three resonances have been demonstrated in [27].

The self-resonant capacitor/inductor hybrids of [28]–[32], are wound-foil structures which exhibit repeatable resonances with low loss. Schematically, these hybrid structures introduce controlled self-inductance in galvanically isolated foil turns (Fig. 4.11b) to implement a series-resonant capacitor, or introduce controlled inter-turn capacitance in a single foil strip (Fig. 4.11c) to produce a parallel-resonant inductor. Such structures can have particularly low loss (a 170 kHz Q of 155 is reported in [32]) because of the low-impedance foil construction and the use of the same conductor for both magnetic and electric storage. Shunt resonators, which for discrete passives were found impractically lossy in the design example, may perhaps be implemented inexpensively with good performance using such hybrid passive elements.

Variations in driving circuitry, manufactured geometry, and temperature (and, importantly, DC magnetizing force and AC flux density when magnetic materials are present) can alter the tuning point of all these resonant structures to such a degree that their filtering properties may be of little benefit in a practical system without tuning. Moreover, distributed models or high-order finite-dimensional approximations are required, in the case of foil resonators at least, to predict resonant frequencies. With active-tuning control for *excitation at resonance*, the full filtering benefits of the above-mentioned structures can be practically realized with minimal design effort. Though not employed for filtering, core-less planar transformers[33] and core-less twisted-coil transformers[34] exhibit maximum-efficiency points characterized by resistive V-I phase relationships at their ports. Effective use of such structures, likewise, is a task ideally suited to phase-sensing control.

Phase-lock controls can be applied to tune a filter resonant frequency rather than a converter switching frequency. An electrically controlled reactance implemented, for instance, with a cross-field reactor (Fig. 4.12, and see [35]–[37]) can shift a filter transmission null as currents are applied to its control winding. As illustrated in Fig. 4.12, a cross-field reactor comprises a magnetic core with two windings that produce perpendicular magnetic fields. The windings are not mutually coupled (i.e., flux from either winding does not link the other) so the device exhibits no “transformer action” in the normal sense. One port, termed the inductance winding, serves as a controlled inductance and carries a ripple current with, possibly, some DC bias. The other port, the control winding, carries a DC current




Figure 4.12: Structural diagram of a cross-field reactor. The magnetic core is wound with two windings (an annular coil and a toroidal coil) that are not coupled in the usual sense.

which drives the core a controlled amount into saturation and adjusts the effective inductance seen from the inductance winding. A nested PLL topology like that presented for the frequency-tuning case can excite the control winding to maintain operation at resonance. The resonance-tuning approach, significantly, can support attenuation of multiple frequencies by independently tuning multiple resonant networks. When an inductance winding is designed to support a DC bias current, the cross-field reactor is less sensitive to control currents (i.e., larger currents are required to adjust the core’s effective permeability). For this reason, an electrically controlled inductance can be implemented with less expensive power-electronic controls when used as a shunt filter element carrying only AC currents. In series-tuned networks considered for the prototype tuning system (Sec. 3.2.2), air-core shunt inductances were found to have unacceptably large losses. The necessity of core loss accompanying a tunable shunt reactance was deemed impractical for the power levels considered, so a resonance-tuning system was never implemented.

Magnetic tuning may also be valuable in ripple-current-steering structures ([18]–[23]) where control of coupling can improve performance. The discussion of magnetically coupled shunt resonator in Sec. 3.3 described a “zero-ripple” condition in which a transformer turns ratio just large enough to make up for imperfect coupling nullifies voltages which drive ripple current toward a port of interest. This process of ripple-current steering — the “zero-ripple” condition — is classically expressed as a condition on *coupling* rather than turns ratio, and early papers describe mechanical means of tuning coupling for maximum filter attenuation [18]. Control over permeability at some region of a magnetic circuit (by means similar to those employed in a cross-field structure) may provide an effective electrical means of controlling coupling. A phase-sensing tuning system could be naturally extended to drive coupling toward the “zero-ripple” condition.

Integrated Filter Elements

THE DESIGN trends of magnetically coupled shunt resonators — particularly attempts to reduce their shunt impedance across frequency — lead to a elegant extension of the filtering applications of the transformer network shown in Fig. 5.1a. Sec. 5.1 details the theory behind this technique: a passive inductance cancellation method implemented with magnetically coupled windings. The experimental results presented in Sec. 5.2 demonstrate the EMI performance achievable with this new method. The inductance-nulling magnetics can be implemented in several fashions, including conventional windings and printed PCB windings, or incorporation directly into a film-wound power-capacitor structure, requiring little additional volume. The resulting capacitor/transformer — an integrated filter element — shows a dramatic and repeatable performance increase over conventional filter networks. Sec. 5.3 addresses the construction of integrated filter elements, arguing that their film-wound structure permits inexpensive and reliable manufacturing. Finally, Sec. 5.4 suggests directions for further development of the inductance-nulling method and the integrated filter element.




Figure 5.1: (a) Model of the integrated filter element with parasitics removed. (b) The coupled windings of the integrated element, with magnetizing and leakage inductances shown. (c) A Δ -Y transformation of the impedances measured between each terminal pair (with the third terminal open-circuited, in each case) results in a T-model for the transformer.

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \begin{bmatrix} \frac{N_1^2}{\mathfrak{R}_{\ell 1}} + \frac{N_1^2}{\mathfrak{R}_M} & \frac{N_1 N_2}{\mathfrak{R}_M} \\ \frac{N_1 N_2}{\mathfrak{R}_M} & \frac{N_2^2}{\mathfrak{R}_{\ell 2}} + \frac{N_2^2}{\mathfrak{R}_M} \end{bmatrix} \begin{bmatrix} i_1 \\ i_2 \end{bmatrix}$$

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \begin{bmatrix} L_{11} & L_M \\ L_M & L_{22} \end{bmatrix} \begin{bmatrix} i_1 \\ i_2 \end{bmatrix}$$




Figure 5.2: Coupled windings, showing the paths of mutual flux Φ_M and leakage fluxes Φ_{11} and Φ_{22} . The inductance matrix which relates flux linkages to coil currents can have a “physical” form which reflects the actual paths seen by the magnetic flux, or a “phenomenological” form like that shown on the lower left. In this inductance matrix, L_{11} and L_{22} are the self-inductances measured from either winding, and L_M is the windings’ mutual inductance. Without further information about the magnetic coupling (e.g., the turns ratio), such terminal measurements do not determine a unique physical model of the magnetic circuit.

5.1 Principle of Operation

A passive inductance matrix is, from energy considerations, positive semidefinite. The entries of an inductance matrix, or the inductances used to model the terminal relations of a multi-port inductive network, however, can be negative. Consider the physical model in Fig. 5.1b of the transformer of Fig. 5.1a, including parasitics. The model includes the magnetizing inductance L_μ , reflecting finite core permeability, and the leakage inductances $L_{\ell 1}$ and $L_{\ell 2}$, modelling the imperfect coupling between transformer windings. $L_{\ell 1}$ and $L_{\ell 2}$ are large insofar as flux from one winding does not link turns on the other winding (cf. Φ_{11} and Φ_{22} in Fig. 5.2). The leakage and magnetizing inductances are “physical” — hence positive — because they correspond to energy storage within the magnetic structure. Note, however, that the physical model has four parameters (the turns ratio and the three inductances already mentioned) but can be modelled by a two-port network characterized by *three* impedances. As suggested by the equivalent inductance matrix formulation in Fig. 5.2, other “non-physical” inductances can preserve the terminal V-I (λ -I) relationships of a magnetic structure. By including, for instance, an inaccessible internal node in the transformer model (Fig. 5.1c), a branch inductance can be negative while preserving the positive inductances seen from each port.

Shunt-path inductance cancellation is possible when a transformer’s mutual induc-




Figure 5.3: (a) The integrated-element schematic from Fig. 5.1a redrawn, including important parasitics, but otherwise leaving terminal I-V relations unchanged. (b) When $L_{11} - L_M$ is chosen to be close to the capacitor's ESL, the shunt network reduces to a capacitance with small ESL, i.e. $\Delta L = -L_M + L_{11} + L_{ESL} \approx 0$.

tance L_M exceeds the self-inductances L_{11} or L_{22} of either winding. The maximum mutual inductance between two coils is limited by the positive semidefinite condition of the inductance matrix:

$$L_{11}L_{22} - L_M^2 \geq 0 \quad \rightarrow \quad L_M \leq \sqrt{L_{11}L_{22}}$$

I.e., L_M is constrained to be less than or equal to the geometric mean between the self inductances. The coupling coefficient k is defined by the extent to which L_M achieves this maximum value:

$$k = \frac{L_M}{\sqrt{L_{11}L_{22}}}$$

The designer is free to choose $L_{22} \gg L_{11}$, which even for moderate coupling can result in a difference $L_{22} - L_M$ less than zero (Fig. 5.1a). Note that the impedance seen across the the N_2 winding (with the N_1 winding open-circuited opposite Node A) is still just $L_{22} - L_M + L_M = L_{22}$, the winding self-inductance.

Fig. 5.3a shows the application of the T-model from Fig. 5.1c to a capacitor whose equivalent series inductance (ESL) the transformer is intended to cancel. When $L_{22} - L_M$ is chosen to be negative and close in magnitude to L_{ESL} , $\Delta L = -L_M + L_{11} + L_{ESL} \approx 0$, and the integrated element can be modelled by the T-network of Fig. 5.3b with an (almost) purely capacitive and resistive shunt impedance. L_{11} is necessarily larger than L_M whenever

$L_{22} - L_M < 0$, so the network adds a series inductance which can react with a capacitance shunted across the output port (the “quiet” port opposite Node A in Fig. 5.1a) to provide roll-off equivalent to two cascaded L-section filters.




Figure 5.4: Construction of an integrated filter element, with windings added outside the capacitor package for a proof-of-concept design. A foil layer is wound first, and kept short to provide a low-impedance shunt path for ripple current. The DC winding is placed on top of the foil turns, such that its self-inductance (and the coupling of between the windings) is large enough to increase the mutual inductance L_M above the AC-winding self-inductance.

5.1.1 Implementations of an Integrated Filter Element

In order to develop low shunt impedance over a broad frequency range, a practical integrated-element design must provide reliable shunt inductance cancellation with as little additional shunt resistance as possible. The inductance $L_{22} - L_M$ is negative but still lossy, and increases by its equivalent series resistance the minimum shunt impedance of the integrated element. Short sections of Litz wire or short, broad foils are therefore the best construction alternatives for the AC winding (the winding with self-inductance L_{22}) because they contribute minimally to the overall AC shunt resistance.

To achieve an L_M sufficiently large for effective inductance cancellation requires a DC-winding self-inductance L_{11} larger than L_{22} and a coupling coefficient equal to:

$$k = \frac{L_{22} - L_{ESL}}{\sqrt{L_{11}L_{22}}}$$

In experiments with various construction techniques, foil AC windings minimized additional shunt resistance and achieved the coupling necessary for ESL cancellation.

The experimental inductance-cancelling transformers were wound on the outside of United Chemi-Con U767D 2200 μF , 35V electrolytic capacitors as shown in Fig. 5.4. The capacitor body was used as a coil form to demonstrate the viability of incorporating similar windings inside the package along with the capacitor plates. In the best-performing design, the shunt-leg inductance (the AC winding) was wound from two layers of 1 mil copper foil 1 inch wide, separated by 1 mil Mylar adhesive tape. The AC-winding length was selected for the minimum length necessary to achieve the coupling (hence the negative inductance) capable of cancelling the electrolytic capacitor's ESL. The capacitor package's circumference was 7.1 cm, and one and three-quarters turns were found just sufficient for inductance cancellation, and limited the total AC-winding resistance to 7 m Ω at 100 kHz. A DC winding comprising several turns of 18-gauge magnet wire was coiled tightly over the AC winding and glued in place. The three-terminal integrated element was completed by soldering the the DC and AC windings together at one end (Node A in Fig. 5.1a) then soldering the free end of the AC winding to the capacitor's positive terminal.

