FAST, ACCURATE POWER MEASUREMENT AND OPTIMIZATION FOR
MICROPROCESSOR PLATFORMS

BY

MATTHEW ROBERT JOHNSON

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2015

Urbana, Illinois

Doctoral Committee:

Professor Sanjay J. Patel, Chair
Associate Professor Deming Chen
Adjunct Assistant Professor Matthew I. Frank
Associate Professor Steven S. Lumetta

ABSTRACT

Power and energy consumption have become important for all computers, but the tools used to measure and optimize power on physical hardware lag far behind performance focused tools. Existing measurement apparata have low analog bandwidth, do not explicitly correlate power data with processor activity, and are not explained in sufficient detail to quantify uncertainty in their data. We present the design, implementation, and application of Jouler's Loupe, a measurement device that overcomes these obstacles and enables a new generation of fast, fundamentally sound energy-efficiency-focused tools. We demonstrate substantial opportunity for energy-focused software optimizations on a mobile CPU core.

To my parents, for their love and support.

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

ADC	Analog-to-Digital Converter
CMRR	Common Mode Rejection Ratio
DUT	Device Under Test
ESL	Equivalent Series Inductance
ESR	Equivalent Series Resistance
FET	Field Effect Transistor
FIR	Finite Impulse Response
FPGA	Field-Programmable Gate Array
GBP	Gain-Bandwidth Product
IIR	Infinite Impulse Response
LC	Inductor-Capacitor
LSB	Least Significant Bit
LUT	Lookup Table
PCB	Printed Circuit Board
PCIe	Peripheral Component Interconnect Express
PDN	Power Delivery Network
PLL	Phase-Locked Loop
PMIC	Power Management Integrated Circuit
ppm	Parts Per Million
PSRR	Power Supply Rejection Ratio

RC	Resistor-Capacitor
RF	Radio Frequency
RL	Resistor-Inductor
RLC	Resistor-Inductor-Capacitor
RSS	Root Sum Square
RTI	Referred-To-Input
RTO	Referred-To-Output
SPICE	Simulation Program with Integrated Circuit Emphasis
TCR	Temperature Coefficient of Resistance
VRM	Voltage Regulator Module

CHAPTER 1

INTRODUCTION

Power and energy consumption have become extremely important for designers, software developers, and end users of nearly all modern computers as a proxy for mobile device battery life, server electricity cost, or peak performance within a fixed power or thermal envelope. Despite this importance, the tools available to measure and optimize the power consumption of real hardware are primitive relative to their counterparts in the performance domain such as cycle-accurate timers, profilers, profile-guided software optimization, and auto-tuners.

Existing power measurement techniques lack the following three desirable attributes:

- An analog-to-digital signal path with low noise, high analog bandwidth, and high sampling rate for voltage and current signals;

- A precise temporal correspondence between power measurements and device-under-test (DUT) activity;

- A rigorous derivation of power measurement uncertainty and error sources that can be used to construct robust measurement and optimization algorithms.

As a result, these techniques can only reliably measure power at a coarse granularity. Without uncertainty analysis, the power consumption of two pieces of hardware or software cannot be compared in a rigorous way. In this thesis, we present the design, implementation, and characterization of the Jouler's Loupe, or Loupe, a measurement apparatus that overcomes these challenges; it serves as an ideal building block for the next generation of much faster, more ubiquitous power measurement and optimization tools for a wide range of processor systems.

The contributions of this work include:

- An enumeration and analysis of the key design parameters and concerns for fast, accurate power measurement;

- The design and implementation of the Jouler's Loupe, a power measurement device that can reliably sample $10\times$-$100\times$ faster than state of the art methods;

- A case study showing how Loupe enables developers to build practical tools to measure and optimize the energy efficiency of software.

CHAPTER 2

RELATED WORK

2.1   Power Measurement Methodology

Measuring the power of an RF or microwave signal is a common problem in communications and electromagnetic compatibility, and a wealth of commercial equipment is available for this purpose. RF power sensors must handle signals up to tens of GHz and do not digitize their input signals directly, but rather integrate them in the analog domain using a thermistor, thermocouple, or diode, then send the integrated signal to an RF power analyzer to be digitized [1]. RF power sensors' frequency responses do not extend down to DC and power analyzers, despite their high sampling frequencies up to 100MHz, only provide postprocessed average or peak power readings at up to $\approx 1.5\text{kHz}$, owing to the low analog bandwidth of the underlying sensing mechanisms. For instance, since thermistor- and thermocouple-based sensors integrate the RF signal by dissipating it across a resistor and measuring its temperature, typical 10–90% rise times measure in seconds. Therefore, RF power measurement techniques are not directly applicable to high-speed processor power measurements, which require very high analog bandwidth extending all the way down to DC.

The four major current sensing modalities are the current sense (or *shunt*) resistor, the current transformer, the Hall effect sensor, and the Rogowski coil [2]. The latter three modalities can be implemented non-intrusively by wrapping a loop or coil around the conductor being measured. This feature allows for the measurement of very high currents up to thousands of amperes, provides galvanic isolation between DUT and measurement system, and avoids the self-heating and power dissipation inherent in the shunt resistor method. The downside of non-intrusive methods is that they are more suited to measuring power through a wire or cable and are inconvenient for

PCB-level measurements, since commercial designs rarely allow for a loop to be wrapped around the power trace or plane leading to a processor. Furthermore, the three coil-based modalities typically have bandwidth of a few MHz or less and variously have issues with linearity, hysteresis, temperature dependence, and measuring DC current. Despite the relatively high power consumption of sense resistors and the attendant self-heating and temperature dependence issues, sense resistors provide the best foundation for a high speed processor power measurement platform due to their simplicity; their excellent frequency response and linearity; their ability to measure DC current; their suitability to measuring currents both on a wire or cable and on a PCB; and the ability to easily use a common power measurement platform to measure different current ranges by using different resistances in a common form factor.

Commercial current probes for oscilloscopes have many desirable characteristics, but are ultimately unsuitable for our purposes except as a useful cross-validation tool. The highest-performance probes, such as the Agilent N2783B, use both a Hall effect sensor and a current transformer to measure from DC to a -3dB bandwidth of 100MHz. However, these probes have the disadvantages of non-intrusive methods discussed previously, and higher-frequency probes are usually limited to measuring smaller conductors. For example, the N2783B can only measure conductors up to 5mm in diameter. For PCB-level power rail measurements, the parasitic impedance of a 5mm wire long enough to fit through the probe would cause substantial ringing problems, as discussed in Section 4.1.1. In addition to the added impedance of the wire itself, even non-contact current probes also incur an insertion impedance. While this impedance is on the order of $1\text{m}\Omega$ at low frequencies, it quickly grows to tens or hundreds of $\text{m}\Omega$ at 10MHz, far higher than the $33\text{m}\Omega$ impedance of a $10\text{m}\Omega$ sense resistor with 0.5nH parasitic inductance. While off-the-shelf current probes are specified with fairly flat gain responses out to their specified bandwidth, no concrete flatness guarantees beyond the -3dB bandwidth are provided and no information is provided as to phase linearity/group delay flatness; a sensor with flat gain and nonlinear phase may produce significant distortions in the time-domain waveform, causing error in power and energy estimates over short time periods. Current probes require periodic degaussing and offset voltage nulling due to residual magnetization of the magnetic core, and are thus more suited for controlled laboratory

4

measurements than continuous, automated in-system use. Current probes also have limited sensitivity in the μA–mA range; for example, the minimum current the N2783B can reliably measure is 5mA. One technique to improve sensitivity is to wind multiple turns of the wire through the probe opening, multiplying the probe's output voltage by the number of turns. This technique has the downside of further increasing inductance along the wire by coiling it and potentially needing to increase its length, and is limited by the conductor size limit of the probe. Further limitations of commercial current probes for our purpose include their large physical size which precludes their use in space-constrained production systems, their current cost of thousands of dollars, and the lack of insight into their error sources and behavior due to their proprietary design. We opted to design our own sense resistor-based power sensor that addresses all of the above issues.

A large body of literature focuses on the accurate measurement of AC mains power, with applications ranging from individual electrical appliances to large industrial equipment, entire buildings, and large-scale electrical grid design. The signals of interest in this domain have low fundamental frequencies from 50Hz up to perhaps 1kHz. One important feature of mains power measurement devices is the ability to synchronize the sample clock to the fundamental frequency for fast convergence of peak and RMS power estimates on periodic signals. Processor DC power rails have no such periodicity, and as such the speed of power estimate convergence is purely dependent on analog bandwidth, noise, and sample rate. Svensson designed and characterized an accurate digital watt meter for AC loads which may have significant harmonic content up to several kHz [3]. Like Svensson, we use a shunt resistance and an ADC in our power measurement apparatus to digitize current and voltage waveforms, and we perform an uncertainty analysis to enable the informed use of the apparatus in statistically rigorous measurement and optimization applications. However, our apparatus has several orders of magnitude more bandwidth and is suited for PCB-level or system-level measurement, and our error analysis is more focused on thorough characterization of the device than on the uncertainty of particular derived quantities from the mains power domain like apparent, active, and reactive power.

Chang et al. [4] introduced a switched-capacitor-based methodology for measuring the dynamic energy of simple ARM microcontrollers at a cycle granularity. The methodology was later applied to FPGAs by Lee et al. [5].

Under this methodology, two or more capacitors are connected across the processor's supply pins, and the measurement system controls digital switches that connect one of the capacitors to the processor at a time. The processor then draws current from that capacitor for a short period of time, rather than from the power supply, while the power supply recharges the other capacitors which are disconnected from the processor. While the capacitor is powering the processor, its voltage drops from V_{DD} to V_2, and the total energy drawn during the time period can be calculated as:

$$E = \int_{t_1}^{t_2} I(t)V(t)\mathrm{d}t = \frac{1}{2}C\left(V_{DD}^2 - V_2^2\right) \tag{2.1}$$

The switched capacitor method has an important advantage over sense resistor-based approaches in that it uses the capacitor as a high speed *analog* integrator, seemingly sidestepping the need for high analog bandwidth to the ADC, high ADC sampling rate, and the computational expense of integrating measured current and voltage data in the digital domain. However, three key downsides to this approach motivate a high speed sense resistor-based system like Loupe. First, physical capacitors deviate in many ways from their ideal specification, including leakage, parasitic inductance, and capacitance dependent on applied voltage and temperature. These deviations limit the utility of the naïve energy computation in Equation 2.1, particularly when absolute accuracy is important to compare results across processors or measurement methods. Sense resistors also have nonidealities such as parasitic inductance and thermoelectric voltage, but they are characterized much more thoroughly than those of capacitors, and the designer can compensate for them more readily. Second, in order to get cycle-level energy measurements, Chang et al. and Lee et al. provide an external clock to the system under test so that each capacitor can power the processor — and the voltage drop can be measured — for a single cycle. As a practical concern, nearly all processors of interest today derive their clock from an on-chip PLL, not external pins. Most commercial products do not allow for single-stepping the processor from an external clock at all, so the automatic cycle-level synchronization of power measurements and processor activity assumed in the literature could not be maintained. Even if single-stepping were possible, leakage has come to represent a far larger fraction of total power consumption in modern processors than in previous generations, so accurate total power

measurements require running the processor at realistic speeds to generate realistic thermal conditions and to capture a realistic balance between static and dynamic power. The third limitation of the switched capacitor approach is that transistors and other on-die structures are highly non-linear, and the switching current and total energy is dependent on supply voltage over time. On one hand, it is desirable to size the capacitor and measurement period to generate a large voltage drop that can be measured precisely. On the other, it is desirable to keep the supply voltage very close to nominal V_{DD} to avoid altering the operating point of the DUT and causing the measured energy consumption to differ from that under realistic operating conditions. This fundamental tension is analogous to the burden voltage for a sense resistor, but the solution in this case — carefully designed amplifiers to allow small voltage drops to be measured precisely — eliminates much of the simplicity advantage of the switched capacitor approach, and does not address the other two downsides mentioned earlier. While switched capacitors provide an interesting way to cross-validate sense resistor-based energy measurements, the internal clocking of modern processors and the need to run the DUT at its normal operating point for best results preclude the cycle-level measurement granularity achieved in the literature and limit any accuracy benefit of the former technique.

Even if the bandwidth, accuracy, and simplicity characteristics of resistance-based current sensing are desirable, the power dissipated across the sense resistor may be untenable in some low-power production systems. Previous studies have evaluated *lossless* current sensing, where instead of a sense resistor, the voltage drop associated with current is measured across parasitic resistances present in the unmodified power supply. The resistances used are the DC resistance (DCR) of the switching FETs [6] or inductors [6, 7] in a single- or multi-phase buck regulator, or even the copper trace or plane between the regulator and the processor [8]. In the case of the inductor, the DC component of the current is measured by using a single-pole RC filter to cancel out the reactance of the inductor. While the resulting measurements are sufficiently accurate to detect and mitigate overcurrent conditions or share load current between multiple phases[9], lossless sensing approaches have three downsides for fast, accurate processor power measurement. First, the notion of a single resistance value belies the complexity of physical power inductors; manufacturer simulation models do include such a

resistance, but also have a parallel combination of frequency-dependent R, RC, and RL circuits [10] that makes the real frequency response considerably more complex and limits the accuracy of RC-filtered current data above DC. Second, current sense resistors are designed with great attention detail in areas like thermoelectric voltage and stability over temperature, lifetime, and environmental conditions that are not important for power inductors, and thus provide superior stability and characterizability. Finally, since the inductor is placed physically and electrically "far" from the processor, it will see less of the high frequency processor current content than a sense resistor near the PCB-mounted decoupling capacitors, limiting the effective time granularity of energy and power measurements. MOSFET- or copper-based lossless sensing suffer from largely the same drawbacks. In the case of copper-based sensing, placement of the low-side voltage endpoint is an additional issue, since this placement fully determines the measured current, but the voltage will vary substantially, both spatially across the power plane or planes and in time across different use cases as the current draw distribution among the process balls changes. Nevertheless, characterizing the temporal resolution and accuracy upper bounds of DCR sensing systems is an interesting area for future study, since lossless sensing may make the measurement and optimization techniques developed in this dissertation applicable to an even broader range of systems and use cases.

2.2 Power Measurement Applications

High-side shunt resistors both generate a small voltage in response to the current passing through them. This microvolt- or millivolt-level signal must be amplified significantly to occupy a significant fraction of the 1–4 V input range of most high-speed ADCs. Amplification is also necessary to reduce the effect of noise on the signal between the current sensor and the ADC. As shown in Table 2.1, most existing power measurement setups do not amplify the current signal at all; those that do amplify use specialized current sense amplifiers or amplifiers integrated into the Hall effect sensor, both of which have a typical 3dB bandwidth of 10–100kHz. Applications with absolute accuracy requirements like power measurement require a flat amplitude response and linear phase response in the frequency range of interest; the

Table 2.1: A small sample of existing power measurement techniques. Most papers do not describe the measurement setup in significant detail. Most techniques do not explicitly amplify the current signal, leading to low dynamic range at the ADC; those that do amplify it have relatively low bandwidth. The commercial AC power measurement devices used by Do et al. and others are based on a sense resistor [11] or current transformer [12]. Jouler's Loupe uses a sense resistor, has a 3dB bandwidth of 60–200MHz, and uses 64Msps 12-bit ADCs.

Citation	Method	Amplifier BW (3dB)	ADC Specs	Correlated w/ Proc. Activity?
Bedard et al.[13]	Sense Resistor	N/A	6.6kHz, 12b	No
Weissel and Bellosa[14]	Sense Resistor	N/A	$\leq$3Msps	No
Rajamani et al.[15]	Sense Resistor	$\leq$10kHz	333ksps, 16b	No
Zhu et al.[16]	Sense Resistor	$\leq$80kHz	200ksps, 16b	No
Carroll and Heiser[17]	Sense Resistor	N/A	250ksps, 16b	No
Contreras and Martonosi[18]	Sense Resistor	N/A	100sps	Yes
Govindan et al.[19]	Hall effect	100kHz	10ksps, 14b	No
Esmaeilzadeh et al.[20]	Hall effect	80kHz	50sps, $\approx$7b	No
Do et al.[21]	AC Power	N/A	1Hz readings	No

0.1dB bandwidth is more applicable and is about 1/6th the 3dB bandwidth for voltage-feedback amplifiers, or about 14–17kHz for the amplifiers in Table 2.1 [22]. This low bandwidth makes it impossible to accurately measure code executions shorter than 1–5ms (millions of cycles on modern processors), and increases the required run time for a given level of accuracy and confidence [23].

Of the methods in Table 2.1, only Contreras and Martonosi [18] can derive a relationship between a power sample and processor activity, but they do not exploit this capability for fine-grained measurement. As observed by Jung et al. [24], misalignment or uncertain alignment in time between when the code of interest is executed and power sampling instants leads to error in the estimate of that code's energy consumption. Minimizing this error by increasing run length further decreases the effective throughput of the measurement system. These two limitations of existing measurement schemes make them ineffective for executions shorter than about a million cycles, stymying the development of energy-aware compilers, profilers, and auto-tuners. In Chapter 4, we will show how Loupe improves bandwidth by several orders of magnitude and extends effective power measurements down to the microsecond timescale.

Also related to this work are relatively new commercial devices for measuring power consumption of smartphones [25] or embedded CPUs [26]. While

the latter device does collect processor activity data along with power, the two data streams are not explicitly correlated using a common timebase and the processor activity data is not immediately useful to automated software tools. Furthermore, these devices suffer from the same bandwidth and dynamic range limitations as the methods in the academic literature.

CHAPTER 3

MOTIVATION

A primary challenge in high speed current measurement for processors is the decoupling capacitance and parasitic impedance between transistors drawing current on the processor die and a current measurement point on the PCB. Power delivery networks (PDNs) are designed to supply current to on-die transistors at nearly constant supply voltage, regardless of time-varying load characteristics. The parasitic inductances of the on-die supply grid, package substrate, and PCB develop a differential voltage in response to changing load current according to the equation $V = L\frac{dI}{dt}$. Thus, the primary PDN design goal is to minimize impedance $|Z| = \frac{|V|}{|I|}$ over a broad frequency range, where I is the current drawn by the processor and V is the voltage drop caused by the parasitic inductance and resistance. A PCB-mounted voltage regulator by itself can only maintain sufficiently low impedance up to the kHz to low MHz range, while processor current draw has significant spectral content up to several GHz. To bridge this gap, a hierarchy of decoupling capacitors are placed on the PCB, package, and processor die. Each successive level of capacitors has smaller capacitance, but a lower-inductance path to the load, and supplies current at low impedance for successively higher frequency ranges. Thus, the charge for high-frequency current transients is supplied by on-die decoupling capacitance, mid-frequency transients are handled by package-mounted capacitors, and lower-frequency content is handled by PCB-mounted capacitors and the regulator. The highest-frequency information in the current signal is significantly attenuated at the PCB level, and the observed signal is essentially a time-averaged version of the true current draw on the chip.

We quantify the potential benefit of higher power sampling frequencies using a bottom-up approach. We construct a relatively detailed SPICE model of a complete processor PDN, from the VRM through the on-die power distribution grid. We then drive the SPICE model using current consumption

data from a cycle-level simulator running representative workloads. We can then determine the resolution and accuracy bounds of PCB sense resistor measurement approaches by comparing PCB-level current and voltage waveforms to the true on-die load.

