On-Chip Current Sensing for Monolithic Buck LED Driver

by

Alexander A. Penn

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

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

ARCHIVES

©Linear Technology Corp, MMX. All rights reserved.

Author .....
Department of Electrical Engineering and Computer Science
May 24, 2013

Certified by .....
Michael G. Negrete
Design Engineering Section Leader, Linear Technology
Thesis Supervisor

Certified by .....
David Perreault
Professor of Electrical Engineering
Thesis Supervisor

Accepted by .....
Dennis M. Freeman
Chairman, Masters of Engineering Thesis Committee

On-Chip Current Sensing for Monolithic Buck LED Driver

by

Alexander A. Penn

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2013, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

This thesis explores and analyzes two different options for on-chip sensing and regulating current for a buck LED driver, as compared to a common external solution. With the rise of LED use in common lighting applications, it becomes increasingly necessary to develop solutions that are simple to use and accurate. In this paper, much of the focus is on reducing the external pin and component count while maintaining output current accuracy and power efficiency.

Thesis Supervisor: Michael G. Negrete
Title: Design Engineering Section Leader, Linear Technology

Thesis Supervisor: David Perreault
Title: Professor of Electrical Engineering

Acknowledgments

I would first like to thank John Tilly, Bryan Legates, Linear Technology, and the VI-A program for providing me with this great opportunity to learn more about circuit design and power electronics. Special thanks to Mike Negrete and Professor Perreault for taking time out of their busy schedules to supervise my thesis. Also, thanks to Josh Caldwell, Zeke Randolph, and Jay Lin for being really helpful and patient with explaining circuits and letting me bounce ideas off of you. Thank you to following individuals who went out of their way to make me feel welcome at Linear: Leo Chen, Jonny Paolucci, Jeremy Wong, Cuyler Latorraca, Alice Pitzer, Yin Wu, Karl Peterson, and Manny Sanchez.

Last but not least, I would especially like to thank my family and friends for supporting me up to this point and throughout the M.Eng program. I really could not have done it without you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Thesis Organization

Chapter 1 provides some background and motivation behind this thesis project. Chapter 2 analyzes the buck switching regulator as well as the control schemes and current measurement methods. Chapter 3 describes the proposed methods for solving the current problems. Chapter 4 goes in depth into specific challenges of each method, as well as challenges that are shared by both methods. Some important circuits are also described. Chapter 5 includes results from system level simulations and some analysis. Chapter6 is the final chapter with conclusions drawn and some future work to be explored.

1.2 Background and Motivation

As light emitting diodes (LEDs) become more widespread, they are beginning to replace other types of lighting in popular applications. Three key metrics that LEDs excel in are power efficiency, color accuracy, and intensity. While the power efficiency is a given in most use cases, the color accuracy and intensity are factors of the control method and result from careful current and voltage regulation. Because LEDs are run off direct current, it is possible to control them digitally by modulating on or off the current flowing through, thus retaining the correct color temperature while

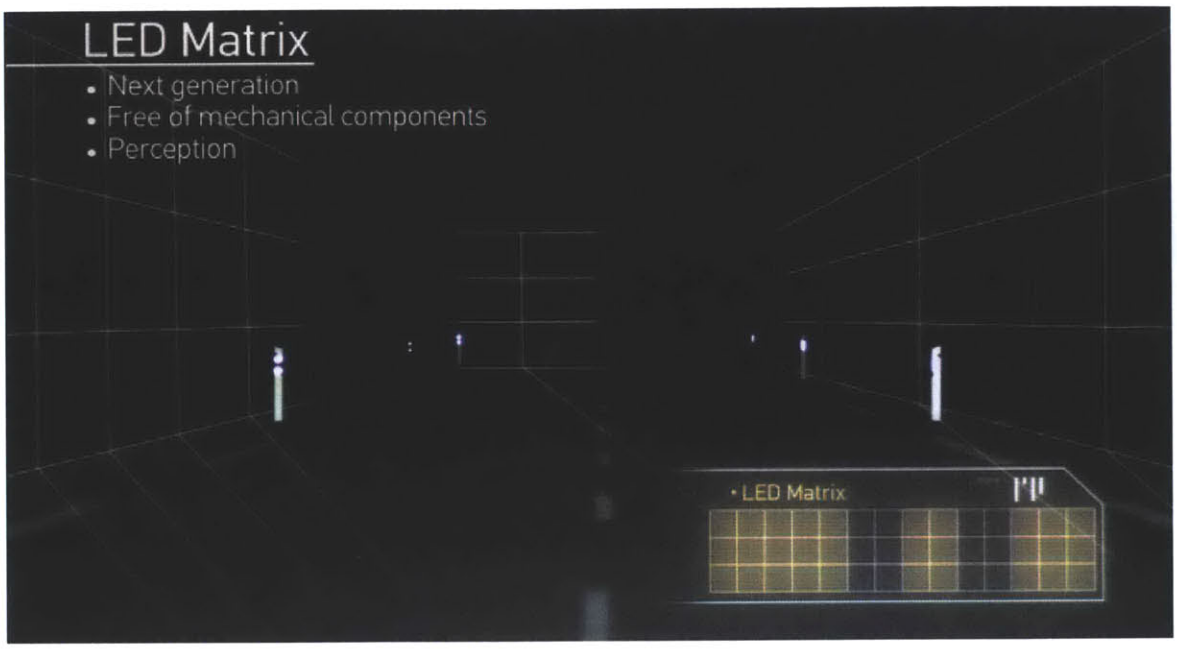

Figure 1-1: Automotive lighting application for LED drivers[1].

changing the brightness.

One area that is beginning to make use of LEDs is car headlamps, like shown in figure 1-1. Traditionally, they have been powered by halogen lamps and more recently xenon lamps on the high end. In addition to the previously stated benefits, one huge advantage of LEDs is the operating color spectrum, which includes colors that more closely resemble daylight. Furthermore, in the past decade, headlamp manufacturers have been attempting to greatly improve visibility through solutions such as the Adaptive Frontlighting System (AFS) , which changes the area of illumination based on factors such as other cars, weather, and speed (Figure 1-2).

Existing methods rely on a mechanical means to direct the light usually by controlling a lens or mirror; using this method, it is generally difficult to generate a precise illumination. While LEDs currently present a significantly more costly alternative, the much improved degree of control in which they can be used to illuminate a scene justifies the switch. With an LED array, headlamp manufacturers can build a much more reliable solution that does away with moving components and offers much more flexibility in the presence of multiple drivers on the road. Another benefit of LEDs that is often overlooked is the added style and customization options that

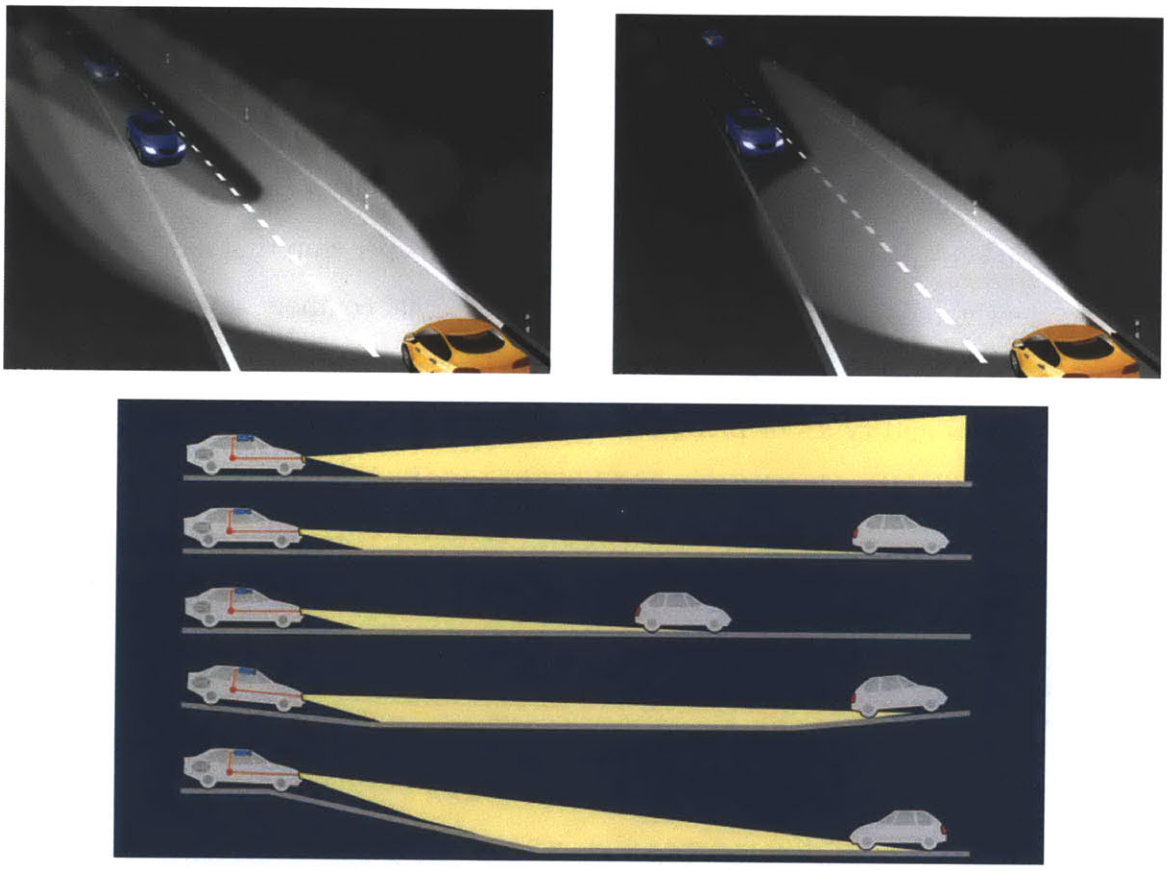

Figure 1-2: Illustration of different Advanced Frontlighting Systems scenarios[1].

result from being more compact and easy to power.

1.3 Major Applications

The main benefit for the proposed chip applies to any application in which a constant current source is needed over which the output voltage varies with time for multiple channels. This includes but is not limited to the automotive lighting application mentioned previously. A constant current source with no output capacitor can ensure safe operation in the event of a change in the output; for instance, if an LED shorts out, there is no stored charge on a capacitor to allow a large current spike to damage the entire LED string. The techniques described here allow for a compact, simple solution that fits the maximum number of output channels per chip that can be operated independently, reducing the number of required external components and simplifying the overall setup. Different channels are required for automotive applications because different colors and types of LEDs are required for different functions, such as turn signaling, day lighting, low-beams, and high-beams (Figure 1-3).

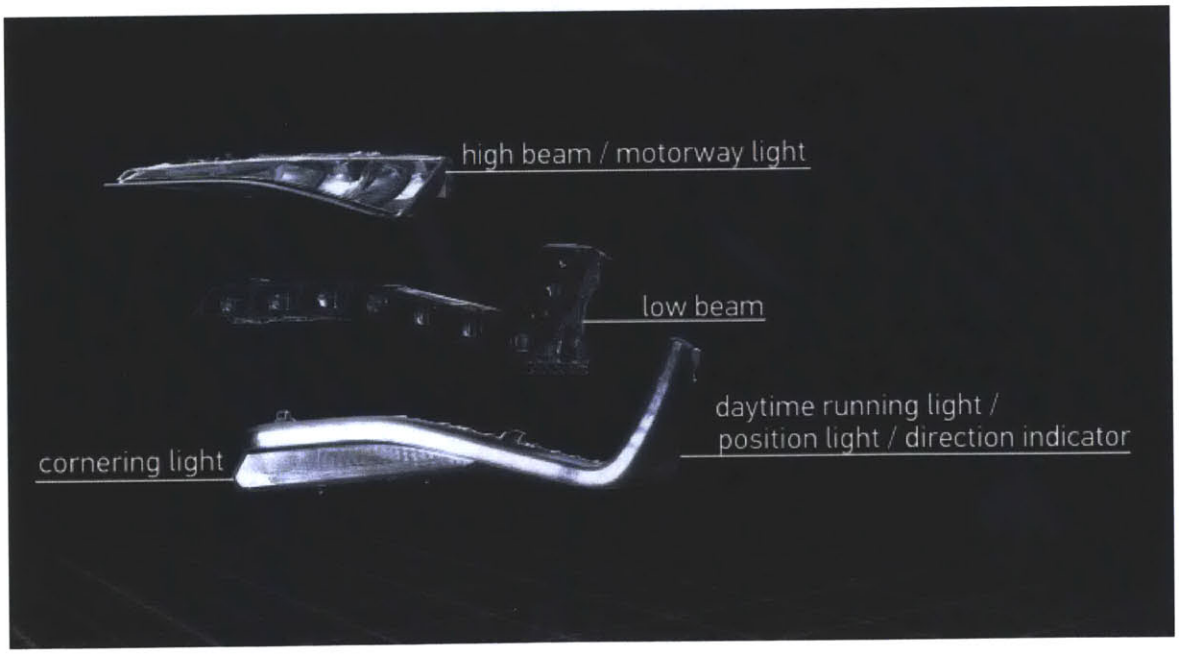

Figure 1-3: Why multiple channels are necessary[1].

While Linear Technology already has multi-channel LED drivers available for pur-

Table 1.1: Proposed table of specifications.

Parameter	LT3595	Proposed Part
Analog Dimming Ratio	10:1	100:1
PWM Dimming Ratio (@200Hz)	5000:1	1000:1
Accuracy	90%	95%
Switching Frequency	2MHz	2MHz
Independent Channels	16	4
Max LED Current/Channel	Up to 50mA	Up to 2A

chase, none of the current products fit the use case exactly. The following image (Figure 1-4) shows a general application of the most similar existing product. While it does away with the sense resistor, its performance does not meet the requirements for aforementioned applications, mainly in the areas of accuracy and current output. Figure 1-5 shows a single channel - note the topology is not the traditional buck.

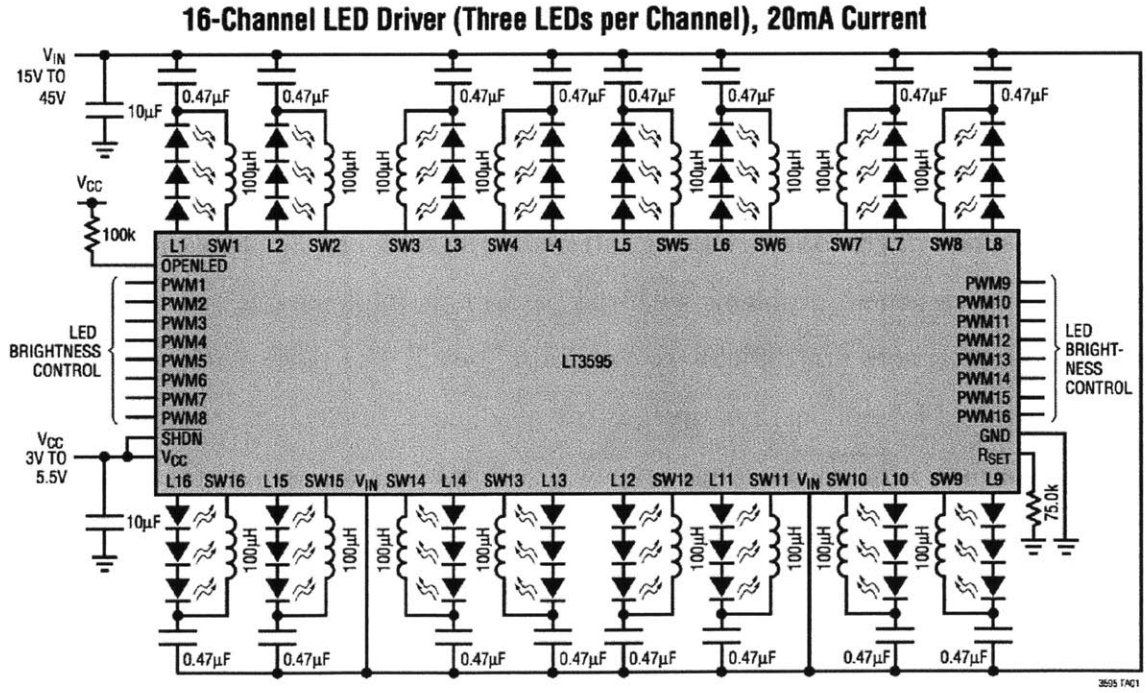

Figure 1-4: Closest existing Linear part, the LT3595[2].

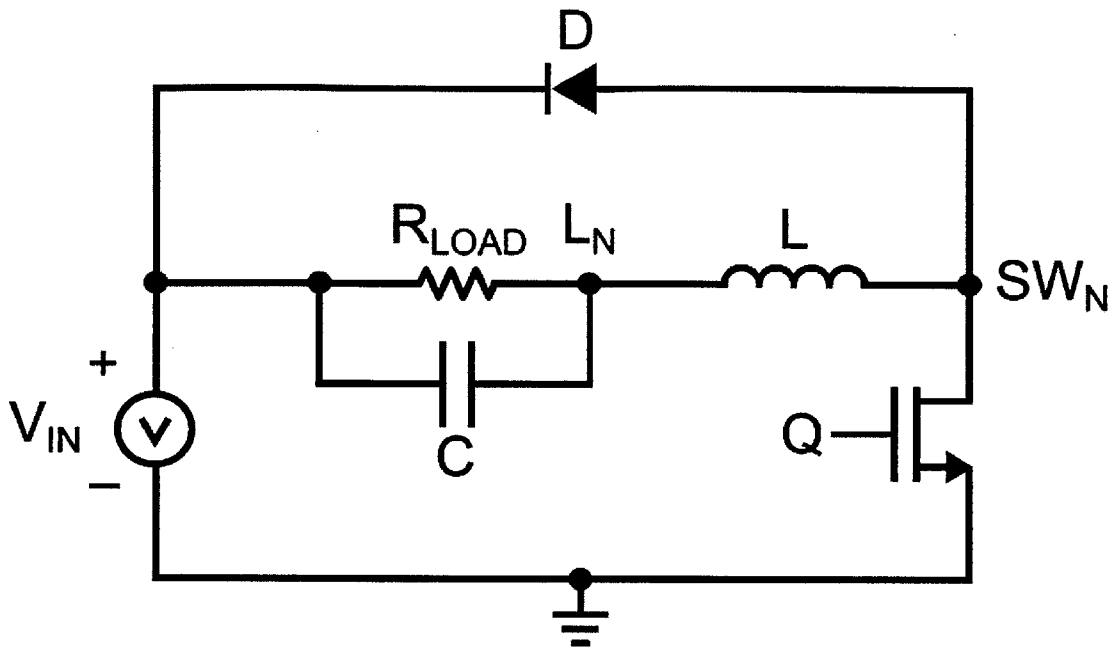

Figure 1-5: Simplified diagram of a single LT3595 channel, which uses the buck-mode topology.

1.4 Proposed Table of Specifications

The idea behind this project is based on Zhen Lis work on a single pin, buck-mode current source. In his thesis, Zhen proposes a method of sensing and regulating current in a buck-mode topology using internal sense resistors [3]. This project aims to address the two main problems with the buck-mode implementation: 1) accurate on-chip sense resistors are difficult to fabricate, and 2) the buck-mode topology is not ideally suited for applications such as car headlamps. Instead, a traditional buck topology is a better fit, which opens up new possibilities as well as different problems in terms of current sensing and regulation.