Because the coupling between the windings was not known *a priori*, the enamel on the DC winding was removed at various tap points to permit adjustment of the self-inductance L_{11} . The EMI experiments detailed below were conducted with contacts soldered at the various taps: corresponding to each is a different mutual inductance and hence a different negative inductance $L_{22} - L_M$. The transformer impedances were measured beforehand to find the tap point corresponding to the best possible cancellation of L_{ESL} . In practice, the measurement uncertainty of small winding impedances required each experiment be repeated with slight tap adjustments to improve filtering performance.

The construction methods for the experimental integrated elements were indeed crude, but the tap-point adjustment — for fixed winding geometry — only needs to be determined once for a given magnetics design. The viability of repeated inductance cancellation, its dependence on repeatable geometry rather than material or contact properties, is explored further in Sec. 5.3. The U767D has an especially high self-resonant frequency, with a mean ESL of about 17.6 nH and mean ESR of 22.1 m Ω . Allowing for a residual inductance ΔL equal to 5% of the uncanceled ESL, the upward inflection of the U767D's impedance could be increased 20 times (beyond 1 MHz), and so be useful at switching *and* EMI frequencies. Where α is the fraction residual inductance (0.05, in this case), the upper impedance corner where $X_{ESL} = R_{ESR}$ is at an angular frequency

$$\omega = \frac{R_{ESR}}{\alpha L_{ESL}}$$

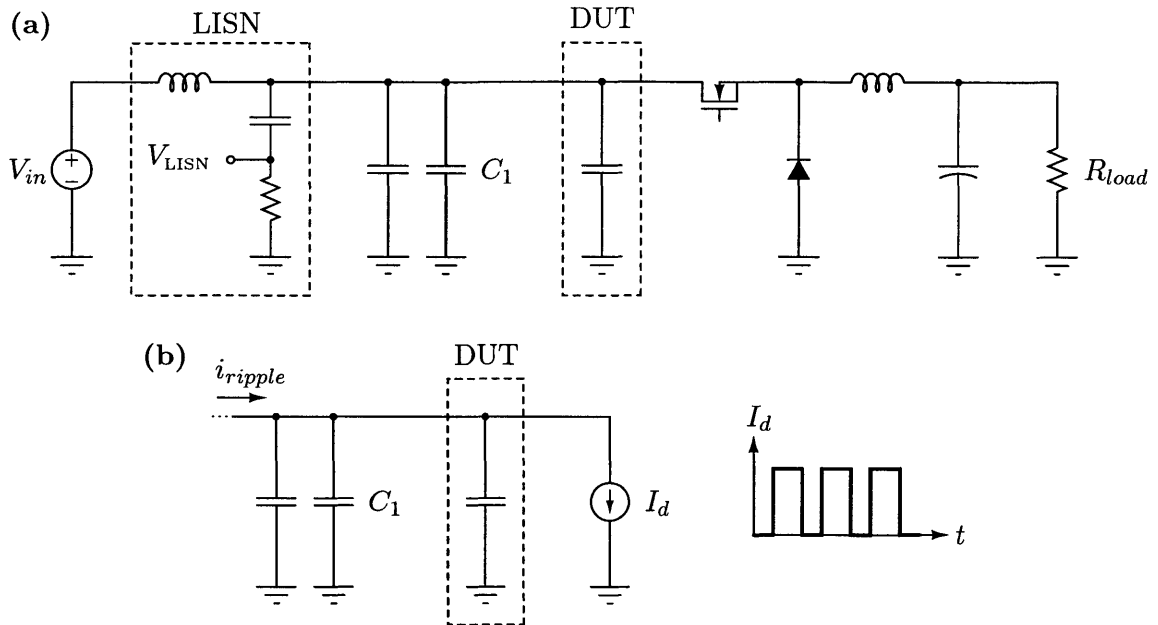


Figure 5.5: (a) Full schematic of the conducted-EMI test setup. C_1 denotes two $22\ \mu\text{F}$ monolithic ceramic capacitors (United Chemi-Con TCD51E1E226M). The devices under test are all based on the the $2200\ \mu\text{F}$ U767D capacitor: the integrated element described in Sec. 5.1.1, the “full-wound” capacitor of Fig. 5.7, and the capacitor alone. LISN data: Solar Electronics 8309-5-TS-100N, shunt capacitance $0.10\ \mu\text{F}$, series inductance $5\ \mu\text{H}$, $50\ \Omega$ termination by instrument. (b) Functional diagram of the EMI test stand. The buck converter provides a pulsed-current stimulus to the input-filter network. The input filter diverts AC currents away from the DC source V_{in} , and its effectiveness is measured by the LISN voltage V_{LISN} , itself a measure of i_{ripple} .

5.2 Experimental results

Conducted EMI performance of the integrated filter element was measured with the test setup of Fig. 5.5. The experimental shunt elements were employed as the principle low-impedance elements in the input filter of a buck converter. As is typical in converter input filters, a high-frequency capacitor (in this case the two capacitors C_1), were added in parallel with the device under test. Attenuation improvement was measured at the LISN port in A-B comparisons between a simple capacitor, an integrated filter element, and “full-wound” L-section filter (Fig. 5.7). This latter element was introduced to distinguish between attenuation due to series reactance rolling off with C_1 , and extra attenuation due to inductance cancellation. The “full-wound” element makes no use of transformer action, but, with three times the turns, has a much higher series inductance than the integrated filter element.

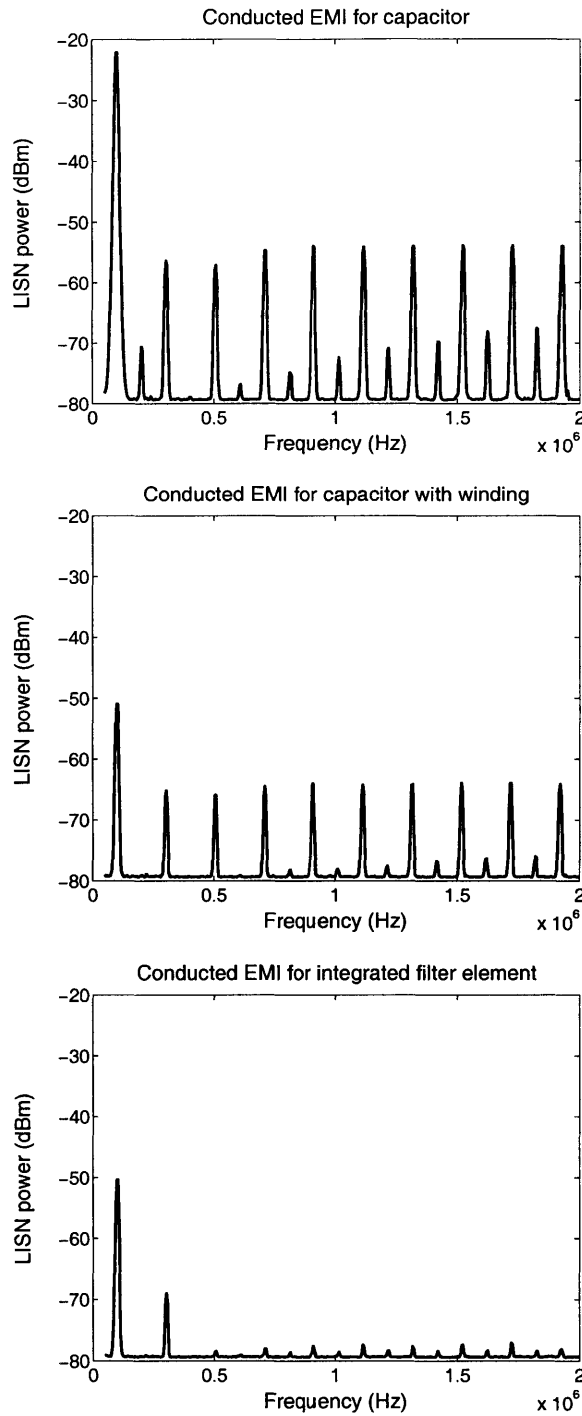


Figure 5.6: LISN power spectra for the case of (a) the capacitor alone, (b) the capacitor and inductor of Fig. 5.7, and (c) the integrated filter element shown in Fig. 5.4. The converter was operated at 50% duty cycle, at 10.04 A DC output current.




Figure 5.7: A “full-wound” inductor-capacitor structure used to distinguish, by comparison with the integrated filter element of Fig. 5.4, the effects series reactance and shunt inductance cancellation. The inductor-capacitor structure shown here adds series inductance L_s (wound on the capacitor) but does not alter the shunt impedance of the capacitor C .

5.2.1 Results from the switching converter

LISN-power spectra obtained for the capacitor, inductor-capacitor, and integrated filter element are shown in Fig. 5.6a–c, respectively. A spectrum analyzer set at 9 kHz resolution bandwidth measured the LISN voltage through a 20 dB attenuator. The converter was operated at 50% duty cycle with 10.04 A DC output current, so the MOSFET drain drew a 20 A p-p square wave of current from the input supply and filter. The results show that while series inductance aids in attenuation around the switching frequency (100 kHz in this example), the low-frequency rise of the electrolytic-capacitor impedance (i.e., when its inductance is not reduced) severely limits EMI performance. The integrated element adds 27 dBm attenuation at the switching fundamental over the simple capacitor, and around 25 dBm attenuation at EMI frequencies. While attenuation at the fundamental is about equal for the “full-wound” inductor-capacitor and the integrated element, the inductance-cancellation technique improves EMI-frequency attenuation by about 15 dBm.

5.2.2 Series inductance

The power spectra of Fig. 5.6 were measured in the presence of common-mode currents unintentionally conducted through the control power supply. Extra common-mode filtering decreased the performance improvement of the integrated filter element over the inductor-capacitor, and suggests that part of the latter element’s performance limits might be due




Figure 5.8: Frequency response of the input filter (transfer function from input current to LISN voltage) for the converter of Fig. 5.5. This curve was obtained by varying the switching frequency of the test converter across a $-50\%/+100\%$ range and observing the attenuation of fundamental and harmonic ripple components.

to coupling past its conduction path. The inductor-capacitor and integrated element have different impedances seen from their output ports, so the higher series inductance of the “full-wound” element may have resulted in an unfair comparison.

To resolve this ambiguity and characterize the integrated-element impedances as a function of frequency, two experiments were carried out to approximate a high-current sweep with a sinusoidal source. In the first experiment, the converter switching frequency (using the setup of Fig. 5.5) was adjusted to permit A-B attenuation comparisons over a continuum of frequencies. The frequency response (transfer function from input current to LISN voltage) of the capacitor and integrated element, calculated from power measurements at the fundamental and second-harmonic frequencies, is shown in Fig. 5.8. The frequency response shows the filtering effect of integrated element’s series inductance: the integrated filter element benefits substantially, not surprisingly, from roll-off of C_1 with its additional series inductance. After about 300 kHz, the integrated-element attenuation computed with this method became unreliable. High di/dt in the switching stage coupled enough power through parasitic paths to obscure the the conducted EMI performance of the DUT. Loss of precision in this same frequency range also obscured the comparison between the inductor-

capacitor and the integrated element in Sec. 5.2.