3.1   PDN Modeling

Our SPICE model of a PDN is similar at a high level to the lumped model used by Gupta et al. [27]. However, our model improves accuracy and allows the model to be more easily adapted to new designs by adding a 4-element lumped model of the VRM, using individual lumped RLC models for each decoupling capacitor and its parasitics, and deriving its parasitic component values from first principles when possible rather than by empirically modifying the values to approximate the shape of a measured PDN impedance curve. For the purposes of evaluating the benefit of high frequency power sampling, we model the PDN of an existing system. The Hardkernel ODROID-XU+E single board computer houses a Samsung Exynos 5410 SoC with 4 high-performance ARM A15 cores and 4 energy-efficient ARM A7 cores. Here, we model the voltage rail powering the A15 cores exclusively, which includes several values of board-level decoupling capacitors and a 10mΩ sense resistor used for a built-in low-frequency current sensor. The structure of the model and the particular component values used in our simulations can be found in Figure 3.1 and Table 3.1. We set V_{DD} at the VRM to compensate for the DC IR drop due to the sense resistor and PCB and provide $1.0V$ at the package. Figure 3.2 shows the impedance of the PDN for current stimuli from 10kHz to 10GHz, with and without the 10mΩ sense resistor. The sense resistor substantially increases impedance, and thus supply voltage noise, at low frequencies, providing another reason to prefer very small sense resistances.

3.1.1   Power Plane Impedance Modeling

Most high-current IC power rails are connected to the PCB-mounted decoupling capacitors and voltage regulator via large power and ground copper pours, or *planes*, on adjacent PCB layers. The two planes separated by a dielectric form a parallel plate capacitor, and this capacitor supplies charge
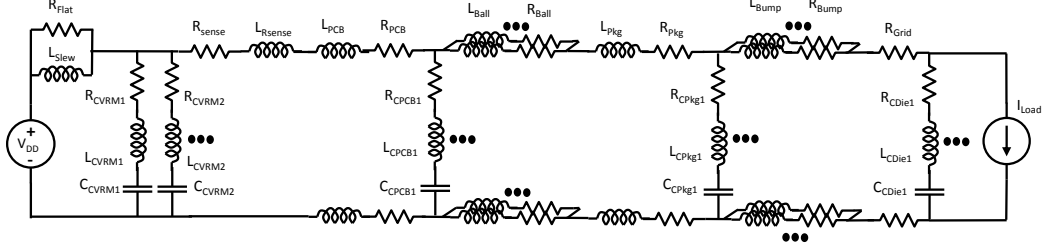
Figure 3.1: Conceptual PDN, modeling VRM responsiveness, decoupling capacitors, plane and package parasitics, and the on-die power grid.

Table 3.1: Component values used in SPICE model of ODROID-XU+E A15 power rail.

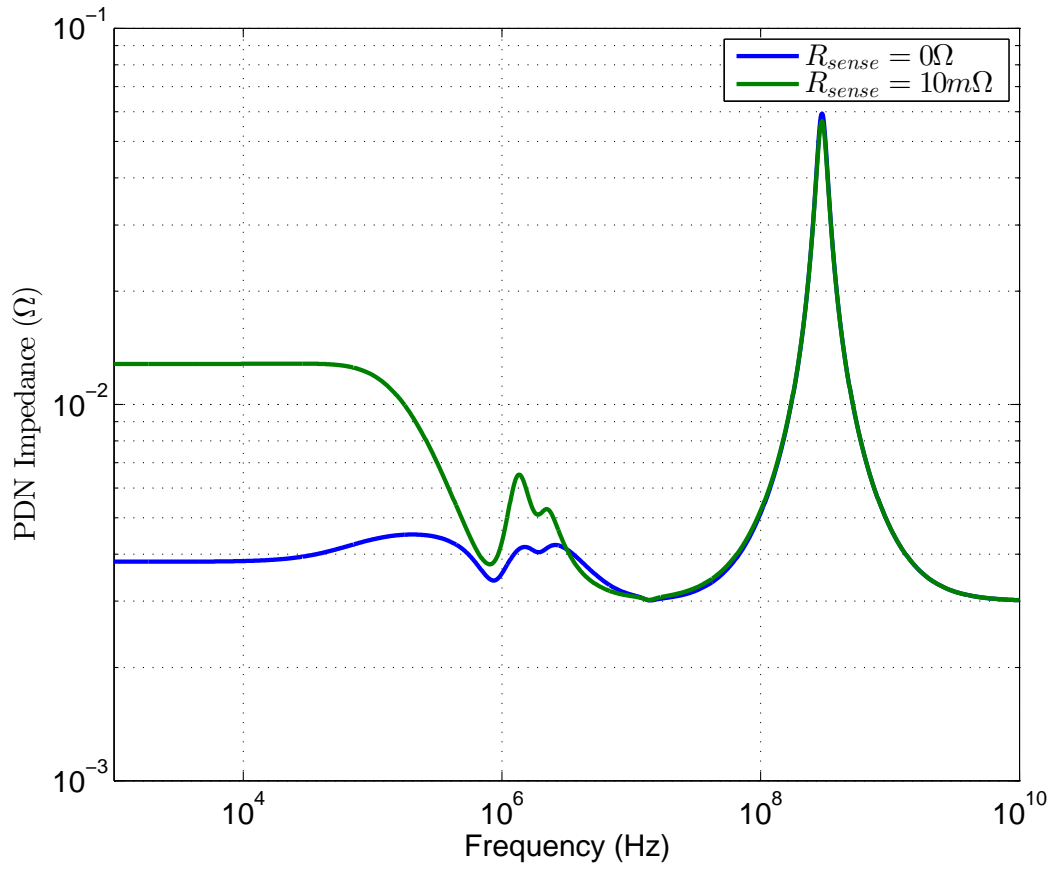| Component	Value
VRM	R_{flat} =0.75mΩ, $L_{slew} = 1.\bar{6}$nH
C_{VRM}	2×(22μF / 3mΩ / 1.4751nH)
R_{sense}	10mΩ
PCB	R_{pcb}=0.1mΩ, L_{pcb}=21pH
C_{PCB1}	2×(22μF / 3mΩ / 1.4751nH)
C_{PCB2}	1×(4.7μF / 7mΩ / 1.4256nH)
C_{PCB3}	12×(0.1μF / 26mΩ / 1.339nH)
Pkg. Balls	45×(424$\mu\Omega$ / 32pH)
C_{Pkg1}	26μF / 541.5$\mu\Omega$ / 5.61pH
Die Bumps	0.3mΩ / 0.5pH Lumped
On-die	C_{Die1}=50nF / 1mΩ, R_{grid}=1mΩ

Figure 3.2: Modeled impedance of the ODROID-XU+E A15 power rail with and without a 10mΩ current sense resistor.
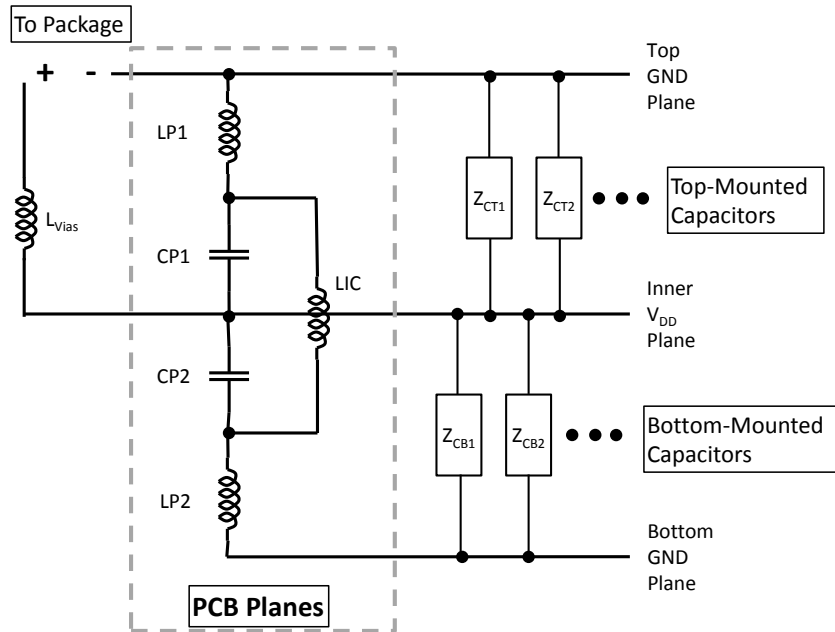
Figure 3.3: We model the PCB planes and capacitors using the topology proposed by Shringarpure et al. [28], where the plane is split into two LC elements and a mutual inductance, and top- and bottom-mounted decoupling capacitors are connected to separate GND nodes. Since the PMIC and processor are mounted on the top of the PCB, the top GND node in this figure corresponds to the GND node in Figure 3.1.

to the IC in the mid-frequency range between those of the PCB-mounted capacitors and the package-mounted or on-die capacitors. We model the power planes and PCB-mounted capacitors using the split topology presented by Shringarpure et al. [28], as shown in Figure 3.3. The planes are modeled as an LC circuit split into two halves, $LP1/CP1$ and $LP2/CP2$, connected to the top- and bottom-mounted decoupling capacitors, respectively. The mutual inductance between the two LC circuits is represented by a third inductor, LIC. To estimate the values of the three inductors and two capacitors, we use the methodology introduced by Jones [29]. First, the planes are modeled as solid rectangles, and their self-impedance is calculated using the Modal-Expansion Cavity Model for a rectangular patch given by Carver and Mink [30]. For simplicity, we assume that the top and bottom ground planes are equidistant from the V_{DD} plane so that $LP1 = LP2$ and $CP1 = CP2$, and that the processor package and die are approximately centered on the planes. We then iteratively calculate the LC circuit's impedance using SPICE and modify the component values until the LC circuit closely matches the first resonance and anti-resonance of the more detailed model.

We estimate that the A15 power planes for the ODROID-XU+E measure 1.5 by 1.0 inches and are separated by a 2.0-mil sheet of FR-4 ($\epsilon_r = 4.7$). The cavity model shows first resonance and anti-resonance at 834MHz and 3.63GHz, respectively. We find that component values of $LP1 = LP2 =$43pH, $CP1 = CP2 =$400pF, and $LIC =$9.6pH closely match the more detailed model up to 3.63GHz, an upper bound which is more than sufficient for evaluating voltage and current waveforms at the sense resistor at ≈100MHz and below.

3.1.2   Decoupling Capacitor Modeling

We model each decoupling capacitor as the series combination of the nominal capacitance value, equivalent series resistance (ESR), and two inductances: the equivalent series inductance and a mounting inductance L_{mount}. The mounting inductance is a result of the current loop formed by the capacitor, the power and ground planes, and the PCB tracks and vias that connect the capacitor to the planes. We calculate these individual component values using estimated PCB stackup, capacitor placement, and via geometry

information from PCB layout data and simplified physically-based models provided by Altera[31], rather than using an electromagnetic simulator.

3.1.3 Power Consumption Simulation Data

We use version 5.3 of the Sniper simulator [32], which can simulate single- or multi-threaded workloads running on parameterized x86 core models. We model a single core similar to Intel's Silvermont running at 1GHz. This relatively low-power, low-frequency processor will tend to have lower frequency content and lower dynamic range in its power consumption than faster or wider cores; thus, our results here provide a conservative estimate of the benefit of high power sampling frequency. We simulate representative 100M-instruction sections of the SPEC2006 benchmarks [33], derived using a SimPoint [34]-like methodology, for a total workload corpus of several billion instructions. We use McPAT 1.0[35] to estimate current consumption on a cycle-by-cycle basis using the simulator's block activity data. We drive I_{load} in the SPICE simulations using this cycle-granularity current trace, generating observed current traces at the die, package, board, and VRM levels along with the true load current.

3.1.4 Results

Each successive layer of decoupling capacitance from the die out to the VRM handles progressively lower-frequency current demand, so the observed sense resistor current waveform will certainly not reflect all the high-frequency content present in the load current, and thus the observed current can only accurately predict the load current when averaged over relatively long time periods. Typical averaging periods in the literature range from several milliseconds to several seconds, or even minutes in the case of full-system AC power measurement. We can quantify the accuracy of much higher sampling rates and averaging over shorter time intervals by analyzing the corpus of load and observed current for our modeled PDN and simulated power consumption data. Sample results for the `astar` benchmark are shown in Figures 3.4 and 3.5. Figure 3.5 compares, for a given sampling frequency, the average observed current over a sample period to the actual average load current, to
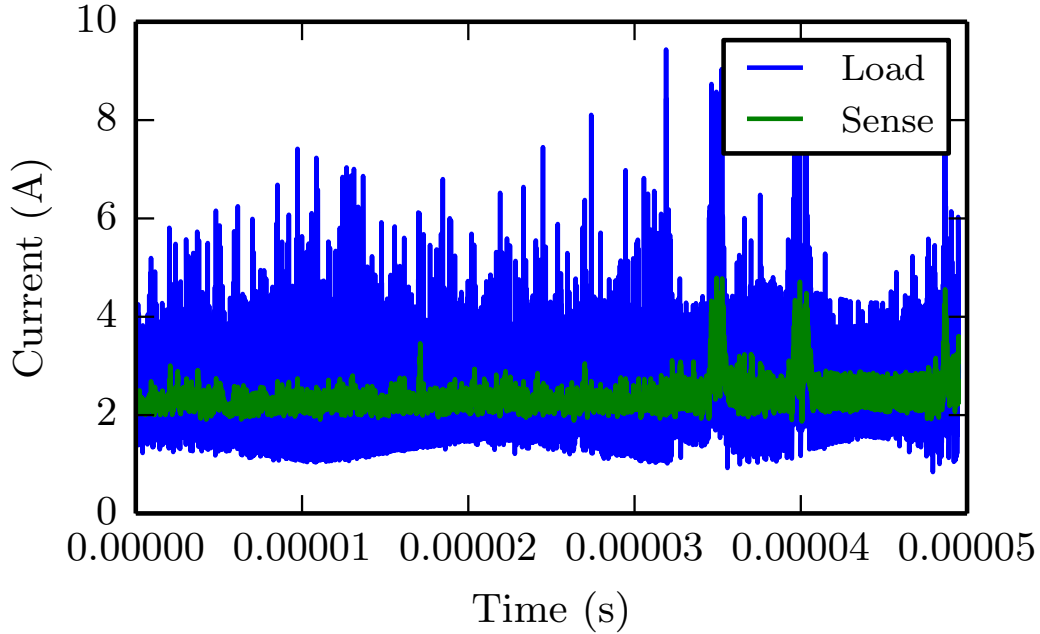
Figure 3.4: Sample load and sense resistor current trace from the `astar` benchmark. The decoupling capacitors on the die, package, and board handle high-frequency current transients, so the sensed current waveform does not capture the rapid variations in load current, but accurately captures the average over longer time periods.

evaluate how well the sense resistor waveform represents the true load. Note that to achieve the stated error percentiles, the measurement system must have analog bandwidth matching or exceeding the sampling rate. The data show that, while high-frequency load transients cannot be fully observed at the current sense resistor, sampling rates in the MHz range are justified and can provide insight into the current consumption of much shorter pieces of code than existing methods. The relatively substantial 99.9th and 99.99th percentile errors at high sampling rates suggest that power measurement and optimization algorithms at these rates must take these statistical properties into account to make correct decisions in the face of uncertainty.
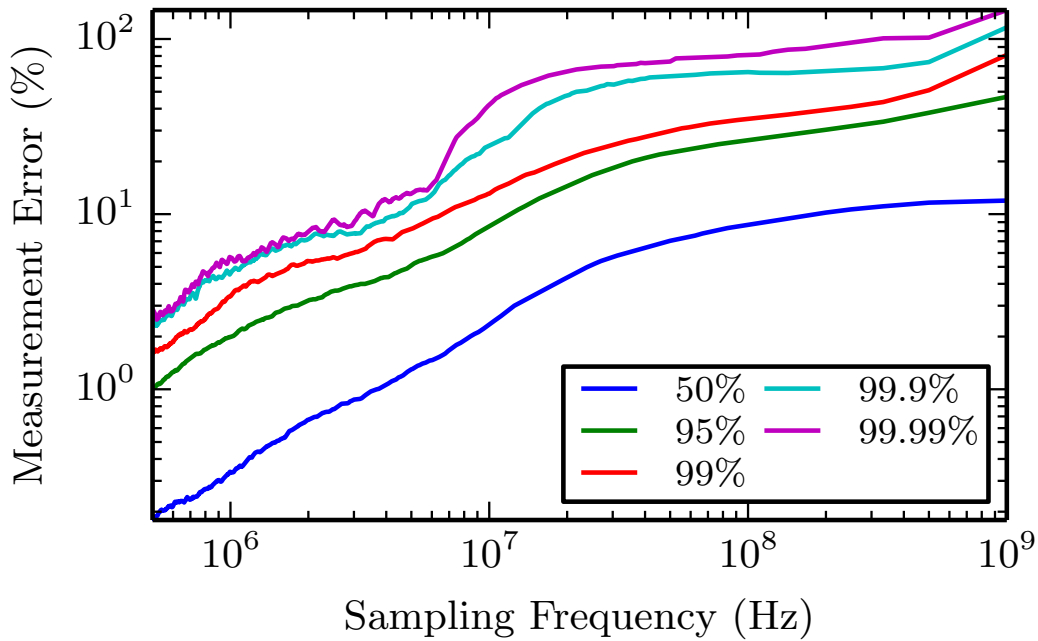
Figure 3.5: Percentiles of the error between observed and actual load current vs. sampling frequency/interval length for the `astar` benchmark. Low sampling rates (up to $\approx$ 100kHz, not shown) can capture average current consumption of long intervals with $\ll$1% accuracy. Median error ranges from 0.2% at 1MHz to 2% at 10MHz. 99th percentile error ranges from 3% at 1MHz to 10% at 10MHz.

CHAPTER 4

LOUPE DESIGN AND IMPLEMENTATION

4.1   Design Goals

To achieve maximum bandwidth, accuracy, and flexibility, Loupe is based on current sense resistors; high-bandwidth amplifiers shift and amplify the voltage and current signals so that their anticipated value ranges fully occupy a ±1V ADC input range. To minimize the effect on the DUT's operating voltage, the sense resistor should have low enough resistance that the worst-case voltage drop across it is very small compared to the rail voltage. The divide between the small current sense signal (0–10mV/0–120mV for 1V/12V power rails) and the large ADC input range leads to a high required gain, up to 200× for low-voltage rails feeding processors directly. An additional factor of 2 gain is required if wideband current and voltage signals are to be transmitted to the ADC via a terminated transmission line, since the termination resistors act as a 2:1 voltage divider. Achieving high bandwidth *and* high gain is one of the biggest challenges in designing Loupe, due to the fundamental tradeoff between the two in amplifier design; in fact, amplifiers are often specified in terms of a fixed gain-bandwidth product. We achieve multi-MHz current amplifier bandwidth and gains of 36–134 in our two proof of concept implementations using two cascaded stages of very high-speed amplifiers; the first is an AD8129 differential amplifier in all cases with a fixed gain of 10, and the second stage is an OPA843 or OPA847 depending on the required gain. For voltage amplifiers, we used a single stage of LM7171A or LMH6609 depending on required gain.

As mentioned in Section 2, a flat gain response and linear phase response (flat group delay) for the frequencies of interest is necessary to recover a clean time-domain amplified signal without complex and error-prone inverse filtering. Our target for Loupe is flat group delay and 0.1dB gain response

flatness from DC to at least 10MHz for all channels. We verified using SPICE simulations that all amplifier channels had a 0.1dB bandwidth of 10–50MHz; channel group delays ranged from 1.5ns and 4ns, and varied less than 90ps over the 0–10MHz band. These small, flat group delays can easily be calibrated out in the digital domain.

The sense resistor itself must be selected carefully for high-frequency measurements, paying special attention to parasitic inductance and temperature coefficient (tempco) as discussed in Section 4.1.1. The amplifier power supplies should be designed with the amplifier power supply rejection ratio (PSRR) in mind so that voltage ripple on the amplifier power supply has a minimal impact on the output signal; our limit for Loupe was less than $\frac{1}{4}$ LSB for a 12-bit ADC ($< \frac{1}{4} * \frac{4V}{2^{12}} = 81.4\mu V$ on the amplifier output). Likewise, variations in DUT rail voltage due to supply noise or DVFS will manifest as a common-mode input signal at the amplifiers, but the amplifiers should have sufficiently high CMRR to avoid large erroneous swings in the output current waveform. For Loupe, we specify that the maximum variation in the DUT rail voltage should cause less than $\frac{1}{4}$ 12-bit LSB variation at the output due to amplifier CMRR. Finally, besides having sufficient gain-bandwidth product, the amplifiers' slew rate should be sufficient to drive full scale output sinusoids throughout the frequency band of interest with the load capacitance presented by the Loupe PCB, coaxial cable, and ADC-side buffer.