Chapter 2

System Overview

2.1 Buck Switching Regulator

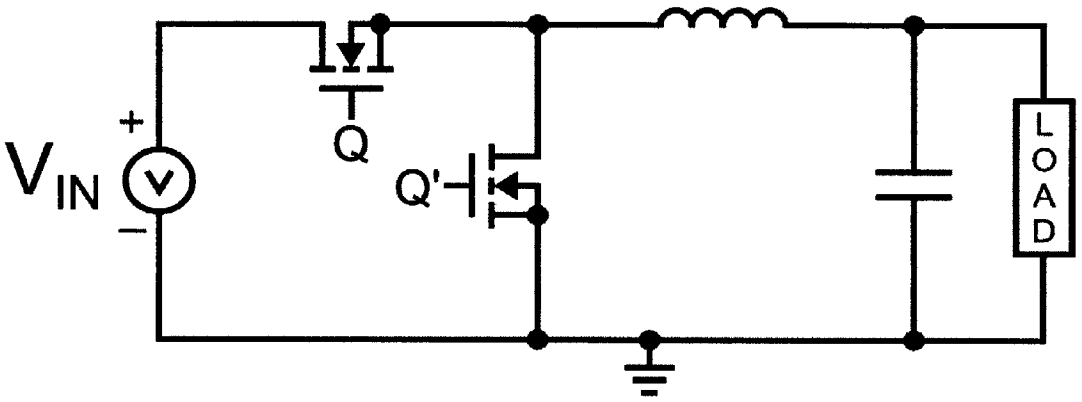

Figure 2-1: Simple buck switching regulator topology.

The topology of a synchronous buck switching converter is shown here in Figure 2-1 [4]. A regulator includes the converter, as well as the control circuitry necessary to generate the duty cycle function, Q. A more in-depth look at the buck regulator system is provided later in the chapter.

The DC voltage source is connected to a pair of transistors that alternate switching on and off to produce a stepped-down voltage at the load. In this simple open-loop case, for a given input voltage and load, the output voltage is a simple function of

the switching duty cycle. The loop can then be closed to regulate to a certain output voltage or current. In the case of the vehicle headlighting application, it is more useful to regulate to a specific output current because the relationship between LED light output and current is much more linear than it is with voltage. Thus, a current source behavior for the buck converter, in which output voltage is not regulated, is desired. This is achieved by setting the feedback so the duty cycle is dependent on the current flowing through the load.

In the next sections, methods of sensing and regulating the current are compared.

2.1.1 Single Output Pin Buck Current Source

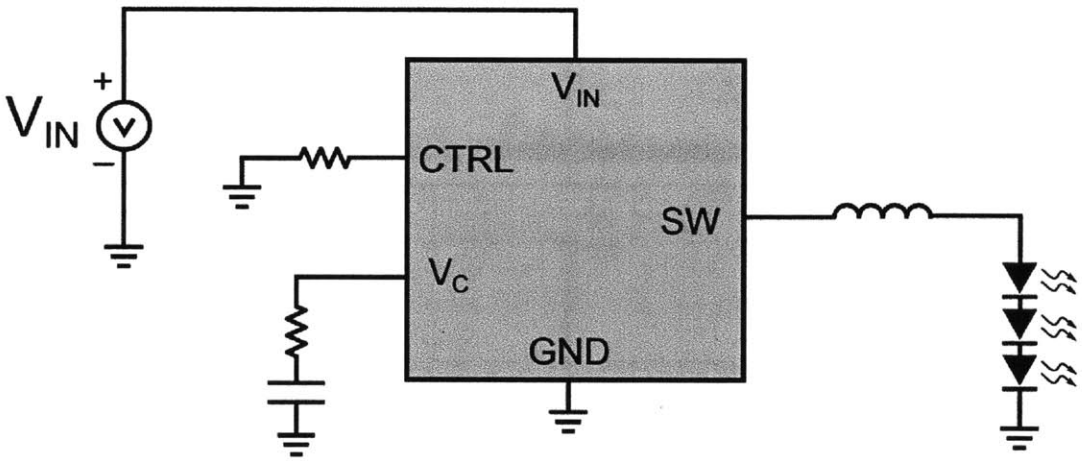

Figure 2-2: Typical application for a single buck LED driver channel.

The results of this work are intended for use in a multi-channel current source chip. Above (Figure 2-2) is a single-channel implementation of such an application. Expanding it to more channels simply involves more switch, control, and compensation pins per channel, with the compensation pins possibly being optional, as discussed later. The single pin aspect refers to the switch pin, which is the only external connection to the load.

2.2 Previous Work

2.2.1 Current Sensing

In this section, a few methods of current sensing that are regularly used are described and illustrated (Figure 2-3) [5, 6, 7, 8].

External Sense Resistor: This method requires the use of an externally located sense resistor, usually located in series with the inductor or load. Through measuring the voltage across the voltage and applying Ohms law, one can determine the current flowing through the resistor. The value of the sense resistor is usually small in order to minimize the power lost in the resistor. The main disadvantage of this method is it requires two extra external pins to measure the voltage across the resistor. Further reducing the appeal of this method is the high relative cost of the actual sense resistor. The more accurate the resistor, the more expensive it is. In cases where accuracy is absolutely essential and there is only one channel, the benefits greatly outweigh the cost. But in applications with multiple channels and a looser requirement on accuracy, the external sense resistor loses appeal.

FET On-Resistance: Expanding on the idea of a sense resistor is the approach of using the actual power switch on-resistance. The difference here is that the voltage is measured across the transistor and using an average resistance value of the transistor, the current flowing through it can be estimated. The goal of this method is to address the two concerns listed previously cost and simplicity. By using the existing power switches, the system becomes simpler and cost is reduced because no additional sense resistors are required. Although this method does address the simplicity and cost tradeoffs, the accuracy is greatly diminished. Unfortunately the transistors are subject to process variations that can lead to a drop in accuracy of at least 20%! While the cost savings on this method are great, the accuracy tradeoff is unacceptable for an application like car head lighting, in which the color accuracy of the headlamp is relatively important.

Inductor Voltage: In the same vein as the using the FET on-resistance, another method takes the approach of measuring the voltage across the inductor. This method

takes advantage of the current-voltage relationship of an inductor. The voltage across an inductor is equal to the inductance times the change in current over time. By integrating the voltage and knowing the inductor value, one can then determine the current flowing through the inductor. The difficult aspect is knowing the inductor value. In some easy cases, the inductor value can be specified by the manufacturer, but that leads to inflexibility in the system. With automotive lighting, much of the inductance is caused by wiring that runs from the electronics out to the actual components. In such a case, the inductance is not so easy to take into account. The alternative is to use a self-calibrating circuit that senses the inductance. Further complicating the system is the fact that cables run to and from the load, so the voltage across the inductor is not so simple to measure either.

Sense-FET: This last method is relatively new but is gaining popularity. It makes use of a transistor, now referred to as sense-FET, sized much smaller than the power switch at a ratio of approximately 1000:1 or greater. By forcing the same drain to source voltage as the power switch across the sense-FET, the current flowing through it is thus proportionally smaller by the same ratio. The main requirement of this method is that the power switch must be located on-chip to ensure good matching with the sense-FET. Even so, the matching between two transistors cannot be guaranteed above 95%. Further work, such as trimming or averaging, must be done to increase accuracy.

2.3 Control Scheme

There are two main methods of control for a buck switching regulator: Current-mode and Voltage-mode. Each has its advantages and disadvantages, which are described below.

2.3.1 Voltage-mode Control

In voltage-mode control, the average current information is compared to a reference voltage through an error amplifier. The output of that error amplifier is then com-

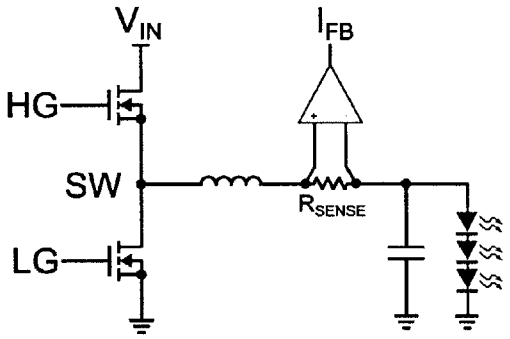

(a) External sense resistor

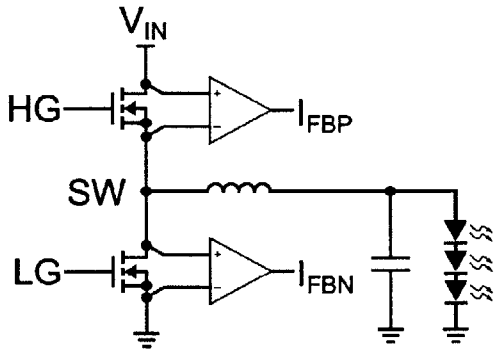

(b) FET on-resistance

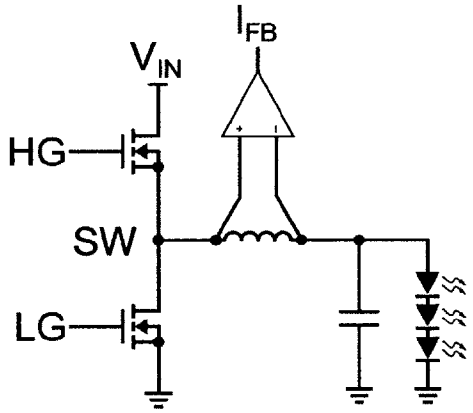

(c) Inductor voltage integration

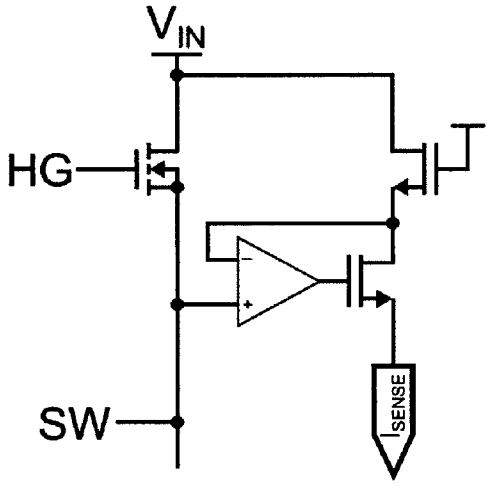

(d) Sense-FET

Figure 2-3: Different current sensing methods.

pared with a voltage ramp signal, which directly controls the duty cycle. Overall, voltage-mode control is simpler to implement.

2.3.2 Current-mode Control

In current-mode control, the average current information is compared to a reference voltage through an error amplifier, just like the voltage-mode control. After that, the output of the error amplifier is compared with the peak current information summed with a compensating ramp. Typical system representations are shown in Figures 2-4 and 2-5.

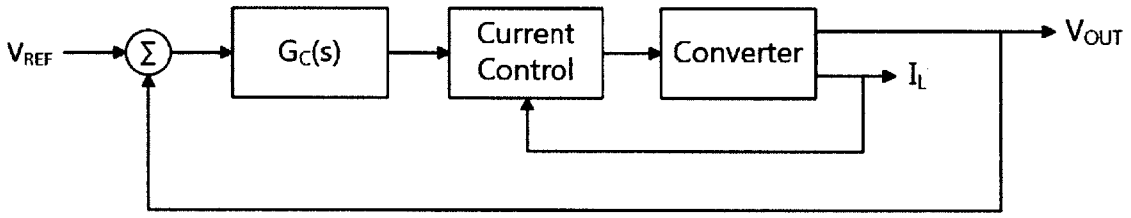

Figure 2-4: Traditional current-mode control block diagram.

The type of control system used (Figure 2-6) is a slightly modified version of the traditional peak current-mode control implementation. In the traditional system, since the loop is regulating both voltage and current, it is usually necessary to obtain a sensed current value. However, for a current source, the system is somewhat simpler in that it just needs to regulate to the set current value.

2.3.3 Loop Analysis

Buck Converter Transfer Function

In this analysis, the output load is assumed to be a string of LEDs, which is approximated as a voltage source for simplicity (Figure 2-7). Compared to the traditional buck converter transfer function analysis[10], this modified configuration ends up simpler because the voltage source replaces a resistor-capacitor network. Instead of two

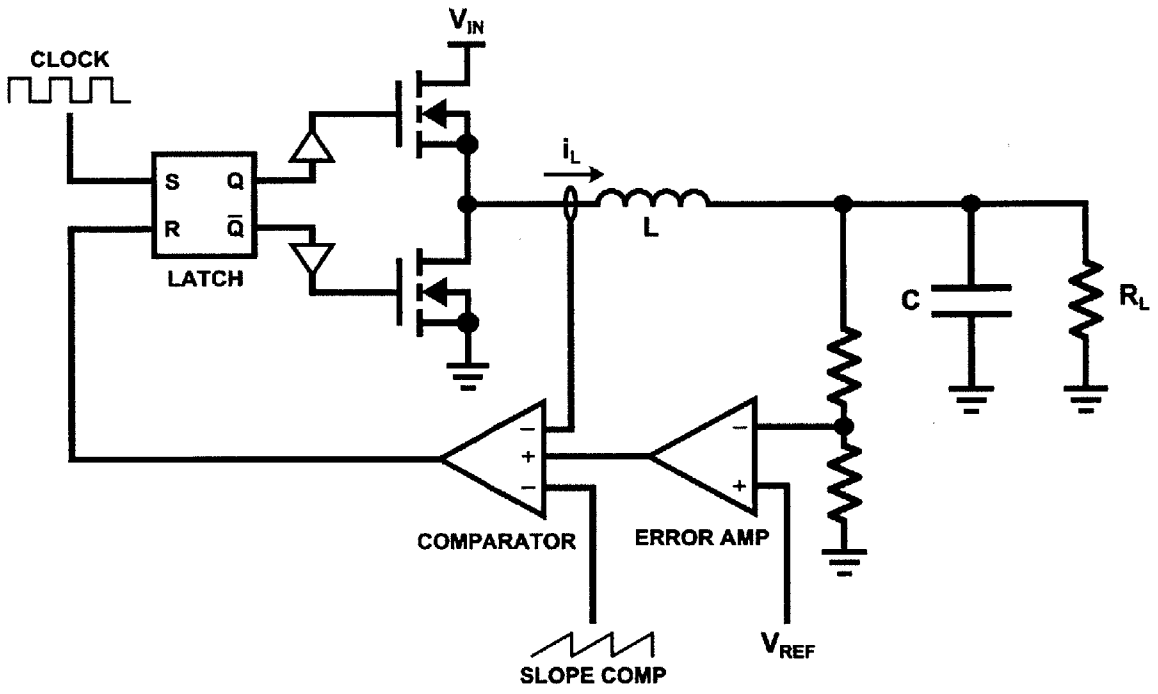

Figure 2-5: Typical implementation of current-mode control. Redrawn from Unitrode[9].

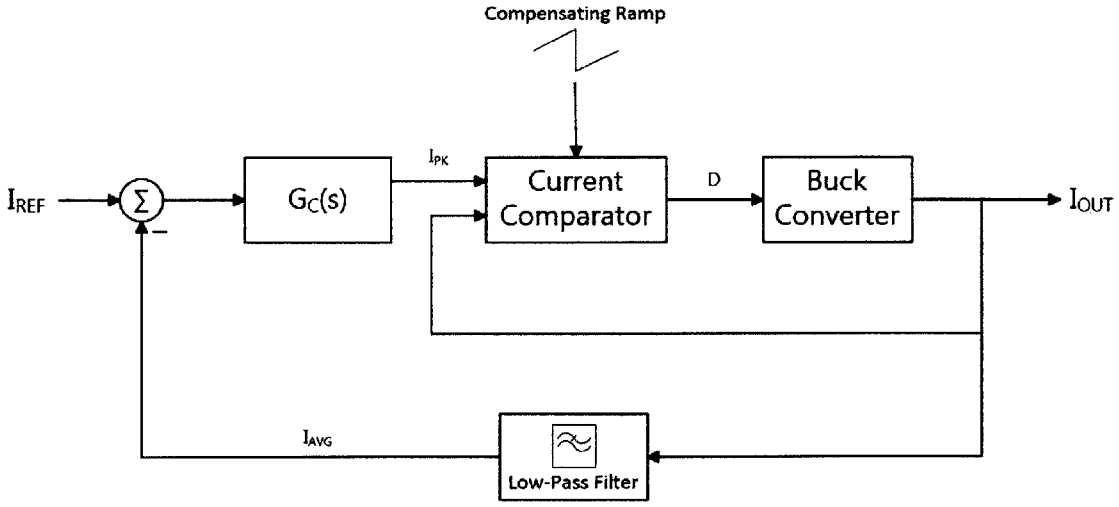

Figure 2-6: Modified peak current-mode control for regulating only current.

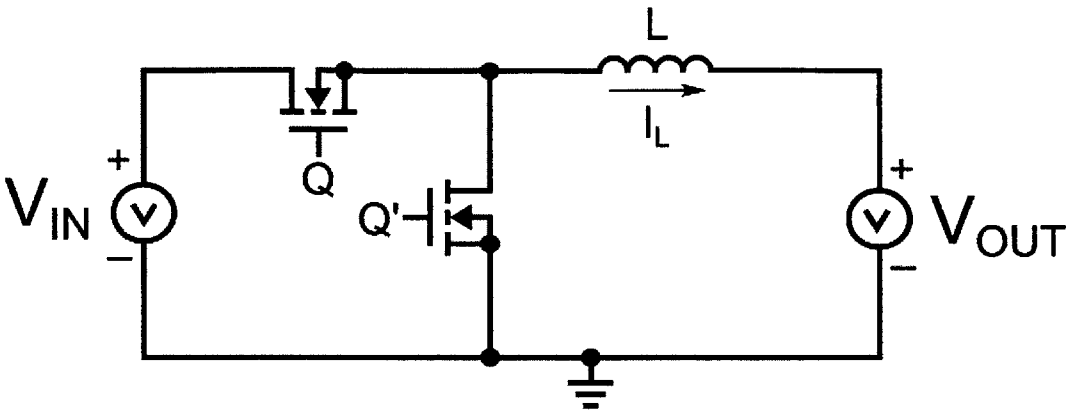

Figure 2-7: Simplification of load for analysis purposes.

switched equations of state, since there is no output capacitor, only the equation for the inductor voltage remains.

$$V_L = L \cdot \frac{dI_L}{dt}$$

$$\frac{dI_L}{dt} = \frac{V_{IN} - V_{OUT}}{L} \cdot Q(t) + \frac{-V_{OUT}}{L} \cdot (1 - Q(t))$$

$$\frac{dI_L}{dt} = \frac{V_{IN}}{L} \cdot Q(t) - \frac{V_{OUT}}{L}$$

$$\frac{d\bar{I_L}}{dt} = \frac{V_{IN}}{L} \cdot \bar{d} - \frac{V_{OUT}}{L}$$

Laplace Transform:

*V_{OUT} is constant

$$s \cdot I_L(s) = \frac{V_{IN}}{L} \cdot d(s)$$

$$H(s) = \frac{I_L(s)}{d(s)} = \frac{V_{IN}}{L} \cdot \frac{1}{s}$$

The resulting transfer function of the buck converter without output capacitor as a function of the duty cycle and using a voltage source at the output is a single pole located at zero. The transfer function of the closed inner loop is shown below:

$$G(s) = \frac{I_L(s)}{I_p(s)} = \frac{k \cdot H(s)}{1 + k \cdot H(s)}$$

But first the behavior of the current comparator, represented as the gain variable k, must be derived. Finding the slopes of the inductor current (Figure 2-8) is useful for deriving k.