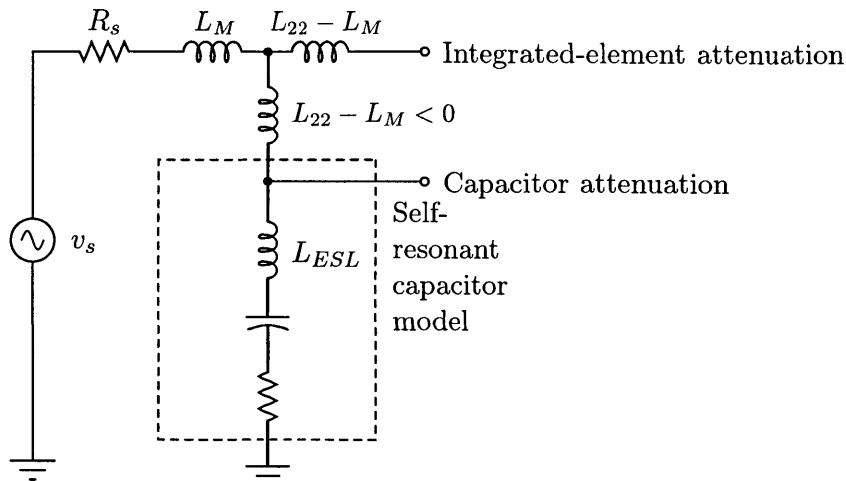


Figure 5.9: Experimental setup for the measurement of inductance cancellation. The reactances jL_M and $j(L_{22} - L_M)$ are small for frequencies below 1 MHz compared to the source impedance R_s .

5.2.3 Shunt inductance cancellation

Passive inductance cancellation only becomes apparent, for experimental elements considered, at frequencies in excess of 100 kHz. To circumvent stray coupling at these frequencies and determine the degree of cancellation, attenuation performance was measured at low current with the experimental setup of Fig. 5.9. A network analyzer was used to drive ripple current into the device under test, and to measure the resulting ripple at different locations. The 50Ω source impedance R_s was high enough (given the magnitude of other network impedances) for v_s to appear as a current source. I.e., the reactance magnitudes ωL_M and $\omega(L_{22} - L_M)$ were small compared to R_s . The attenuating performance of the capacitor alone was measured by driving the integrated filter element at its input node (the common node of the AC and DC windings) and measuring the voltage at the capacitor positive terminal (the point labeled “capacitor attenuation” in Fig. 5.9). A voltage measurement at the integrated element’s output port was a measure of its attenuation *from low shunt impedance alone*. Because the high-impedance measurement at the node labeled “integrated-element attenuation” in Fig. 5.9 drew so little current, any attenuation measured at this node in excess of that seen for the capacitor was due to a decrease in shunt inductive reactance, and *not* due to filtering by the series inductance $L_{22} - L_M$.

The capacitor and integrated-element attenuation are shown for frequencies to 1 MHz




Figure 5.10: Gain from driving voltage to voltage measured at the nodes labeled “capacitor attenuation” and “integrated-element attenuation” in Fig. 5.9.

in Fig. 5.10. The low-frequency attenuation of either element is roughly identical at low frequencies. At frequencies above 100 kHz, the capacitor impedance is primarily inductive, so its attenuation performance levels off to the constant performance of an inductive divider. The high-frequency divider ratio for the integrated element is about 12 dB better — after rolling off about a decade longer — than the capacitor alone. This performance represents a cancellation of the electrolytic capacitor’s ESL to approximately 15% of its original value. Note that the ordinate of Fig. 5.10, labeled “ripple throughput,” has not been corrected for source magnitude changes under variable driving-point impedance. Because of the capacitor and integrated-element attenuation measurements were made simultaneously, however, the two curves in Fig. 5.10 are a viable *comparison* of attenuation, even if they do not precisely measure *absolute* filtering performance.

5.3 Manufacturing

The value of the proposed technique relies on the repeatable cancellation of capacitor self-inductance. The histograms of Fig. 5.11 summarize the ESR and ESL measurements for 30 electrolytic capacitors of a type typically used in EMI filters for automotive applications.

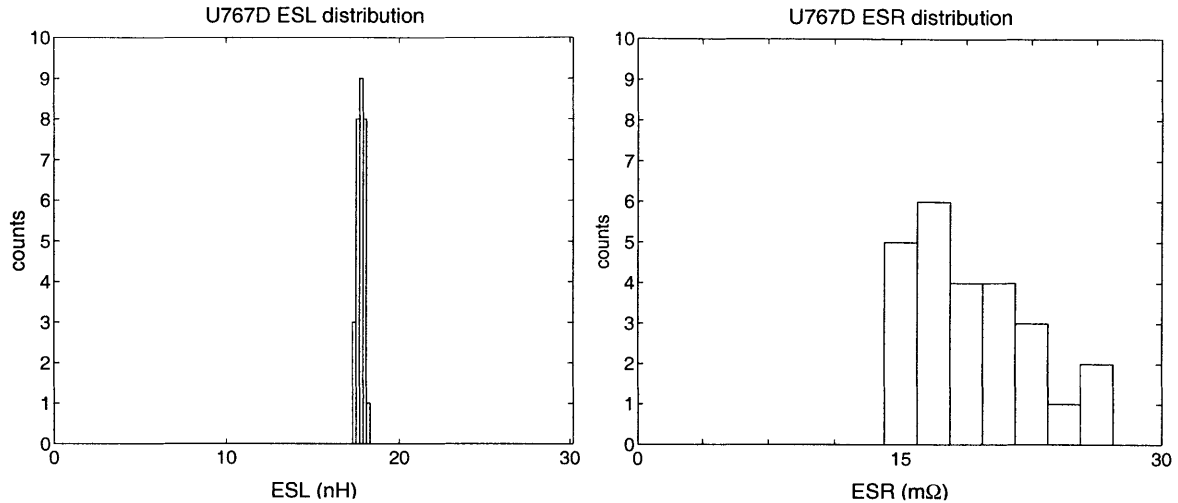


Figure 5.11: ESR and ESL histograms for 30 United Chemi-Con U767D 1200 μF capacitors. ESL range: 17.29 to 18.13 nH, $\sigma = 44.6$ pH. ESR range: 14.2 m Ω to 60.9 m Ω (outlier not shown).

The data show a remarkably tight clustering of ESL, $\pm 2.4\%$ for all the units measured, compared to a deviation in the ESR for some units of 175% from the mean value. There is no strong correlation between ESL and ESR for these capacitors, and even the ESR outliers exhibited a “normal” ESL. Inductance depends on geometry rather than material or contact properties which can strongly impact capacitor losses. These data suggest that repeatable inductance cancellation to within a few percent the base ESL should be practical.

Because of the repeatability of wound geometry and inductance, the coupled windings tested outside of the capacitor package — tapped as needed for best performance — could be manufactured with a set tap point and reliable inductance cancellation. A typical four-layer electrolytic capacitor is shown in Fig. 5.12a. One of the plate electrodes is etched to produce a porous film with high surface area. The spacer material is not a dielectric, but an electrolyte-impregnated layer which brings one plate into intimate electrical contact with the porous surface, where a high value of capacitance is achieved by a thin oxide coating of the interstices. Coupled windings — the AC and DC windings in Fig. 5.12b — could be added over the basic capacitor structure or made from extensions of the plate windings.

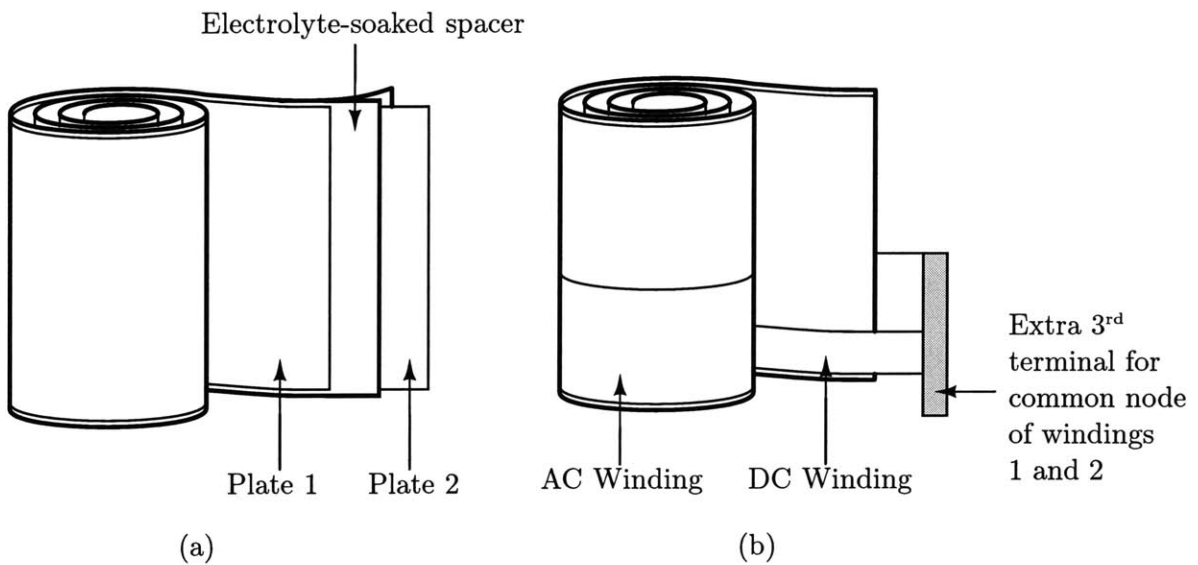


Figure 5.12: Incorporation of the coupled windings into the structure of a power capacitor.

5.4 Further work

The preliminary measurements presented here illustrate the large potential advantage of inductance-cancellation methods. Nevertheless, much development is required. Winding construction methods (including external windings, printed PCB windings, and integrated windings) need to be fully evaluated. Design trade-offs for integrated elements, including thermal design, eddy-current effects, and low AC losses must also be explored. The potential benefits of the technology fully warrant such investigation.

Conclusions

LARGE passive components in filters or power stages are detrimental to the transient performance of a converter and can contribute significantly to its volume and cost. This work has introduced topology- and component-level techniques which reduce the volume of passive filter elements required for a given level of ripple performance.

The first technique — a topological approach — employs filter networks with resonant branch impedances to provide high impedance mismatch at discrete frequencies. Component tolerances and operating conditions can cause uncontrolled resonators to miss their design frequency and fail to attenuate ripple. Chapter IV therefore introduces this thesis' major contribution: a phase-sensing tuning system which drives a converter's switching frequency (or any suitable tuning actuator) toward a resonant network's maximum-impedance point. The phase-lock system tuning is a general means of controlling the phase relationship among arbitrary frequency components of periodic signals, and has broad utility beyond power-filtering applications.

The component-level innovation introduced by this thesis is the integrated filter element, a three-terminal, passive inductance-cancellation structure inspired by trends in zero-ripple filter topologies. The magnetically coupled windings of an integrated element increase, by shunt inductance cancellation, the frequency of the impedance rise of a power capacitor, making the capacitor useful for filtering at switching *and* EMI frequencies. The windings themselves can be patterned from foil layers in capacitor packages, a construction technique which requires little additional volume, is compatible with existing capacitor construction, and offers repeatable magnetic coupling.