For Loupe, we wanted to decouple the ADCs from the amplifier and have them on separate boards connected by coaxial cable, to enable upgrading the ADC independently and using an off-the-shelf ADC board to avoid the time-consuming task of designing, assembling, validating, and calibrating a high-speed ADC PCB. The ADCs should have at least 12 bits of precision, should sample at upwards of twice the analog bandwidth of the power sensor ($\geq$20MHz), and should be able to stream data back to a wide range of DUTs or analysis systems over a high-bandwidth standard interface. The ADCs should also ideally be on the same clock domain as the hardware that receives processor activity data, so that the power and activity data can be precisely aligned using a common timebase. Our design goal for the coaxial cable was to have less than 0.1dB maximum insertion loss from 0–10MHz; if using RG174 cable, this yields a maximum cable length of 3 feet, which is sufficient in most cases to run a cable from the possibly space-constrained DUT enclosure to the external ADC board.

4.1.1 Current Sense Resistor

The current sense resistor, while a conceptually simple component, plays a key role in the ultimate accuracy of current measurements, especially at high speed, and must be chosen carefully. Nominal resistance varies part-to-part, but this variation can be calibrated out. Other important properties of sense resistors cannot easily be calibrated out, including parasitic inductance, temperature coefficient, and thermoelectric voltage. For long-lived power measurement systems, long term resistance stability over thermal and mechanical shock, vibration, and humidity are also important, but we do not consider them in this work due to our relatively controlled measurement environment and inability to meaningfully test these parameters.

All real components have some amount of parasitic inductance and capacitance. We consider only surface mount resistors, as the high lead inductance of through-hole inductors renders them unsuitable. Typical thick- or thin-film surface mount resistors have parasitic capacitance on the order of tens of femtofarads in parallel with the series RL circuit; we ignore this capacitance in this work, as it only has a significant impact for RF or microwave circuits in the GHz range. The parasitic inductance can range from picohenries to tens of nanohenries for the large resistors used in current sensing, and resists changes in current by producing a differential voltage across the resistor according to the differential equation $V = L\frac{dI}{dt}$. In a current sensing application, this differential voltage is added to the "real" voltage caused by IR drop, and amounts to error in the measured current. The inductance adds reactance proportional to frequency producing an overall impedance $Z = R + j2\pi f L$, artificially increasing the measured amplitude of high-frequency current changes and rendering the amplifiers' flat gain/linear phase response useless. The resistor's resistance R, its inductance L, and any capacitance C to which the resistor is connected form an RLC circuit, also called a damped resonator. At the circuit's resonant frequency, calculated as $\frac{1}{2\pi\sqrt{LC}}$, energy transfers repeatedly between the inductor and capacitor, and is slowly dissipated (damped) as it travels through the resistor. The oscillating voltage across the resulting RLC circuit after a rapid current change is known as *ringing*, and is depicted in Figure 4.1. This ringing has two negative effects. First, while the error in the sensed current due to ringing will likely average out over time, this is not necessarily true in the short
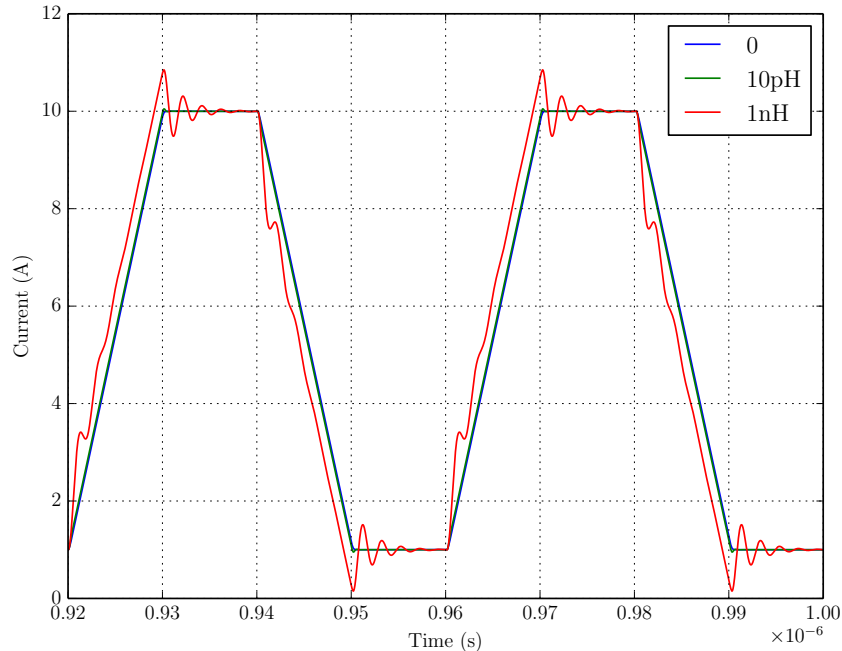
Figure 4.1: Measured current using a 10mΩ current sense resistor with various parasitic inductances. The true current draw switches between 1A and 10A with 10ns rise and fall times. Inductances up to several pH do not contribute significant error, but inductances of 1nH and above produce significant ringing. The maximum allowable parasitic inductance value varies with resistor value, error budget, and load current frequency content.

term, so power and energy estimates for short time intervals will be greatly affected. Second, ringing produces large peaks in the current sense voltage during load transients. These peaks can be handled either by reducing the gain in the power sensor so that the peak still fits in the input range of the ADC, or by keeping higher gain and saturating the ADC during the peaks. The first approach adds quantization error by reducing ADC resolution for all current data, and the second adds clipping error. For our purposes a parasitic inductance under 1nH gives sufficiently low error. If additional correction is necessary due to higher frequency content, more stringent error requirements, or a highly inductive resistor, a single-pole RC filter can be used between the sense resistor and the first-stage amplifier to cancel out the parasitic inductance with capacitance [36].

A resistor's value changes with temperature. While the true temperature-resistance curve is usually nonlinear and varies based on material composition, the effect of temperature is usually approximated using a linear coefficient called the temperature coefficient of resistance, also known as TCR or

23

tempco [37]. The TCR is a conservative measure of the fractional change in resistance value (relative to nominal value) for every degree change in temperature and is usually expressed in $\frac{ppm}{^\circ C}$. Given an estimate of how much the resistor's temperature can be expected to vary over a period of time, one can use the TCR to compute an upper bound on the error in measured current due to temperature. The resistor's temperature may vary not just due to ambient temperature, but also due to self-heating from dissipating power. Larger sense resistance values lead to larger current sense voltages and easier current sensing, but also dissipate more power, heat the resistor more, and cause more temperature-related error. A similar tradeoff exists between physically large resistors with low thermal resistance and small resistors with low parasitic inductance. The designer should choose a resistor with a TCR as small as possible given the design constraints, and the effects of TCR can be mitigated by keeping the resistor's temperature more constant. The temperature can be kept more constant by designing the system to use resistors with small value and/or small thermal resistance, or by mechanical means such as a heatsink, forced air, or a temperature-controlled oven like those used for precision oscillators and voltage references. Copper PCB traces have very poor TCR of $3930\frac{ppm}{^\circ C}$; indeed, high TCR and the small current handling capability of thin copper layers are the primary reasons for using special current sense resistors rather than carefully sized PCB traces. When integrating a sense resistor into a power sensor, it is therefore important that the PCB traces from the resistor to the first stage amplifier be short, geometrically symmetric, and thermally symmetric to minimize their contribution to sensed current error.

Wherever two dissimilar metals meet, as in the resistive element and the substrate of a surface mount resistor or the substrate and the copper PCB pads, a thermocouple is formed. That is, if the two metals are at different temperatures, they develop a differential voltage proportional to the temperature difference; this thermoelectric voltage is specified for a pair of materials in $\frac{V}{^\circ C}$. While this effect is used to great effect for measuring temperature, in the case of a current sense resistor the thermocouples formed are considered parasitic, since the thermoelectric voltage manifests as sensed current error. The thermoelectric voltage is analogous to TCR in terms of system design considerations; these considerations are broken into two broad categories, avoidance and mitigation. Thermoelectric voltage minimization is

mostly undertaken by component manufacturers by selecting thermoelectrically compatible materials; its impact can be mitigated by minimizing the temperature difference across all parasitic thermocouples. These mitigation techniques are mostly at the PCB layout level; the easiest technique is to align a resistor such that its terminals are isothermal based on a first-order model of major heat sources in the system, such as high-power components. Another technique is to implement a single resistance using two series half-value resistors next to one another; even if the terminals of a given resistor are not isothermal, producing a thermoelectric voltage, the voltages of the two adjacent resistors will cancel [38].

4.1.2   Analog-to-Digital Converter

The voltage and current signals must ultimately be digitized in order to store and make decisions based on power measurements. The first design choice is whether to digitize voltage and current on two separate ADCs and multiply in the digital domain, or use an analog multiplier and a single ADC. While analog multipliers have simplicity, power, and cost benefits over a high-speed ADC, there is no commercially available multiplier with the combination of high bandwidth and high accuracy matching high-end ADCs. In our proof of concept, we digitize voltage and current separately to achieve the highest bandwidth and accuracy possible and gain the ability to examine each separately.[1] The most obviously important specifications of the ADC in this context are sampling rate and resolution, but several other specifications also impact the uncertainty of the resulting power measurements. Integral nonlinearity (INL) and differential nonlinearity (DNL) characterize the nonideality of the mapping from ADC input voltages to output codes. The system designer can use a wide variety of techniques to compensate for the DC and AC incarnations of these nonlinearities [39]. In this work, we use a simple one-dimensional lookup table (LUT) to compensate for ADC nonlinearities and amplifier gain, offset, and linearity errors. While our technique maps a single ADC sample to an output current or voltage value considering no other information, it could be extended to consider several previous samples and internal ADC error sources.

---

[1]For example, to characterize voltage droop in the PDN or capture current stimulus data to drive simulations using alternative PDNs.

Besides maintaining a correct code-to-input-voltage mapping, the other primary requirement for obtaining an accurate digital waveform is well controlled sample timing; that is, the position of the sample instants should be well known in order to properly reconstruct a digital waveform with the same shape as its analog counterpart. In this work, we consider only uniform sampling regimes, in which the sampling instants are ideally regularly spaced and sample jitter should be minimized. For externally clocked ADCs, sample jitter consists of external clock jitter, which can be improved at the system level, and internal aperture jitter, which is intrinsic to the ADC selected. Nonuniform or periodic nonuniform sampling, where samples are taken at irregularly spaced but known times, are also promising techniques in a power measurement context. Nonuniform sampling allows perfect reconstruction of certain signals that do not strictly meet the classical Nyquist criterion of being bandlimited to half the average sample rate [40]. In offline power measurement applications like a power optimizing compiler, the system has the ability to run the same piece of code repeatedly, thus presenting a periodic signal to the ADC. An importance sampling algorithm like VEGAS [41] could be used to vary sample positions between runs to get more information about quickly varying portions of the sampled waveforms; in this way, the overall power and energy waveforms could be reconstructed accurately even if the bandwidth of these waveforms greatly exceeded average ADC sample rate, as long as the system has fine-grained control over sample position as in the system proposed by Papenfuss et al. [42].

4.1.3   Amplifiers

Amplifiers are needed to apply gain and offset to the raw voltage and current signals from the current sense resistor so that they occupy the entire ADC input voltage range and maximize effective resolution. The μV- to mV-level sense current signals and multi-Volt ADC input ranges lead to relatively high required gain, up to $100\times$ or more for low-voltage rails with low burden voltage tolerance. The amplifiers must also have very high bandwidth to be able to accurately measure the power and energy consumption of very short pieces of code. Amplifiers are often specified in terms of their gain-bandwidth product (GBP); our application requires very high

GBP, particularly for the current channel, since the differential current sense voltage is necessarily much smaller than the rail voltage. In this work, we only consider operational amplifiers (op-amps), as opposed to standalone FET amplifiers, to take advantage of their complex integrated matching and compensation mechanisms. Op-amps often have fairly complex closed-loop gain characteristics which vary according to load capacitance, and are often characterized using a scalar bandwidth figure, the -3dB bandwidth. The -3dB bandwidth is the frequency $f_{-3\text{dB}}$ for which $gain(f) \geq gain(0$ Hz$) - 3\text{dB} \ \forall \ f \leq f_{-3\text{dB}}$. For applications demanding absolute accuracy across a broad frequency range, a more demanding specification is in order; in this work, we use the -0.1dB flatness specification commonly used in analog video. The -0.1dB flatness bandwidth is the frequency $f_{-0.1dB}$ for which $|gain(f_1) - gain(f_2)| \leq 0.1dB \ \forall \ f_1, f_2 \leq f_{-0.1dB}$. That is, given two equal-power sinusoidal inputs of arbitrary frequencies up to the -0.1dB flatness point, the two output waveforms will have power within 0.1dB of one another [43]. This 0.1dB power difference corresponds to a maximum amplitude difference of just over 1%. For first-order op-amp systems, the -0.1dB flatness bandwidth will be about 6.55× lower than the -3dB bandwidth due to the more stringent requirement [22].

We set out to design one power sensor architecture that will work for any voltage rail in commonly used processors and memory systems, so the amplifiers connected to the sense resistor must be able to handle a common-mode voltage of up to 12V with full accuracy.

For our application, high DC and AC accuracy are also important, so we look for amplifiers with high common mode rejection ratio (CMRR) and power supply rejection ratio (PSRR), and low bias current, offset current, offset voltage, and gain error.

4.2 Implementation

We now describe a proof of concept implementation of Loupe, including two separate power sensor PCBs for different categories of DUT, satisfying the design constraints and goals described in Section 4.1. Figure 4.2 shows a diagram of the complete system.
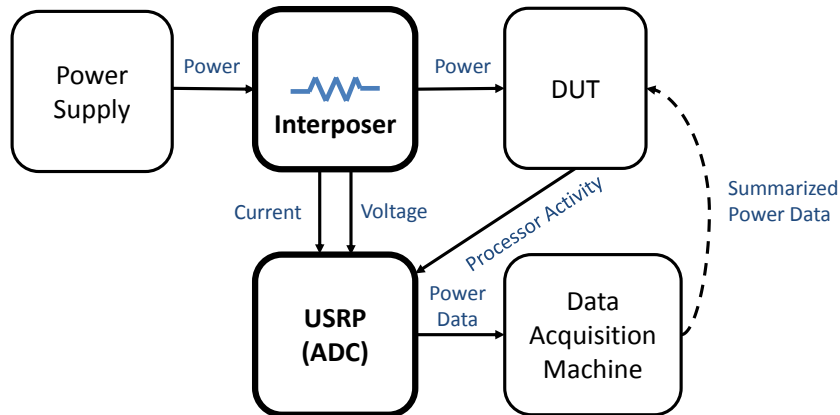
Figure 4.2: In our proof-of-concept experimental setup, the interposer provides amplified voltage and current signals to the high-speed ADCs on the USRP. The USRP's FPGA combines the power data with processor activity data sent directly from the DUT and streams the combined data to a second measurement computer over USB 2.0. The combined data stream could also be sent, directly from the USRP or via the measurement computer, back to the DUT to aid self-hosted power optimization tools like compilers and auto-tuners.

4.2.1 Power Sensor PCB

Bus-based accelerators like PCIe GPUs are an easy candidate to instrument with Loupe, since power can be intercepted with no soldering or other system modifications required. Our first power sensor implementation was thus designed as an interposer between a PCIe card and a motherboard. Motherboards use three rails to provide up to 75W to PCIe cards — 12V, 3.3V, and a lower-current 3.3VAUX — and our power sensor has three voltage and channels to measure each separately.

Amplifier Circuit Design

The full amplifier schematics can be found in Figures A.1–A.2.

Noise Analysis In addition to the systematic biases in output voltage due to component tolerances and repeatable ambient temperature effects, we must consider the effect of random noise. Noise is one of the two major sources of non-systematic, uncalibratable noise in our system, along with self-heating-related thermal effects in the sense resistor and ICs. Thus, though we include the computed noise parameters in our overall uncertainty analyses,

we also examine it separately here to gain insight into the fundamental limits of accuracy in our system.

In op-amp circuits, the main sources of error are input noise current and voltage associated with the op-amp itself and Johnson noise, also known as thermal noise, from discrete offset and gain resistors [44]. The noise current and voltage op-amp specifications are aggregates of multiple underlying noise mechanisms such as Johnson noise from integrated resistors, shot noise, and flicker noise. The Johnson noise spectral density observed across a resistance R, measured in $\frac{\text{V}}{\sqrt{\text{Hz}}}$, is proportional to absolute temperature T and is calculated as $4k_BTR$, where k_B is Boltzmann's constant. Johnson noise is spectrally white, meaning that it has equal power at all frequencies; therefore, to convert the noise density to an RMS noise voltage, a frequency band of interest with bandwidth B must be defined. With this constraint, the RMS noise voltage over the frequencies of interest is $\sqrt{4k_BTRB}$, and the distribution of noise voltage over time is assumed to be normal.

In the following analysis, we denote Johnson noise associated with a resistor R_x as simply R_x; input noise current on the inverting and non-inverting inputs as I_{in-} and I_{in+}, respectively; and input noise voltage as V_{nin}. Table 4.1 lists the relevant component values, parameters, and noise analysis for the baseline design of our 0–9A, 0.3–1.7V mobile SoC power sensor. In our analysis, noise voltages are referred to the output of the relevant amplifier stage or stages (RTO noise), because we are ultimately interested in the average- and worst-case noise produced at the ADC, rather than comparing the noise magnitude against the magnitude of any particular input signal. Total output noise is $\approx 795\mu$V RMS over a 10MHz bandwidth, corresponding to 0.815 LSB at the 12-bit ADC. RMS noise of $\frac{1}{3}$–1 LSB actually has some benefit in our application since it enables dithering, where oversampling can be used to average out quantization error and improve resolution [45, 46]. Stage 1 contributes 2.26× as much gain and has 4.57× as much standalone RTO noise as stage 2. As in many multi-stage amplifier systems, we find that the noise of the first stage dominates. Specifically, stage 1's input voltage noise alone, when amplified by stage 2's noise gain, is 91.3% as large as the overall RSS noise.

By quantitatively evaluating noise in the design stage, we can quickly evaluate several alternative designs or environmental conditions. Figure 4.3 shows the noise components of the stage 1 amplifier for the baseline design

Table 4.1: Noise analysis parameters and results for a two-stage 0–9A current amplifier circuit.

Parameter	Value	Noise Contribution
Bandwidth	10MHz	
Temperature	25°C	
Stage1 Part Number	AD8129	
Stage 1 $I_{n+/-}$	2.79nA	1.4μV / 1.4μV
Stage 1 $I_{n\text{FB}}$	1.4$\frac{\text{pA}}{\sqrt{\text{Hz}}}$	8.85μV
Stage 1 $V_{n\text{in}}$	133μV	133μV
Stage 1 R_{filt}	50Ω	2×8.85μV
Stage 1 R_f	2kΩ	18.1μV
Stage 1 R_g	221Ω	54.5μV
Total Stage 1 RTO Noise		**146.13μV**
Stage 2 Part Number	OPA847	
Stage 2 $I_{n+/-}$	2.5$\frac{\text{pA}}{\sqrt{\text{Hz}}}$	5.51μV / 1.38μV
Stage 2 $V_{n\text{in}}$	0.85$\frac{\text{nV}}{\sqrt{\text{Hz}}}$	14.6μV
Stage 2 R_f	174Ω	5.34μV
Stage 2 R_g	39.2Ω	11.26μV
Stage 2 $R_{\text{Offset1/2}}$	137Ω / 2kΩ	24.9μV
Total Stage 2 RTO Noise		**32.00μV**
Total Stage 1 + Stage 2 RTO Noise		**795.42μV**

and 4 alternatives. The four alternative design decisions are: reducing the feedback resistor values by 4× to 499Ω/54.9Ω, using 1kΩ input filter resistors, eliminating the input filter resistors, and using an AD8130 amplifier with lower noise and a gain of 2× versus the baseline AD8129's 10×. When using the alternative part, the stage 2 gain is increased to compensate for the lost stage 1 gain. While reducing the feedback resistor values is an attractive option to reduce Johnson noise and the impact of the FB input current noise, it only reduces stage 1 RSS noise by 5.8% due to the dominance of input voltage noise. The benefit may not be worth the 4× increase in drive current required by the lower-valued resistors, especially considering the resulting self heating not modeled in this analysis. Similarly, using a larger input filter resistor or no filter at all impacts stage 1 RSS noise by +4.9% and -6.2%, respectively.