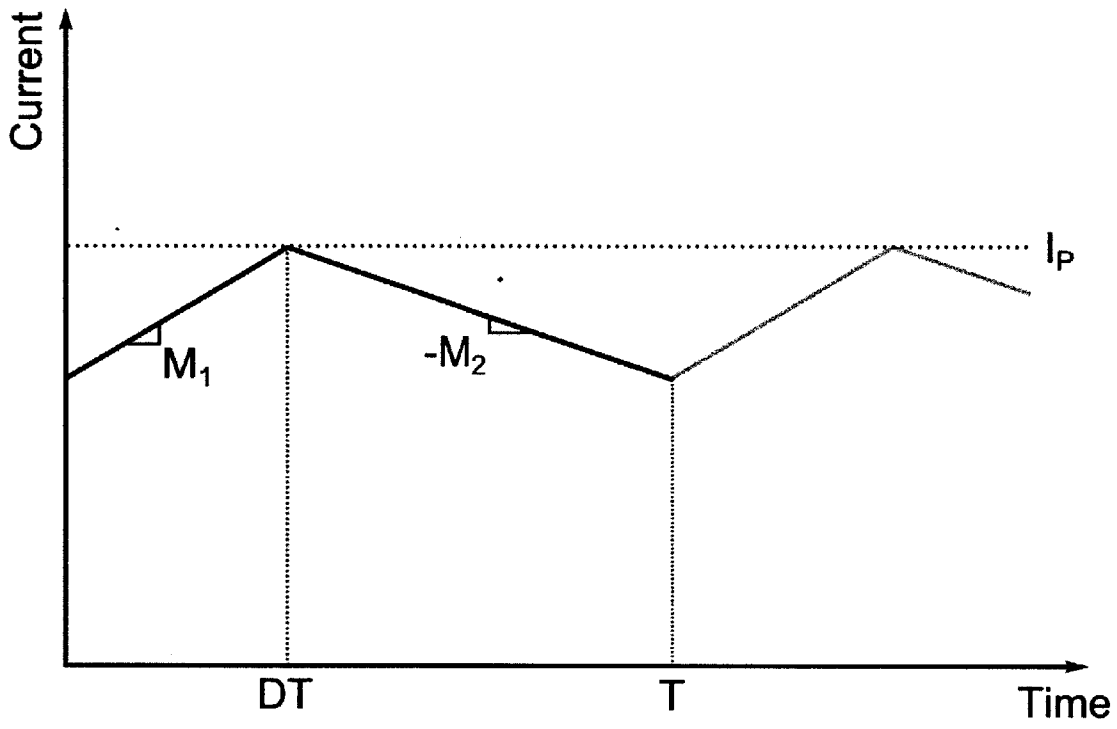

Figure 2-8: Illustration of current during one period along with the corresponding slopes.

$$M_1 = \frac{V_{IN} - V_{OUT}}{L}$$

$$M_2 = -\frac{V_{OUT}}{L}$$

$$i_L = i_p - \frac{1}{T}[\frac{1}{2}M_1 d^2 T^2 + \frac{1}{2}M_2 d'^2 T^2]$$

$$i_L = i_p - \frac{V_{IN}}{2L}d^2 T + \frac{V_{OUT}}{2L}T[d^2 + d'^2]$$

$$\tilde{i_L} = \tilde{i_p} - \frac{V_{IN}T}{L}\tilde{d} - \frac{D^2T}{2L}\widetilde{v_{IN}} + \frac{V_{OUT}T}{L}\tilde{d} + \frac{D^2T}{2L}\widetilde{v_{OUT}} - \frac{V_{OUT}T}{L}\tilde{d} + \frac{D'^2T}{2L}\widetilde{v_{OUT}}$$

Assuming $v_{IN} = V_{IN}$, $v_{OUT} = V_{OUT}$ (i.e. no perturbation in V_{IN} or V_{OUT}):

$$\tilde{i_L} = \tilde{i_p} - \frac{V_{IN}T}{L}\tilde{d}$$

$$\tilde{d} = \frac{L}{V_{IN}T}(\tilde{i_p} - \tilde{i_L})$$

Thus: $k = \frac{L \cdot f_{SW}}{V_{IN}}$

Closing that inner loop to get a transfer function as a function of the peak current results in a system with a single pole located at a frequency related to the switching frequency and some scaling factor.

$$G(s) = \frac{I_L(s)}{I_p(s)} = \frac{f_{SW}}{f_{SW} + s} = \frac{1}{\frac{s}{f_{SW}} + 1}$$

The inductor value actually cancels out from the pole location, meaning its sizing only affects the ripple. The same goes for the input voltage. The only remaining variable that determines this pole location is the switching frequency. Note that the pole is not located at the switching frequency; due to the conversion from radians, the pole location actually happens to be $2 * \pi$ times smaller than the switching frequency. This result is important to take into account when determining the compensation capacitor.

Compensation Function

The compensation will likely only require a capacitor as a single pole for stability purposes instead of a resistor-capacitor network as a pole-zero. The transfer function for this system (Figure 2-9) can be derived as follows:

$$V_{OUT} = \frac{i_{OUT}}{sC} = \frac{G_M V_{IN}}{sC}$$

$$\frac{V_{OUT}}{V_{IN}} = \frac{G_M}{sC}$$

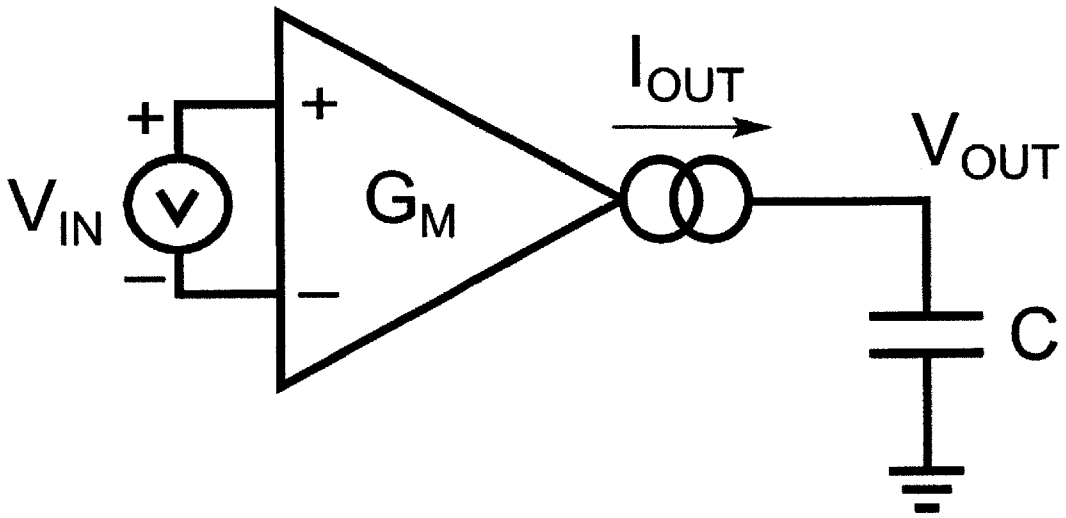

Figure 2-9: Operational transconductance amplifier feeding a capacitor.

The pole is then located at zero, with the crossover frequency being at:

$$f_c = \frac{G_M}{C}$$

Given that pole location, the transconductance and capacitor size must be optimized to provide the minimum power consumption and die area, maximum noise immunity.

Overall System

The resulting system has a maximum of three important poles: one each from the compensation, buck converter, and low-pass filter. The low-pass filter has minimal influence due to its high frequency pole location. In fact, it may even be possible to completely remove the low-pass filter and use the compensating pole to low-pass filter the error signal and end up with the same peak current value. In doing so, the system is simplified into a two-pole system, in which the compensation value as well as the gain can be easily adjusted for stability. Because the other pole is located at $f_{sw}/2\pi$,

the compensation must act as a dominant pole and ensure proper phase margin of around 60 degrees. At first order, that means the pole must be placed approximately 5 times smaller than $f_{sw}/2\pi$ for adequate phase margin. An optimized dominant pole location is computed here in Figure 2-10 to achieve the quickest loop response.

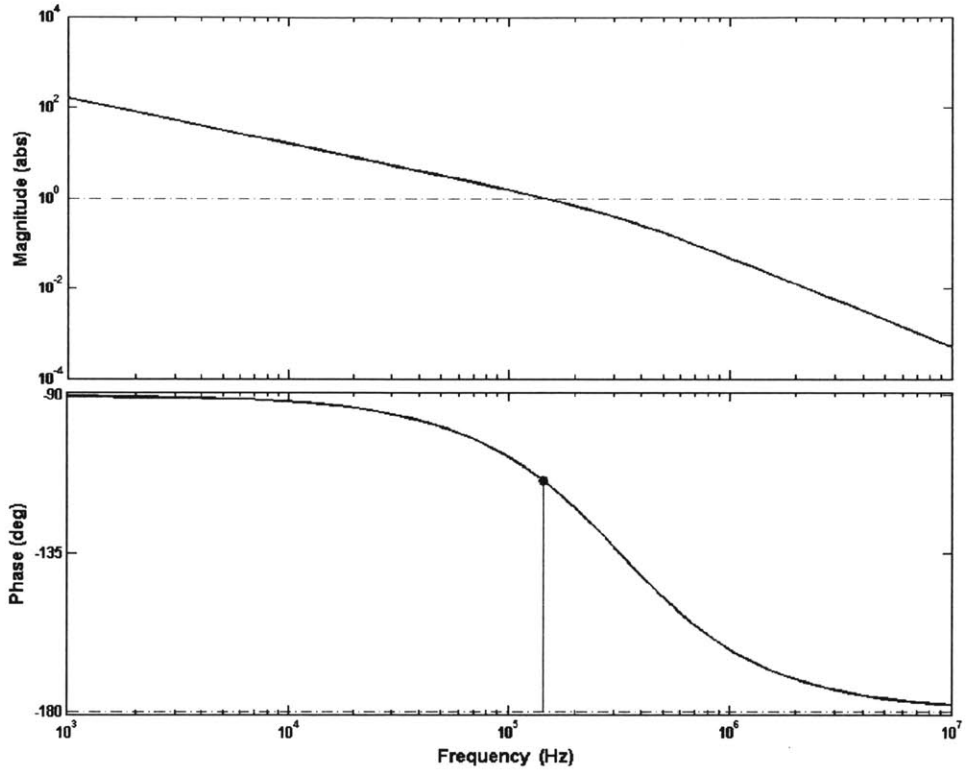

Figure 2-10: Bode plot of optimally compensated control loop. Phase margin is 65 degrees with the crossover frequency occurring at 145 kHz.

Chapter 3

Proposed Solutions

3.1 Current Sensing

With the benefits and drawbacks in mind, the current sensing method most suited for the intended application is the sense-FET. Through the use of averaging or trimming, the inaccuracies presented from poor matching can be overcome to provide an acceptable level of accuracy without increasing the external component complexity. While using the FET on-resistance also allows for a single-output pin solution, it is much more difficult to trim the on-resistance to a specific value than it is to trim two FETs to match. Out of the listed methods for sensing currents, the senseFET provides the best balance between external circuit simplicity and output accuracy.

3.1.1 SenseFET Accuracy

The main issue with using a Sense-FET is the low accuracy resulting from poor matching. If the sense-FET is accurate only to 10 or 20 percent, the resulting system can only then be accurate at a maximum of the same amount. The poor matching between the sense-FET and power switch is due to the ratio of sizes being so large. That means no matter where the sense-FET is placed, a large variation in matching within the large area taken up by the switch is unavoidable.

3.1.2 Averaging

One way to mitigate the effect is to place sense-FETs at different points within the area of the switch and then use a different FET for each switching cycle[11]. The average measurement ends up being very close to the desired ratio between the switch and sense-FETs. The advantage of this method is no additional testing time is required to guarantee some base level of accuracy, since the average of the sense-FETs is designed to be already close to the actual value. They are more likely to be better matched over temperature and load due to their locations being spread throughout the switch device area.

3.1.3 Trimming

The other way to increase accuracy is to use a traditional trimming scheme to create a parallel combination of FETs that are at the exact sensing ratio desired. This entails the use of binary sized FETs and a trimming circuit to be operated at testing. The advantage of this method is an absolute accuracy value can be achieved, with no additional circuitry needed during operation. Of course a trimming circuitry is required during testing to burn in the trim bits.

3.1.4 SenseFET Ratio

Ideally, the senseFET is set to mirror the actual current at the exact ratio set by the size, which typically runs in the range of 1000:1 to 50,000:1 for this specific process[12]. However, in reality, the drain-source voltage at small currents can become so small that accuracy becomes an issue. If the buck converter is meant to be run at a wide output current range, this problem must be addressed. One remedy is to change the ratio between the senseFET and power FET. By decreasing the effective size of the power device, the ratio decreases and a larger signal is generated. Although the tradeoff for this is additional circuit complexity, the accuracy and possibly power efficiency gained is well worth it. Furthermore, an optimal balance between resolution (steps) and overall complexity must be determined. It may also be necessary to

implement some sort of hysteresis to prevent the ratio from switching back and forth if the control voltage is noisy. To obtain a balance between accuracy and efficiency, ratios of 10,000:1 and 1,000:1 are tested.

3.2 Regulation Schemes

Taking all of the previous work into account, three solutions for current sensing and regulation in the previously described system are proposed in this paper.

The most similar solution to the traditional current-mode control scheme is to use a sense-FET to generate a sensed current for the top gate and bottom gate, which are then summed to provide the total current information (Figure 3-1). By sending the current through a resistor, the summed sense currents can be converted into a sense voltage that is proportional to the actual current. This voltage is then low-pass filtered to obtain an average current value, which is compared to a control voltage and used to regulate the system. A current-summing method such as this is the subject of many previous works, but its use in a purely current-mode control system is less well known. Most of the previous work focuses on special techniques, such as averaging multiple sense-FETs, to increase the accuracy given poorly matched sense-FETs[11]. However, with the proper trimming circuits, the matching can be trimmed to the desired accuracy. This method from this point on is referred to as the current summing method.

Instead of regulating using a reference control voltage, one way of simplifying the loop is to regulate directly using a reference current. To do this requires a careful feedback loop that first sets a control current through a sense-FET connected to the switch pin. A transconductance amplifier connected to the other side of the sense-FET measures the drain-source voltage as referenced to ground and outputs to the compensation node. The system in Figure 3-2 regulates until the drain-source voltage across the bottom switch is on average equal to the voltage across the sense-FET.

A slight modification of this approach is to use the compensation capacitor to low-pass filter the signal instead of a separate sampling capacitor. Instead of sampling on

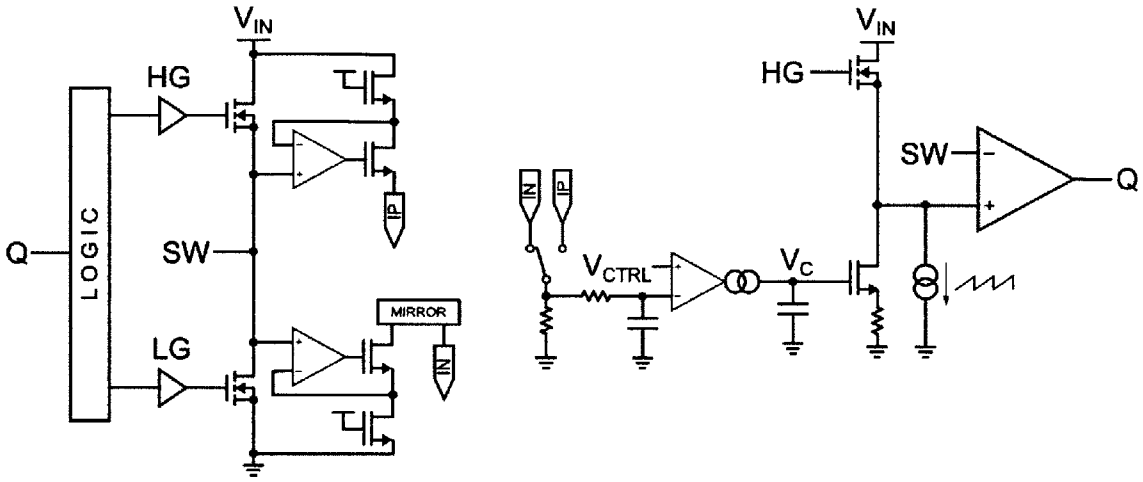

Figure 3-1: Current summing method.

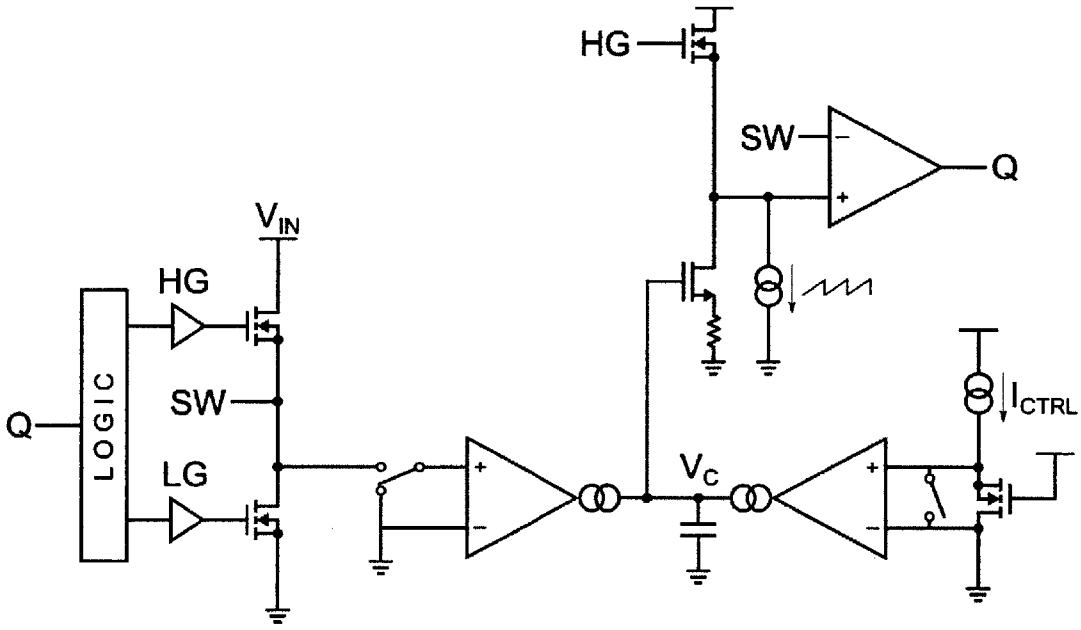

Figure 3-2: Opposing transconductance amplifier method.

the inputs, the amplifiers are fed a signal that alternates between the actual switch voltage and ground. When the switch pin is connected, the amplifiers will be sourcing or sinking the appropriate currents. When the inputs are grounded, the amplifiers shall not output any current and the compensation node shall stay at a constant voltage. The main advantage of this approach is it allows for a much larger capacitor value since the compensation capacitor is usually located externally. As a side benefit, it also reduces the number of circuit components used and somewhat simplifies the circuit. The caveat to this approach is since the input to the transconductance amplifiers is changing every clock cycle, the amplifiers have to respond more quickly to keep up. A higher bandwidth requirement is then placed on these amplifiers.

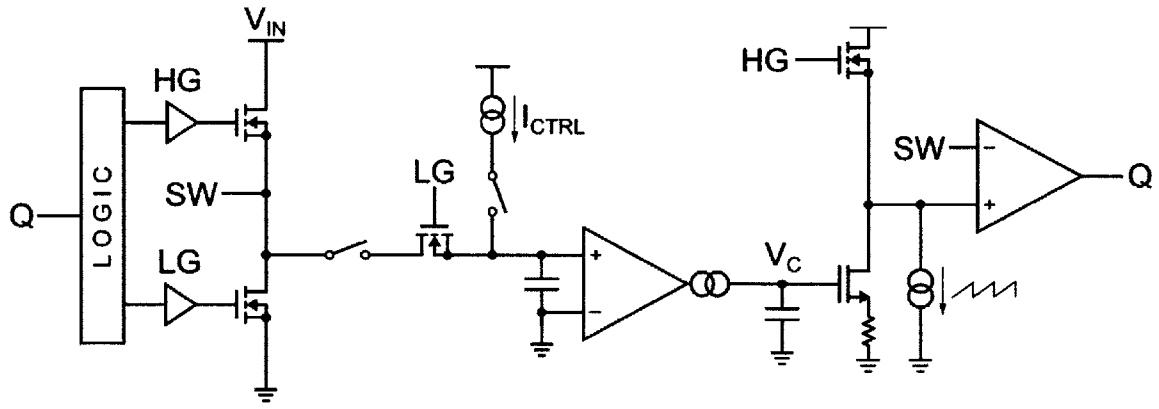

Figure 3-3: (Single) Transconductance amplifier method.

The next approach (Figure 3-3) is similar to the previous approach but attempts to remedy the bandwidth problems. The idea is to use a sample and hold scheme on the bottom power FET to sense the average drain-source voltage. This voltage is then fed to a transconductance amplifier, which pulls current due to the negative voltage of the switch node. A matched transconductance amplifier also feeds into the output of the first, which means at equilibrium, the two amplifiers shall provide the same amount of current. Since they are matched, the voltage at the inputs shall be exactly opposite. A sense-FET placed across the inputs of the second amplifier is fed a control current. The system then regulates to a point in which the voltages across both amplifier inputs are exactly opposite. This method is now referred to as the transconductance amplifier method. Ultimately, the latter two methods are similar

enough that for comparison purposes, only the first and third systems are actually evaluated.