6.1 Conclusions: actively tuned filters

By ensuring effective attenuation at the ripple fundamental or a ripple harmonic frequency, active tuning eases the filtering requirement — and so lowers the volume and energy storage — of an accompanying network. For both an input filter and a power stage employing

Conclusions

tuned excitation of parallel resonances, savings of about a factor of three in volume, mass, and inductance have been demonstrated. Such reductions need not be realized as volume savings: a designer can, for instance, maintain passive-component volume and performance at a lower switching frequency and higher efficiency, or maintain the volume of a conventional filter while achieving better ripple performance at a constant switching frequency. Designs which do opt for smaller passive elements enjoy improved transient performance and, possibly, the replacement of large capacitors with less expensive, lower-value elements.

The tuning control system, with its low-level analog signals in close proximity to large switched currents, was difficult to implement properly. Careful attention to common impedances, especially those seen by power-stage return currents, was required to obtain lock. EMI also affected performance, and may have been responsible for lock-out in several experiments, though this connection was never demonstrated conclusively. Once low-impedance grounding techniques were employed, with careful separation of power and analog returns, the phase-lock system achieved reliable lock with no lock-out observed during trials over the course of a week.

The real design challenge for tuned filters seems not to lie in the tuning system, but in the resonators themselves. In particular, series-tuned shunt resonators offer the best ripple attenuation with capacitor-heavy, low-Q designs, a trend opposed to the need for tuning. A difficult three-parameter optimization of volume, Q, and attenuation performance is likely necessary to arrive at suitable series-tuned networks. For the switching frequency and power level considered in the example converter, series networks of substantial Q and low-tuning-point loss were always too large to compare favorably shunt capacitors. Applications restricted to relatively small capacitances, however, could benefit from resonant shunts.

As noted for the example filters, parallel resonators placed in series with a quiet port can reduce the filter magnetics volume by about a factor of three. Resonator Q and ripple attenuation increase together, so resonator design amounts to a two-parameter trade-off of performance and volume. For parallel-resonant designs, however, low ripple ratios favor resonant filters. Circulating tank currents are Q times larger than ripple currents at the resonator terminals, and can produce large peak AC flux density (hence high loss) in a resonant inductor. Should the designer choose to operate with higher power-stage ripple, a resonant inductor using the same core material may have to be made larger to support circulating currents with acceptable Q. Nevertheless, resonant filter designs with typical ripple ratios of 5–20% can still be dramatically smaller than conventional low-pass designs of comparable performance.

Highlighting the need for lower ripple ratios in parallel-tuned designs, consider the resonant power-stage example, in which the conventional filter mass necessary to match the deep continuous-conduction performance of the resonant filter was impractically large. The resonant filter, at one-third the size, was perhaps of acceptable size, and would enable operation at lower ripple ratios than might otherwise be practical. This power-stage example, then, demonstrates the use of resonators to achieve a *performance improvement* in constant volume, rather than a *volume decrease* for constant performance. Note also that low ripple ratios (less than 1%) are common in input and output filters for switching converters: in such applications, resonant networks offer clear volume and mass savings.

Magnetically coupled shunt resonators, like their simple series-tuned counterparts, exhibit design trends opposed to the need for tuning. Low Q and large capacitance provide the best ripple attenuation, and once again a three-parameter search may be necessary to find suitable components. A limit on capacitor values imposed by an application constraint, however, can favor high-Q magnetically coupled resonators.

6.2 Conclusions: integrated filter elements

The performance trends of magnetically coupled shunt resonators favor shunt paths with small impedance at as many frequencies as possible. Such small impedance can be achieved with a low-Q resonance, but is realized even more effectively by the attempted *elimination* of resonance. This strategy is employed by integrated filter elements, which use self-inductances in excess of winding mutual inductance to develop a negative inductance in the shunt branch of their equivalent T-model. Shunt ESL cancellation is always accompanied by the introduction of positive inductance in series with the other two terminals, an extra reactance with which the designer can realize higher-order filters than with a capacitor alone.

In proof-of-concept experiments, ESL cancellation and the accompanying introduction of series reactance were measured. A high-frequency comparison between the ripple attenuation of a capacitor and an integrated element demonstrated a cancellation of the electrolytic capacitor's ESL to approximately 15% of its original value. The series inductance was significant enough to roll off at 50 kHz with 40 μF , providing a ripple gain 20 dB better than a filter employing only capacitors. These results were measured for film windings on the outside of a capacitor package, demonstrating the compatibility of the approach with eddy-current losses in the electrolytic element. The windings comprised a foil winding

Conclusions

under a single layer of 18 gauge wire, and the capacitor canister, with a 2 mm gap between the inside package surface and plate windings, could have accommodated this winding build internally.

The repeatability of magnetic coupling was addressed by impedance measurements of collection of capacitors with the same part number. ESL varied $\pm 2.4\%$ for all the units measured, compared to a deviation in the ESR for some units of 175% from the mean value. The inductance of wound foils depends on geometry rather than material or contact properties which can strongly impact losses, suggesting that inductance cancellation to within a few percent of initial ESL should be practical.

6.3 Further work

Active-tuning controls need to be developed with a manufacturer's attention to detail and economy. Unnecessary complexity in the control circuitry of the prototype system should be streamlined, and long-term reliability and performance assessed. Integration of tuning and PWM controls indeed seems possible, but further experience with the pathologies of the tuning controller would be a reasonable prerequisite for IC design. Resonators would profit from further general treatment, aimed at resolving, for instance, under what conditions series-tuned resonators become attractive alternatives to capacitors. Resonator design with ferrite materials has hardly been considered, but may offer pleasant results for high ripple-ratio resonators. Active tuning further stands in a unique position to realize the filtering benefits of hybrid reactive structures which introduce high-Q resonances at predictable, but not tightly controlled, frequencies.

The preliminary measurements on integrated elements illustrate the advantages of inductance-cancellation methods. External windings do not demonstrate thoroughly, however, the thermal viability of the integrated filter element. The inclusion of DC conduction paths within a power capacitor may have implications for the plate materials, plate geometry, and the capacitor's power rating, all of which remain unexplored. ESL is relatively constant unit to unit at constant temperature, but the variation in coil dimensions (hence ESL) over a range of operating conditions needs to be characterized. The experiments conducted so far suggest that the attenuation benefit of inductance cancellation is insensitive to eddy-current dissipation within the capacitor. Should ESL cancellation be achieved more precisely, such losses may be found to place a lower bound on repeatable performance. So too, sensitivity to coupling changes in the presence of magnetic materials, while

not detectable in the initial measurements, may prove important as the technology progresses. Most significantly, winding construction methods, whether external, printed on the surrounding PCB, or integrated in the capacitor package, need to be fully evaluated.

Appendix A
MATLAB files

coredata.m	Micrometals iron-powder toroid database	page 103
wiredata.m	Wire database	page 109
coredesign.m	Iterative core-selection algorithm	page 112
convergence.m	Script to check convergence of iterative network design based on exponential loss model for inductors	page 115
coreloss.m	Core loss computation	page 116
acresistance.m	AC resistance calculation (copper loss and core loss)	page 117
dcresistance.m	DC resistance calculation (copper loss)	page 117
permBpk.m	Percent initial permeability vs. peak AC flux density	page 118
permH.m	Percent initial permeability vs. DC magnetizing force	page 119
iripple.m	Phasor analysis to determine ripple	page 119
par.m	parallel-impedance calculation	page 120
general.m	script to determine best balance of resonator Q and tuning-point loss, in normalized coordinates	page 120

A.1 coredata.m

% Micrometals toroidal core data

% awgcol: hash table associating AWG (first column) with a column

% of winding tables (singlewinding and fullwinding below)

`awgcol = [18 3 ; 17 4 ; 16 5 ; 15 6 ; 14 7 ; 13 8 ; 12 9 ; 11 10 ; 10 11];`

MATLAB files

```
% singlewinding: winding table for T-series cores
% single-layer winding
% columns:
%   core   T-series #
%   suffix 0 none  1 A  2 B  3 C  4 D
%   turns  18 to 10 AWG
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, p.64
singlewinding = [ ...
  80 0 30 27 23 20 17 15 13 11 9 ; ...
  80 4 30 27 23 20 17 15 13 11 9 ; ...
  90 0 34 30 26 23 20 17 15 13 11 ; ...
  94 0 35 31 27 24 21 18 15 13 11 ; ...
  106 0 36 31 27 24 21 18 15 13 11 ; ...
  106 1 36 31 27 24 21 18 15 13 11 ; ...
  106 2 36 31 27 24 21 18 15 13 11 ; ...
  124 0 46 40 36 31 27 24 21 18 15 ; ...
  130 0 51 45 40 35 31 27 23 20 17 ; ...
  130 1 51 45 40 35 31 27 23 20 17 ; ...
  131 0 41 36 32 28 24 21 18 16 13 ; ...
  132 0 45 40 35 31 27 23 20 18 15 ; ...
  141 0 59 52 46 40 35 31 27 24 20 ; ...
  150 0 56 49 44 38 34 29 26 22 19 ; ...
  157 0 64 56 50 44 39 34 30 26 23 ; ...
  175 0 73 64 57 50 44 39 34 30 26 ; ...
  184 0 63 56 50 44 38 34 29 26 22 ; ...
  200 0 86 76 67 60 53 46 41 36 31 ; ...
  200 2 86 76 67 60 53 46 41 36 31 ; ...
  201 0 63 56 50 44 38 34 29 26 22 ; ...
% 224 3 0 0 0 0 0 0 0 0 0 ; ...
  225 0 97 86 76 68 60 53 46 41 36 ; ...
% 249 0 0 0 0 0 0 0 0 0 0 ; ...
  250 0 86 76 67 60 53 46 41 36 31 ; ...
  300 0 136 121 108 96 85 75 66 58 52 ; ...
  300 4 136 121 108 96 85 75 66 58 52 ; ...
  400 0 160 142 126 113 100 88 78 69 61 ; ...
  400 4 160 142 126 113 100 88 78 69 61 ; ...
  520 0 221 197 176 156 139 123 109 97 86 ; ...
  520 4 221 197 176 156 139 123 109 97 86 ; ...
  650 0 250 223 199 177 158 139 124 110 98 ; ...
];

% fullwinding: winding table for T-series cores
% 45% of toroid ID remaining
% columns:
```

```

%   core   T-series #
%   suffix 0 none  1 A  2 B  3 C  4 D
%   turns  18 to 10 AWG
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, p.65
fullwinding = [ ...
  80 0 57 45 36 29 23 18 14 11 9 ; ...
  80 4 57 45 36 29 23 18 14 11 9 ; ...
  90 0 70 56 45 36 28 22 18 14 11 ; ...
  94 0 73 58 46 37 29 23 18 14 11 ; ...
  106 0 75 60 48 38 30 24 19 15 12 ; ...
  106 1 75 60 48 38 30 24 19 15 12 ; ...
  106 2 75 60 48 38 30 24 19 15 12 ; ...
  124 0 117 93 75 60 47 37 30 23 19 ; ...
  130 0 142 113 90 72 57 45 36 28 23 ; ...
  130 1 142 113 90 72 57 45 36 28 23 ; ...
  131 0 95 76 61 48 38 30 24 19 15 ; ...
  132 0 114 91 73 58 46 36 29 23 18 ; ...
  141 0 180 144 115 92 73 57 46 36 29 ; ...
  150 0 166 132 106 85 67 53 42 33 27 ; ...
  157 0 210 168 134 107 85 67 53 42 34 ; ...
  175 0 267 213 170 136 108 85 68 54 43 ; ...
  184 0 210 168 134 107 85 67 53 42 34 ; ...
  200 0 365 290 232 186 148 116 93 74 59 ; ...
  200 2 365 290 232 186 148 116 93 74 59 ; ...
  201 0 210 168 134 107 85 67 53 42 34 ; ...
% 224 3 0 0 0 0 0 0 0 0 0 ; ...
  225 0 461 367 294 235 186 147 117 93 74 ; ...
% 249 0 0 0 0 0 0 0 0 0 0 ; ...
  250 0 365 290 232 186 148 116 93 74 59 ; ...
  300 0 870 693 554 443 352 278 221 176 140 ; ...
  300 4 870 693 554 443 352 278 221 176 140 ; ...
  400 0 1182 942 754 602 479 378 301 240 191 ; ...
  400 4 1182 942 754 602 479 378 301 240 191 ; ...
  520 0 2261 1765 1413 1129 898 708 564 450 358 ; ...
  520 4 2261 1765 1413 1129 898 708 564 450 358 ; ...
  650 0 2861 2280 1824 1458 1159 914 729 581 463 ; ...
];
% Cores: geometrical and dissipation data for toroidal cores
% columns:
%   core   T-series #
%   suffix 0 none  1 A  2 B  3 C  4 D
%   OD      (mm)
%   ID      (mm)