Figure 4.3 shows the noise components of the stage 2 amplifier and the cascaded two-stage system for several scenarios, all using the baseline stage 1 design. The first three groups of bars show the typical behavior at 25°C,
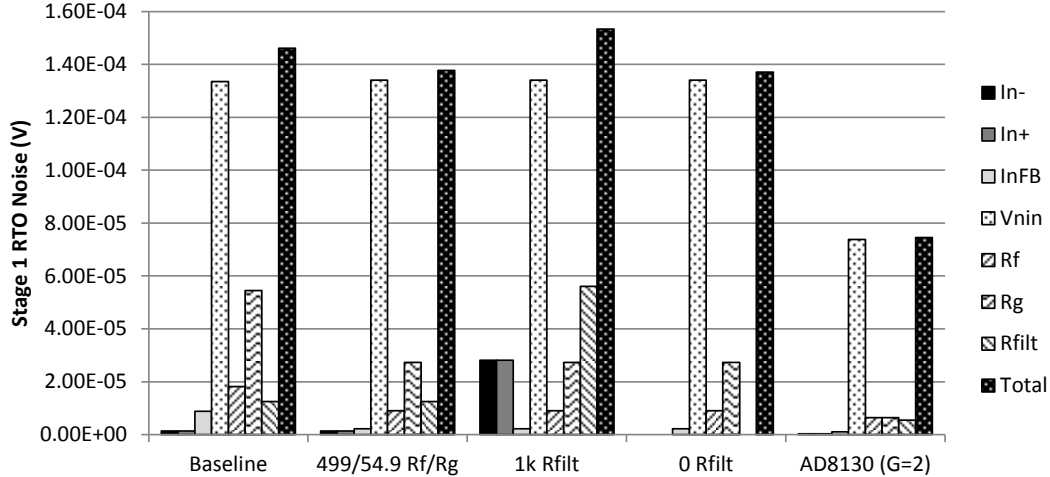
Figure 4.3: Stage 1 current amplifier noise components for the baseline
design and four potential design variants.

worst-case behavior at 25°C, and worst-case behavior over the $0 - -70$°C
temperature range. The results vary by less than 0.3%, since the dominant
$V_{nin,Stage1}$ does not change. The fourth group shows the small effect (0.04%)
of reducing the offset resistor values by 10×; this noise reduction clearly does
not outweigh the 10× static power increase. The fifth group shows the overall
system implications of switching to an AD8130 in stage 1 and increasing the
gain in stage 2. In this case, the stage 1 RTO noise is halved despite the
AD8130's larger V_{nin} due to the decreased noise gain. However, the stage 2
RTO noise is increased by 3.36× due to its correspondingly increased noise
gain. The overall effect is a noise increase of 2.18×, so the change is not
beneficial.

Figure 4.5 shows the noise components for the single stage voltage amplifier
in our 0.3–1.7V power sensor. This amplifier receives a much larger input
signal than the current amplifiers, and its noise gain is thus much lower at
3.85×. The overall RTO voltage noise is 47.81μV, or 0.049 LSBs at the ADC.
The dominant component, as with the current amplifiers, is V_{nin}.

SPICE Characterization The amplifiers on the power sensor have a lin-
ear response on all voltage and current channels through the entire range
specified by the PCIe standard, as shown in Table 4.2. Group delays for
all six channels are small and extremely flat from DC to 20MHz, yielding
minimal time-domain waveform distortion. All six channels also have flat
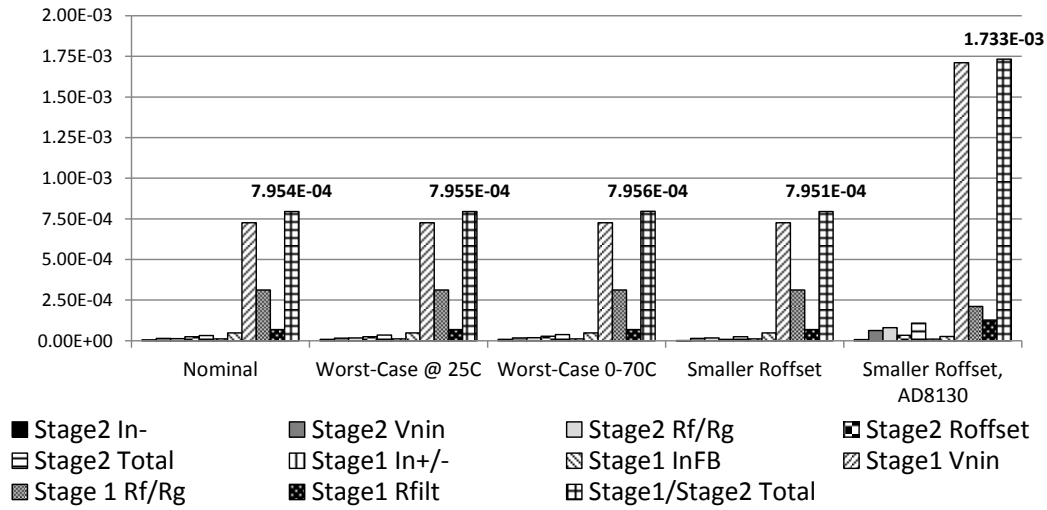AC responses to 10MHz or above, and amplify the signal range of interest to

31

Figure 4.4: Stage 2 amplifier and two-stage system noise components for the typical case, worst case, worst case over temperature, and two potential design variants. All values are in RMS Volts.
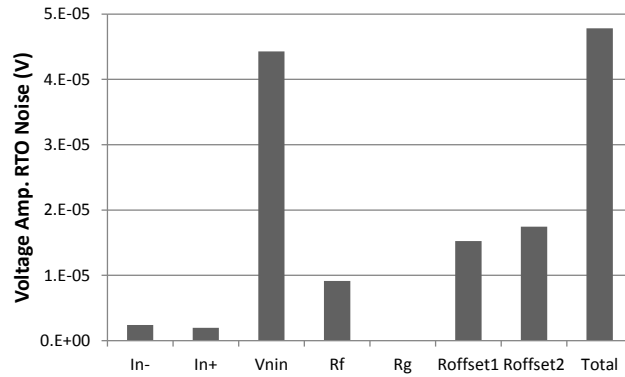

Figure 4.5: Voltage amplifier noise components. Due to the lower noise gain, the RTO voltage noise is much lower than that of the current amplifiers, at 47.81μV RMS.

Table 4.2: Range of input values resulting in a linear DC response at the USRP ADC for each of six channels on the PCIe interposer. All channels have a linear response over the entire range allowed by the PCIe 1.0–3.0 standards. Small, stable group delays from 0–20MHz enable simple time registration of voltage, current, and processor activity in the digital domain.

Channel	Measured ADC Linear Range	PCIe Specified Range	Group Delay
V(12V)	10.1–14.08V	12.8–13.2V	1.54–1.56ns
V(3.3V)	2.73–3.77V	2.97–3.63V	2.51–2.60ns
V(3.3VAUX)	2.73–3.77V	2.97–3.63V	2.51–2.60ns
I(12V)	0–9A	0–130mV (0–5.5A)	3.75–3.79ns
I(3.3V)	0–3.11A	0–27mV (0–3A)	2.63–2.66ns
I(3.3VAUX)	0–0.4A	0–26.25mV (0–0.375A)	2.63–2.66ns

occupy a [-2V, +2V] range, as shown in Figure 4.6. The termination resistors on either end of the coaxial cables act as a voltage divider and attenuate the amplified signal by 2× to fit the USRP ADC's $[-1V, +1V]$ input range.

Power Supply Circuit Design

PCB Design

Loupe's amplifiers are housed in a custom designed 4-layer PCB, shown in the renderings in Figure 4.7 and the photographs in Figure 4.8. The amplifier PCB is in the form factor of a PCIe card, and the DUT input and output connectors are PCIe male and female connectors. The board can fit inside a standard ATX PC case between any PCIe 1.0–3.0 card and the motherboard. We populated one copy of the board with amplifiers for all 3 PCIe power rails in accordance with the PCIe 3.0 specified voltage and current ranges, and another copy with amplifiers for low-voltage, high-current power rails like those powering mobile SoCs, as shown in Figure 4.9. Together, these two PCB variants span devices with two orders of magnitude difference in power consumption with a single design.

(a) V(12V) AC

(b) I(12V) AC

(c) V(12V) AC (zoomed)

(d) I(12V) AC (zoomed)

(e) V(12V) DC

(f) I(12V) DC

(g) V(3.3V) AC

(h) I(3.3V) AC

(i) V(3.3V) AC (zoomed)

(j) I(3.3V) AC (zoomed)
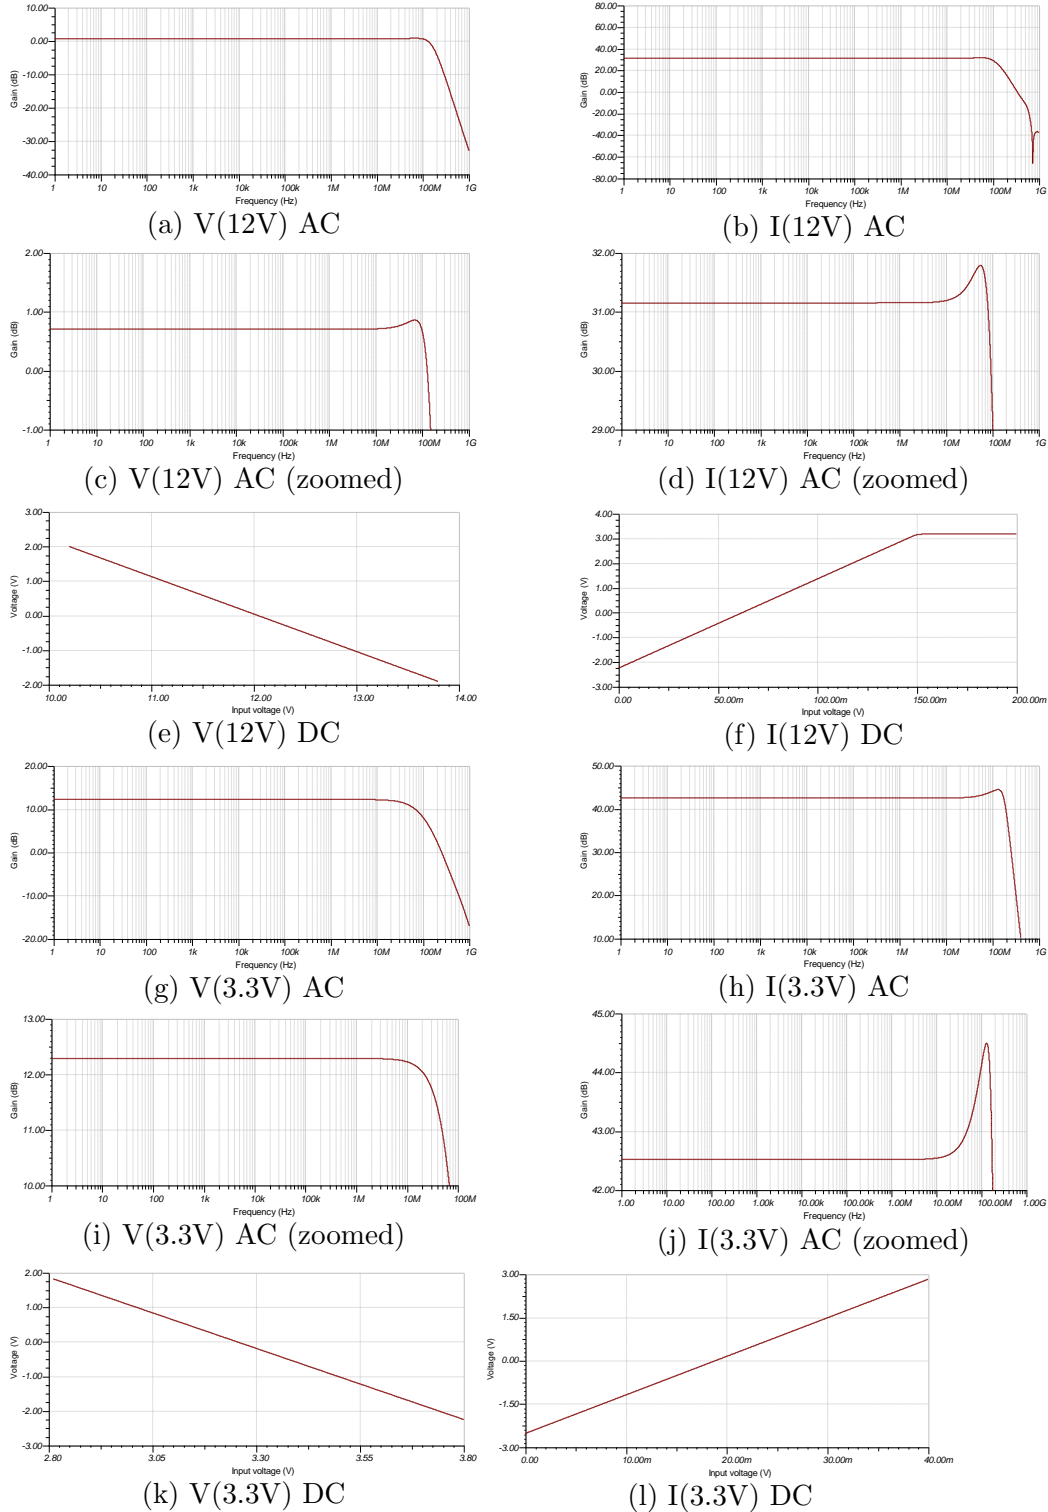

(k) V(3.3V) DC

(l) I(3.3V) DC

Figure 4.6: Signal chain AC and DC responses shown for 12V and 3.3V voltage and current rails on PCIe power sensor. The 3.3V and 3.3VAUX rails share identical amplifier circuits and differ only in sense resistor value. All amplifiers have linear DC responses and amplify the signal values of interest to occupy [-2V, +2V]. All AC responses have better than 0.1dB flatness from 0Hz to at least 10MHz and minimal peaking above 10MHz.
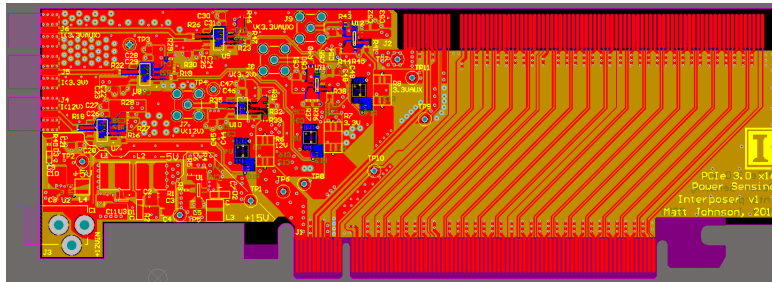
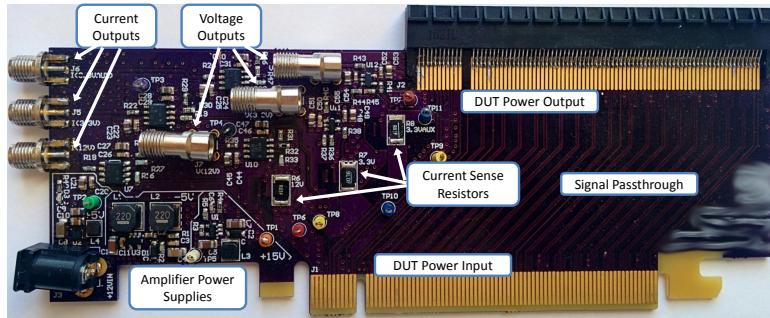(a) Front (2D)
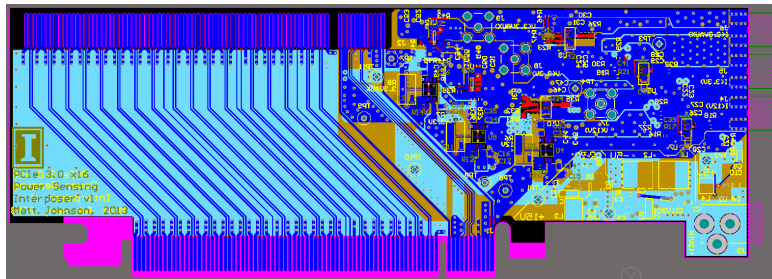
(b) Back (2D)
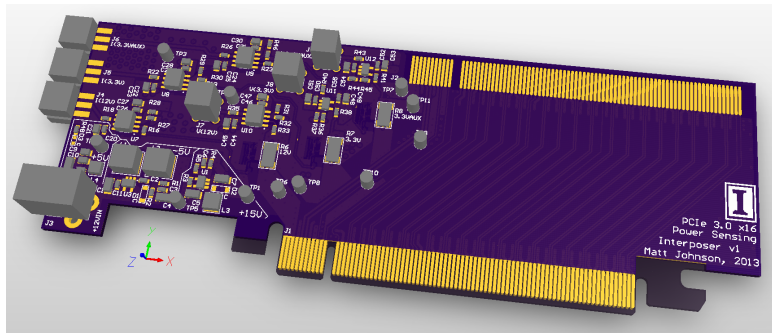
(c) Front (3D)
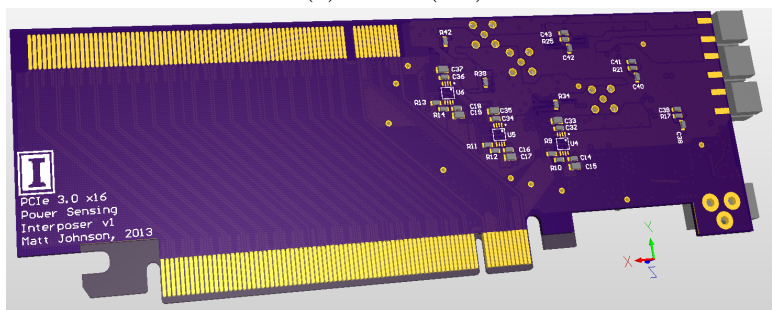
(d) Back (3D)

Figure 4.7: 2D and 3D renderings of the 4-layer PCIe interposer PCB.
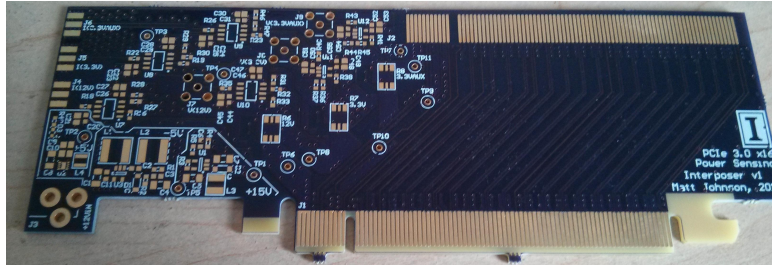
(a) Bare Board

(b) Populated

Figure 4.8: Photographs of the bare and populated 4-layer PCIe interposer PCB. The interposer passes through PCIe data signals undisturbed, and uses 1 and 2 stages of op-amps to amplify the voltage and current signals, respectively, corresponding to the DUT-side voltage and voltage drop across the current sense resistors on each of PCIe's 3 power rails. The amplified current and voltage signals are sent to a high-speed ADC via RG174 coaxial cable.
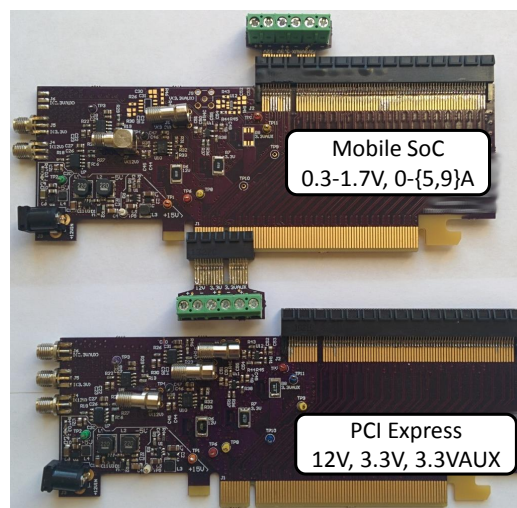

Figure 4.9: The Loupe amplifier PCB design can be populated with different footprint-compatible amplifiers and passive components to measure a wide range of systems, including a 75W PCIe card and a mobile SoC. Breakout daughtercards enable the measurement of non-PCIe power rails.
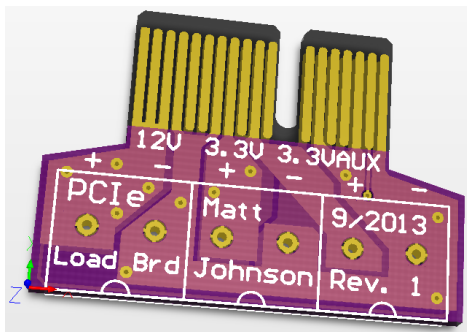
36

Figure 4.10: The Jouler's Loupe adapter PCB allows non-PCIe voltage rails to be measured using the PCIe form factor power sensor by adapting between male or female PCIe connectors and terminal blocks.