Chapter 4

Analysis of Systems

Since the main focus of this work is to compare the different regulation schemes, general circuits like oscillators and bandgap references are not discussed.

4.1 Current Summing

4.1.1 Operational Amplifiers

An operational amplifier is necessary to maintain the voltage between the source and drain of the sense-FET to match the source-drain voltage of the power FET. The simplest implementation is to use a source-follower instead of a full-on differential amplifier. One caveat is the drain-source voltage must be somewhat large or else the amplifier saturates and generates an incorrect sensed current. A common gate differential pair is used to generate the sensed current. A common gate topology is implemented over a common source topology for its speed; because the sensed current is sensitive to the bias currents at low loads, a simple level shift is used to isolate the signal path.

Top Gate Sense

The top gate sense circuit (Figure 4-1) is based on a common gate design, which allows it to be fast. To prevent the circuit from slewing, a sample and hold capacitor

is used at the switch node input to keep the generated current at approximately the correct value while the top gate is off (Figure 4-2). If this is not implemented, the generated current must slew from no current all the way to the correct value and then back down to nothing.

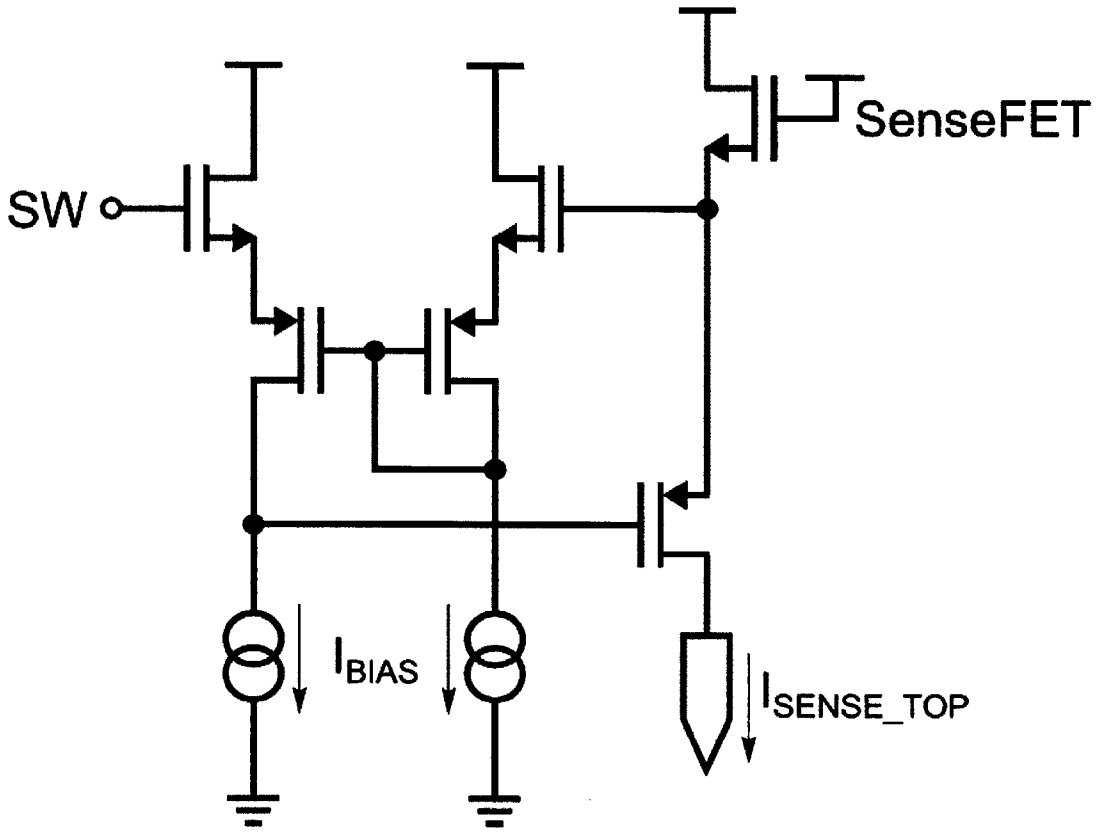

Figure 4-1: Top gate sense amplifier.

One area of caution is in the timing of the sample and hold scheme. As V_{IN} increases, the time it takes for the switch node to pull up to the rail as the top switch turns on increases as well. Thus, if the sampling signal is set to the same as the top gate signal, the error due to the top gate sensing circuit will proportionally increase with VIN. The sampling signal must be delayed a certain amount based on V_{IN}.

When using the current summing approach, the idea is to add the sensed currents through a resistor to obtain a voltage. However, when switching between the sensed top and bottom currents, there are some inaccuracies that follow through.

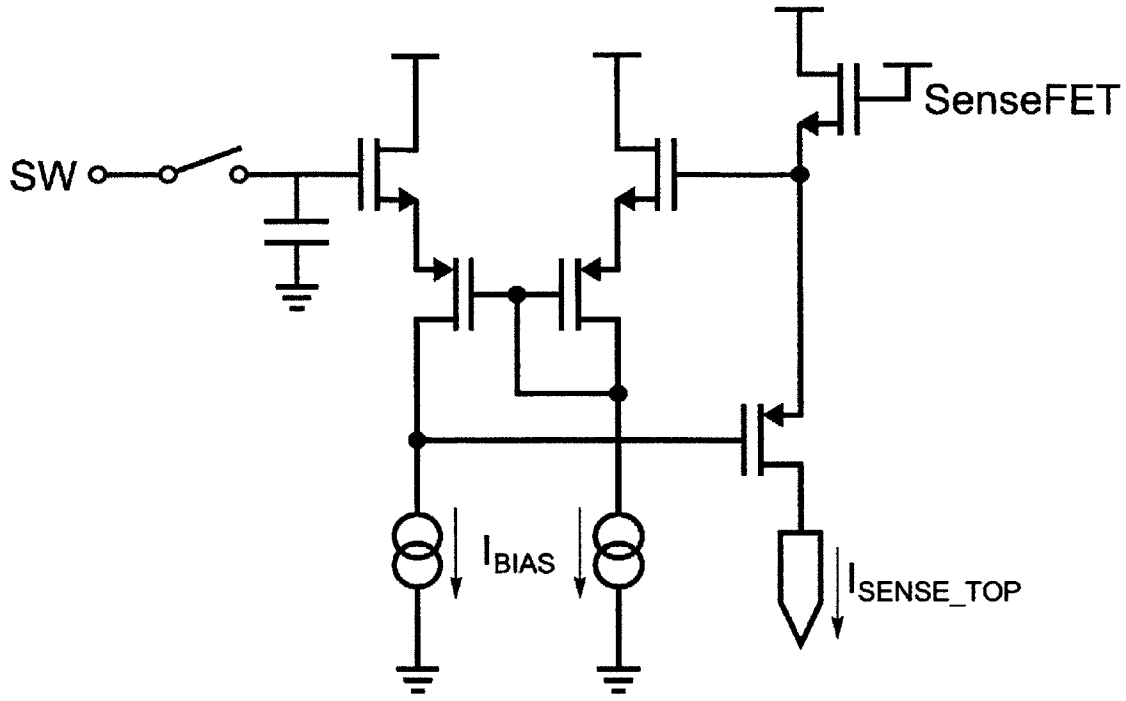

Figure 4-2: Modified top gate sense amplifier.

The op-amps and current mirrors generating those sensed currents do not have infinite bandwidth to keep up with the switching, so there is some time constant for the sensed current to reach the appropriate level. For the bottom gate, because there is zero voltage switching, the error at the switch node is only a diode drop. However, for the top gate, the switch voltage must swing from close to ground all the way to near the rail. This leaves the potential for an enormous amount of error from the top gate current sense. One solution to this is to use a blanking signal at the turn-on transient to block out the error term. Even with that in place, there is some error term for the value that is held during the blanking.

Bottom Gate Sense

The bottom gate sense (Figure 4-3) is a little more complex in implementation than the top gate sense because of the negative switch voltage. That requires both the sense and power devices to be located before the input. Since the current is also flowing in the opposite direction, a current mirror must be generated for a current-to-voltage

conversion. Unfortunately, as the op-amp is switching from its regulation point and the rail, the current mirror is generally in slow moving from triode to saturation, which limits the bandwidth of the loop. However, because the output, which is low impedance, is fed back into the input, a sample and hold cannot be implemented at the input.

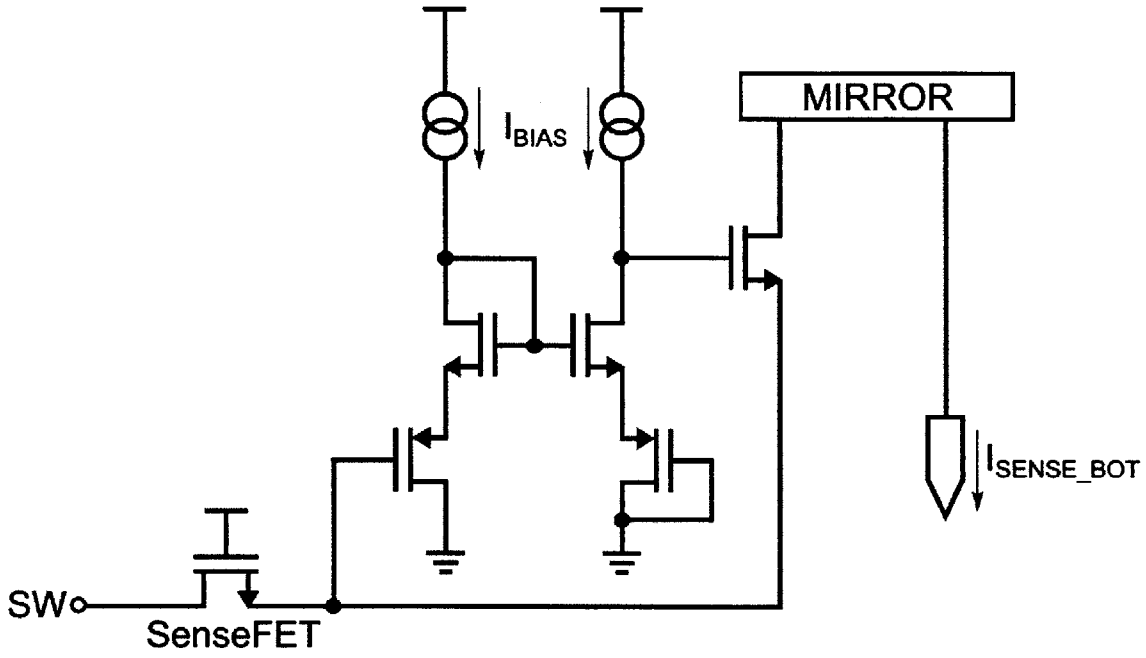

Figure 4-3: Bottom gate sense amplifier.

The approach to solving this problem is to hold the output state by breaking the loop, effectively low-pass filtering the input. Because this bottom-gate sense is mainly used for the averaging scheme - as opposed to the top-gate which is also used for the peak current detection, having an accurate instantaneous current readout for the bottom gate is not vital to the overall system accuracy. The caveat with this if there is any non-zero voltage on the input, the amplifier is railed to one side or the other. When the sampling capacitor is switched on, the railed voltage dominates and the system slowly recovers based on the slew rate of the amplifier. Two methods are explored to alleviate this problem. The first idea is to implement a buffer that is switched in during the off-time of the power device. This helps to maintain the capacitor voltage at the output of the op-amp so it is not completely railed. There

is a transient effect that causes some systematic offset. The other method is to add a current buffer stage after the op-amp output to provide the current necessary to charge the holding capacitor quickly, as shown in Figure 4-4. In doing so, a second pole is created: one from the high gain node p_1, and the other from the holding capacitor at node p_2. Although it settles on the correct value, the resulting loop does not have adequate phase margin and oscillates to the final value. To fix that, a zero must be inserted by placing a resistor (already shown) in series with the holding capacitor at node p_2.

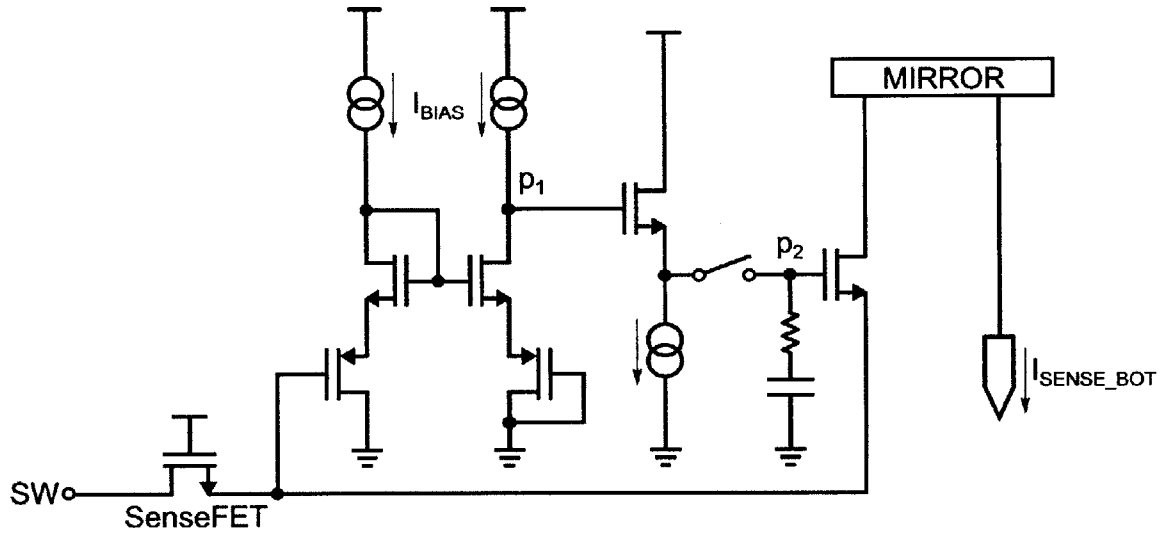

Figure 4-4: Modified bottom gate sense amplifier with current buffer stage.

4.1.2 Charge Injection

Using a sample and hold scheme can lead to larger error due to charge injection[13]. If the sampling switches are only switched on for short amounts of time, the gate capacitance can provide a direct path for charge to flow, skewing the held value. However, if the system uses an averaging scheme, any extra charge becomes averaged out over a relatively long period of time, thus reducing the error observed.

$$q_{ch} = -WLC_{OX}(V_H - V_{IN} - V_T)$$

$$\Delta V = -\frac{(V_H - V_L)C_{OL}}{C_L} - \frac{WLC_{OX}(V_H - V_{IN} - V_T)}{2C_L}$$

$$V_{OS} = -(V_H - V_L)\frac{C_{OL}}{C_L} - (V_H - V_T)\frac{WLC_{OX}}{2C_L}$$

For the current summing topology, the switches must be carefully sized so that settling time due to charge injection is as short as possible. A simple test bench (Figure 4-5) is set up to determine the optimal gate sizing.

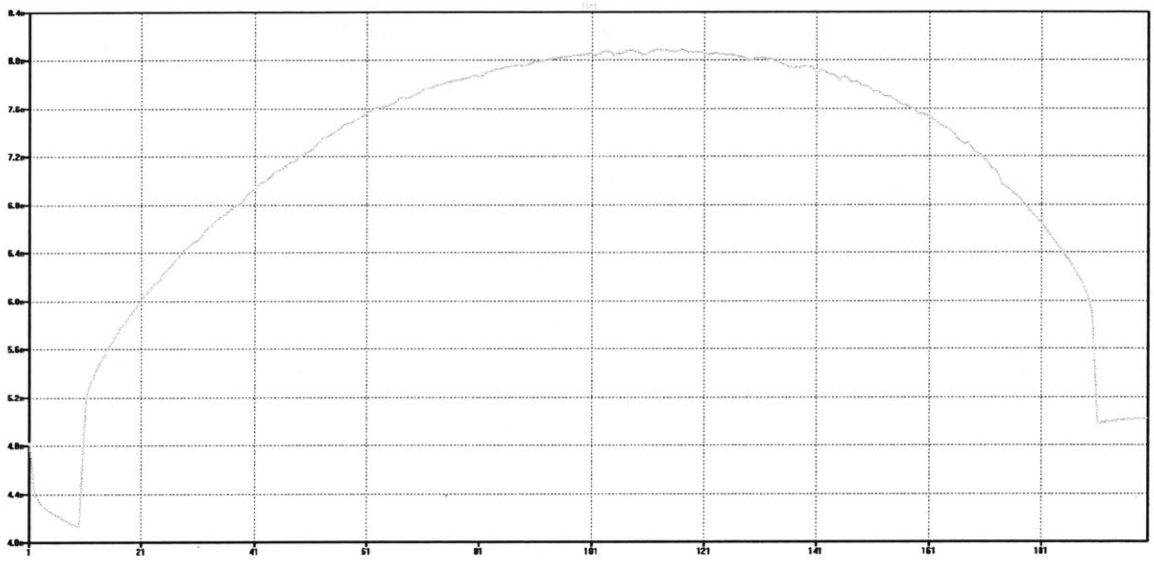

Figure 4-5: Charge injection: switch size vs. settling time to 90% of final value.

4.2 Transconductance Method

4.2.1 Gate to Source Voltage

INSERT diagram and point out relevant transistors

The difficulty in running a bottom-side current sense is that the switch node is a negative voltage when the bottom switch is on. Because there is no negative rail, some workarounds must be used to ensure operation. One method involves applying a proportional current to a ground-referenced senseFET. The drain to source voltage across that device shall be the same in magnitude as the voltage across the bottom switch. However, since the senseFET source is at ground, the drain is positive instead

of negative like the power device. Thus, the on-resistances of the two devices are slightly disproportionate. As the devices are operated in the linear region, they are governed by the equation:

$$R_{ON} = \frac{1}{\mu_n C_{ox} \frac{W}{L}(V_{GS} - V_{TH})} + R_{const}$$

The threshold voltages shall be the same around 1.5V, but the gate to source voltages are slightly different. Because the current is flowing in the opposite direction of the normal device polarity for both devices, the drains act like sources and vice versa. The gates are tied together at 5V, but one source is at ground and the other is a V_{DS} below ground. Thus, they differ by approximately 300mV. The result is an overdrive value that differs by a maximum of approximately 10% at the maximum output current. Consequently, the on-resistance of the power device fluctuates based on the change in drain voltage. The actual effect is less than 10% in the worst case because of the device layout. The actual layout of the power devices is slightly different than a conventional NMOS so there is a constant resistance term that lessens the effect of the overdrive voltage. According to the graph below (Figure 4-6), the effect of a different overdrive voltage ends up being at worst about 2%.

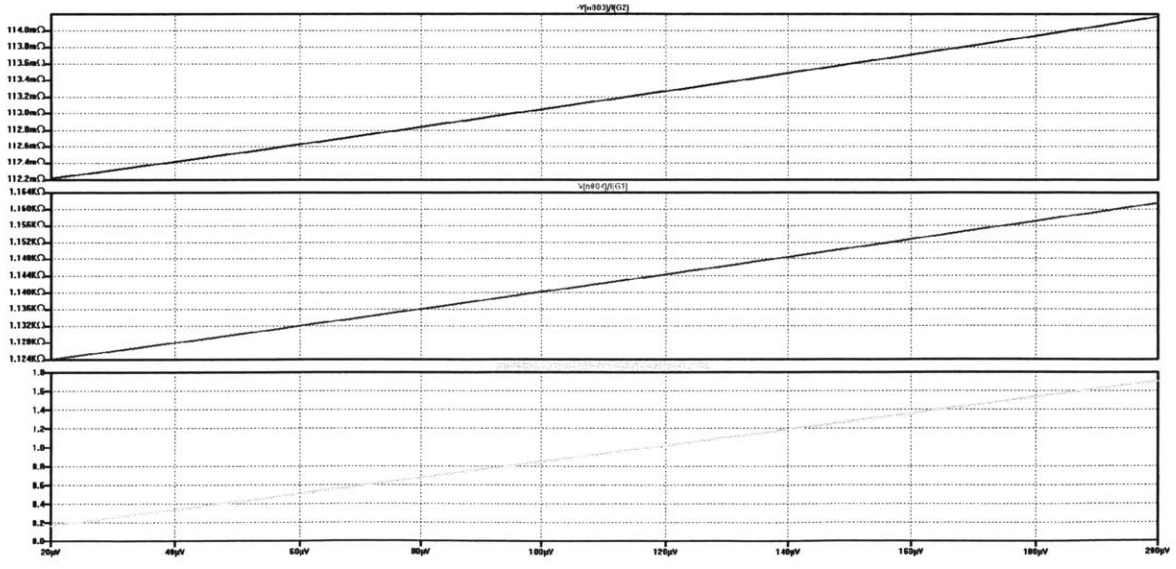

Figure 4-6: Ramped current through NMOS power devices. SenseFET on-resistance (top), Power FET on-resistance (middle), Percentage of difference (bottom).