```

MATLAB files

```
% height (mm)
% l mag path length (cm)
% A mag xsection (cm^2)
% V mag volume (cm^3)
% MLT mean length per turn (cm/turn)
% W40C watts dissipation for 40C temp rise
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, pp.5-13 and 64,65
Cores = [...
  80 0 20.2 12.6 9.53 5.14 .347 1.78 2.80 1.30 ; ...
  80 4 20.2 12.6 12.7 5.14 .453 2.33 4.07 1.84 ; ...
  90 0 22.9 14.0 9.53 5.78 .395 2.28 3.64 1.88 ; ...
  94 0 23.9 14.2 7.92 5.97 .362 2.16 3.44 1.85 ; ...
  106 0 26.9 14.5 11.1 6.49 .659 4.28 4.49 2.59 ; ...
  106 1 26.9 14.5 7.92 6.49 .461 3.00 3.86 2.25 ; ...
  106 2 26.9 14.5 14.6 6.49 .858 5.57 5.19 2.97 ; ...
  124 0 31.6 18.0 7.11 7.75 .459 3.55 3.95 2.79 ; ...
  130 0 33.0 19.8 11.1 8.28 .698 5.78 4.75 3.53 ; ...
  130 1 33.0 19.8 5.72 8.28 .361 2.99 3.67 2.78 ; ...
  131 0 33.0 16.3 11.1 7.72 .885 6.84 5.11 3.52 ; ...
  132 0 33.0 17.8 11.1 7.96 .805 6.41 4.95 3.53 ; ...
  141 0 35.9 22.4 10.5 9.14 .674 6.16 4.75 3.92 ; ...
  150 0 38.4 21.5 11.1 9.38 .887 8.31 5.28 4.45 ; ...
  157 0 39.9 24.1 14.5 10.1 1.06 10.7 5.89 5.29 ; ...
  175 0 44.5 27.2 16.5 11.2 1.34 15.0 6.58 6.16 ; ...
  184 0 46.7 24.1 18.0 11.2 1.88 21.0 7.54 7.47 ; ...
  200 0 50.8 31.8 14.0 13.0 1.27 16.4 6.50 7.61 ; ...
  200 2 50.8 31.8 25.4 13.0 2.32 30.0 8.78 10.1 ; ...
  201 0 50.8 24.1 22.2 11.8 2.81 33.2 8.90 9.28 ; ...
  % 224 3 57.2 31.8 19.1 14.0 2.31 32.2 0 0 ; ...
  225 0 57.2 35.7 25.4 14.6 2.59 37.8 6.93 9.16 ; ...
  % 249 0 63.5 35.7 25.4 15.6 3.36 52.3 0 0 ; ...
  250 0 63.5 31.8 25.4 15.0 3.84 57.4 10.4 13.9 ; ...
  300 0 77.2 49.0 12.7 19.8 1.68 33.4 7.95 14.5 ; ...
  300 4 77.2 49.0 25.4 19.8 3.38 67.0 10.5 18.7 ; ...
  400 0 102 57.2 16.5 25.0 3.46 86.4 11.1 25.2 ; ...
  400 4 102 57.2 33.0 25.0 6.85 171 14.4 32.1 ; ...
  520 0 132 78.2 20.3 33.1 5.24 173 13.7 41.5 ; ...
  520 4 132 78.2 40.6 33.1 10.5 347 17.7 52.7 ; ...
  650 0 165.1 88.9 50.8 39.9 18.4 734 23.1 82.5 ; ...
];

Cores=sortrows(Cores,8); % sort in order of ascending magnetic volume
Cores(:,3:5)=Cores(:,3:5)*0.1; % convert to cgs
```

```

% Materials: basic materials data
% columns:
%   Micrometals mix #
%   relative permeability
%   density (g/cc)
%   relative cost
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, p.1
Materials = [...
    2 10 5.0 2.7 ; ...
    8 35 6.5 5.0 ; ...
    18 55 6.6 3.4 ; ...
    26 75 7.0 1.0 ; ...
    28 22 6.0 1.9 ; ...
    33 33 6.3 1.6 ; ...
    38 85 7.1 1.1 ; ...
    40 60 6.9 1.0 ; ...
    45 100 7.2 2.6 ; ...
    52 75 7.0 1.4 ; ...
];
150

% Permvsh: coefficients for percent permeability vs. DC magnetizing force
% columns:
%   Micrometals mix #
%   a coefficient
%   b coefficient
%   c coefficient
%   d coefficient
%   e coefficient
% %mu_{0} = ((a + c*H + e*H^2)/(1 + b*H + d*H^2))^0.5
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, p.28
Permvsh = [...
    2 10000 -4.99e-3 -49.5 9.16e-6 0.0865 ; ...
    8 10090 4.26e-3 30.9 7.68e-5 -0.0119 ; ...
    18 9990 8.36e-4 14.4 3.92e-4 0.0853 ; ...
    26 10090 5.05e-3 13.1 1.17e-3 0.0212 ; ...
    28 10140 4.68e-4 -30.2 1.45e-5 0.0505 ; ...
    33 10200 5.12e-3 7.39 9.62e-5 0.0298 ; ...
    38 9960 -1.53e-4 5.87 1.17e-3 -0.0256 ; ...
    40 10240 4.32e-3 12.8 6.26e-4 0.0267 ; ...
    45 10014 6.07e-3 45.2 1.79e-3 -0.0578 ; ...
    52 10240 6.71e-3 24.7 7.75e-4 -0.0105 ; ...
];
170
180

```

MATLAB files

```
% Percent permeability vs. peak AC flux density
% columns:
%   Micrometals mix #
%   a coefficient
%   b coefficient
%   c coefficient
%   d coefficient
%   e coefficient
%  $\mu_{0} = ((a + c*B + e*B^2)/(1 + b*B + d*B^2))^{0.5}$  for 2,8,18,26,40
%  $\mu_{0} = a + b*B + c*B^{0.5} + d*B^2$  for 30, 33, 38, 45, 52
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, p.28
PermvBpk = [...
  2  9970  5.77e-4  7.29 -8.96e-8 -1.18e-3 ; ...
  8  9990  4.52e-4  11.4  8.82e-9 -8.29e-4 ; ...
 18 10270  1.01e-4  12.3  2.70e-8 -8.43e-4 ; ...
 26 10600  7.21e-5  37.8 -7.74e-9 -3.56e-3 ; ...
 28  93.4 -2.99e-2  2.08  8.30e-7      0 ; ...
 33  92.6 -2.51e-2  2.36  1.07e-7      0 ; ...
 38  90.7  4.01e-3  3.15 -2.18e-6      0 ; ...
 40 10480  1.62e-4  40.8 -6.51e-9 -3.35e-3 ; ...
 45  88.3  5.78e-3  3.80 -2.72e-6      0 ; ...
 52  92.0  1.34e-2  2.77 -3.66e-6      0 ; ...
];

% Coreloss: Core loss vs. peak AC flux density
% columns:
%   material
%   a coefficient
%   b coefficient
%   c coefficient
%   d coefficient
%  $CL (mW/cm^3) = f / (a*B^{-3} + b*B^{-2.3} + c*B^{-1.65}) + d*f^2*B^2$ 
% cf. Micrometals "Power Conversion and Line Filter Applications" Issue J
% January 2001, p.28
Coreloss = [...
  2  4.0e9  3.0e8  2.7e6  8.0e-15 ; ...
  8  1.9e9  2.0e8  9.0e5  2.5e-14 ; ...
 18  8.0e8  1.7e8  9.0e5  3.1e-14 ; ...
 26  1.0e9  1.1e8  1.9e6  1.9e-13 ; ...
 28  3.0e8  3.2e7  1.9e6  3.1e-13 ; ...
 33  3.4e8  2.0e7  2.0e6  3.7e-13 ; ...
 38  1.2e9  1.3e8  1.9e6  3.2e-13 ; ...
 40  1.1e9  3.3e7  2.5e6  3.1e-13 ; ...
 45  1.2e9  1.3e8  2.4e6  1.2e-13 ; ...
];
```



```
52 1.0e9 1.1e8 2.1e6 6.9e-14 ; ...
];
```

A.2 *wiredata.m*

```
% Wire data in cgs
% KSV p. 582
%
% AWG
% Diameter (cm)
% ohm/cm (75deg C)
% g/cm
% turns/(sq cm)

% AWG 10
% Diameter (mm)
% ohm/km (75deg C)
% kg/km
% turns/(sq cm)

Wire = [...
  0 8.25 0.392 475 0 ; ...
  1 7.35 0.494 377 0 ; ...
  2 6.54 0.624 299 0 ; ...
  3 5.83 0.786 237 0 ; ... 20
  4 5.19 0.991 188 0 ; ...
  5 4.62 1.25 149 0 ; ...
  6 4.12 1.58 118 0 ; ...
  7 3.67 1.99 93.8 0 ; ...
  8 3.26 2.51 74.4 0 ; ...
  9 2.91 3.16 59.0 0 ; ...
  10 2.59 3.99 46.8 14 ; ...
  % 11 2.31 5.03 37.1 17 ; ...
  12 2.05 6.34 29.4 22 ; ...
  % 13 1.83 7.99 23.3 27 ; ... 30
  14 1.63 10.1 18.5 34 ; ...
  % 15 1.45 12.7 14.7 40 ; ...
  16 1.29 16.0 11.6 51 ; ...
  % 17 1.15 20.2 9.23 63 ; ...
```

MATLAB files

```
18 1.02 25.5 7.32 79 ; ...
% 19 0.912 32.1 5.80 98 ; ...
20 0.812 40.5 4.60 123 ; ...
% 21 0.723 51.1 3.65 153 ; ...
22 0.644 64.4 2.89 192 ; ...
% 23 0.573 81.2 2.30 237 ; ...
24 0.511 102 1.82 293 ; ...
% 25 0.455 129 1.44 364 ; ...
26 0.405 163 1.15 454 ; ...
% 27 0.361 205 1.10 575 ; ...
28 0.321 259 1.39 710 ; ...
% 29 0.286 327 1.75 871 ; ...
30 0.255 412 2.21 1090; ...
];

% convert to cgs
Wire(:,2)=Wire(:,2)*0.1;
Wire(:,3)=Wire(:,3)*10e-5;
Wire(:,4)=Wire(:,4)*0.01;

wirearea=pi*Wire(:,2).*Wire(:,2)*0.25; %cm^2
pack=(wirearea.*Wire(:,5))*0.01; %packing factor
```