4.2.2 Connection to DUT

Our original PCIe form factor power sensor can measure unmodified PCIe 1.0–3.0 cards plugging into unmodified PCIe motherboards via its built-in male and female PCIe connectors. To enable the same power sensor to measure current through wires in non-PCIe contexts, we designed an adapter board as shown in Figure 4.10. The same adapter board can plug into a female PCIe connector using the bare edge connectors or a male PCIe connector by soldering a female PCIe connector to the edge connectors on the adapter board. The other side of the adapter contains terminal blocks where DUT power wires can be attached without soldering. Using one male-configured and one female-configured adapter board on each side of the power sensor, we are able to intercept any external DUT power rail for measurement. Such a configuration is best suited for low-frequency power measurement or for measuring power through a wire or connector, since the loop inductance incurred by the ≈4.5 inch span of the adapter boards may cause significant ringing in PCB power rails with rapidly varying current. The small-form-factor power sensor we developed to address this problem has two parallel sets of DUT connections: terminal blocks as in the PCIe adapter boards, and large plated through holes with short, thick gauge wires soldered to the through holes and the relevant positions on the DUT PCB. The sense resistor on the small-form-factor is less than 2mm from the plated through holes for a dramatically reduced current loop area compared to the PCIe sensor plus adapter boards.
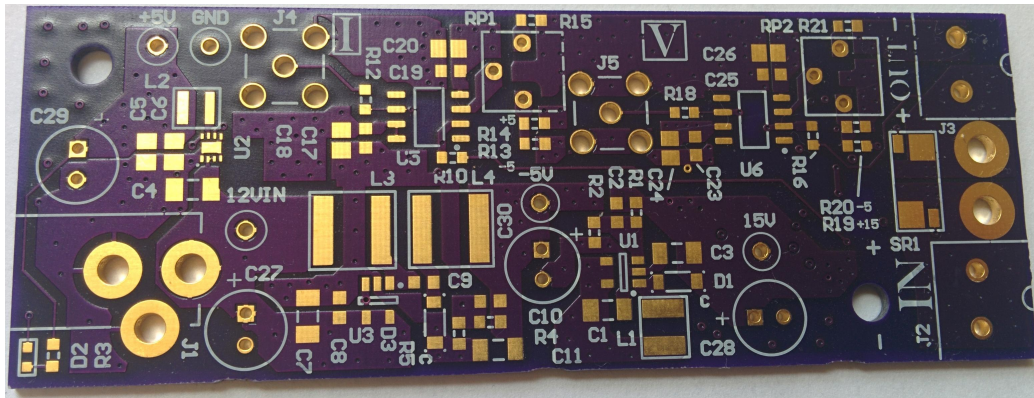
(a) Bare Board
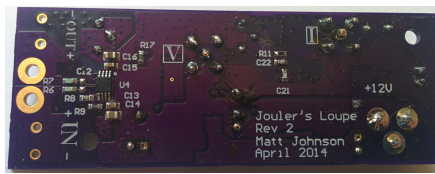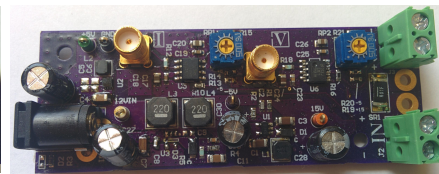
(b) Populated (Back)   (c) Populated (Front)

Figure 4.11: The small form factor single channel power sensor incorporates some lessons learned during the design of the PCIe sensor and enables very low inductance DUT connections for PCB level power measurements.

4.2.3   Small Form Factor Power Sensor

As mentioned in Section 4.2.2, the PCIe power sensor is well suited to measuring currents through wires or connectors, but adds a relatively large parasitic inductance when used to measure non-PCIe PCB-level power rails. Many PCB-level power rails of interest are also part of space-constrained systems where the sheer size of the PCIe power sensor is problematic. We designed a smaller single channel power sensor to address these issues, as shown in Figure 4.11. The smaller sensor uses a circuit topology and components similar to the PCIe sensor, including on-board +5V, +15V, and -5V voltage regulators for the amplifiers. We also opportunistically modified the design in small ways to further improve measurement uncertainty and adjustability. First, we added the ability to select the voltages used in the resistive divider that sets the DC offset for the voltage and stage 2 current amplifiers by including multiple parallel offset resistor footprints. This gives the user additional flexibility in the voltage and current ranges that can be measured, the supported ADC input ranges, and the resistor values that can be used to achieve a given offset. We also added a potentiometer in the offset

voltage dividers to enable some offset adjustability after assembly to compensate for small DC offsets in the amplifiers and ADC. This adjustability is helpful in avoiding clipping at the ADC when the amplifier output range is very close to the ADC input range to maximize resolution. Care must be taken to maintain the good noise performance of the power sensor by using a small value potentiometer with low Johnson noise; in this work, we use a 50Ω potentiometer. We also used a new Kelvin sensing pad layout for the current sense resistor that can accommodate the Vishay WSK2512 series sense resistors in addition to the WSL2512 series supported on the PCIe sensor; the WSK series offers an improved TCR of $\pm 35 \frac{ppm}{^\circ C}$, down from $\pm 50 \frac{ppm}{^\circ C}$ for the WSL.

In this work, we populated the smaller power sensors with components to measure two power rails with relatively low voltage, 0.3–1.7V, and relatively high current, 0–5A and 0–9A. These measurement ranges are suitable for measurement of two CPU, GPU, or DRAM power rails on a mobile SoC. Both channels were designed assuming a $10\text{m}\Omega$ sense resistor, for a full-scale resistor voltage of 50mV and 90mV, respectively. Figure 4.12 shows the DC transfer function and frequency response of these two channels. Due to the smaller rail voltage, higher gain is required and these channels have slightly lower -0.1dB bandwidth than the 3.3–12V PCIe rail amplifiers, but still exceed the 10MHz specification.

4.2.4   Digital Processing

For the high-speed ADCs and incorporation of processor activity data, Loupe uses an Ettus Research USRP1, a device originally designed for software-defined radio. The USRP1 includes 4 64MSPS, 12-bit ADCs, a modest FPGA, and a USB 2.0 interface. The coaxial cables carrying amplified voltage and current signals are connected to the ADCs via low-frequency receiver (LFRX) daughterboards. The daughterboards contain SMA coaxial connectors, a unity-gain single-ended-to-differential amplifier, and headers to breakout GPIO pins on the FPGA. We transmit processor activity data to the FPGA via a custom cable from the DUT's signaling mechanism to the daughterboard's 0.1 inch headers on the daughterboard. When using the PCIe amplifier board, the host x86 PC sends the USRP processor activity
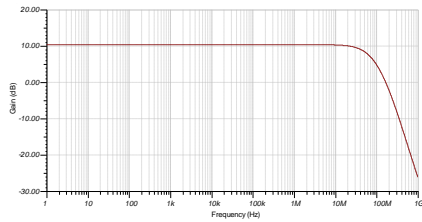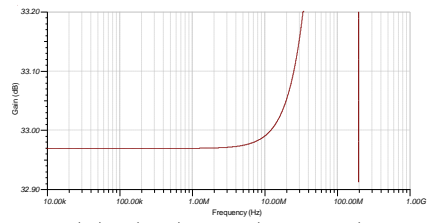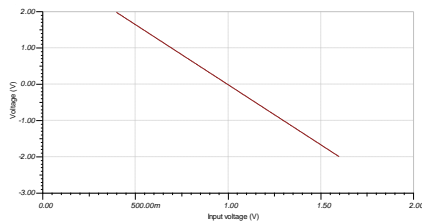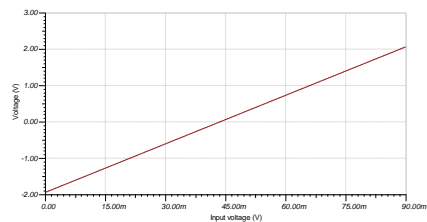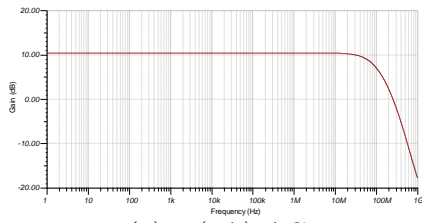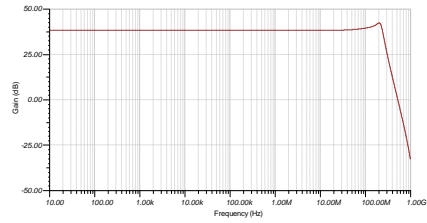
(a) V(9A) AC
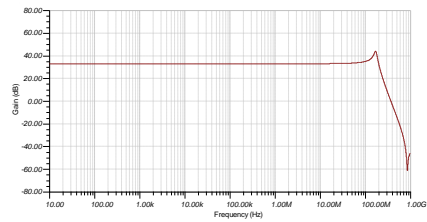
(b) I(9A) AC
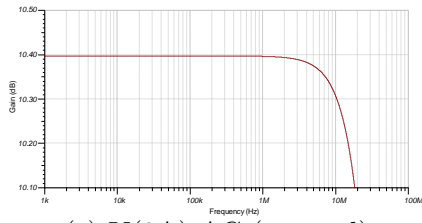
(c) V(9A) AC (zoomed)

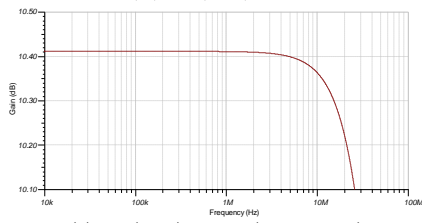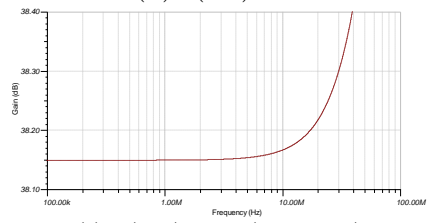(d) I(9A) AC (zoomed)

(e) V(9A) DC
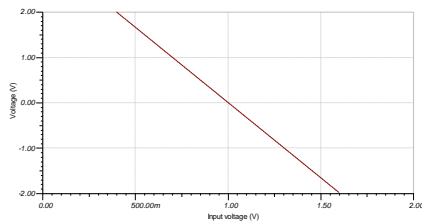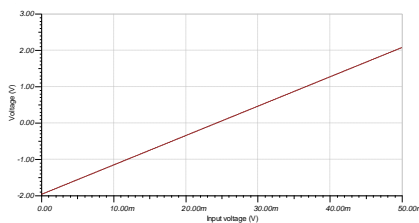
(f) I(9A) DC

(g) V(5A) AC

(h) I(5A) AC

(i) V(5A) AC (zoomed)

(j) I(5A) AC (zoomed)

(k) V(5A) DC

(l) I(5A) DC

Figure 4.12: Signal chain AC and DC responses of the voltage and current channels on the small form factor 5A and 9A sensors.
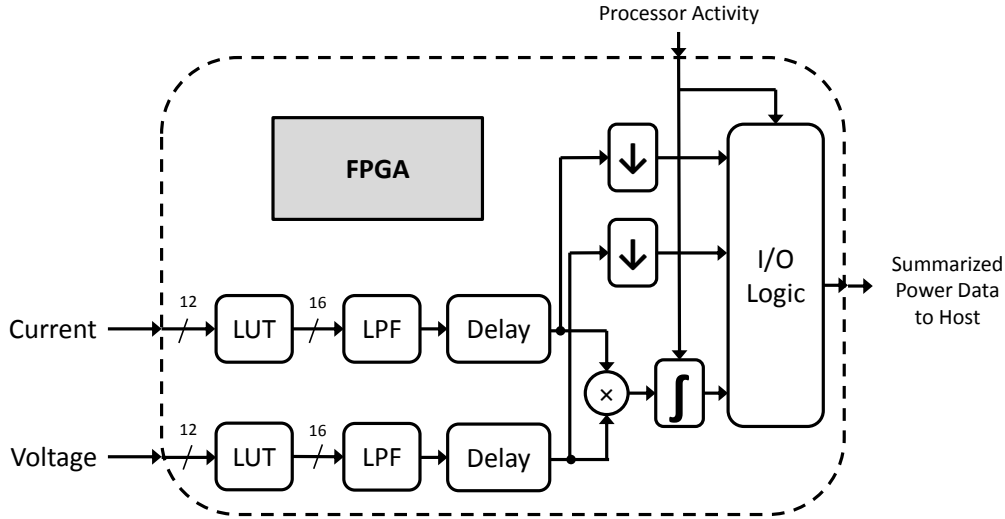
Figure 4.13: The FPGA in Jouler's Loupe post-processes digitized current and voltage data and integrates it with processor activity data to output a stream of temporally registered power samples or energy estimates. The post-processed data is sent to the DUT itself or a separate host machine via USB; other USRP devices are available with higher-bandwidth 1Gb Ethernet, 10Gb Ethernet, or PCIe interfaces.

data via a StarTech PEX1P PCIe-connected parallel port used as bit-bang GPIO; on an Intel i7-2600k processor at 3.4GHz, the PC can write to the parallel port 573000 times per second. The mean write latency is $1.365\mu s$, with a very small standard deviation of 22.274ns. Like the amplifier group delay, then, this skew is very consistent and its mean is calibrated out using the digital delay lines on the FPGA. Our knowledge of the distribution of write latencies is useful in characterizing the uncertainty of aggregate time-interval energy estimates, as discussed in Section 4.3.2. When using the small form factor amplifier board, the mobile SoC DUT sends the USRP processor activity data via memory-mapped GPIO pins, accessed from userspace when necessary using a Linux kernel module.

The ADCs, which reside on the USRP's main board, are directly connected to the FPGA, where all at-speed digital-domain processing takes place. Figure 4.13 shows an overview of the signal processing chain implemented in the FPGA. We heavily modified the USRP's stock FPGA image to remove unused communications-centric functionality, add needed calibration and filtering logic, and improve dynamic range and rounding logic. The incoming 12-bit current and voltage samples from the ADC are first remapped to 16-bit

values via separate current and voltage LUTs. The LUTs are filled with values computed from a per-power-sensor calibration procedure and can correct for any nonlinearities in the overall analog transfer function, including gain resistor component tolerances, amplifier gain error and offset voltage, and ADC INL and DNL. In our implementation, all FPGA logic is fixed-point; that is, it assumes a simple linear relationship between the N-bit digital value and the underlying current, voltage, power, or energy. A floating point or other nonlinear mapping would improve resolution when the measured voltage or current is low, a helpful quality for measuring mobile and embedded devices with a large dynamic range of possible power consumption; we leave an investigation of such mappings to future work.

The ADC analog inputs include a simple first-order RC antialiasing filter to attenuate very high-frequency content, but the ADC sample rate f_s greatly exceeds the analog bandwidth of the power sensor f_a. Therefore, the raw digitized signal may still contain significant frequency content in the interval $[f_a, \frac{f_s}{2}]$ which does not necessarily exhibit the flat group delay characteristics we have shown for the passband. This high frequency content should therefore be attenuated to avoid corrupting the time domain properties of the measured power signal; this attenuation is accomplished by linear-phase FIR low pass filters, labeled "LPF" in Figure 4.13.

The remapped and filtered current and voltage data is fed to a digital delay line (DDL) which delays each signal by a programmable amount of time. Integer cycle delays are implemented with a programmable shift register. The fractional cycle remainder of the desired delay, if desired, is achieved with a fractional delay filter (FDF), which performs bandlimited interpolation to shift its input signal by a fractional number of samples; in the frequency domain, the ideal FDF has linear phase (i.e., flat group delay) and flat amplitude response [47]. The desired delay only depends on the power sensor and DUT characteristics, so we are able to use a fixed FDF for a given power sensor/DUT pair, improving the quality of the filter and reducing its implementation cost. In future implementations , a more expensive variable FDF may be used to enable setting the delay at runtime and avoid regenerating a new FPGA image for each power sensor/DUT pair.

After the DDL, the current and voltage data is sent down two parallel paths: the first multiplies current and voltage to obtain power samples and integrates the power samples to form energy estimates over time intervals

42

defined by the external processor activity data, and the second simply down-samples the current and voltage data to fit the available host bandwidth. Finally, a block of output processing logic sends either the integrated energy estimates or a stream of interleaved current, voltage, and processor activity data to the host, depending on the mode set by the user.

4.3   Analysis

4.3.1   Calibration

Apart from the oscilloscope-based delay calibration, the end-to-end DC current-to-ADC-code and voltage-to-ADC-code transfer functions were measured using a low-noise linear power supply and a calibrated 6.5-digit multimeter. As predicted by SPICE simulations in the design phase, every channel's DC response was extremely linear throughout the ADC's input range ($r^2 > 0.9999998$ for 15-point best fit lines). The calibration phase helps correct for any DC gain and offset deviations caused by component tolerances.

4.3.2   Uncertainty Analysis

Single Measurement Uncertainty

Inspired by Nakutis [48], we used the software tool GUM Workbench [49] to analyze the uncertainty in our current measurements. In an uncertainty analysis, one specifies all the uncertain quantities, such as temperatures, temperature coefficients, and component tolerances, any correlations between these quantities, and the algebraic equations that combine them to form the output variable (measured current in this case). The software tool then analytically computes the contribution of each input quantity's uncertainty to the overall output uncertainty, and can run Monte Carlo simulations to visualize complex output variable distributions.

Our uncertainty analysis involved 39 equations and 99 quantities, including component tolerances, temperature variations, amplifier bias current and offset voltage, and ADC clock jitter. Even with the conservatively tight tolerance for most components on the initial Loupe PCB, when the board is

uncalibrated and actual component values are not known, the uncertainty is as high as 3.4% of full scale; of that uncertainty, most of it is caused by uncertainty in stage 2 amplifier bias current (57%), sense resistor room temperature value (12.6%), and stage 1 amplifier feedback resistor nominal value (10.1%).