Although the 2% error is small, this problem does not exist for the single transconductance amplifier topology because the drains are tied together and the sources are regulated to the same voltage.

This problem also does not exist for the top side devices.

4.2.2 Transconductance Amplifier

For each possibility for implementing the transconductance method, there is a specific requirement for the transconductance amplifier. Generally speaking, a transconductance amplifier takes a voltage at the inputs and turns it into a current at the output. In the case of the single transconductance amplifier method, the requirement is it needs to take both positive and negative input voltages near ground.

For the case of the opposing transconductance amplifiers, they need to take a wider input range centered at ground and they need to be highly linear about that point. Additionally, they need to be well matched; any mismatch between the two amplifiers appears as an error in the output current.

There are a few factors to take into consideration. The first is whether the input pair shall be NMOS or PMOS (Figure 4-7). Generally NMOS transistors are faster but in this application, the input voltages will be near and possibly below ground. Therefore, if the input pair is to be NMOS, a level shift is required to bring the voltages up to the required minimum voltage. With a PMOS input pair, the input voltages can generally be lower, but there is still a limit as to not push the NMOS current mirrors into triode. For the purpose of this application, the switch node shall not run more than a diode drop below ground; otherwise the gate on-resistances are set too high. Thus, the lower bound on the input for a PMOS input pair only needs to be enough to satisfy a drain-source saturation voltage plus diode drop. To further reduce the acceptable input voltage, a folded cascode topology can be utilized. The benefit of this is it reduces the voltage required to support the current mirror down to two drain-source saturation voltages. Unfortunately, both approaches require increased power consumption.

In the scenario of the single transconductance amplifier, both inputs shall sit very

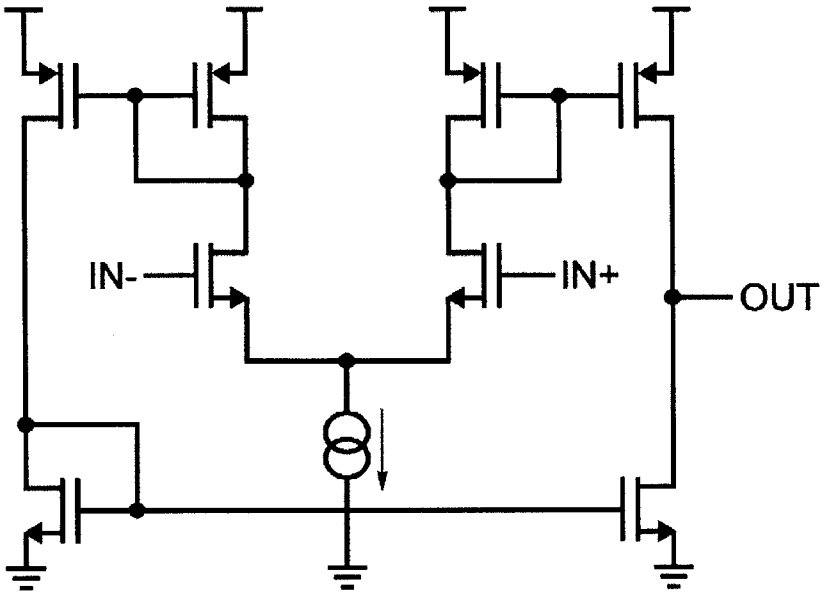

(a) N-input pair

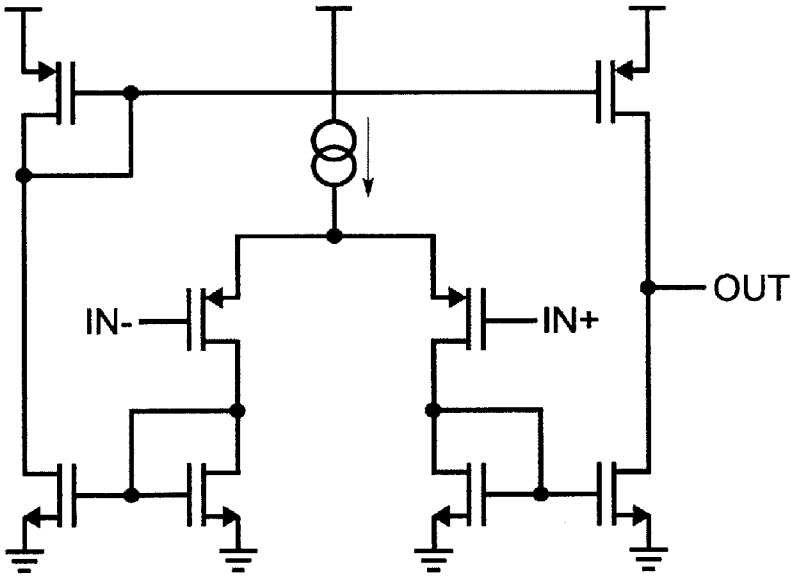

(b) P-input pair

Figure 4-7: Simple transconductance amplifiers.

close to ground, which means a simple P-type input pair with no folded cascode implementation shall be adequate. For the opposing transconductance amplifiers, the inputs are the positive and negative drain-source voltage of the bottom gate, respectively. Because of that, either the level shifted N-type or folded-cascode P-type implementations are required.

In order to maintain the proper transconductance, wide devices with short channels are required for the inputs. The current mirrors must also be wide to reduce the drain-source voltage but have long channels for good matching. Unfortunately, for the folded cascode topology, the output voltage of the amplifier sets the drain-source voltage of the cascode device, leading to channel length modulation. This difference becomes an output voltage-dependent offset that ultimately throws off the average output current. Luckily, if the amplifier output is held to a small enough range, the offset shall be relatively constant over different output loads. Any offset seen in actual implementation can be canceled out by some constant current source. If this is not sufficient, more exotic techniques often used in deep submicron CMOS processes can be employed. For the process used at Linear and the purposes of this application, those techniques are not necessary.

To address some accuracy issues that are discussed later, a transconductance amplifier is designed for the top gate that operates on the V_{IN} voltage and has a ground-referenced output (Figure 4-8). It is determined that a complementary NMOS input pair folded cascode transconductance amplifier best meets the requirements. The requirements are that it has to take inputs close to Vin but with the output being the Vc node, which is an input to a low voltage, ground referenced circuit. Because of the high voltage, protection devices are needed to limit the actual output voltage of the circuit; all other devices used are low voltage for good matching. The additional devices required for this design become the limiting factor when determining the V_C range, which is then used for a voltage to current conversion with the current comparator. As in the case of the low-voltage transconductance amplifier, care is taken to ensure minimal current offset due to the variable output voltage. Note that the input drains are connected before the current source cascode devices. Because the inputs

operate at relatively close voltages while the output voltage operates in the range of approximately $1 - 4$V, the current offset due to channel length modulation reflected from the output voltage outweighs the offset due to channel length modulation from the input pair. Adding any additional cascode devices further reduces the possible output voltage range for lower V_{IN}, which negatively impacts the resolution of the subsequent voltage to current conversion.

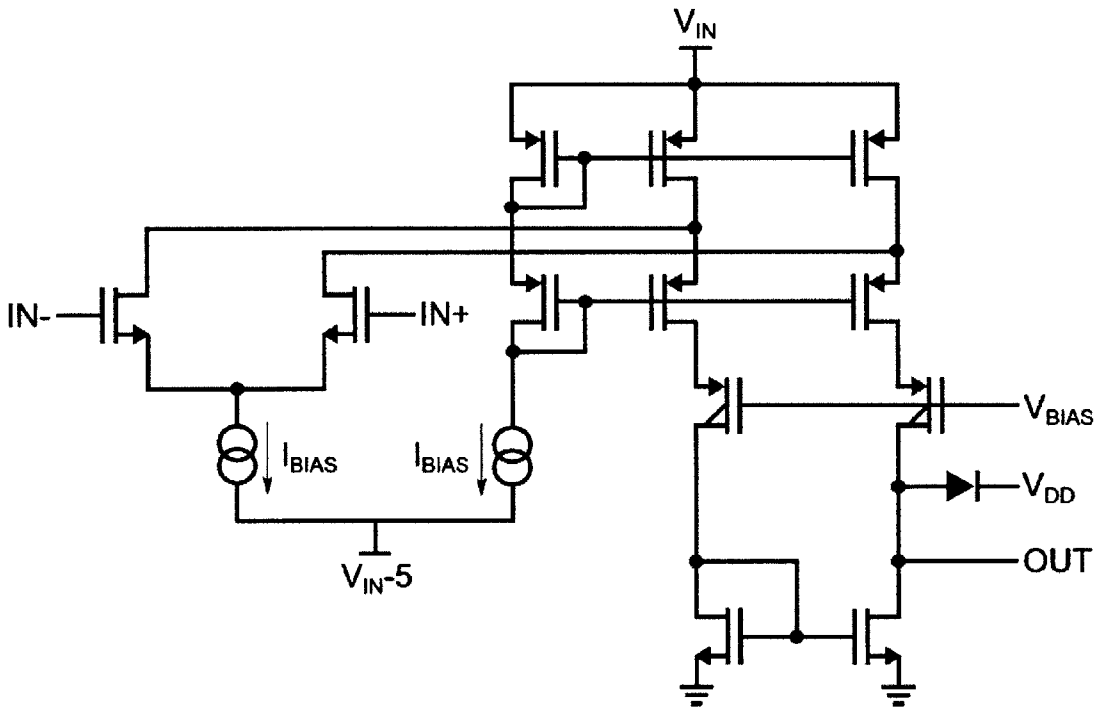

Figure 4-8: Transconductance amplifier for high voltage top gate sensing.

Unfortunately, the one of the most important aspects - matching between two transconductance amplifiers, depends on process variation, which is more difficult to control. The problem is especially difficult with a top gate and bottom gate sense, as they are completely different topologies as opposed to matching two of the same circuit. To overcome this, it may be necessary to perform some trim operations - or else any offsets in the amplifiers can ruin any matching between the power device and senseFET.

Transconductance Amplifier Output Offset

In the opposing transconductance amplifier topology, if the two amplifiers are not well matched, they will have some offset. If there is any inherent offset in the amplifiers, it is reflected directly into the output current. Thus, much care must be taken in layout to ensure the amplifiers are as well matched as possible. Luckily, the mismatch between the two amplifiers can be trimmed out using the senseFET trim circuit.

Transconductance Amplifier Input Offset

The input referred offset of the transconductance amplifier shall also be reduced as much as possible. The methods of doing this include using a large bipolar input pair or some other techniques such as chopper or auto-zeroing[14]. It is also possible to use trimming as well.

Transconductance Matching

Using the switched transconductance amplifiers connected to the same node, it is important to maintain a similar transconductance. If there is any difference, the difference reflects directly into how much each part is weighted. The ideal case is for the weighting to be equal at 50% duty cycle. However, as long as the transconductances are relatively close, the error due to uneven weighting is negligible.

4.3 Common Issues

4.3.1 Sample and Hold vs. Sample and Off

One aspect that plays a factor in determining the output current error in the initial sample and hold scheme is the holding error, which is proportional to the duty cycle. With bottom-gate sensing, the error occurs when the bottom gate is off; the value being held is the minimum value of the inductor current waveform. The calculated average current is thus lower than the actual current, leading to a higher output current. As shown in Figure 4-9, when sampling and holding, the error of the system

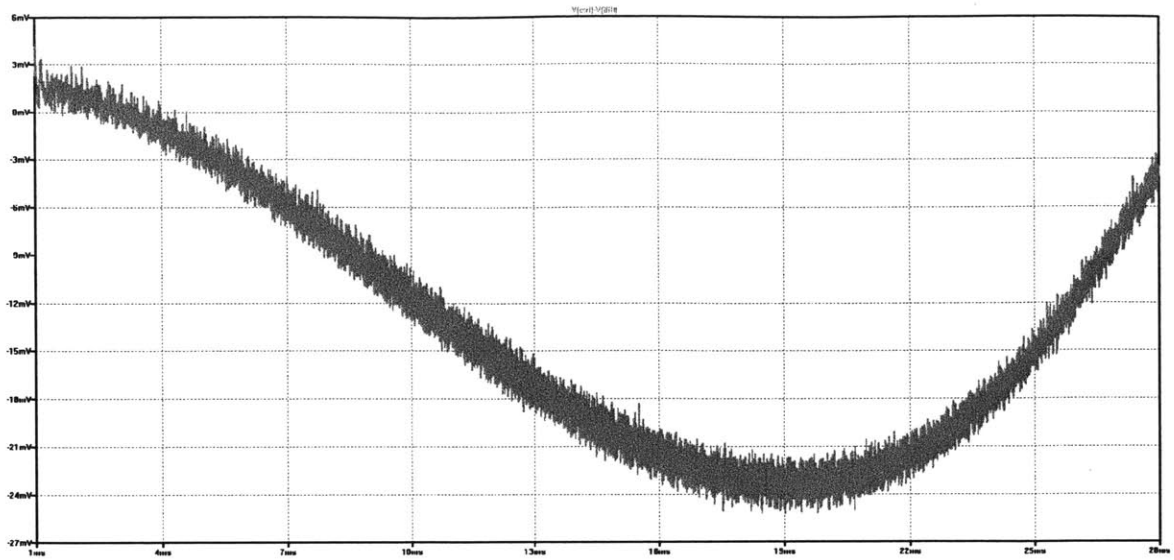

Figure 4-9: Error in output current vs. duty cycle. Note the waveform plotted is the control current minus the output current, so a negative value means the output current is too high.

is greatest when the duty cycle is near 50%. This occurs because the current ripple is greatest at medium duty cycle. The holding error is proportional to the ripple times the off-time of the switch. While a small duty cycle has a large off-time of the switch, the decrease in current ripple causes the overall error to be smaller than the case with large ripple and a medium duty cycle; this explains the bow-like shape of the error waveform.

One way to go about reducing the holding error is to turn off the measuring scheme when the switch is also off. In theory, there is no holding error because nothing is being held. However, this scheme then becomes more susceptible to any non-idealities that would impact the measurement, such as the turn-on spiking. As the duty cycle decreases, factors such as the body diode voltage drop will contribute the same error for a constant period of time while the actual measurement time decreases, thus increasing the effect on the measurement value. This causes the overall error to increase as duty cycle decreases, effectively replacing the holding error with a different error.

In the case of a low-duty cycle application, it may be acceptable to use only bottom-gate sensing. But for a general purpose part, this is undesirable. The short-

comings described are exactly the reasoning behind implementing the following section.

4.3.2 Top Gate vs. Bottom Gate Sensing

To really understand why the problems from the previous section occur, it helps to explore idea of sensing current on the top gate or bottom gate.

The situation for sensing on the top gate is directly analogous to sensing on the bottom gate. The error occurs in top gate sensing when the top gate is off, and the current value being held is the maximum current value. The calculated average current is higher than the actual, thus the actual current is lower than desired. And in contrast with bottom gate sensing, using the top gate could still be acceptable for applications with high duty cycles. The issue becomes slightly more complex, however, when observing the actual error over duty cycle. While it is true that the incorrect value is held for the respective extreme duty cycle, the error observed in the output current is also a function of the error in the held value, which also changes with duty cycle. In fact, the error in the held value is largest for a moderate duty cycle. This actually impacts the end result much more than the amount of time spent holding the incorrect value.

Thus, each method has its benefits and drawbacks. With top gate sensing, the duty cycle is limited on the low-end, since the top gate needs to be on for a certain amount of time to calculate a relatively accurate current. And instead of the dip in output accuracy like shown in Figure 34, there would be a bump. The opposite is true for bottom gate sensing, with the limit being on the high end of the duty cycle. All circuit elements being equal, the advantage goes to bottom side sensing because the circuit can be referenced to ground, while the top-side circuit would have to be floating with the input voltage. For a circuit to float with the input voltage, which can go up to 60V or more, the techniques are not necessarily difficult, but the circuit is inherently more difficult to implement because care must be taken to ensure the devices operate with safe voltages.

Neither top-gate sensing nor bottom gate sensing on its own is ideal. The goal is

to obtain the maximum performance on the widest possible output current range, so it would be advantageous to combine the best of both worlds - that is, sense on the top and bottom gates. Through sensing on both the top and bottom gates, there is little holding error because something is always measuring the actual inductor current. The method of sensing both gates can be clearly demonstrated with the current summing approach. Each circuit is activated for its respective portion of the duty cycle and outputs a current proportional to the inductor current; the resulting currents are then summed together. It is possible that there is some error due to transitioning between the top and bottom gate sensing, but since it is averaged over the entire period, that error should be negligible. With the transconductance amplifier-based design, it is also possible to apply a top and bottom gate sensing method, as shown in Figure 35. Of course, some of the disadvantages remain from top and bottom gate sensing, such as high voltage and complex circuits.

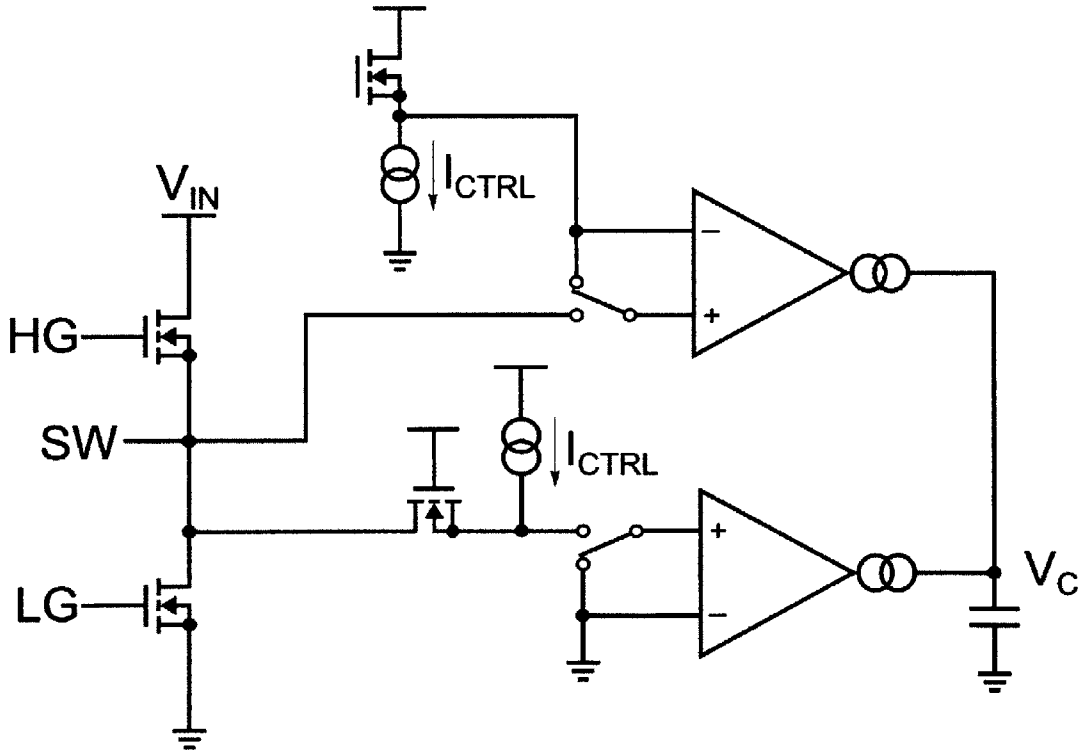

Figure 4-10: How to implement sensing on both top and bottom gates.