A.3 inductance.m

```
% inductance.m
% returns inductance for a given material, core index, turns, irpp, idc, L, fsw
% and awg index

function [ind,Q,Feloss]=inductance(matind,coreind,N,Irpp,Idc,Iacrms,L,fsw,awgind)
global Cores Coreloss Permvsh Permvsbpk Materials Wire wirearea;

mu0=4*pi*1e-7; %H/m

% compute Hdc (DC magnetizing force)
l=Cores(coreind,6); %mean magnetic path length
Hdc=0.4*pi*N*Idc/l; %H (oersted)

% find % initial permeability due to DC saturation
percpermHdc=permH(matind,Hdc);

% compute Bpk (peak AC flux density)
A=Cores(coreind,7);
```

```

Bpk=L*Irpp*1e8/2/A/N;
                                                                    20

% find % initial permeability due to peak AC flux
percpermBpk=permBpk(matind,Bpk);

% compute effective permeability
mur=Materials(matind,2);
mu=mu0*mur*(percpermHdc/100)*(percpermBpk/100);
percperm=mu/mu0/mur*100;

% compute L with this core, material, and winding
Amks=A*0.0001;           % 10000cm^2 = 1m^2
lmks=1*0.01;           % 100cm = 1m
ind=mu*Amks*N*N/lmks;
                                                                    30

%now iterate AC-flux loss computation

% compute Bpk (peak AC flux density)
Bpk=ind*Irpp*1e8/2/A/N;

% find % initial permeability due to peak AC flux
percpermBpk=permBpk(matind,Bpk);
                                                                    40

% compute effective permeability
mu=mu0*mur*(percpermHdc/100)*(percpermBpk/100);
percperm=mu/mu0/mur*100;

% compute L with this core, material, and winding (iteration)
ind=mu*Amks*N*N/lmks;

% compute core loss
                                                                    50
Feloss=coreloss(matind,Bpk,fsw)*Cores(coreind,8)*0.001; %watts
      %0.001 mW to W
MLT=Cores(coreind,9);
Rac=acresistance(MLT*N+10,awgind,fsw);
Rdc=dccresistance(MLT*N+10,awgind);

acdispipation=Feloss+Iacrms*Iacrms*Rac;
%dispipation=Feloss+Iacrms*Iacrms*Rac+Idc*Idc*Rdc
eqR=acdispipation/(Iacrms)^2; % +Rdc;
Q=2*pi*fsw*ind/eqR;
                                                                    60

%Q=2*pi*0.5*ind*Imax*Imax/(T*(Feloss+Iacrms*Iacrms*Rac+Idc*Idc*Rdc));

```

A.4 `coredesign.m`

```
% coredesign.m  
% design a toroidal inductor for minimum size (lowest T series)
```

```
% inputs:  
Irpp=0.565; %amps  
Idc=25; %amps  
Iacrms=0.125 %amps - used to compute AC Cu loss  
L=170e-6; %henry  
fsw=100e3; %Hz
```

```
Jmax=500; %A/cm^2 maximum current density
```

```
global Coreloss Permvsh PermvshBpk Materials Wire Cores wirearea;
```

```
load wiredata;  
load coredata;
```

```
nonefound=NaN;
```

```
Lbest=nonefound*ones(length(Materials(:,1)),1);  
corebest=Lbest;  
Nbest=Lbest;  
Qbest=Lbest;  
Felossbest=Lbest;
```

```
% .....  
% choose wire
```

```
Jwire=Idc./wirearea;
```

```

for k=1:length(wirearea),
    if Jwire(k)>Jmax
        break
    end
    awgind=k;
end
awg=Wire(awgind,1);

% .....
% choose material

for matind=1:10,
    material=Materials(matind,1);
    strikes=0;

% .....
% design inductor for smallest T-series
% core

% choose core (start big and work down)

for coreind=length(Cores(:,1))-1:1;

    for k=1:length(awgcol(:,1)),           %get column index to
        if awgcol(k,1)==awg               %winding table
            c=awgcol(k,2);
        end
    end

    % N = single-winding turns
    N=singlewinding(coreind,c);

    % N = "full-wound" turns (45% ID remaining)
    %N=fullwinding(coreind,c);

    [Lnew,Qnew,Feloss]=inductance(matind,coreind,N,Irpp,Idc,Iacrms,L,fsw,awgind);

    % throw away designs with high dissipation
    if Feloss<Cores(coreind,10)
        if Lnew>L
            Lbest(matind)=Lnew;
            Nbest(matind)=N;
            corebest(matind)=coreind;
        end
    end
end
end

```

MATLAB files

```
    Qbest(matind)=Qnew;
    Felossbest(matind)=Feloss;
else
    strikes=strikes+1;
end
else
    strikes=strikes+1;
end

if strikes>=5

    if isnan(Lbest(matind))==1 % we never found a good core
        break
    end

    % now take off turns to match target L and reduce losses
    bestcoreind=corebest(matind);
    % for easy reading, output T-series number, not index
    corebest(matind)=Cores(bestcoreind,1)+Cores(bestcoreind,2)/10;
    for Nnew=Nbest(matind)-1:-1:1,
        [Lnew,Qnew,Feloss]=inductance(matind,bestcoreind,Nnew,Irpp,Idc,Iacrms,L,fsw,awgind);
        if Lnew<L
            break % keep last: we took off too many turns
        end
        if Qnew<Cores(coreind,10)
            break
        end
        Lbest(matind)=Lnew;
        Nbest(matind)=Nnew;
        Qbest(matind)=Qnew;
        Felossbest(matind)=Feloss;
    end
    break
end

end

end

[Materials(:,1) corebest Nbest 1e6*Lbest Qbest Felossbest]
```

A.5 convergence.m

```

% convergence.m
% uses simple loss model,  $Q=4.4025 \cdot I_{rpp}^{-2.387}$ , regardless of inductance,
% for  $I_{dc}=25A$  and  $f_{sw}=100kHz$ 

fsw=100e3;
T=1/fsw;
fres=fsw;
wres=2*3.14159*fres;
Pout=300;
Vout=12;
Iout=Pout/Vout;
RL=Vout*Vout/Pout;
Ltotal=200e-6;

QL1=20;
QL2=20;

ripp3=0;
ripptarget=0
irms3=0;
irmstarget=3;

p=0;
for D=0.2:0.1:0.8,
p=p+1;
m=0;
%for  $L1=Ltotal*0.1:Ltotal*0.1:Ltotal*0.9$ ,
L1=120e-6

QL1=20;
QL2=20;

m=m+1;
L2=Ltotal-L1;
C=1/wres/wres./L2;
CR=0.003;

iterations=10
for n=1:iterations,

L1R=wres*L1/QL1;
L2R=wres*L2/QL2;

```

MATLAB files

```
ZL1=tf([L1 0],1)+L1R;
ZL2=tf([L2 0],1)+L2R;
ZC=tf(1,[C 0])+CR;

%eq source to L1 current
G1=1/(ZL1+par(ZL2,ZC)+RL);
%eq source to L2 current
G2=ZC/(ZC+ZL2)*G1;

[rippestimate1(p,m,n),irms]=ioutripple(fsw,Vout/D,G1,D,Iout);
QL1=4.4025*rippestimate1(p,m,n)^(-2.387);
QL1history(p,m,n)=QL1;

[rippestimate2(p,m,n),irms]=ioutripple(fsw,Vout/D,G2,D,Iout);
QL2=4.4025*rippestimate2(p,m,n)^(-2.387);
QL2history(p,m,n)=QL2;

end

%end

end

iteration=1:iterations;
plot(iteration,QL1history(1,6,:), 'k-', iteration,QL2history(1,6,:), 'k- . ')
plot(iteration,rippestimate1(1,6,:), 'k-', iteration,rippestimate2(1,6,:), 'k- . ')
```

50

60

70

A.6 coreloss.m

```
% coreloss.m
% returns core loss in mW/cm^3 for a given peak AC flux density Bpk (gauss)
% and frequency (Hz)

function [loss]=coreloss(matindex,B,f)
global Coreloss

a=Coreloss(matindex,2);
b=Coreloss(matindex,3);
c=Coreloss(matindex,4);
```

10

```
d=Coreloss(matindex,5);

loss=f./(a./(B.^3)+b./(B.^2.3)+c./(B.^1.65)) + d*f.*B.*B;
```

A.7 *acresistance.m*

```
% acresistance.m
% compute resistance of a length of wire (cm) at frequency f (Hz)
% length      wire length (cm)
% awg        wire gauge index
% fsw        frequency for AC resistance evaluation (Hz)

function [acres]=acresistance(length,awgind,f)
global Wire wirearea

rho=1.56e-8;      %Cu resistivity ohm-m
rho=rho*100;     %ohm-cm
mur=1;           %relative permeability
mu0=4*pi*1e-7*0.01; %H/cm

skind=sqrt(2*rho/2/pi/f/mu0/mur); %skin depth (cm)

DCarea=wirearea(awgind);
DCres=length/DCarea*rho;
ACradius=Wire(awgind,2)/2-skind;
if ACradius<0
    ACradius=0;
end
ACarea=DCarea-(3.14159*ACradius^2);

acres=DCarea/ACarea*DCres;
```

A.8 *dcreistance.m*

```
% dcreistance.m
% compute resistance of a length of wire (cm)
% length      wire length (cm)
% awg        wire gauge index

function [DCres]=dcreistance(length,awgind)
```

MATLAB files

```
global Wire wirearea
```

```
rho=1.56e-8;      %Cu resistivity ohm-m  
rho=rho*100;     %ohm-cm
```

10

```
DCarea=wirearea(awgind);  
DCres=length/DCarea*rho;
```

A.9 permBpk.m

```
% permBpk.m  
% returns percent initial permeability vs. peak AC flux density (gauss)
```

```
function [percent]=permBpk(matindex,B)  
global PermvBpk Materials;
```

```
a=PermvBpk(matindex,2);  
b=PermvBpk(matindex,3);  
c=PermvBpk(matindex,4);  
d=PermvBpk(matindex,5);  
e=PermvBpk(matindex,6);
```

10

```
m=Materials(matindex,1);
```

```
if m==8 | m==18 | m==26 | m==40  
    percent=sqrt(real((a+c*B+e*B.*B)./(1+b*B+d*B.*B)));  
elseif m==28 | m==33 | m==38 | m==45 | m==52  
    percent=a+b*B+c*B.^0.5+d*B.*B;  
else % a hack for the -2 curve  
    percent=sqrt(((a+c*B+e*B.*B)./(1+b*B+d*B.*B)));  
    for i=1:length(percent),  
        if i>0.93*length(percent)  
            percent(i)=0;  
        end  
    end  
end  
end
```

20

30

A.10 permH.m

```
% permH.m
% returns percent initial permeability vs. H (oersted)

function [percent]=permH(matindex,H)
global PermvH

a=PermvH(matindex,2);
b=PermvH(matindex,3);
c=PermvH(matindex,4);
d=PermvH(matindex,5);
e=PermvH(matindex,6);

percent=sqrt((a+c*H+e*H.*H)/(1+b*H+d*H.*H));
```

10

A.11 iripple.m

```
function [irpp,irms]=iripple(fsw,amp,G,D,dc)
% Phasor analysis: evaluate response of G to a periodic gate function
% fsw   source freq (Hz)
% amp   source amplitude
% G     transfer function to evaluate
% D     source positive duty ratio (0 to 1)
% dc    DC current (A)

% for stand-alone testing
%fsw=100e3;
%D=0.2;
%amp=12/D;

harmonics=100;
ppp=100;           %points per period
n=0:1:harmonics;
T=1/fsw;
t=0:T/ppp:T;
wsw=2*pi*fsw;
```