Once a board has been carefully calibrated, the uncertainty in many quantities, including nominal passive values and amplifier offset/bias, can be reduced to nearly 0. Calibration drops the uncertainty by more than 20× to 0.16% of full scale; by far the dominant source of remaining uncertainty is the sense resistor's tempco and associated resistance change due to self-heating by up to 25°C. To further reduce the uncertainty of a calibrated Loupe board, the results indicate that a heatsink or fan for the sense resistors would be most effective.

Application-Level Uncertainty

In most cases, the interesting outputs of a power measurement apparatus at the application level will be estimates of power or energy derived from multiple current and voltage samples. Additional statistical analysis can be applied to understand the properties of these higher-level quantities in order to inform a new class of more rigorous power-focused characterization and optimization algorithms.

For the simple case of measuring a constant steady-state power, we can treat individual power samples as independent random samples from a normal distribution $\mathcal{N}(\mu, \sigma)$. The uncertainty or standard deviation of the average of n samples from such a distribution is $\frac{\sigma}{\sqrt{N}}$. Thus, to achieve any desired uncertainty σ_d, we must take the average of $(\frac{\sigma}{\sigma_d})^2$ samples.

A more common case for power measurement in processor systems is the estimation of the energy consumption of a piece of code using processor activity data to select which of the power samples are relevant. In this case, there are not only amplitude errors for each current and voltage measurement, but also timing errors due to the uncertainty in the temporal alignment between the "start" and "stop" signals at the power measurement device and the true starting and stopping points of the code's execution on the processor. The mean value of the processor activity signal delay can be calibrated out by the digital delay lines on the FPGA, but any remaining jitter will cause
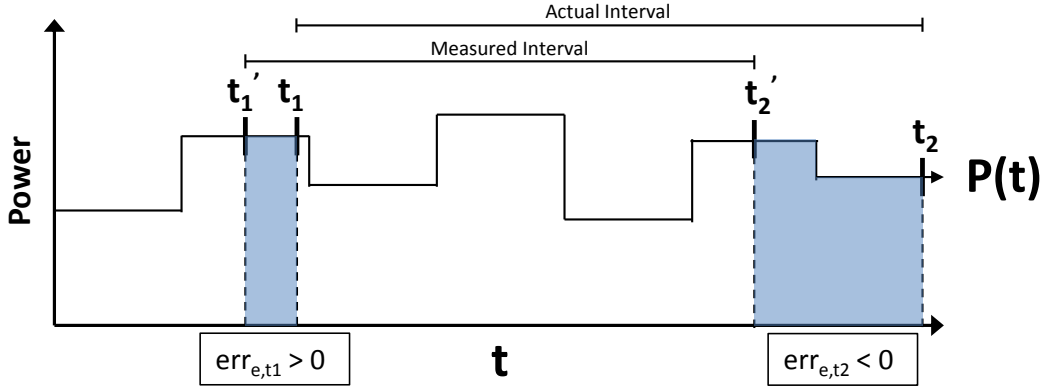
Figure 4.14: An illustration of energy estimation errors resulting from timing errors. Errors in estimating the true start and stop times of a measurement interval $[t_1, t_2]$ cause energy over time periods at the beginning and end of the interval to erroneously be included or excluded when computing energy estimate. We provide a quantitative analysis of the distribution of energy errors given the distribution of timing errors and measured power values.

additional uncertainty in the resulting energy estimate. The amount of total estimation error depends not only on the amount and direction of timing errors on the "start" and "stop" signals, but also on the measured power during the erroneous time periods in question. For example, if power consumption dropped to zero before and after executing a piece of code, then "start" signals arriving too early or "stop" signals arriving too late would not cause any error in the energy estimate. Quantifying the impact of timing uncertainty on energy uncertainty can be handled in two ways: the measured power signal's behavior during the erroneous time periods can be treated statistically as an ensemble of average behavior, yielding a single energy uncertainty number for a given hardware platform, or the error distribution can be more precisely calculated for a given pair of "start" and "stop" signals by looking at the surrounding power measurements.

Assume we have a measured power signal $P(t)$ and we would like to estimate the energy consumed by the processor over the interval $[t_1, t_2]$. Due to timing uncertainty in the system, the start and stop signals arrive at the FPGA at times t_1' and t_2'. The timing errors $err_{t,t_1} = t_1 - t_1'$ and $err_{t,t_2} = t_2 - t_2'$ are independent, identically distributed random variables with a PDF $f(t)$. We assume that the timing errors have zero mean, since in a real system jitter will be relative to the mean measured delay for that hardware platform. The

45

relationship between timing errors and energy estimation errors is shown in Figure 4.14. In the figure, the real start time t_1 is after the estimated start time t_1', causing a positive energy error; that is, the shaded region is included in the energy estimate when it shouldn't be. The real stop time t_2 is also after the estimated stop time t_2', but this causes a negative energy error because the shaded region is not included in the energy estimate when it should be.

In the following equations, we use the slightly modified definite integral notation:

$$\int_x^y f(t)\mathrm{d}t = \begin{cases} \int_x^y f(t)\mathrm{d}t, x \leq y \\ -\int_y^x f(t)\mathrm{d}t, x > y \end{cases}$$

The expected value and variance of the energy error due to the timing error τ in t_1' are given by:

$$\mathrm{E}[err_{e,t_1}] = \int_{-\infty}^{\infty} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau \tag{4.1}$$

$$\sigma^2(err_{e,t_1}) = \int_{-\infty}^{\infty} \left[f(\tau) * \left(\int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right)^2 \right] \mathrm{d}\tau \tag{4.2}$$

While Equations 4.1–4.2 are too computationally expensive for real-time use in the general continuous-time case, a much simpler solution can be found if $P(t)$ is piecewise-constant as in a sampled waveform with sample period t_s using zero-order hold reconstruction. If we register the processor activity data using the ADC sample clock, in the absence of other information we assume that the signals arrive in the center of the clock period. In this case, the inner integrals are very simple when evaluated over a sample period, and we can recast the outer integrals as summations where each term covers one sample period. $P[0]$ is a special case in that half the sample period is before t_1' yielding a negative inner integral, and half is after t_1' yielding a positive inner integral; the integrands for all other sample periods are positive everywhere or negative everywhere in the domain. To calculate expected error more efficiently, we first split the outer integral into 4 pieces:

$$\mathrm{E}[err_{e,t_1}] = \int_{t_1'}^{t_1' - \frac{t_s}{2}} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau + \int_{t_1'-\frac{t_s}{2}}^{-\infty} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau$$

$$+ \int_{t_1'}^{t_1' + \frac{t_s}{2}} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau + \int_{t_1'+\frac{t_s}{2}}^{\infty} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau$$

The first and third terms' inner integrals have domains only covering a single power sample, and their integrands can be replaced by $P[0]$, a constant. The inner integrals thus evaluate to $P[0]\tau$, yielding:

$$\mathrm{E}[err_{e,t_1}] = P[0] \int_{t_1'}^{t_1' - \frac{t_s}{2}} \tau f(\tau)\mathrm{d}\tau + \int_{t_1'-\frac{t_s}{2}}^{-\infty} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau$$

$$+ P[0] \int_{t_1'}^{t_1' + \frac{t_s}{2}} \tau f(\tau)\mathrm{d}\tau + \int_{t_1'+\frac{t_s}{2}}^{\infty} \left[f(\tau) * \int_{t_1'}^{t_1'+\tau} P(t)\mathrm{d}t \right] \mathrm{d}\tau$$

For the second and fourth terms, we split the inner integrals into three parts: the first half sample period from t_1' towards $t_1'+\tau$, zero or more aligned full sample periods, and the last partial sample period to get to $t_1' + \tau$. We will restate τ as the sum of $\frac{t_s}{2}$, an integer number of sample periods it_s, and a sample period offset $o = \tau - it_s - \frac{t_s}{2}$, all multiplied by $sgn(\tau)$. This split yields:

47

$$\mathrm{E}[err_{e,t_1}] = P[0] \int_{t_1'}^{t_1'-\frac{t_s}{2}} \tau f(\tau)\mathrm{d}\tau$$

$$+ \int_{t_1'-\frac{t_s}{2}}^{-\infty} \left[f(\tau) * \left(\int_{t_1'}^{t_1'-\frac{t_s}{2}} P(t)\mathrm{d}t + \int_{t_1'-\frac{t_s}{2}}^{t_1'-\frac{t_s}{2}-it_s} P(t)\mathrm{d}t + \int_{t_1'+\tau-o}^{t_1'+\tau} P(t)\mathrm{d}t \right) \right] \mathrm{d}\tau$$

$$+ P[0] \int_{t_1'}^{t_1'+\frac{t_s}{2}} \tau f(\tau)\mathrm{d}\tau$$

$$+ \int_{t_1'+\frac{t_s}{2}}^{\infty} \left[f(\tau) * \left(\int_{t_1'}^{t_1'+\frac{t_s}{2}} P(t)\mathrm{d}t + \int_{t_1'+\frac{t_s}{2}}^{t_1'+\frac{t_s}{2}+it_s} P(t)\mathrm{d}t + \int_{t_1'+\tau-o}^{t_1'+\tau} P(t)\mathrm{d}t \right) \right] \mathrm{d}\tau$$

Of the three split integral components, the first simply evaluates to $\frac{P[0]t_s}{2}$, the second is a summation of similar terms over the full sample periods between t_1' and $t_1'+\tau$, and the third contains τ and is analogous to the inner integrals of the first and third terms. The outer integral of the second and fourth terms can be expanded into a sum of integrals over sample periods. Performing these two operations leads to:

$$\mathrm{E}[err_{e,t_1}] = P[0] \int_{t_1'}^{t_1'-\frac{t_s}{2}} \tau f(\tau)\mathrm{d}\tau$$

$$+ \sum_{i=-\infty}^{0} \left(\int_{t_1'-\frac{t_s}{2}-(i+1)t_s}^{t_1'-\frac{t_s}{2}-it_s} \left[f(\tau) * \left(-\frac{P[0]t_s}{2} - \sum_{j=i}^{-1}[P[j]t_s] + P[-i-1]*(\tau + it_s + \frac{t_s}{2}) \right) \right] \mathrm{d}\tau \right)$$

$$+ P[0] \int_{t_1'}^{t_1'+\frac{t_s}{2}} \tau f(\tau)\mathrm{d}\tau$$

$$+ \sum_{i=0}^{\infty} \left(\int_{t_1'+\frac{t_s}{2}+it_s}^{t_1'+\frac{t_s}{2}+(i+1)t_s} \left[f(\tau) * \left(\frac{P[0]t_s}{2} + \sum_{j=1}^{i}[P[j]t_s] + P[i+1]*(\tau - it_s - \frac{t_s}{2}) \right) \right] \mathrm{d}\tau \right) \qquad (4.3)$$

Equation 4.3 contains two types of integral terms of the form $\int f(\tau)\mathrm{d}\tau$ and $\int \tau f(\tau)\mathrm{d}\tau$, both with domains of a half or full aligned sample period. These integral terms are all constants that can be precomputed once for a given platform and jitter PDF. For brevity, we adopt the notation $F[i]$ and $G[i]$ for these two sets of coefficients, where i is the index of the relevant P sample; coefficients covering the negative and positive halves of $P[0]$ will use the

subscripts 0^- and 0^+, respectively. We can now rewrite Equation 4.3 as:

$$\mathrm{E}[err_{e,t_1}] = -P[0]G[0^-] - \sum_{i=-\infty}^{-1}\left(t_s\left[\frac{P[0]}{2} + \sum_{j=i+1}^{-1}P[j]\right]F[i] + P[i]G[i]\right)$$
$$+ P[0]G[0^+] + \sum_{i=1}^{\infty}\left(t_s\left[\frac{P[0]}{2} + \sum_{j=1}^{i-1}P[j]\right]F[i] + P[i]G[i]\right)$$

$$(4.4)$$

The analogous equation for variance is:

$$\sigma^2(err_{e,t_1}) = P[0]^2G[0^-] + \sum_{i=-\infty}^{-1}\left(t_s\left[\frac{P[0]^2}{2} + \sum_{j=i+1}^{-1}P[j]^2\right]F[i] + P[i]^2G[i]\right)$$
$$+ P[0]^2G[0^+] + \sum_{i=1}^{\infty}\left(t_s\left[\frac{P[0]^2}{2} + \sum_{j=1}^{i-1}P[j]^2\right]F[i] + P[i]^2G[i]\right)$$

$$(4.5)$$

As a practical matter, the jitter PDF will have bounded support for real systems, and the user may choose to truncate the PDF at a certain point to avoid excess computation to account for extremely rare cases of large timing error. Equations 4.4–4.5 hold for any finite or infinite sets of coefficients $F[i]$ and $G[i]$, and the structure of the inner and outer summations suggests an efficient iterative algorithm to compute the expected value or variance in linear time with respect to $|F[i]|$. Listing 4.1 shows C-like pseudocode for such an algorithm using a PDF supported on the interval $[a\ldots b]$ and, without loss of generality, using floating point values throughout and negative indexing on the P, F, and G arrays.

```
float  ev = P[0]*(G[ZERO_PLUS]  - G[ZERO_MINUS]);
float  sum_p = 0.0f;
for(int  i = -1; i >= a; i--) {
        sum_p += P[i];
        ev -= ts * (P[0]/2 + sum_P) * F[i] + (P[i] * G[i]);
}
sum_p = 0.0f;
for(int  i = 1; i <= b; i++) {
        sum_p += P[i];
        ev += ts * (P[0]/2 + sum_P) * F[i] + (P[i] * G[i]);
}
```

```
return ev;
```

Listing 4.1: .]Iterative algorithm for computing expected energy error for a timing jitter PDF supported on [a,b].

It may seem unintuitive that $\mathrm{E}[err_{e,t_1}]$ is not necessarily zero; that is, even though mean delay has been calibrated out, for a given measurement the mean energy error may be nonzero, meaning that the this expected value should be added to the nominal energy estimate returned to the application. Suppose $t_s = 1$, $f(t)$ is a uniform distribution over $[-1, 2]$, and $P[\text{-}1 \ldots 1]=[0\mathrm{W}, 1\mathrm{W}, 2\mathrm{W}]$. If we want to compute the expected energy error for a run starting during sample period 0 (t'_1 assumed to be 0.5), we first compute the three relevant coefficients $G[-1, 0^-, 0^+, 1]$. For a uniform distribution with domain of size 3, $f(t) = \frac{1}{3}$ everywhere, so $G[-1, 0^-, 0^+, 1] = [\frac{1}{3}, \frac{1}{24}, \frac{1}{24}, \frac{1}{3}]$; F is not used in this example since we are only looking at one sample on either side of t'_1. The summation of Equation 4.4 evaluates to $-1 * \frac{1}{24} - (1 * \frac{1}{2} * \frac{1}{3} + 0 * \frac{1}{6}) + 1 * \frac{1}{24} + (1 * \frac{1}{2} * \frac{1}{3} + 2 * \frac{1}{6}) = \frac{1}{3}J$ expected error. Many terms cancel due to symmetries in the uniform jitter PDF, but the overall expected error is positive, meaning that, on average, our measurements overestimate the energy during $[t_1, 2]$ by $\frac{1}{3}J$. This overestimation is due to the asymmetry in the measured power waveform with respect to t'_1; though the true t_1 is equally likely to be before or after the measured t'_1 given our jitter PDF, the cases where t_1 is after t'_1 have a disproportionate impact on the expected error since the erroneously attributed power is higher.

The astute reader will note that the foregoing analysis has focused on t_1, the beginning of a measurement interval on the processor, exclusively. Equations 4.4 – 4.5 can be used to quantify end-of-interval error around t_2 by replacing $t_1\prime$ with t'_2 and inverting the signs of the four terms in each equation, since $t'_2 < t_2$ leads to an *under*estimation of interval energy, as opposed to *over*estimation when $t'_1 < t_1$.

CHAPTER 5

EXPERIMENTAL METHODOLOGY

Our experimental platform is an ODROID-XU development board running version 3.4.75 of the Linux kernel. The ODROID-XU+E powers the A7 cores off their own rail and has its own 10mΩ sense resistor on-board connected to an INA231 I^2C-connected power monitor. We removed the current sense resistor and connected its pads to the input and output of the Loupe measurement board via two short wires. We were then able to use Loupe's current sense resistor, which has far better thermal dissipation, parasitic inductance, and tempco properties. We communicate processor activity data to the USRP via a memory-mapped GPIO pin using a very simple protocol; we toggle the GPIO at the beginning and end of each measured run. More sophisticated protocols might use multiple GPIOs to do fine-grained accounting of execution time between multiple layers of the software stack. The GPIO is connected to the USRP via a 1.8V-to-3.3V level shifter and a custom-built cable. Its latency mean and standard deviation are even lower than the parallel port solution for x86 machines due to tighter coupling between the processor core and the GPIO.

We demonstrate Loupe's capabilities by using it to modify FFTW, a performance-focused FFT auto-tuner, to measure power for the transforms it runs as well. FFTW can decompose large, multidimensional transforms into combinations of many different algorithms, data layouts, loop nest orders, and so on, and is thus fertile ground for testing power-performance tradeoffs at a software level, since it provides many different ways of performing the same computation.

To ensure that all FFT variants were tested under the same conditions, we turned off DVFS and kept all 4 A7 cores at 1.2GHz, 1.2375V nominal. Precise performance data was gathered using the ARM cycle counter; we used a kernel module to allow userspace access to the appropriate registers. We modified FFTW version 3.3.3 to toggle the GPIO before and after each

timed run. We found through experimentation that power measurements converged to a high degree of repeatability and stability when runs took at least 10000 processor cycles (8.3μs), just over one cycle of the buck regulator powering the A7s. Clear features can be seen in the raw data at a much finer granularity, but nonlinearities arising from the regulator itself and the power delivery network require slightly longer to average out. Further experimentation is needed to develop a robust procedure for obtaining high-confidence measurements of even shorter runs. To reduce interference effects from other processes and the operating system, we run each FFT variant 8 times and use the minimum runtime and minimum measured energy.

CHAPTER 6

RESULTS AND DISCUSSION

Analog Bandwidth   Figure 6.1 shows the benefit of Loupe's high analog bandwidth; it is able to detect short-lived events that lower-bandwidth systems would miss entirely, and can perform other measurements much faster than existing apparata.

FFTW   We used the FFTW planner to perform 1.32 million runs spanning thousands of different FFT problems in about 5 minutes. The problems include real and complex, in-place and out-of-place, and sizes from 1 up to 4.2 million points. The null hypothesis when considering energy-specific optimizations is that optimizing for high performance is the same as optimizing for low energy; that is, static power and baseline dynamic power are great enough to override whatever dynamic power differences exist between multiple strategies to perform the same computation, and the most energy efficient configuration will be the one that finishes first (and thus incurs the least of this fixed cost). If this null hypothesis were true universally, there would be no need for energy-focused optimizations or tools; performance-centric compilers would already be optimal.