4.3.3 Balancing Top and Bottom Sense Circuits

The simplest is to use only one gates sensing mechanism at a time and switch off around 50% duty cycle. Hysteresis can be employed to ensure there is no switching back and forth due to noise, but with two discrete levels, it is likely there will be a discontinuous jump in output current at the transition regardless of where it is. Instead, the best way to avoid the duty cycle dependent error is to just combine the top and bottom gate sensing circuits and turn off each respective circuit off when it is not in use. This in effect weights the top or bottom sensing circuits based on the duty cycle so the gate that is on for longer affects the result more.

For the current summing method, it is already taken care of because the instantaneous current is calculated and then low-pass filtered to obtain the average from both gates. On the other hand, for best accuracy, the two transconductance amplifier-based methods require a method of switching between the top and bottom sensing circuits. The ideal result is a completely linear relation between duty cycle and output current. But due to mismatches and offsets, the top and bottom average sensing circuits likely do not match up, leading to an undesirable discontinuity if they are individually used. Any discontinuity is undesirable because that leads to either a missing code or a code overlap. For better continuity and consistency, it may make sense to have a range in which both sense circuits are regulating. In the end, the simplest method is to connect each circuit for its respective portion of the period, effectively weighting each circuit using the duty cycle signal. The average shall then be a continuous function for all duty cycles.

However, there still remains some error in this implementation! The drain to source capacitance on the top and bottom gates introduces more error into the calculation. Between the two, the top gate error significantly outweighs the bottom gate error. This is because the bottom gate has zero-voltage switching from the body diode conducting before the switch turns on. Thus, the maximum voltage drop at the switch node will be a diode drop. However, with the top gate, there is no body diode to start conducting current, so the switch node starts at 0V and then swings

up to the rail. This introduces a huge error into the system because instead of seeing just a few hundred millivolts, the current sense circuit is led to believe there are tens or hundreds of times more current flowing through the top switch. The solution to this is described in the next section.

4.3.4 Sample Blanking Circuit

When the top switch turns on, its drain-source capacitance is charged all the way up to the input voltage. In order for the channel to properly form, that capacitance must be first discharged, which takes the form of a large current spike through the top switch. Much of that current goes into charging the bottom switch capacitance. So for any current measuring circuits connected, it appears like a lot more current is flowing through the inductor than it actually is. This throws off the accuracy of the current averaging circuit, and thus, a blanking signal shall be used for the top gate transconductance amplifier to blank out that initial spike.

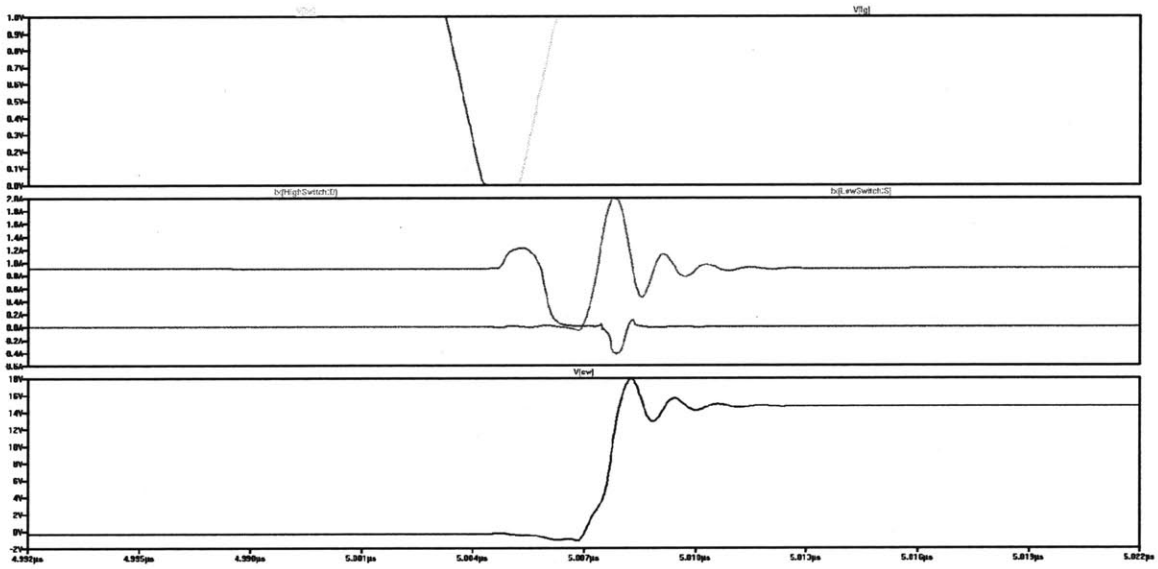

Figure 4-11: Voltages do not rise up immediately when the switch is turned on.

Just blanking out that initial spike is not enough to fix the problem. If the remaining duration of the period is fully sampled, then the period sampled becomes asymmetric. One simple method to block out the largest spikes is to attach an inverter to the switch node and or that signal to the gate control signal. Thus, once the switch

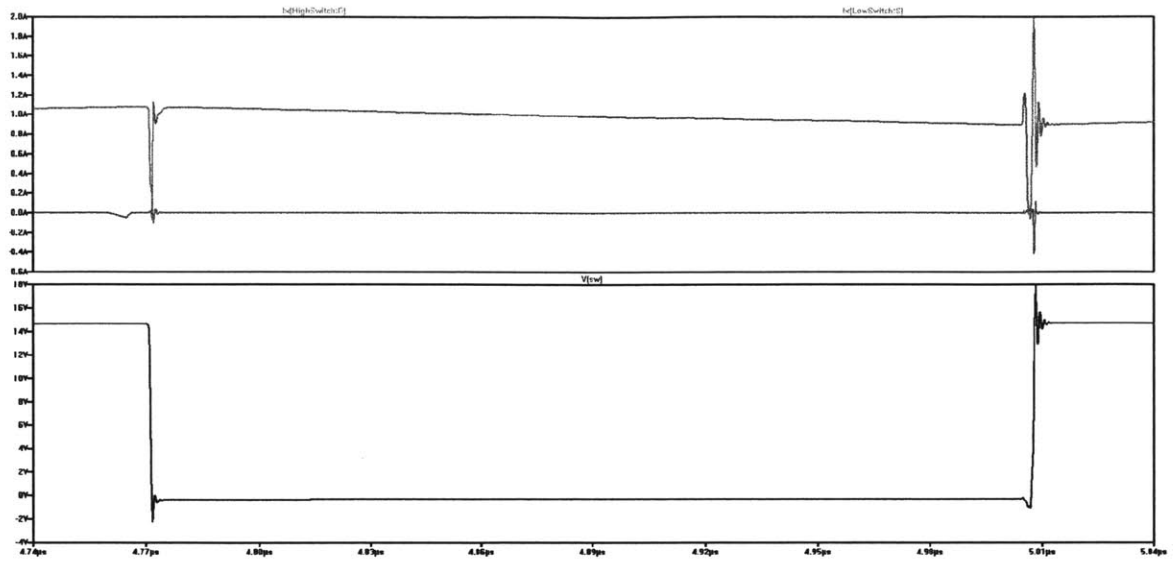

Figure 4-12: Voltages do not rise up immediately when the switch is turned on.

voltage passes a certain threshold, most of the spike is already settled. Then as the top gate turns off, the switch voltage once again decreases and the sampling stops there as well.

A similar problem exists for the bottom switch, but as it is soft-switched, the inaccuracy that results is far less than what results from the top switch. Thus, it is left as is.

In practice, it is not critical to achieve extremely precise sampling times and worry about matching delays because in this implementation, both the top and bottom switches are being sampled. In the case where only a single gate is being sampled, at one of the extremes of duty cycle, a large current spike will throw off the current estimation greatly. But when it is being averaged over an entire duty cycle, its effect is greatly reduced and becomes a constant over all duty cycles. As simulated, this error term becomes negligible and a blanking circuit is not actually needed to achieve the desired accuracy levels.

4.3.5 Sampling Error and Blanking

Due to the current spike on the high gate, blanking is implemented. However, the most simple implementation of using just a inverter to determine when the switch

node has passed a certain threshold still generates some error to the sampled current. Most of the current spike is eliminated, but during the turn-off, since there is a slight delay from the inverter, there is a short time in which the switch samples a low switch voltage, believing it to be a large top gate current. Since the bottom gate is still measuring during this time, the error is a constant throughout duty cycle, as the duration of the error remains a constant regardless of how long the top gate is on.

4.3.6 Voltage to Current Converter

In any case, it is important to keep track of the output voltage range at the Vc node. Both top and bottom sensing circuits must be able to regulate for the designed range of the Vc node, which is then translated into a current that feeds into the peak current comparator. Setting a large voltage range gives higher overall resolution but puts tighter constraints on the output stages of the circuits that precede it. It also means the gain of the error amplifiers does not have to be as high.

4.3.7 Peak Current Detection

Since the feedback method used is peak current-mode, there must be a way to determine the peak current quickly and somewhat accurately. High accuracy is not a major concern because the feedback loop is capable of compensating for any peak current inaccuracies, as it is regulating the average current using circuits that are focusing on high accuracy. Untrimmed, it is possible to achieve approximately 10% accuracy between the peak sensing and power devices. In the current summing method, the existing average current detection is just a low-pass filtered version of the sensed inductor current. Thus, no additional circuitry is necessary. For the other two methods, the average current sensing circuits hold only that - averaged information. The peak detection circuit uses a reference current generated from the compensation node voltage to set the desired peak current. Then, a comparator trips when the inductor current reaches that preset value. Effectively, this method controls the duty cycle.

For the comparator topology, it is simplest to use a common base differential pair

as shown in Figure 4-13. The reason for having a sense-FET on both sides is to match the offset due to the bias currents. This extra bias current does not have much effect on the voltage across the power device, but since the sense-FET is thousands of times smaller, the bias current causes a significant offset. Thus, to achieve the desired accuracy, it is necessary to account for the offset. It is also important to implement some hysteresis (not shown) so the output does not fluctuate with noise. This circuit is only necessary on the top side gate, as the inductor current only increases when the top gate is on.

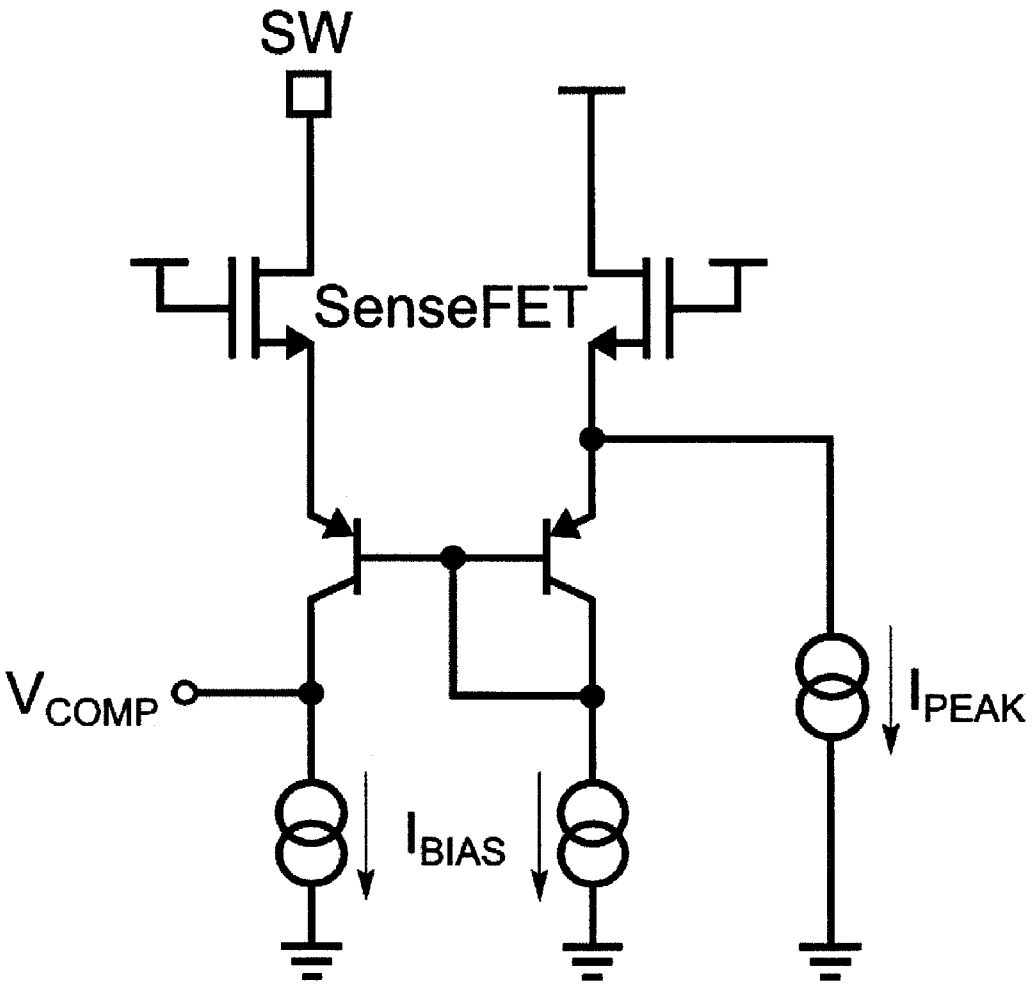

Figure 4-13: Peak current sense comparator.

4.3.8 Overcurrent Protection

It is necessary to implement an overcurrent limit to prevent damage to components. While the external loop is designed to be able to handle this, there is no guarantee of catching a large transient spike. Thus, a separate comparator with a hard-programmed limit is also included. The design is exactly the same as the peak current comparator but uses a constant current reference instead of a variable current sink based on the compensation node voltage. This hard limit adjusts based on the sense ratio - although in the actual circuit implementation, no additional circuitry is needed for the adjustment because the power device change already takes care of that. For large currents, the current limit will be X. For small loads, the hard current limit will be Y. These limits generally are determined by practical values for the inductor, which factors into the current ripple. For a monolithic chip, the power devices are known, so it is much easier to set the current limits. Like the peak current sensing circuit, the overcurrent detection does not require exact matching between the sense-FETs and the power device.

4.3.9 Peak/Over-Current Blanking Circuit

For the over current comparator described above, it would not function properly if implemented on its own. Unfortunately, the switch node experiences ringing due to various parasitic inductances and capacitances, which might lead to false trips. The useful information is only available once the ringing dies down, so a blanking circuit is necessary to block out the ringing. According to a model with 1nH bond-wires and 30ohm parallel dampening resistance from radiated energy, the ringing goes away after approximately 10ns, which corresponds to 5% minimum duty cycle. A simple blanking circuit of a constant 10ns delay should generally be sufficient for the current comparators for the V_{IN} voltage range used, although that may change if a larger range is desired. The simplest implementation is just an inversion of the clock signal, as long as the clock pulse width is set to the necessary value.

One issue to keep in mind is that combined with the W-switching, the minimum

blanking time required changes due to the larger gate capacitance of the larger power switch and as well as the need for extra driver stages. Due to this, it is necessary to either set the blanking time based on the large switch or to change blanking times depending on which switch is being used.

Doing a valley current regulation scheme places a similar constraint on the high end of the duty cycle. There is not much benefit, if any, because the bottom switch is soft-switched, which means the period of time in which the body diode is conducting also has to be blanked.

4.3.10 Shoot-Through Protection

A shoot-through protection circuit (Figure 4-14) is implemented using some logic gates and delays to prevent both switches from being on at the same time. In essence, the circuit simply ensures there is a delay before each gate is turned on so the other gate is able to first shut off completely. Using ideal switches, this is not a factor. But left unaccounted for, the shoot-through current will become the main power loss factor - if the chip does not blow up! It also has a large impact on accuracy; implementing the shoot-through protection circuit decreases error by 2-3%. This likely occurs because during the time in which both switches are conducting, the assumption that the switch current is equal to the inductor current is no longer true.

4.3.11 Sense Ratio Adjusting Circuit

The sense ratio adjusting circuit (Figure 4-15) employs the concepts of a basic flash ADC to choose the correct sense ratio based on the user-defined control voltage. A larger control voltage shall increase the ratio between the switch and sense-FET, while a smaller control voltage shall decrease the ratio. This circuit is designed with a non-ideal control signal in mind, with noise being the main issue. Thus, each of the comparators in the ADC has a bit of hysteresis programmed in to prevent the circuit from switching rapidly between sense ratios at the transition voltages in the presence of noise. Theoretically, this circuit can be used with any number of step

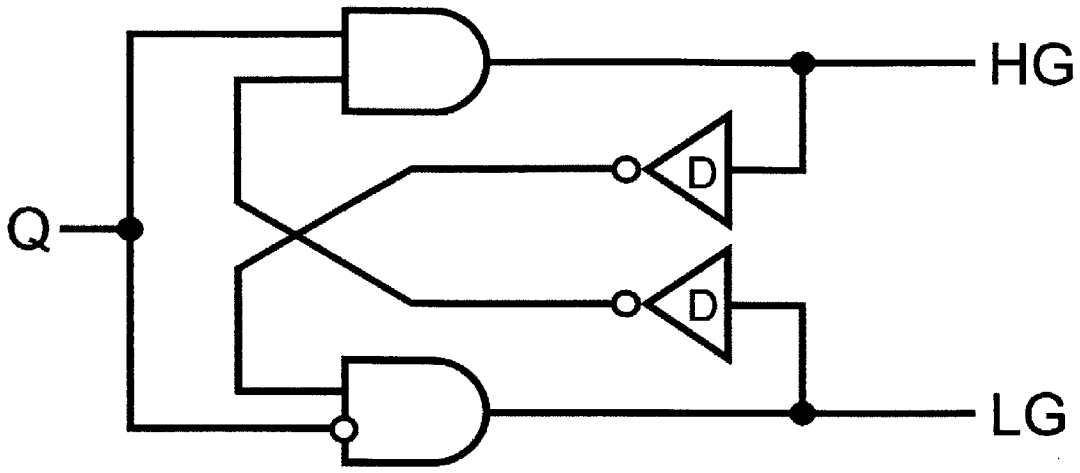

Figure 4-14: Shoot-through prevention circuit.

sizes to achieve the sense ratio desired for any control voltage. However, considering space constraints as well as the marginal benefits of adding more steps, it is much more practical to stick with the fewest steps possible. In this case of a 100:1 analog dimming ratio, having only one step - two ratios, provides the most benefit for accuracy; any marginal benefit of additional steps is outweighed by the cost in size and complexity.