10

MATLAB files

```
F=amp*2*D*sin(n*3.14159*D)./n/3.14159/D;
F(1)=D*amp;
%src=0;
%Gsrc=tf(1);
out=0;
for r=1:harmonics,
    out=out+abs(F(r+1)*evalfr(G,j*r*sw))*cos(r*sw*t+angle(F(r+1)*evalfr(G,j*r*sw)));
%   src=src+abs(F(r+1)*evalfr(Gsrc,j*r*sw))*cos(r*sw*t+angle(F(r+1)*evalfr(Gsrc,j*r*sw)));
end

% plot(out);
% pause

% for stand-alone testing
%plot(out)
%hold on
%plot(src./amp*2)
%hold off
%axis([0,ppp,-1.8,1.8])
irpp=2*max(abs(out));
out=out+idc; % add DC to ripple: note that we don't compute
              % DC from F(1) and tf's evaluated at 0 rps because
              % tf's contain an AC resistance

%for stand-alone testing
%plot(out)
%hold on
%plot(src./amp*2)
%hold off
%axis([0,ppp,-1.8,1.8])
irms=sqrt(out*out'*(T/ppp)/T);
```

A.12 par.m

```
function [parallel] = par(arg1,arg2)
parallel = 1/(1/arg1+1/arg2);
```

A.13 general.m

```

% Generalized low-pass design
% For different values of resonator ESR and Q,
% find locus where resonant network matches simple low-pass
% p-p ripple performance at D=0.3

C1=1;
ZC1=tf(1,[C1 0]);

L=100;
ZL=tf([L 0],1);
H1=ZC1/(ZC1+ZL);

wres=1;

[pptarget,rmstarget]=iripple(wres/2/pi,1,H1,0.3,0.3)

capneeded=zeros(8,16);

m=1;
n=1;
for r=0.05:0.05:0.75,
    for Q=8:2:45,

        Lr=r*Q/wres;
        Cr=1/wres/wres/Lr;
        ZCr=tf(1,[Cr 0]);
        ZLr=tf([Lr 0],1);
        Zres=ZCr+ZLr+r;

        stepfact=0.9;
        C2=C1;
        steps=0
        while steps<30,
            ZC2=tf(1,[C2 0]);
            H2=par(ZC2,Zres)/(par(ZC2,Zres)+ZL);
            [thispp,thisrms]=iripple(wres/2/pi,1,H2,0.3,0.3);
            error=-thispp+pptarget;
            if error/pptarget>.5
                C2=C2*0.7;
            elseif error/pptarget>.2
                C2=C2*0.8;
            elseif error/pptarget>.1
                C2=C2*0.9
            elseif error/pptarget>.01
                C2=C2*0.95

```

MATLAB files

```
elseif abs(error/pptarget)<0.01
    capneeded(m,n)=C2;
    break;
elseif error/pptarget>-0.05
    C2=C2*1.01;
elseif error/pptarget<-0.05
    C2=C2*1.04;
end
steps=steps+1;
end

n=n+1;
end
n=1
m=m+1
end
```

50

60

Appendix B

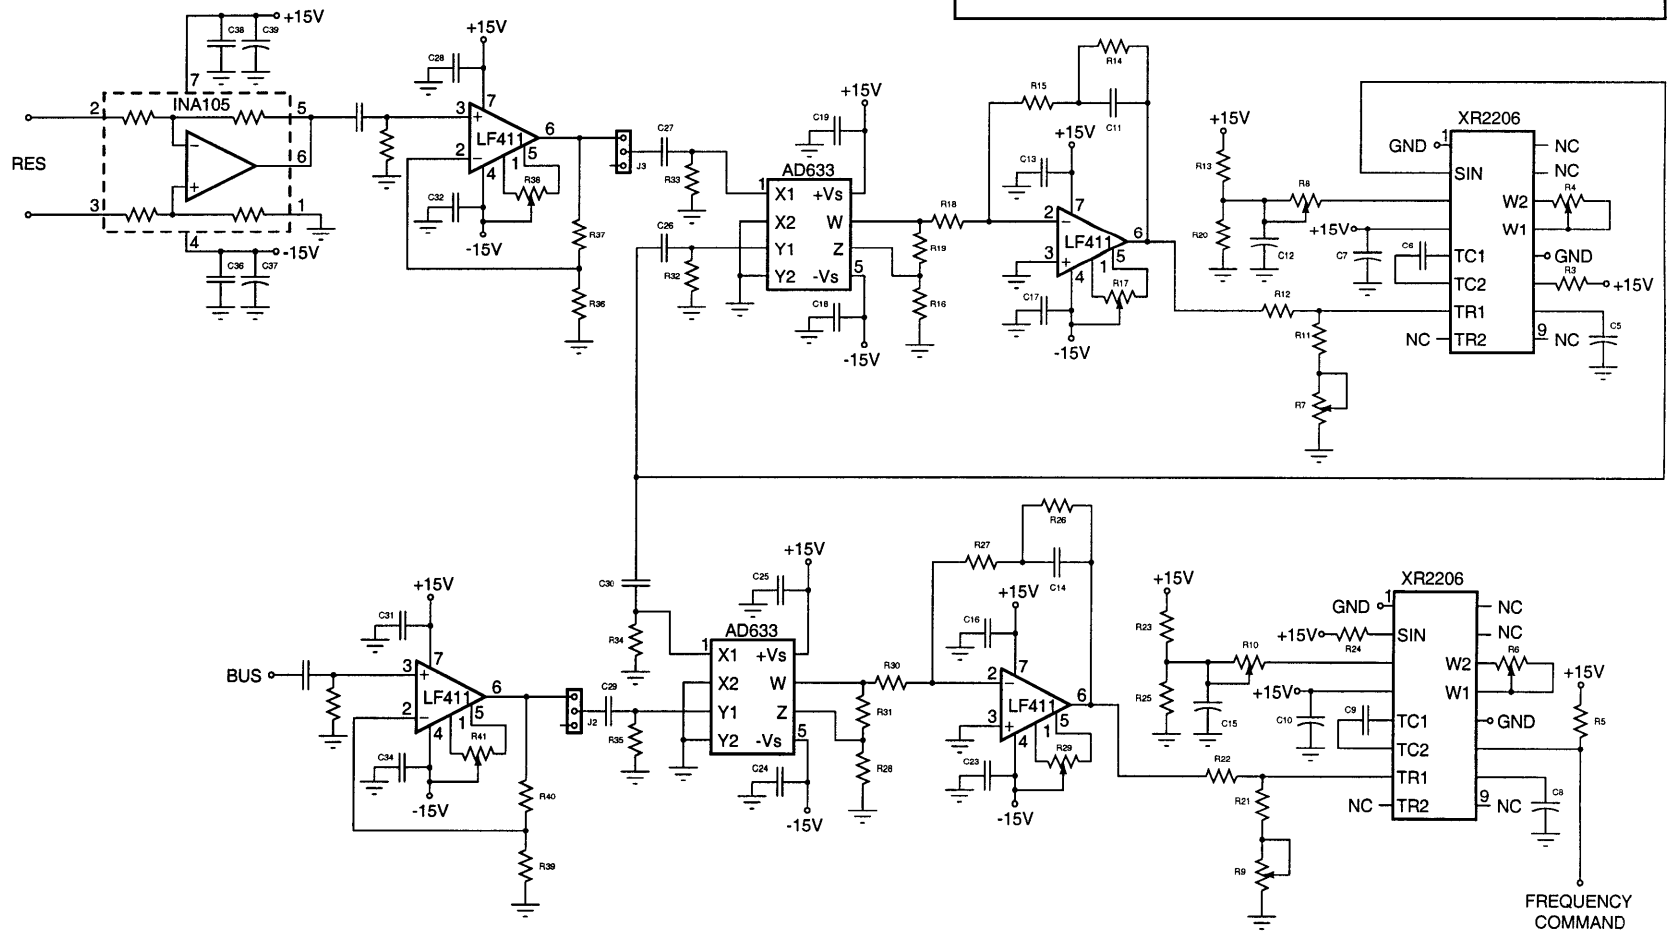
Phase-lock tuning circuit

Phase-lock tuning controller schematic	page 124
PCB top copper layer	page 126
PCB bottom copper layer	page 127
Silkscreen	page 128
Parts list	page 129