Of the 5230 problems where more than one solution was attempted, many indeed supported the null hypothesis; when plotting energy vs. runtime for all solutions attempted, a single solution, the most performant one, would lie on the pareto-optimal frontier, and the frontier would be said to be degenerate. However, a substantial fraction, 433/5230 or 8.3%, yielded interesting energy-performance tradeoffs and a pareto frontier with multiple options. For example, Figure 6.2 shows a configuration with 6 pareto-optimal solutions. These results include the static and baseline dynamic power of the other 3 A7 cores which remain idle during the benchmark but are powered off the same rail as the active core. Subtracting $\frac{3}{4}$ of 30.047mA, the measured median baseline power, from all power results to more fairly reflect the
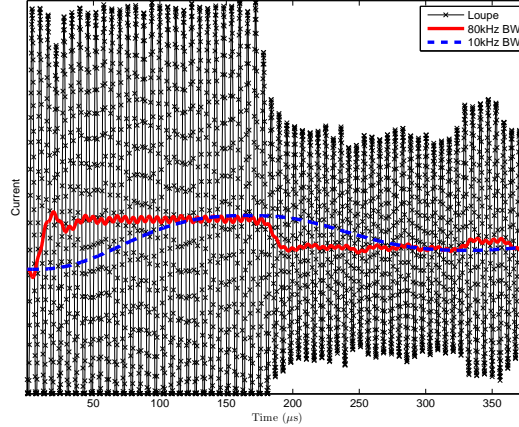
Figure 6.1: Real downsampled 8MSPS current data from the ODROID-XU mobile SoC board and simulated data from existing measurement approaches from the literature with 3dB bandwidths of 10kHz and 80kHz. Loupe captures many samples per cycle of the voltage regulator. Loupe's 10MHz analog bandwidth enables higher accuracy and a much shorter minimum run time.

power consumed on the active core yields 476/5230 problems with interesting energy-performance tradeoffs, or 9.08%. Such a high fraction of problems with interesting energy-performance tradeoffs suggests that there is substantial opportunity for a new generation of tools to optimize software for any weighted combination of energy and performance on modern hardware.
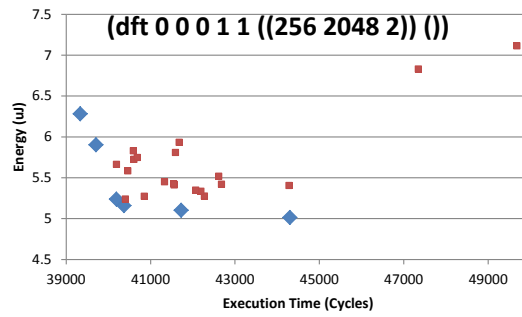
(dft 0 0 0 1 1 ((256 2048 2)) ())

Figure 6.2: Scatter plot of energy consumption vs. execution time for a 256-point complex out-of-place FFT. Of the 27 ways we tried to execute this FFT, the 6 highlighted in blue are on the pareto-optimal frontier. The most energy efficient option is 12.6% slower but uses 20.2% less energy than the fastest option. An energy-focused autotuner could choose among these 6 options according to the user's preference for high performance or energy efficiency.

CHAPTER 7

CONCLUDING REMARKS

In this dissertation, we motivated a faster, more accurate, more rigorously developed power measurement platform. We presented Jouler's Loupe, a power measurement system addressing the shortcomings of previous techniques. Finally, we showed that Loupe can indeed be used to gather power-performance data very quickly on real hardware, and that there exists considerable opportunity for tools, enabled by Loupe, to optimize software for energy efficiency rather than performance.

APPENDIX A

POWER SENSOR DESIGN

A.1 Power Sensor Schematic

A.1.1 PCIe Power Sensor

Figures A.1 – A.5 show the schematics for the three-channel PCIe 3.0 form factor power sensor. All passives are 0603 or larger to facilitate manual assembly. Fewer than three channels can be populated if necessary; indeed, we used a PCIe form factor board as the prototype for a two-channel sensor for a mobile SoC. The gain and offset of the amplifiers can be changed within a wide range, to accommodate systems with widely varying voltage and current ranges of interest, by varying passive component values.

Mobile CPU Power Sensor

The single channel power sensor uses the same power supply configuration as the PCIe power sensor, as depicted in Figure A.2. The amplifiers use nearly the same topology other than the small enhancements listed in Section 4.2.3, but have different passive component values due to their different voltage and current measurement ranges. The amplifier schematics are shown in Figures A.6–A.7.

A.2 Amplifier Error Analysis

In this section, we list the equations used to model the sources of error or uncertainty in the power measurements obtained by our power sensor. These equations model the most relevant specifications of the passive components
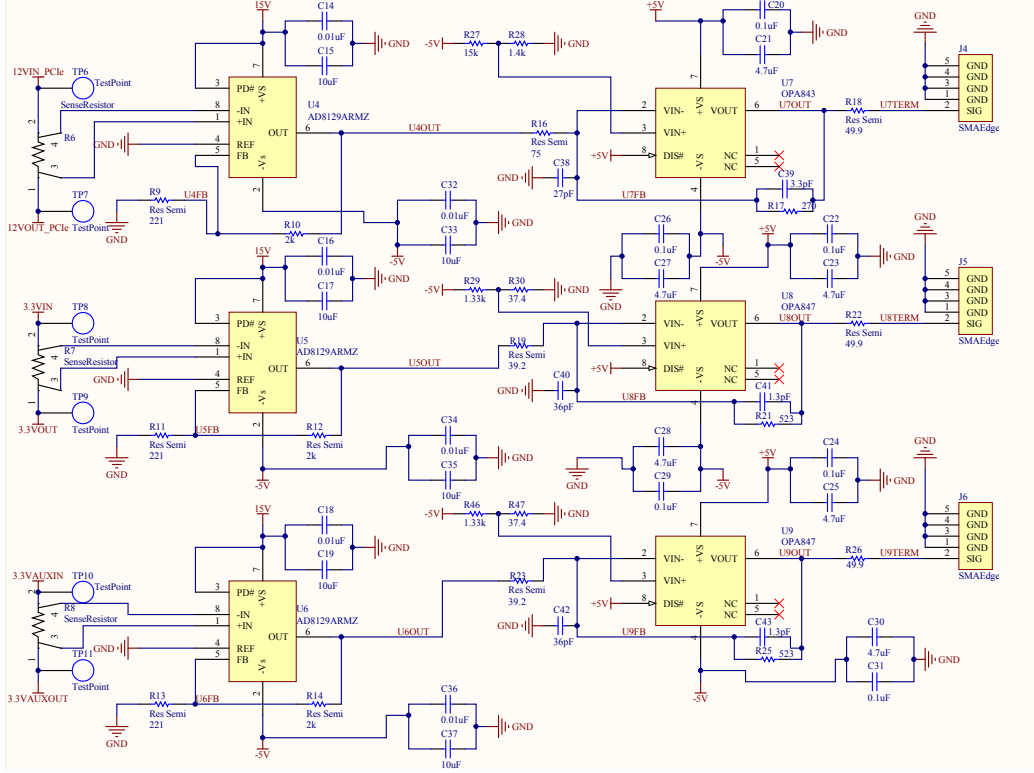
Figure A.1: Current sense resistors and current amplifiers.

and integrated circuits that influence current and voltage measurements. We show equations for a two-stage amplifier, as in the current channel of our power sensors; the single-stage voltage channel is modeled in much the same way, omitting the equations pertaining to the sense resistor and stage 2 amplifier. Note that each of the quantities listed in the equations below may itself be represented as a distribution of values, rather than a single nominal value. In this way, we can capture error sources like resistor value tolerances in our analysis by modeling them implicitly.

A.2.1 Sense Resistor

The following equations compute the actual resistance value R_s and differential voltage V_{R_s} of the sense resistor based on the resistor's nominal value R_{sn}, load current I, TCR TCR_{R_s}, thermoelectric voltage V_{t,R_s}, and the temperature differences from the resistor to nominal ambient temperature (usually $20°C$ or $25°C$) $(\Delta T_{R_s,A})$ and between the two terminals of the resistor
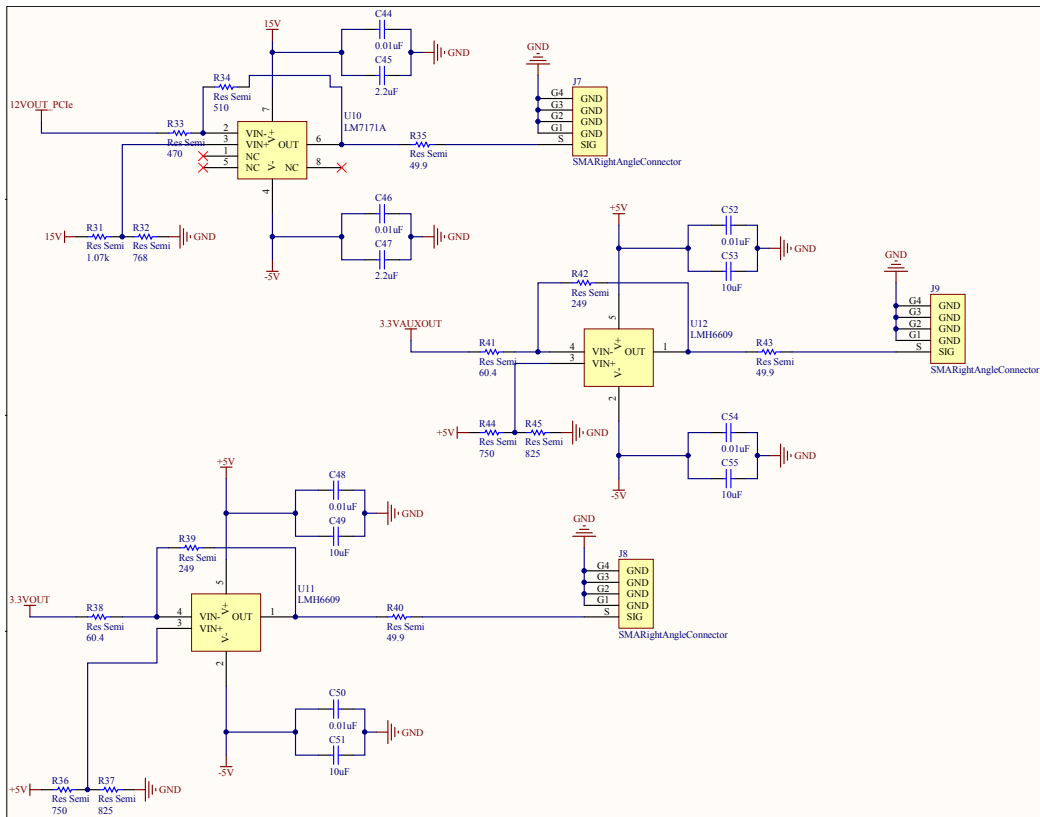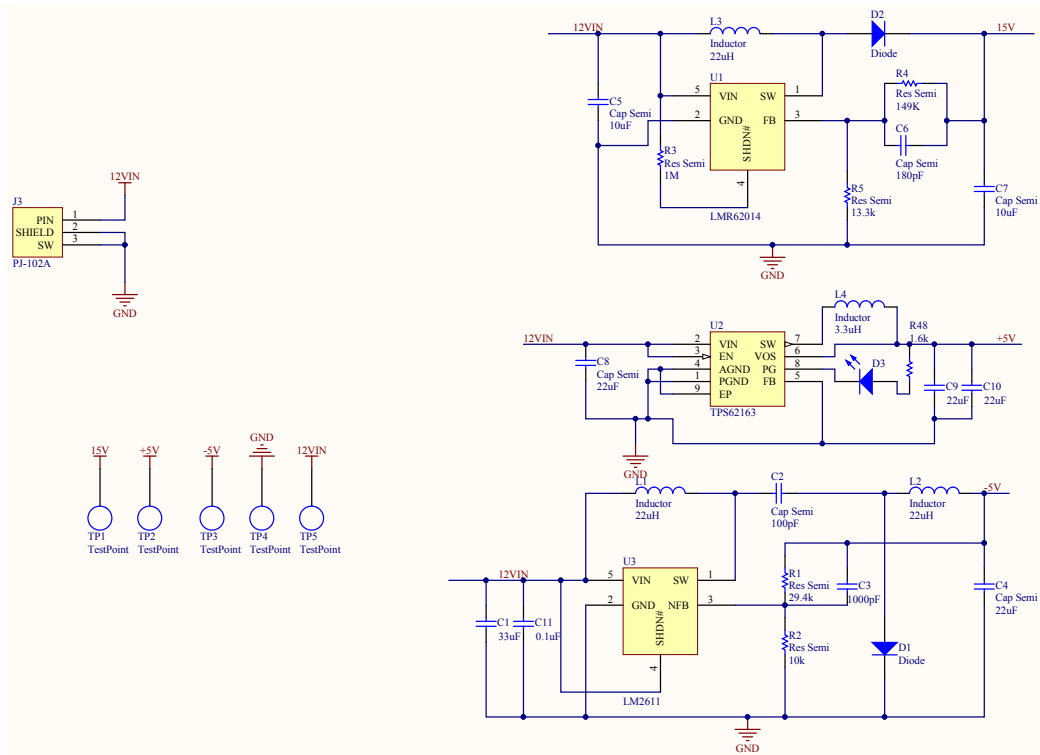
Figure A.2: Voltage amplifiers.
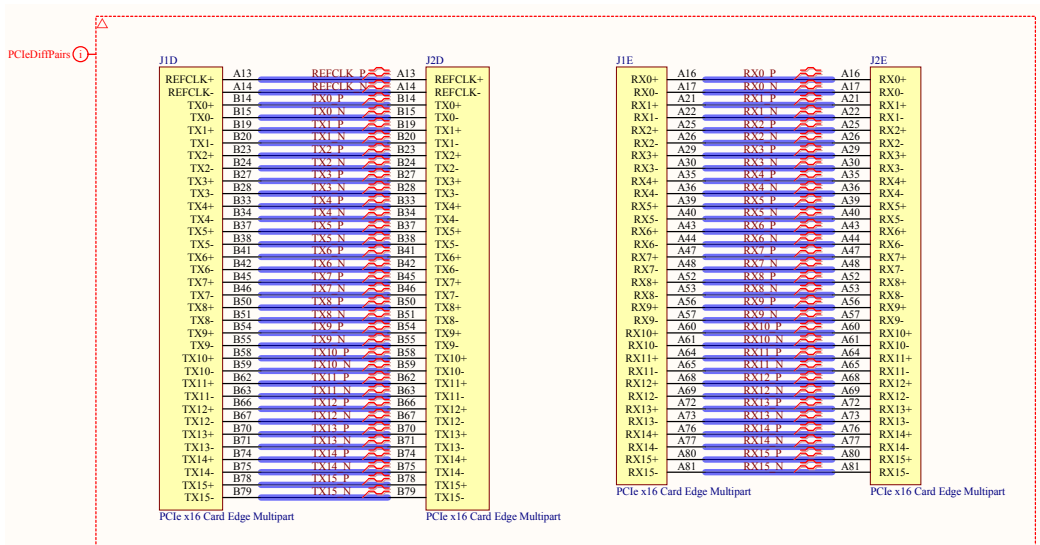
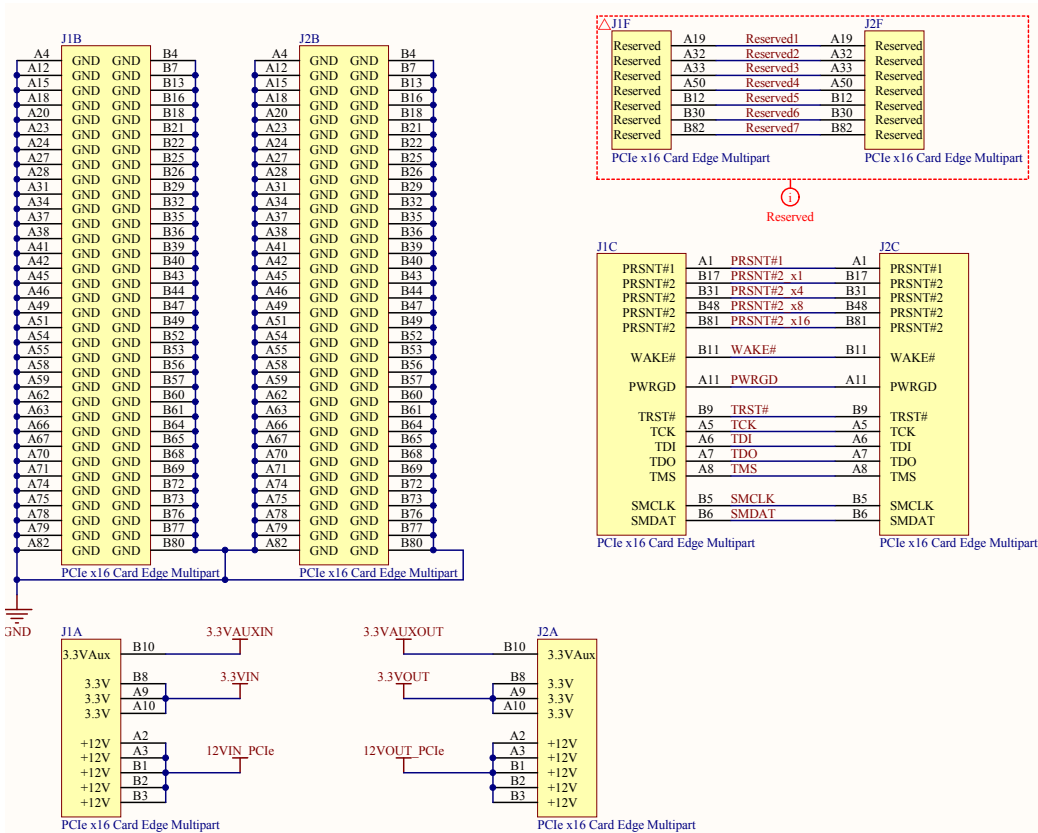Figure A.3: Voltage regulators.