As implemented with any of the proposed topologies, changing the sense ratio by adding or reducing senseFETs does not give any real benefit to the measured voltage, which is the limiting factor at lower currents. While it does boost the signal by increasing the sensed current, the boosted signal happens to be after the point of greatest impact; the drain to source voltage on the power FET remains the same regardless of the ratio, which means the amplifier must maintain the same drain-source voltage on the sense-FET as well. A way to get around this is instead of changing the senseFET, to adjust the actual power FET, which effectively changes the on-resistance and provides a larger or smaller drain-source voltage that actually helps to reduce the error (Figure 4-16). This method is called W switching[15]. Since the drain to source voltage is smaller for low currents, increasing the on-resistance boosts that signal, making it easier for the amplifiers in the regulation circuit. The original motivation behind W switching is to increase the efficiency at low output

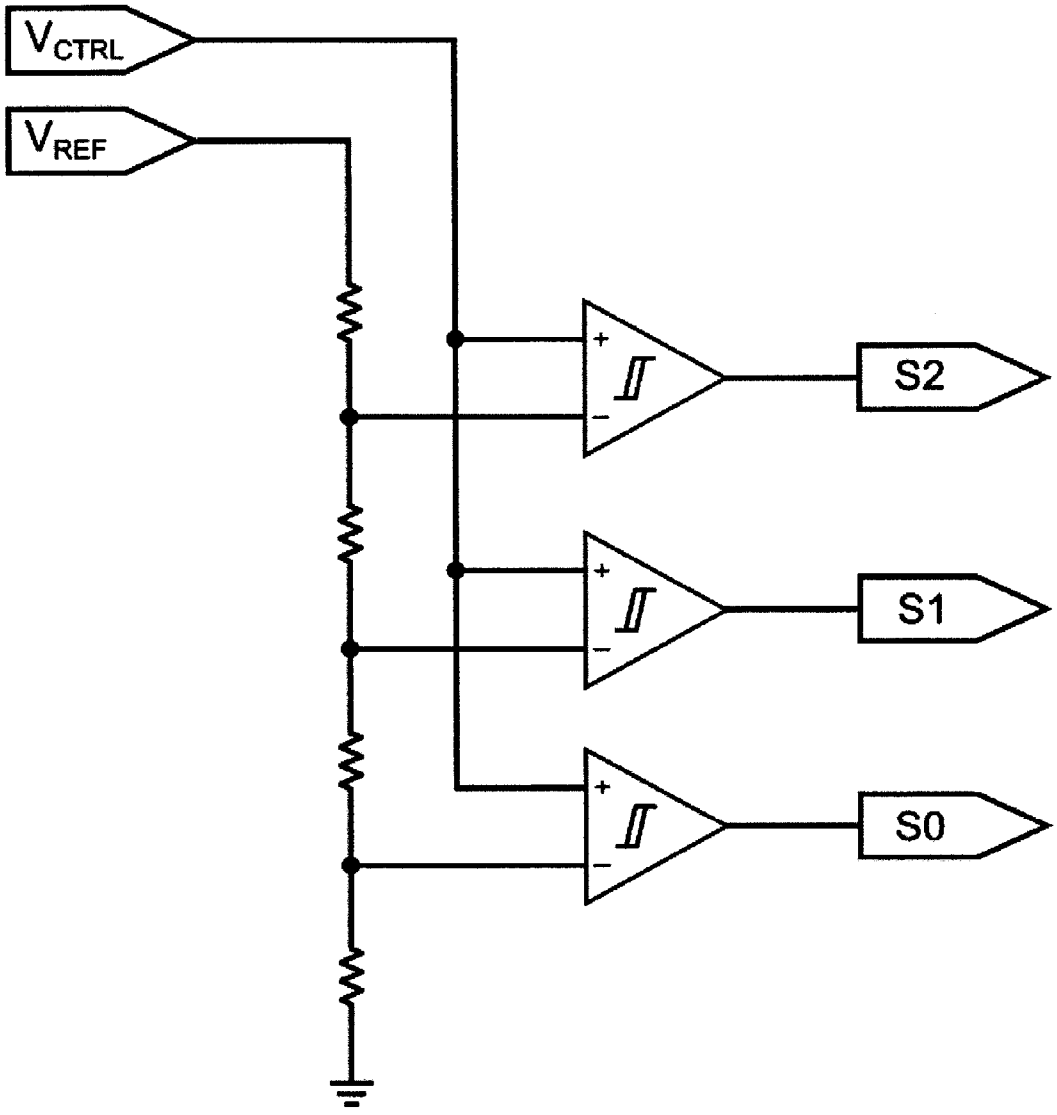

Figure 4-15: Sense ratio adjusting circuit implemented with 3 bits.

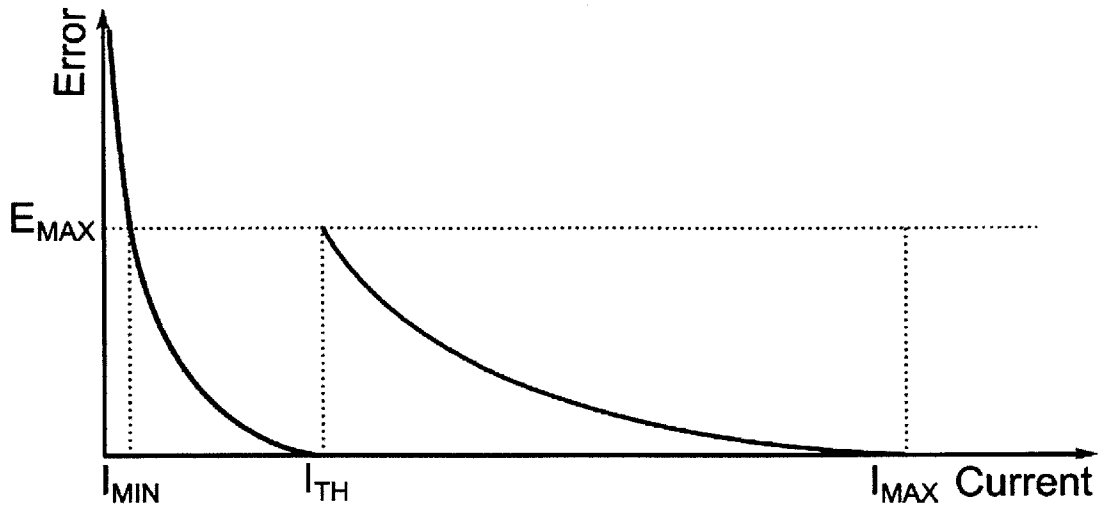

Figure 4-16: Expected error of 2-ratio system with relation to output current.

currents, which is a positive side effect as well. The three main losses in a buck converter result from switching loss (Figure 4-17), gate charge loss, and conduction loss. Listed below are the equations for each:

$$P_{COND} = I_{OUT}^2 R_{DS(ON)} \frac{V_{OUT}}{V_{IN}}$$

$$P_{SW} = \frac{1}{2} V_{DS} I_D t_{SW} f_{SW}$$

$$P_{GATE} = Q_G V_{DD} f_{SW}$$

As the output current decreases, the conduction loss reduces by a square factor and the switching loss decreases linearly. However, the gate charge loss remains constant, so the gate charge loss becomes a bigger factor compared to the other losses. While more complicated solutions exist, such as resonance or gate drive reduction, the simplest way to reduce the gate charge loss is to reduce the gate capacitance, given a set internal rail and switching frequency. By using two separate power FETs of different sizes, the smaller one can be used for lower output currents without the penalty of a large gate capacitance. While the conduction loss increases with the smaller device due to a larger on-resistance, since the power loss is a square factor

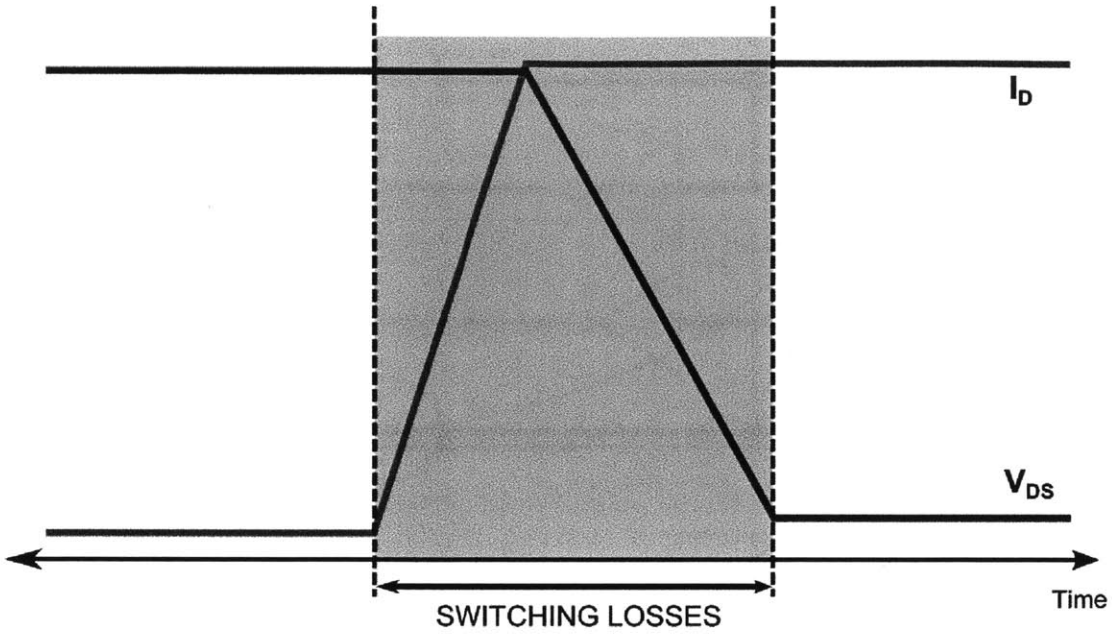

Figure 4-17: Timing diagram of high-side switching losses. Adapted from Klein[16].

of the current, the increase in conduction loss shall still be less than the decrease in switching loss. In practice, changing the gate size only shifts the current peak efficiency curve, as shown in Figure 2 of Williams et al [15].

In terms of practical implementation, the largest gains are made with two separate power FETs; any additional FETs for increased resolution will at best only provide a small marginal benefit and at worst significantly require more space. For layout purposes, the smaller power FET can be placed separately from the larger one, with no regards to matching between them. Instead, either two senseFETs can be used with their own trim circuits or one senseFET with two separate trim circuits. One important note to keep in mind is the when the sense ratio changes, the voltage-to-current gain from the compensation node must be adjusted as well, or else the peak current value becomes set for the wrong ratio. Left on its own, the compensation node voltage adjusts to the new regulation point, resulting in an undesired transient. To account for the new regulation point, the safest method is to reset the compensation node and allow the loop to bring it back up.

For the purpose of this specific application, if the maximum output current is

aimed at 2A, the larger power FET M2 shall be sized accordingly for maximum power efficiency, with the smaller power FET M1 sized approximately ten times smaller (Figure 4-18).

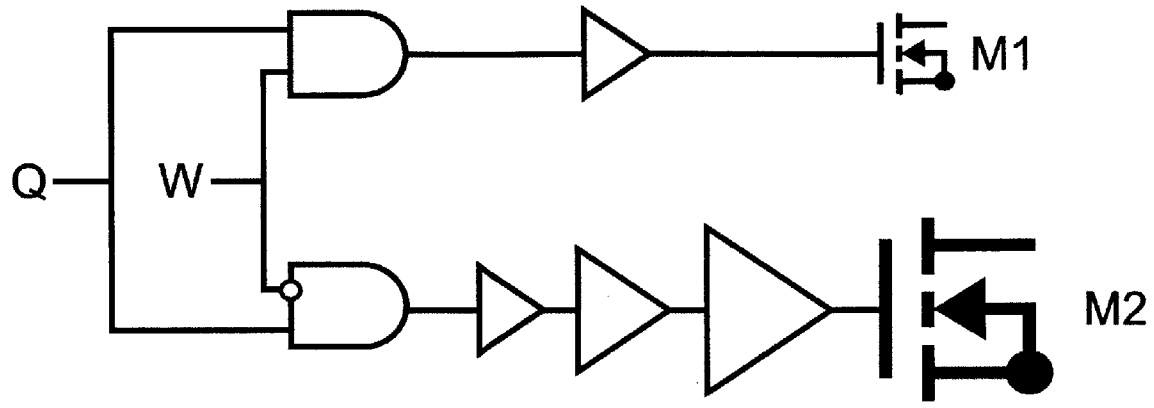

Figure 4-18: W-switching mechanism.

4.3.12 PWM Control

The PWM control circuitry runs open-loop, as the PWM frequency is relatively slow - around 200Hz. For a 1000:1 PWM dimming ratio, that requires a 5μs minimum on-time. Because of the inductor, the ramping down of current takes some time. But because there is no output capacitor to charge up the next cycle, there is also no worry about just shorting the output to ground. The most straightforward implementation of PWM control is to AND the PWM signal with the clock.

The V_C node voltage must be held during PWM off to allow for a fast turn-on time. If it is allowed to run freely to any voltage, since both gates are off, the regulation loop tries to run more current through the switches, pushing the node voltage higher and higher. The node voltage sticks high until the PWM on signals and the peak current is subsequently set too high, leading to a large output current overshoot and slow correction back to the intended value. For something like a 1000:1 PWM dimming ratio, this is unacceptable, as the circuit shuts down again before even having a chance to settle to the correct output value.

4.3.13 Trim Circuit

There are two main points where trimming is crucial. The first is the actual senseFET to ensure the sense ratio is as expected. Trimming is necessary for the senseFETs in both cases to ensure an accurate match of the drain-source voltages because as shown in Figure 4-19, the standard deviation of the FET drain to source voltage is approximately 10% of the mean. It would be futile to aim for less than 5% error if the sense ratio is going to be off by more than 10%.

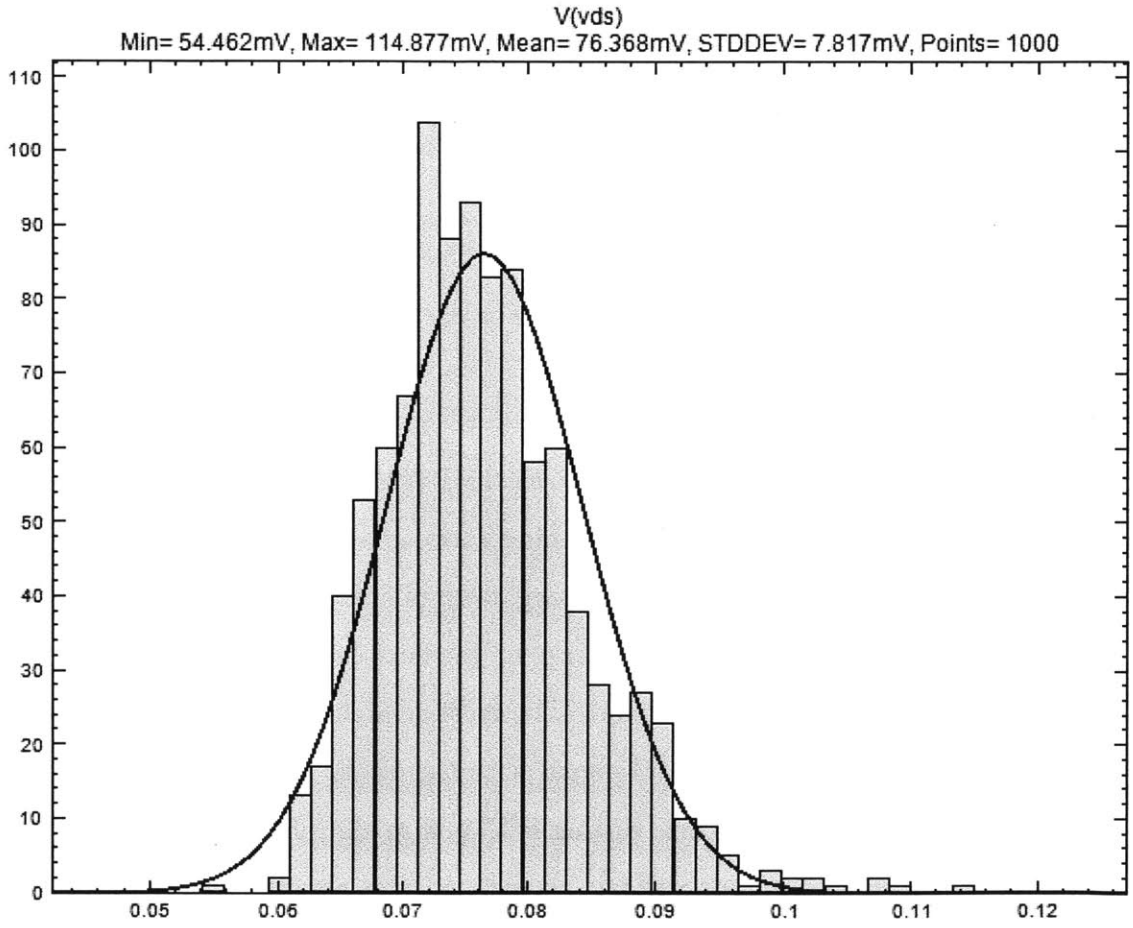

Figure 4-19: Histogram of drain-source voltage of randomized sense-FET over 1000 runs.

The second critical trimming component is the control current source. If that current is not accurate, then having the correct sense ratio will be negated.

Elements in which extreme accuracy is not necessary include the peak current de-

tection circuit, which include the corresponding senseFET and the voltage-to-current resistor. If the previously mentioned elements are trimmed, the control loop can compensate for any inaccuracies presented in the peak detection circuit.

Trimming will also help for the transconductance amplifiers; if they are unmatched, the top and bottom sense circuits will be weighted unevenly. However, assuming the transconductances are close in value, the effect of uneven weighting will be negligible.

However, trimming is not required in all parts of each scheme. For example, in the case of the opposing transconductance amplifiers, the average current is controlled through a set of trimmed senseFETs. Since all methods are still employing peak-current detection, a separate senseFET is required for that purpose as well. But because the average current is already accurately measured and stored on the compensation node, the peak current sensing does not have to be extremely accurate. Thus, the senseFET for that need not be trimmed as carefully as the others - or possibly at all.

Chapter 5

Circuit Testing

5.1 Results

Since experimental circuits are not fabricated, the circuits designed are tested through simulation in LTSPICE with various test setups. For ease of testing, most circuits are tested under specific conditions with some ideal models to isolate errors due to that specific block only. Full, system-level simulations are also run, although the results can only verify some general specifications.

Figure 5-1 illustrates the averaged output current as the output voltage is ramped from 0V to 40V, which is also the value of V_{IN}, the supply voltage. The output current holds within the limits of 5% error from approximately 2V to 39.7V, which corresponds to a range of 5% to 99% of the supply voltage. The upper limit is key, as it allows the supply voltage to just slightly higher than the voltage needed to drive the load.

From Figures 5-2 and 5-3 above, note for the 200mA case, the error grows much larger than 5%. This is because the output current is right at the boundary for W-switching, with the sense ratio set at the larger value to maintain consistency with the 1A and 500mA cases. The error reduces down within acceptable limits with a smaller sense ratio.

The waveform of Figure 5-4 shows the ramp-up of a 100Hz PWM signal at 1000:1 dimming ratio. The initial ramp-up time is dependent on the inductor value, but

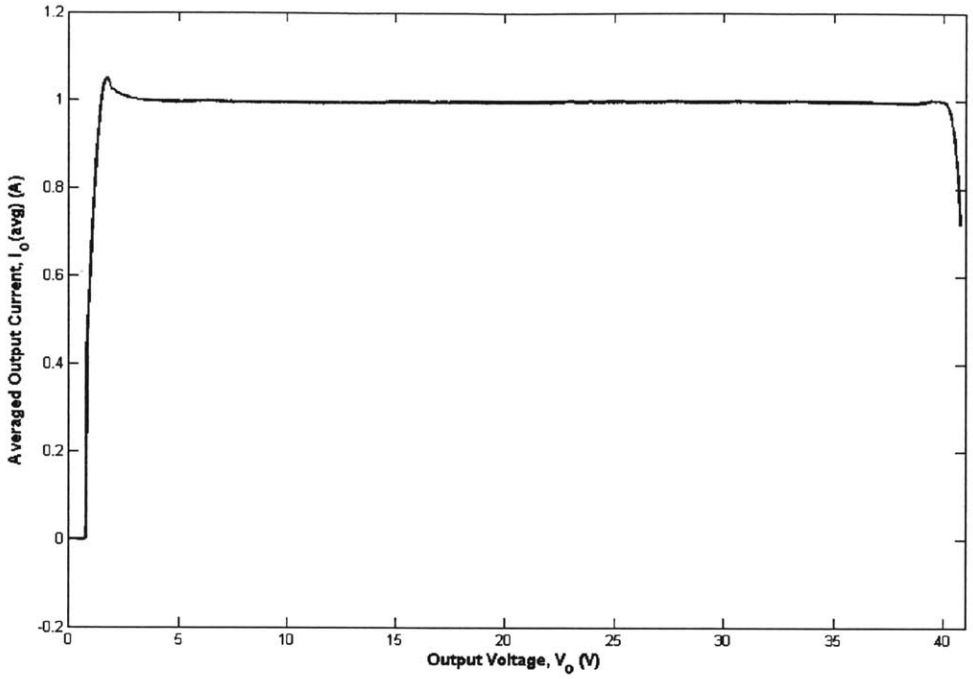

Figure 5-1: Averaged Output Current vs. Output Voltage for transconductance method.

note how the output current quickly reaches and stabilizes at the correct value.