REV 12.26.2000

JOSHUA PHINNEY


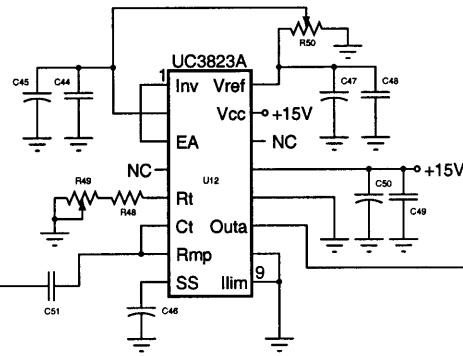



PHASELOCK ACTIVE-TUNING CONTROLLER
FREQUENCY CONTROL




EXTRA INVERTING AMPLIFIER




REV 12.26.2000	JOSHUA PHINNEY
PHASELOCK ACTIVE-TUNING CONTROLLER PWM CONTROLLER & SWITCHING STAGE	

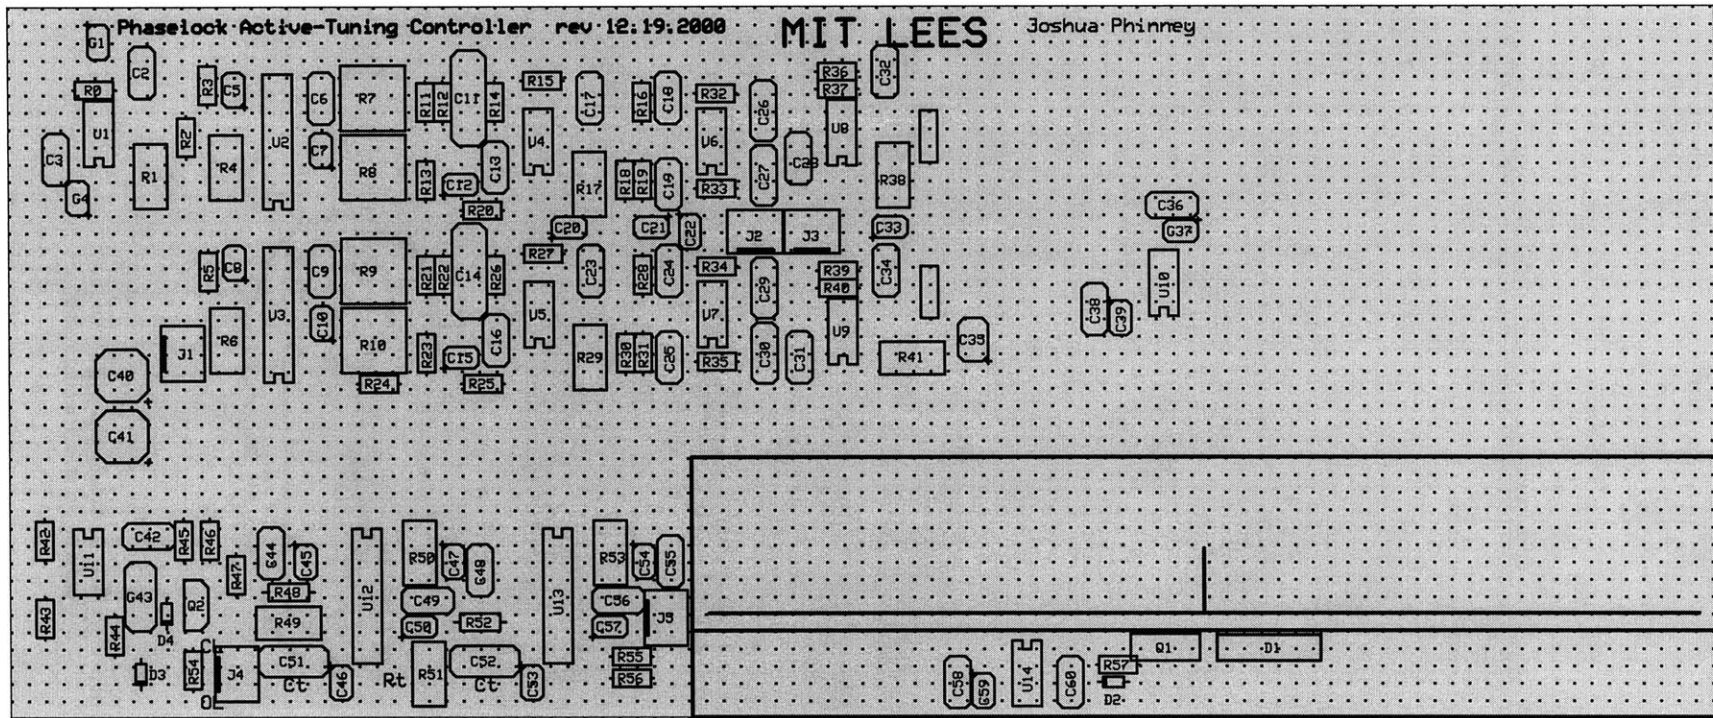




Top copper layer



Bottom copper layer



Slikscreen

Component	Value	Notes
$C_2, C_3, C_{13}, C_{16}, C_{17}, C_{18}, C_{19}, C_{23}, C_{24}, C_{25}, C_{28}, C_{31}, C_{32}, C_{34}, C_{36}, C_{38}, C_{42}, C_{44}, C_{48}, C_{49}, C_{52}, C_{55}, C_{56}, C_{58}$	0.1 μF	monolithic ceramic
$C_{37}, C_{39}, C_{44}, C_{47}, C_{50}, C_{54}, C_{57}, C_{59}$	10 μF	16 V tantalum
$C_5, C_6, C_7, C_8, C_9, C_{10}, C_{46}, C_{53}$	1 μ	16 V tantalum
C_{40}, C_{41}	22 μF	16 V tantalum
$C_{20}, C_{21}, C_{22}, C_{33}$	4.7 μF	16 V tantalum
$C_{26}, C_{27}, C_{29}, C_{30}$	0.01 μF	Z5U ceramic disc
C_6, C_9	0.0022 μF	polyester
C_{11}, C_{14}	0.22 μF	
C_{51}, C_{52}	select	C_t PWM timing
C_{43}	100 pF	ceramic disc
$R_1, R_{17}, R_{29}, R_{38}, R_{41}$		LF411 offset adjust pot (debug only)
R_4, R_6	500 Ω	THD adjust pot or fixed
R_8, R_{10}	50 k Ω	amplitude adjust pot or fixed
R_7, R_9	1 k Ω	VCO base-frequency adjust pot or fixed
R_{49}, R_{51}	2 k Ω	PWM base-frequency adjust pot or fixed

Phase-lock tuning circuit

Component	Value	Notes
R_{50}, R_{53}	1 k Ω	duty-ratio adjust
$R_0, R_2, R_3, R_{24}, R_{32}, R_{33}, R_{34}, R_{35}, R_{42}, R_{43}$	10 k Ω	
$R_5, R_{19}, R_{31}, R_{45}$	1.0 k Ω	
$R_{13}, R_{20}, R_{23}, R_{25}$	5.1 k Ω	
R_{16}, R_{28}	100 k Ω	
R_{18}, R_{30}	180 k Ω	
R_{15}, R_{27}	68 k Ω	
R_{14}, R_{26}	open	
R_{12}, R_{22}	75 k Ω	
R_{11}, R_{21}	4.7 k Ω	
R_{44}	6.8 k Ω	
R_{46}	680 Ω	
R_{47}	330 Ω	
R_{54}	22 Ω	
R_{48}, R_{52}	select	R_t PWM timing
R_{57}	10 Ω	
D_1	40CPQ080	
D_2, D_3, D_4, D_5	1N4148	
D_6	MBR0540T1 40 V Schottky	
Q_1	SUP75N08-10	
Q_2	2N3906	

Bibliography

- [1] S. Upadhye and Y. Atre, "Determination of the design parameters of passive harmonic filters using nonlinear optimization," *1998 IEEE Industrial and Commercial Power Systems Technical Conference*, pp. 155–164, 1998.
- [2] K. Lin, M. Lin, and T. Lin, "An advanced computer code for single-tuned harmonic filter design," *IEEE Transactions on Industry Applications*, vol. 34, pp. 640–648, July/August 1998.
- [3] M. Zhu, D. Perreault, V. Caliskan, T. Neugebauer, S. Guttowski, and J. Kassakian, "Design and evaluation of an active ripple filter with Rogoswki-coil current sensing," *IEEE Power Electronics Specialists Conference*, pp. 874–880, 1999.
- [4] J. Walker, "Design of practical and effective active EMI filters," *Proceedings of Powercon 11*, pp. 1–8 I–3, 1984.
- [5] L. LaWhite and M. Schlecht, "Active filters for 1-MHz power circuits with strict input/output ripple requirements," *IEEE Trans. Pow. Elec.*, vol. PE-2, pp. 282–290, May 1994.
- [6] T. Farkas and M. Schlecht, "Viability of active EMI filters for utility applications," *IEEE Trans. Pow. Elec.*, vol. 9, pp. 328–337, May 1994.
- [7] P. Midya and P. Krein, "Feed-forward active filter for output ripple cancellation," *Int. J. Elec.*, vol. 77, no. 5, pp. 805–818, 1994.
- [8] D. Hamill, "An efficient active ripple filter for output ripple cancellation," *IEEE Trans. Aero. and Electron. Sys.*, vol. 32, pp. 1077–1084, July 1996.
- [9] N. P. *et al*, "Techniques for input ripple current cancellation: Classification and implementation," *IEEE Trans. Pow. Elec.*, vol. 15, pp. 1144–1152, Nov 2000.
- [10] H. Baltes, A. Koll, and D. Lange, "The CMOS MEMS nose — fact or fiction?," *Proceedings of the IEEE International Symposium on Industrial Electronics*, vol. 1, pp. SS152–SS157 vol.1, 1997.
- [11] D. Lange, C. Hagleitner, O. Brand, and H. Baltes, "CMOS resonant beam gas sensing system with on-chip self excitation," *The 14th IEEE International Conference on Micro Electro Mechanical Systems*, pp. 547–552, 2001.
- [12] F. Gardner, *Phaselock Techniques*. New York: John Wiley & Sones, Inc., 1966.
- [13] R. Best, *Phase-Locked Loops*. New York: McGraw-Hill, 3rd ed., 1997.

BIBLIOGRAPHY

- [14] D. Wolaver, *Phase-Locked Loop Circuit Design*. Englewood Cliffs, New Jersey: Prentice Hall, 1991.
- [15] R. Sanneman and J. Rowbotham, "Unlock characteristics of the optimum type II phase-locked loop," *IEEE Trans. Aerospace Navigation Electronics*, vol. ANE-11, pp. 15–24, March 1964.
- [16] S. Feng, W. Sander III, and T. Wilson, "Small-capacitance nondissipative ripple filters for DC supplies," *IEEE Transactions on Magnetics*, vol. MAG-6, pp. 137–142, March 1970.
- [17] Micrometals, "Q curves for iron powder toroids," 1997. Issue G.
- [18] S. Cuk, "A new zero-ripple switching DC-DC converter and integrated magnetics," *IEEE Transactions on Magnetics*, vol. MAG-19, pp. 57–75, March 1983.
- [19] G. Crouse, "Electrical filter," Jan. 27 1928. US patent No. 1920948.
- [20] S. Senini and P. Wolfs, "The coupled inductor filter: Analysis and design for AC systems," *IEEE Trans. Ind. Elec.*, vol. 45, pp. 574–578, August 1998.
- [21] G. Bloom and R. Severns, "The generalized use of integrated magnetics and zero-ripple techniques in switchmode power converters," *IEEE Power Electronics Specialists Conference*, pp. 15–13, 1984.
- [22] D. Hamill and P. T. Krein, "'zero' ripple technique applicable to any dc converter," *IEEE Power Electronics Specialists Conference*, pp. 1165–1171, 1999.
- [23] J. Kolar, H. Sree, N. Mohan, and F. Zach, "Novel aspects of an application of 'zero'-ripple techniques to basic converter topologies," *IEEE Power Electronics Specialists Conference*, pp. 796–803, 1997.
- [24] D. Logue, *Power electronic building block application in optimization, control, and synthesis*. Ph.D. thesis, Dept. of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Power Affiliates Program, 2000.
- [25] J. Rodriguez, 2001. Private discussions relating to doctoral-thesis research in MIT Laboratory for Electromagnetic and Electronic Systems.
- [26] Y. Midorikawa, S. Hayano, , and Y. Saito, "A new inductor having noise-filtering capability," *IEEE Transactions on Magnetics*, vol. 30, pp. 4761–4763, November 1994.
- [27] Y. Midorikawa, S. Hayano, , and Y. Saito, "A multi-resonant type inductor having notch-filtering capability," *IEEE Transactions on Magnetics*, vol. 32, pp. 4998–5000, September 1996.
- [28] R. Kemp, P. Murgatroyd, and N. Walker, "Self-resonance in foil inductors," *Electronics Letters*, vol. 11, pp. 337–338, 24th July 1975.

- [29] R. Reeves, "Inductor-capacitor hybrid," *Proc. IEE*, vol. 122, pp. 1323–1326, November 1975.
- [30] P. Murgatroyd and N. Walker, "Lumped-circuit model for inductor-capacitor hybrid," *Electronics Letters*, vol. 12, pp. 2–3, January 1976.
- [31] R. Reeves, "Choke-capacitor hybrid as a fluorescent-lamp ballast," *Proc. IEE*, vol. 122, pp. 1151–1152, October 1975.
- [32] M. Ehsani, O. Stielau, and J. van Wyk, "Integrated reactive components in power electronic circuits," *IEEE Trans. Pow. Elec.*, vol. 8, pp. 208–215, April 1993.
- [33] S. Tang, S. Hui, and H.-H. Chung, "Coreless planar printed-circuit-board (PCB) transformers — a fundamental concept for signal and energy transfer," *IEEE Transactions on Power Electronics*, vol. 15, pp. 931–941, September 2000.
- [34] S. Hayano, Y. Nakajima, H. Saotome, and Y. Saito, "A new type high-frequency transformer [*sic*]," *IEEE Transactions on Magnetics*, vol. 27, pp. 5205–5207, November 1991.
- [35] R. Hertz and H. Buelteman, "The application of perpendicularly superposed magnetic fields," *AIEE Transactions*, vol. 74, pp. 655–660, November 1955.
- [36] H. McCreary, "The magnetic cross valve," *AIEE Transactions*, vol. 70, no. 2, pp. 1868–1875, 1951.
- [37] F. Beck and J. Kelly, "Magnetization in perpendicularly superposed direct and alternating fields," *Journal of Applied Physics*, vol. 19, pp. 551–562, June 1948.