5.2 Known Issues

5.2.1 Duty Cycle Limitations

There are a few key aspects involving the regulation scheme and physical limitations that determine the possible duty cycle range.

The main limitation comes from the use of a peak current sensing circuit. Because the regulation loop is dependent on a circuit that compares the top gate current to some regulated current, there is some minimum time required for the top gate to be on. This comparison is done very quickly, but the issue that slows down the entire process is the ringing of the switch node. This ringing is likely to cause the comparator to trip prematurely, before the inductor current has a chance to reach the desired average value. Because of that, a blanking circuit must be implemented for

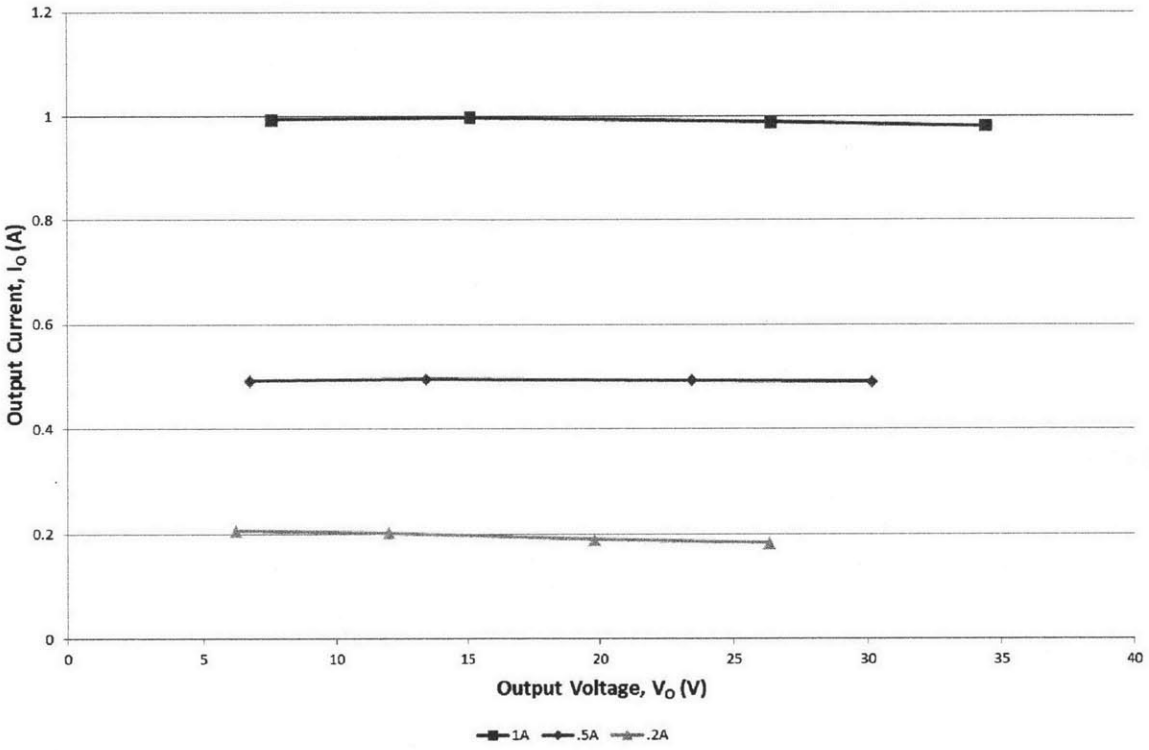

Figure 5-2: Averaged output current vs. output voltage for different current levels with transconductance method.

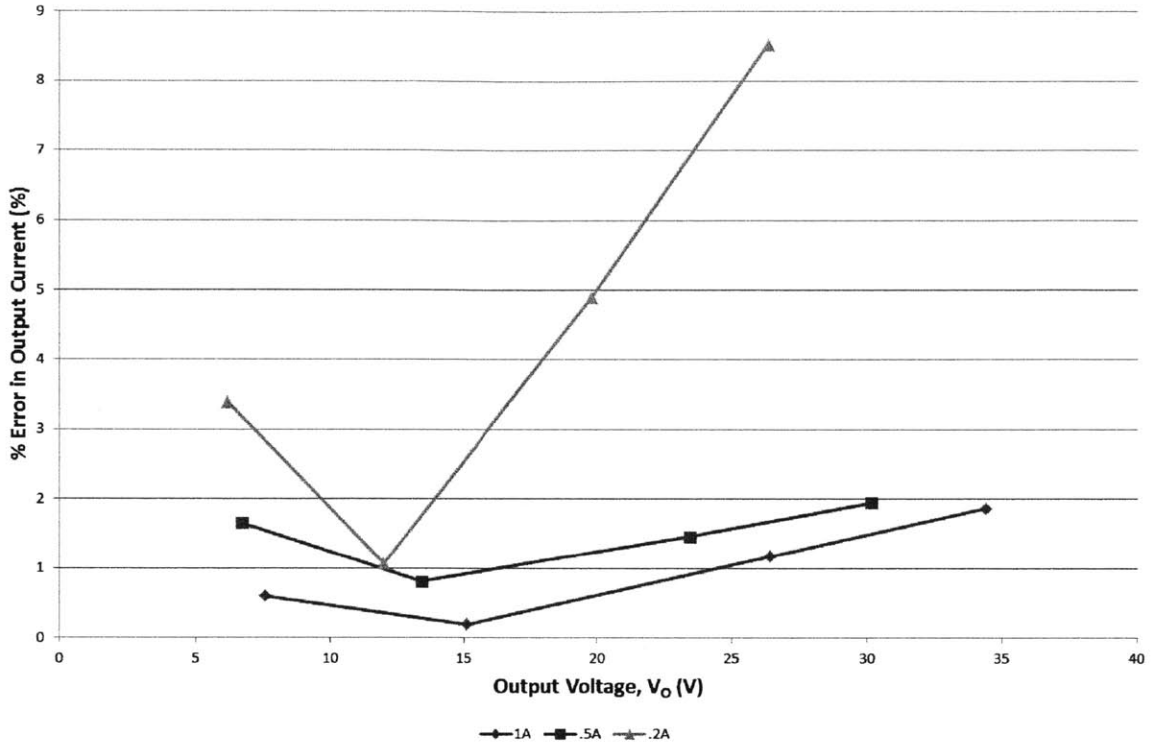

Figure 5-3: Error in averaged output current vs. output voltage for different current levels with transconductance method.

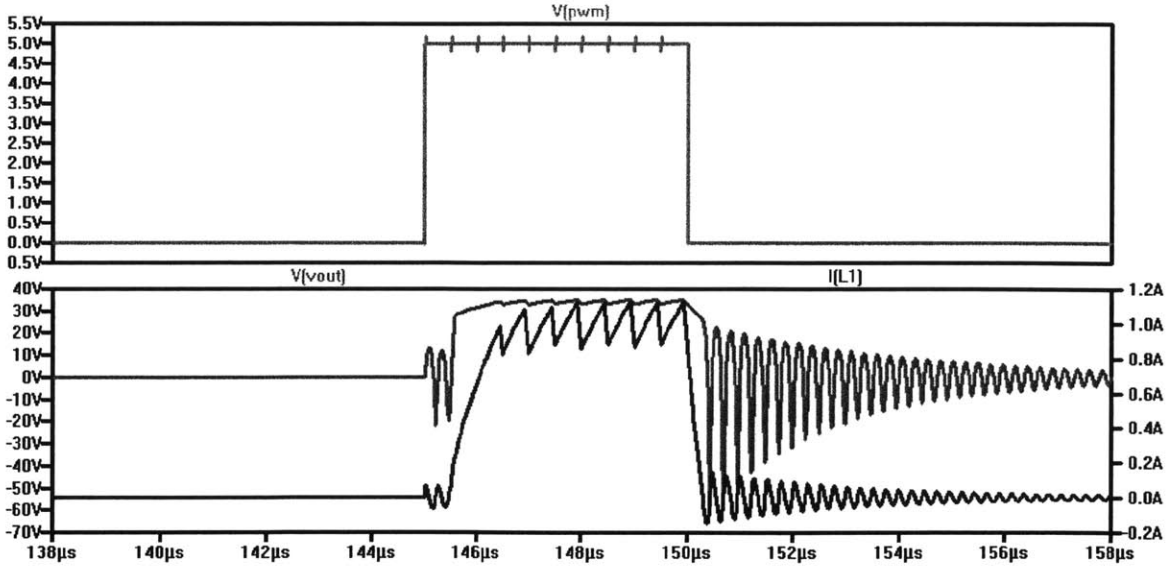

Figure 5-4: Single PWM pulse at 1000:1 dimming ratio. Note the initial rise time is limited by the inductor size.

the duration of the ringing, but that means the top gate must at least be on for that much longer. The duty cycle at the low end is then limited to the shortest possible blanking time that can be used without prematurely tripping the comparator.

The main duty cycle limitations are related to the turn-on times for the power devices. Since the average current is being measured on both the top and bottom gates, accuracy at low or high duty cycles shall remain within acceptable ranges. Thus, the turn-on times relate to the driver delay and device rise times, which shall remain within 10 nanoseconds. At worst case, this presents a 2% boundary at either extreme.

5.2.2 Efficiency at Low Output Currents

While the W-switching method reduces the power loss in the power devices, it does not reduce the quiescent current due to the regulation circuitry. The biggest contributors to power loss at small output currents are the transconductance amplifiers. To achieve adequate bandwidth while minimizing noise requires a relatively large compensation capacitor paired with a corresponding large transconductance. While a large transconductance on its own does not necessarily mean a large operating current, the use of the amplifier for measuring the drain-source voltage of the power devices requires it to operate over a large input range approximately $\pm$ 300mV. That means at the very least the output stage of the transconductance amplifier must be operating with a large current, which does not decrease as the power device size is reduced. Thus, the efficiency of the part also suffers as a result of this tradeoff.

5.3 Layout Considerations

One of the most important aspects of an on-chip current sensing topology is the actual chip layout. Even the best design cannot escape the non-idealities due to the layout; in some cases, layout becomes even more important than the design itself! While this thesis does not explore actual chip layouts for these topologies, some of the major issues are considered.

In each of the methods explored, a senseFET is used to match the drain to source voltage on the power device. It is crucial that the voltage matched is accurate, or else any steps after that cannot do anything to correct for that initial offset. Thus, a trim circuit is necessary to correct for any mismatches in the ratio. However, to ensure high accuracy across operating regions and temperature, the senseFET must also match the temperature coefficient of the power device. In addition to the device itself, the resistance and temperature coefficients of the metal leads must also be matched together.

In implementing the W-switching, it is not essential to match the small and large devices as long as there are two separate trim circuits to take that into account. However, to save space, it makes sense to use a portion of the larger device for the lower loads and then to turn on the entire device for high loads. This is slightly more complex to drive than the case in which two separate devices are used, with only one turned on at a time. One significant difference appears in the gate charge of each device; the smaller device turns on first unless a delay is implemented so both turn on at the same time.

5.4 EMI Considerations

Since the main application for this circuit is automotive lighting, electromagnetic interference (EMI) is of great importance. The circuit must stay outside of the radio bands or else it may interfere with the radio reception. This is the main reason for the high switching frequency of 2MHz to stay out of the AM band. The use of a slope compensation ramp is to prevent any sub-harmonic oscillation, which might interfere with the radio bands. Furthermore, operation must strictly stay within continuous mode, with no frequency modulation.

Since radiated EMI generally is an issue that can be solved externally with grounding planes and metal enclosures, it is not discussed here.

Chapter 6

Conclusion and Future Work

6.1 Conclusion

In the end, all three methods are determined to be viable methods for this specific application of on-chip current sensing and regulating, given ideal devices and processes. Because the intended purpose is just a current source with no voltage feedback, the system is for the most part stable, with only two dominant poles. This leaves many degrees of freedom in which to architect the actual circuit implementations. However, taking the non-idealities into account, there are some methods that are easier to implement than others in reality.

The single transconductance amplifier topology is the preferred method because of its simplicity and accuracy. It is simple because it uses a transconductance amplifier that does not have to be matched with anything. Given an absolute accuracy requirement, it only needs to be paired up with a matched and trimmed senseFET to work properly and accurately. The peak detection circuit, implemented separately, does not have to be extremely accurate because the average will be taken into account already. An added benefit is that the compensation capacitor can be located on-chip, since the circuit blocks can all be implemented with enough bandwidth that the compensating pole can be pushed out to the theoretical maximum; this means another pin can be saved per channel.

In contrast, the opposing transconductance amplifier topology requires two well-

Table 6.1: Summary of Schemes

	Sense Resistor	Current Summing	Transconductance
Extremely High Accuracy (<1% error)	Y	N	N
High Accuracy (<5% error)	Y	Y	Y
Quickly Changing Output Voltage	N	Y	Y
Compatible with Wide V_{IN} Range	Y	N	Y

matched transconductance amplifiers in addition to the matched senseFETs. Making it even more difficult is the fact that in normal operation, they must be sourcing and sinking the exact same current, respectively. That means they must be highly linear with an extremely low offset around ground. While entirely possible, this requires much more work to implement accurately than the single transconductance amplifier method.

Finally, the current summing method presents what seems to be a simple option as well. Because it is generating the exact sensed currents from the top and bottom gates, all that stands between the peak and average currents is a low-pass filter. Unfortunately, the peak sensed current is not so useful for the purpose of the peak current comparator because of the need for a holding capacitor. Thus, it is necessary to also implement the same peak current comparator as the other methods. Although it is simple in theory, this method suffers from bandwidth issues, as the amplifiers must slew at each turn-on transient, leading to high levels of inaccuracy at low output currents. Because of this instability, the compensating capacitor must be larger to lower the unity gain frequency before the extraneous poles and zeros appear. There is no real way to mitigate this issue because any attempts to increase accuracy in either the averaged or peak measured current will lead to a reduction in accuracy in the other.

6.2 Future Work

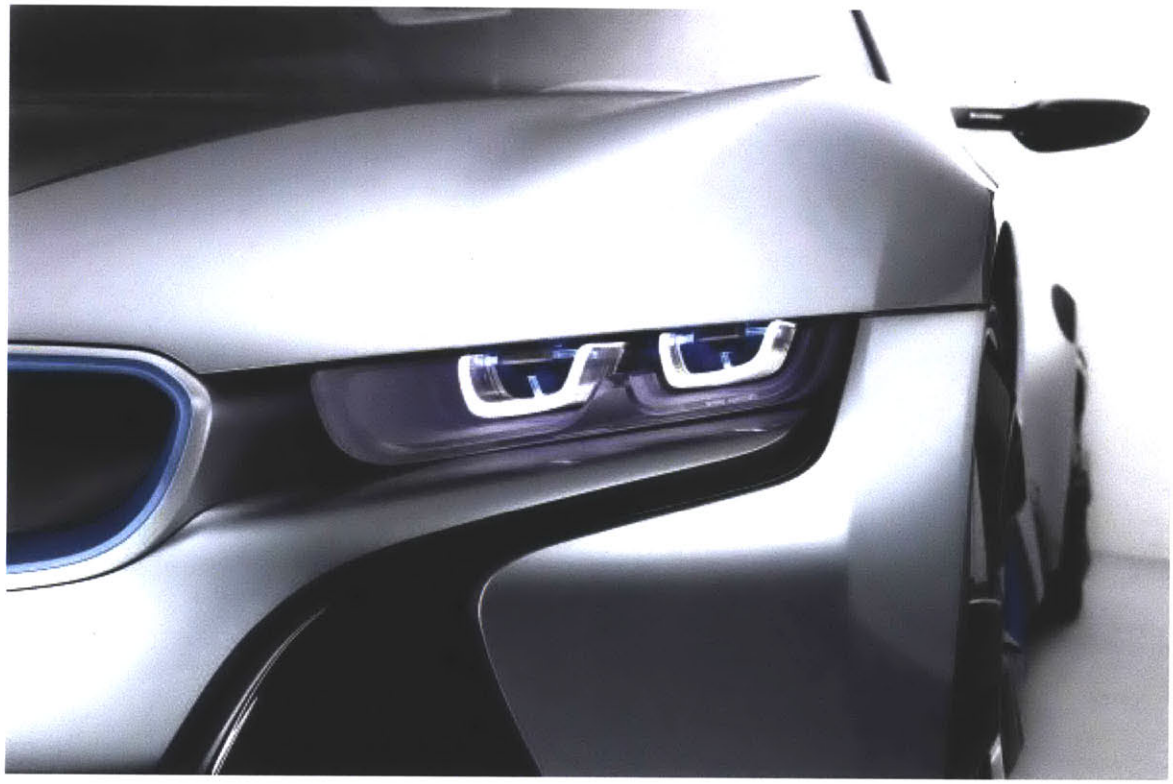

Figure 6-1: BMW concept of laser headlighting[17].

Even though LED technology has not yet fully matured, car manufacturers are already looking ahead to the next generation (Figure 6-1). The benefit of lasers over LEDs is even lower power consumption coupled with a significantly smaller physical footprint. Since LEDs and lasers operate on similar principles, they should both be operable using the methods proposed in this paper. In fact, since the methods proposed are all for accurate constant-current sources, they are extremely well-suited for powering laser diodes.

Bibliography

[1] Hella KGaA Hueck & Co. Headlamps, September 2012. http://www.hella.com/hella-com/620.html.

[2] Linear Technology Corporation. Lt3595 datasheet, July 2012. http://cds.linear.com/docs/en/datasheet/3595f.pdf.

[3] Z. Li. Design and implementation of a truecolor wide dimming single-pin LED driver. M.Eng thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 2010.

[4] J. Kassakian, M. Schlecht, and G. Verghese. *Principles of Power Electronics.* Addison-Wesley, Reading, Massachusetts, second edition, 1992.

[5] K. Cheng, C. Su, and H. Ko. A high-accuracy and high-efficiency on-chip current sensing for current-mode control CMOS DC-DC buck converter. *International Conference on Electronics, Circuits, and Systems (ICECS)*, pages 458–461, August 2008.

[6] R. Pilawa-Podgurski et al. Low-power maximum power point tracker with digital control for thermophotovoltaic generators. *Applied Power Electronics Conference and Exposition (APEC)*, pages 961–967, February 2010.

[7] Y. Lee and C. Hsu. High accuracy CMOS current sensing circuit for current mode control buck converter. *Power Electronics and Drive Systems (PEDS)*, pages 44–48, November 2010.

[8] J. Chen, F. Yang, and C. Chen. A new monolithic fast-response buck converter using spike-reduction current-sensing circuits. *IEEE Transactions on Inductrial Electronics*, 55(3):1101–1111, March 2008.

[9] Unitrode. *Modelling, Analysis and compensation of the current-mode converter.* Texas Instruments Incorporated, 1999.

[10] F. Tan and R. Middlebrook. A unified model for current-programmed converters. *IEEE Transactions on Power Electronics*, 10(7):397–408, July 1995.

[11] S. Rao et al. A 1.2-a buck-boost led driver with on-chip error averaged sensefet-based current sensing technique. *IEEE Journal of Solid-State Circuits*, 46(12).

[12] John Gardner. *SenseFET Matching.* Private communications, August 2012.

[13] G. Wegmann, E. Vittoz, and F. Rahali. Charge injection in analog MOS switches. *IEEE Journal of Solid-State Circuits*, (6):1091–1097, December 1987.

[14] C. Enz and G. Temes. Circuit techniques for reducing the effects of op-amp imperfections: autozeroing, correlated double sampling, and chopper stabilization. *Proceedings of the IEEE*, 84(11):1584–1614, November 1996.

[15] R. Williams et al. The dual-gate W-switched power MOSFET: A new concept for improving light load efficiency in DC/DC converters. *IEEE International Symposium on Power Semiconductor Devices and IC's*, pages 193–196, May 1997.

[16] J. Klein. *Synchronous buck MOSFET loss calculations with Excel model.* Fairchild Semiconductor, 2006.

[17] BMW Group. BMW develops laser light for the car [Press Release], January 2011.