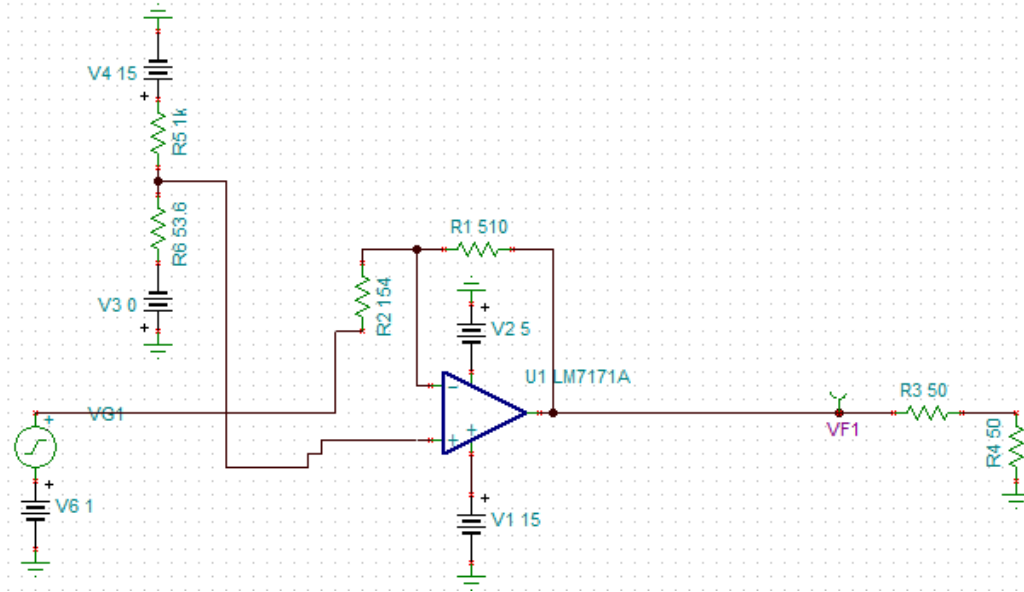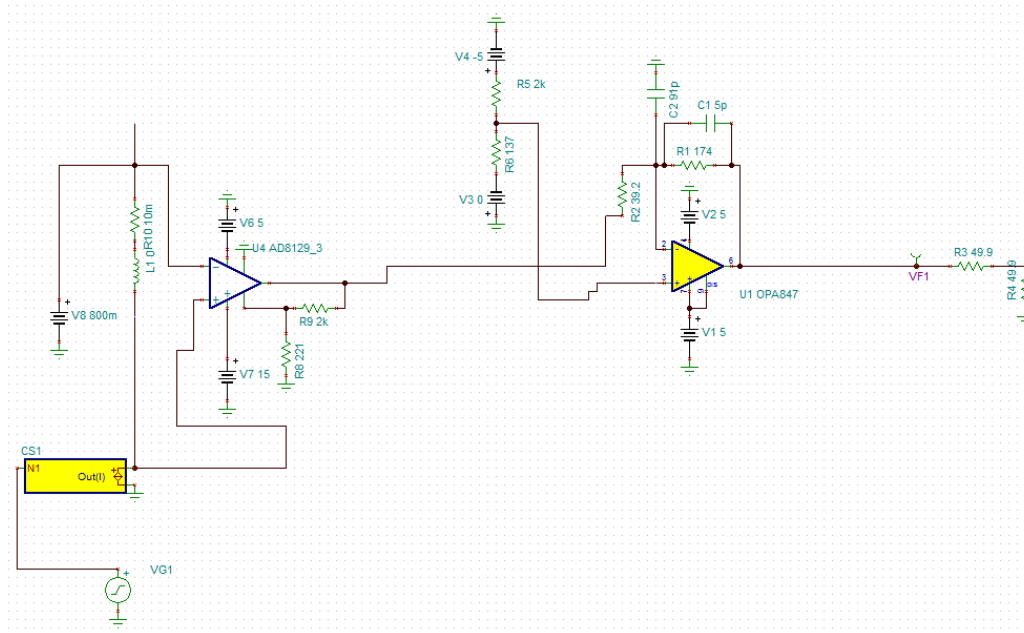Figure A.4: PCIe differential pairs.
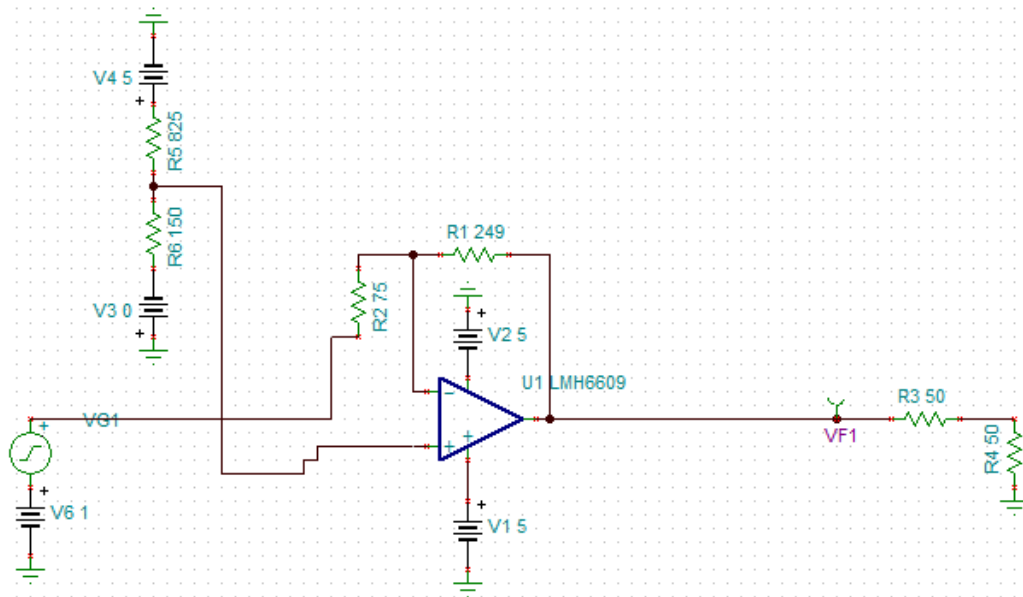
Figure A.5: Miscellaneous PCIe signals.
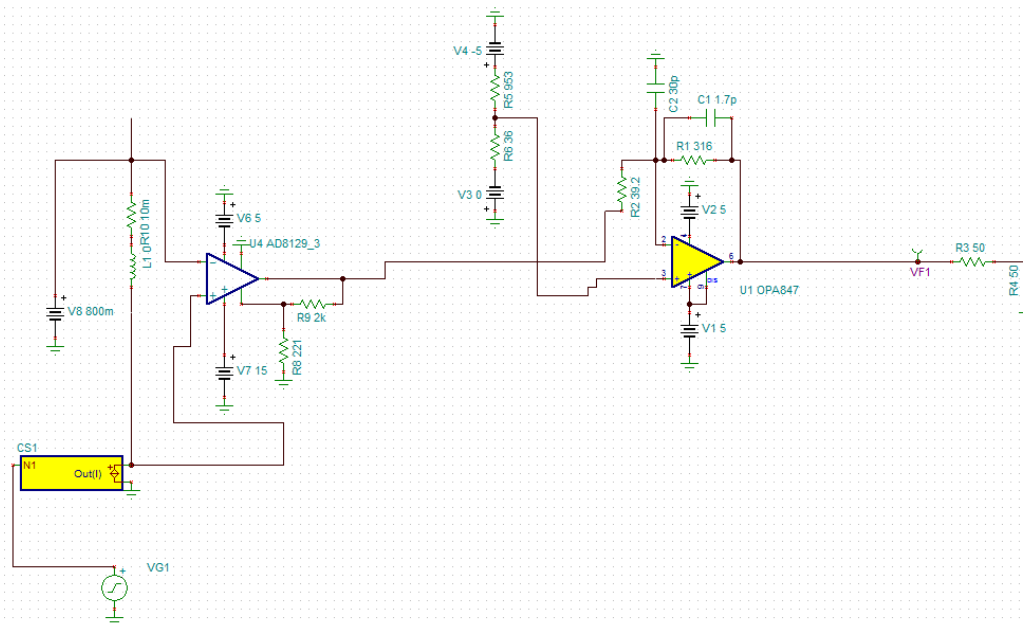
(a) Voltage Amplifier

(b) Current Amplifiers

Figure A.6: 0–9A channel amplifier schematics.

(a) 0–5A Voltage Amplifier

(b) 0–5A Current Amplifiers

Figure A.7: 0–5A channel amplifier schematics.

$(\Delta T_{R_s,S})$:

$$R_s = R_{sn} \times (1 + \Delta T_{R_s,A} \times TCR_{R_s}) \tag{A.1}$$

$$V_{R_s} = I \times R_s + \Delta T_{R_s,S} \times V_{t,R_s} \tag{A.2}$$

A.2.2 Stage 1 Amplifier

The stage 1 amplifier is a non-inverting amplifier whose gain is set by a feedback resistor R_{s1f} and gain resistor R_{s1g}. We model value tolerances, temperature difference from ambient, and TCR of the resistors. We also model the effects of the amplifier's CMRR, PSRR, offset voltage, and offset voltage temperature drift on the ultimate input voltage being amplified. Finally, we model the amplifier's DC gain error, gain error temperature drift, and gain nonlinearity.

$$R_{s1f} = R_{s1fn} \times (1 + \Delta T_{R_{s1f},A} \times TCR_{R_{s1f}}) \tag{A.3}$$

$$R_{s1g} = R_{s1gn} \times (1 + \Delta T_{R_{s1g},A} \times TCR_{R_{s1g}}) \tag{A.4}$$

$$V_{cm,Stage1} = (V_{Rail} + (V_{Rail} - V_{R_s}))/2 \tag{A.5}$$

$$\begin{aligned} V_{in,Stage1} = &- V_{R_s} + V_{os,Stage1} + TC_{V_{os,Stage1}} \times \Delta T_{Stage1} \\ &+ \frac{V_{cm,Stage1}}{CMRR_{Stage1}} + \frac{Ripple_{V_{dd},Stage1}}{PSRR_{Stage1}} \end{aligned} \tag{A.6}$$

$$\begin{aligned} Gain_{Stage1} = &(1 + \frac{R_{s1f}}{R_{s1g}}) \times (1 + DCGainError_{Stage1} \\ &+ \Delta T_{Stage1} \times TC_{Stage1,GainError} + GainNonlinearity_{Stage1}) \end{aligned} \tag{A.7}$$

$$V_{out,Stage1} = V_{in,Stage1} \times Gain_{Stage1} \tag{A.8}$$

A.2.3 Stage 2 Amplifier and Transmission Line Termination Resistors

The stage 2 amplifier is an inverting amplifier whose gain is set by a feedback resistor R_{s2f} and gain resistor R_{s2g}. The amplifier's offset is set by a voltage divider of two resistors R_{s2o1} and R_{s2o2}. We model value tolerances, temperature difference from ambient, and TCR of all four resistors. We also model the effects of the amplifier's bias currents (I_{b-} and I_{b+} for the inverting and

noninverting inputs, respectively), CMRR, PSRR, offset voltage, and offset voltage temperature drift on the ultimate input voltage being amplified.

$$R_{s2f} = R_{s2fn} \times (1 + \Delta T_{R_{s2f},A} \times TCR_{R_{s2f}}) \tag{A.9}$$

$$R_{s2g} = R_{s2gn} \times (1 + \Delta T_{R_{s2g},A} \times TCR_{R_{s2g}}) \tag{A.10}$$

$$R_{s2o1} = R_{s2o1n} \times (1 + \Delta T_{R_{s2o1},A} \times TCR_{R_{s2o1}}) \tag{A.11}$$

$$R_{s2o2} = R_{s2o2n} \times (1 + \Delta T_{R_{s2o2},A} \times TCR_{R_{s2o2}}) \tag{A.12}$$

$$I_{b+,Stage2} = I_{b,Stage2} + \frac{I_{offset,Stage2}}{2} \tag{A.13}$$

$$I_{b-,Stage2} = I_{b,Stage2} - \frac{I_{offset,Stage2}}{2} \tag{A.14}$$

$$V_{I_{b+,Stage2}} = \frac{I_{b+,Stage2}}{\frac{1}{R_{s2o1}} + \frac{1}{R_{s2o2}}} \tag{A.15}$$

$$V_{I_{b-,Stage2}} = \frac{I_{b-,Stage2}}{\frac{1}{R_{s2f}} + \frac{1}{R_{s2g}}} \tag{A.16}$$

$$V_{VoltageDivider} = V_{dd,VoltageDivider} \times \frac{R_{s2o1}}{R_{s2o1} + R_{s2o2}} \tag{A.17}$$

$$V_{cm,Stage2} = \frac{V_{VoltageDivider} + V_{out,Stage1}}{2} \tag{A.18}$$

$$\begin{aligned} V_{in,Stage2} = & V_{out,Stage1} - V_{VoltageDivider} + Vos, Stage2 \\ & + TC_{V_{os,Stage2}} \times \Delta T_{Stage2} + \frac{V_{cm,Stage2}}{CMRR_{Stage2}} \\ & + \frac{Ripple_{V_{dd},Stage2}}{PSRR_{Stage2}} + (V_{I_{b+,Stage2}} - V_{I_{b-,Stage2}}) \end{aligned} \tag{A.19}$$

$$Gain_{Stage2} = -\frac{R_{s2f}}{R_{s2g}} \tag{A.20}$$

$$V_{out,Stage2} = V_{VoltageDivider} + (V_{in,Stage2} \times Gain_{Stage2}) \tag{A.21}$$

A.2.4 Transmission Line Termination Resistors and ADC Buffer

The stage 2 amplifier is connected to the ADC buffer via a transmission line, with series and parallel termination resistors R_{t1} and R_{t2} on the amplifier and ADC buffer side, respectively. Since the transmission line is properly terminated and the frequencies and currents of interest do not result in significant attenuation over short coaxial cables, we do not model the cable itself. The ADC buffer is a single-ended to differential amplifier whose gain is set by

two pairs of resistors, R_{abf-} and R_{abg-} on the negative input and R_{abf+} and R_{abg+} on the positive input. We model value tolerances, temperature difference from ambient, and TCR of all six resistors. We model the effects of varying resistor values on the buffer's gain and the differential output voltage presented to the ADC.

$$R_{t1} = R_{tn} \times (1 + \Delta T_{R_{t1},A} \times TCR_{R_{t1}}) \tag{A.22}$$

$$R_{t2} = R_{tn} \times (1 + \Delta T_{R_{t2},A} \times TCR_{R_{t2}}) \tag{A.23}$$

$$V_{in,ab} = V_{out,Stage2} \times \frac{R_{t2}}{R_{t1} + R_{t2}} \tag{A.24}$$

$$R_{abg+} = R_{abg+n} \times (1 + \Delta T_{R_{abg+},A} \times TC_{R_{abg+}}) \tag{A.25}$$

$$R_{abg-} = R_{abg-n} \times (1 + \Delta T_{R_{abg-},A} \times TC_{R_{abg-}}) \tag{A.26}$$

$$R_{abf+} = R_{abf+n} \times (1 + \Delta T_{R_{abf+},A} \times TC_{R_{abf+}}) \tag{A.27}$$

$$R_{abf-} = R_{abf-n} \times (1 + \Delta T_{R_{abf-},A} \times TC_{R_{abf-}}) \tag{A.28}$$

$$\beta_{ab1} = \frac{R_{abg+}}{R_{abg+} + R_{abf+}} \tag{A.29}$$

$$\beta_{ab2} = \frac{R_{abg-}}{R_{abg-} + R_{abf-}} \tag{A.30}$$

$$Gain_{ab} = 2 \times \frac{1 - \beta_{ab1}}{\beta_{ab1} + \beta_{ab2}} \tag{A.31}$$

$$V_{out,ab} = V_{in,ab} \times Gain_{ab} \tag{A.32}$$

A.2.5 ADC

Sample Jitter

In a real uniform sampling ADC, the sampling instants will deviate slightly from their ideal positions at $t = N \times \frac{1}{f_s}$. This deviation is known as sample jitter, and has two main components: the jitter in the triggering edges of the external ADC sample clock, and the jitter in the internal delay between the triggering clock edge and the sample capture. These two components are known as clock jitter and aperture delay jitter, respectively. The modeling

equations used in our uncertainty analysis are as follows:

$$\frac{dV_{in,ADC}}{dt} = \sqrt{2} \times \pi \times f_{in} \times Amplitude_{in} \tag{A.33}$$

$$t_{SampleJitter} = t_{ClockJitter} + t_{ApertureDelayJitter} \tag{A.34}$$

$$V_{in,ADCWithJitter} = V_{in,ADC} + \frac{dV_{in,ADC}}{dt} \times t_{SampleJitter} \tag{A.35}$$

$$\tag{A.36}$$

Note that we only model the aperture delay *jitter* and not the mean aperture delay itself, as this delay can be calibrated out in the full system along with the average delays of the amplifiers, coaxial cable, and processor activity signaling mechanism.

The ADC in the USRP1 used for this work is clocked using a temperature-compensated crystal oscillator (TCXO) with ≤1ps RMS phase jitter [50]. The ADC itself has 1.2ps RMS aperture delay [51], for a maximum root sum square (RSS) sample jitter of 1.56ps. The TCXO may be replaced by hand with a more expensive pin-compatible model to reduce clock jitter [52]. Its timing characteristics are expressed in terms of its phase noise spectrum. Translating these to a scalar jitter figure for uncertainty analysis [53] yields an RMS jitter of 0.347ps, reducing the total RSS sample jitter of 1.25ps. If sample jitter is small enough that the input waveform maps to the same ADC code at the ideal and actual sampling instants, the jitter has no practical effect. For an N-bit ADC and a full-scale input sinusoid with frequency f_{in}, the jitter must be less than $\frac{1}{\pi \times f_{in} \times 2^{N+1}}$ to cause a sampled voltage error of less than $\frac{1}{2}$ LSB in the worst case [54]. For a 10MHz sinusoid matching the input bandwidth used in our prototype power sensor and a 12-bit ADC, the critical jitter value is 3.886ps. Assuming a normal jitter distribution, the actual jitter will exceed the critical value 1.29% of the time using the stock TCXO, and 0.19% of the time using the upgraded TCXO. Note that a full-scale 10MHz sinusoid is the worst case for slew rate and thus jitter-induced error, and jitter is low enough using the stock TCXO that jitter will seldom cause a different ADC code to be read.

The slew rate can either be modeled using empirical data for a particular system, or by modeling the analog waveform as an ensemble of one or more sinusoids with given frequencies and amplitudes. For a given sinusoid with amplitude A and frequency f, the sinusoid is described by $f(t) = A \times sin(2 \times$

$\pi \times f \times t$) and its slew rate by $\frac{df(t)}{dt} = 2 \times \pi \times f \times A \times cos(2 \times \pi \times f \times t)$. The distribution of slew rates over a period is U-shaped, with a mean value of 0, a mean absolute value of $4f$, and peaks at $\frac{df(t)}{dt} = \pm(2 \times \pi \times f \times A)$. The standard uncertainty of this distribution is $\sqrt{2} \times \pi \times f \times A$; GUM Workbench and other uncertainty analysis software can handle U-shaped distributions directly.

We model the ADC's input noise, quantization error, and INL. We also model the effect of sample jitter on the digitized signal by multiplying the jitter by the slew rate of the signal in the neighborhood of the sample instant. Note that we only model aperture delay jitter, not the mean aperture delay value, because constant aperture delay can be calibrated out in the full system along with the average delays of the amplifiers, coaxial cable, and processor activity signaling mechanism.

$$V_{in,ADC} = V_{out,ab} + InputNoise_{ADC} \tag{A.37}$$

$$V_{LSB} = \frac{V_{InputRange,ADC}}{2^{N_{bits,ADC}}} \tag{A.38}$$

$$V_{QuantizationError,ADC} = N_{LSB_{QuantizationError}} \times V_{LSB} \tag{A.39}$$

$$V_{INL,ADC} = N_{LSB_{INL}} \times V_{LSB} \tag{A.40}$$

$$V_{quantized,ADC} = V_{in,ADCWithJitter} + V_{INL,ADC} + V_{QuantizationError,ADC} \tag{A.41}$$

REFERENCES

[1] "Fundamentals of RF and Microwave Power Measurements," Agilent Technologies, Tech. Rep. 64-1B, 2000.

[2] W. Koon, "Current sensing for energy metering," in *Proceedings of IIC-China*, 2002, pp. 321–324.

[3] S. Svensson, "Power measurement techniques for non-sinusoidal conditions," Ph.D. dissertation, Chalmers University of Technology, 1999.

[4] N. Chang, K. Kim, and H. G. Lee, "Cycle-accurate energy consumption measurement and analysis: Case study of arm7tdmi," in *Low Power Electronics and Design, 2000. ISLPED '00. Proceedings of the 2000 International Symposium on*, 2000, pp. 185–190.

[5] H. G. Lee, K. Lee, Y. Choi, and N. Chang, "Cycle-accurate energy measurement and characterization of fpgas," *Analog Integr. Circuits Signal Process.*, vol. 42, no. 3, pp. 239–251, Mar. 2005. [Online]. Available: http://dx.doi.org/10.1007/s10470-005-6758-5

[6] X. Zhou, X. Peng, and F. Lee, "A high power density, high efficiency and fast transient voltage regulator module with a novel current sensing and current sharing technique," in *Applied Power Electronics Conference and Exposition, 1999. APEC '99. Fourteenth Annual*, vol. 1, Mar 1999, pp. 289–294 vol.1.

[7] D. Goder, "System to protect switch mode DC/DC converters against overload current," Patent US 6 127 814 A, 10 03, 2000. [Online]. Available: http://www.google.com/patents/US6127814

[8] S. Ziegler, H.-C. Iu, R. Woodward, and L. Borle, "Theoretical and practical analysis of a current sensing principle that exploits the resistance of the copper trace," in *Power Electronics Specialists Conference, 2008. PESC 2008. IEEE*, June 2008, pp. 4790–4796.

[9] A. Ikriannikov and O. Djekic, "Investigation of DCR Current Sensing in Multiphase Voltage Regulators," in *IBM Power and Cooling Technology Symposium 2007*, 2007.

[10] "SPICE Model – 0603PS," Coilcraft, Inc., Tech. Rep. 267-1, August 2007.

[11] *P4400 Kill A Watt Operation Manual*, P3 International.

[12] *Watts Up? Power Analyzer, Watt Meter, and Electricity Monitor*, Electronic Educational Devices, February 2008.

[13] D. Bedard, M. Lim, R. Fowler, and A. Porterfield, "Powermon: Fine-grained and integrated power monitoring for commodity computer systems," in *IEEE SoutheastCon 2010 (SoutheastCon), Proceedings of the*, 2010, pp. 479–484.

[14] A. Weissel and F. Bellosa, "Process cruise control: event-driven clock scaling for dynamic power management," in *Proceedings of the 2002 international conference on Compilers, architecture, and synthesis for embedded systems*, ser. CASES '02. New York, NY, USA: ACM, 2002. [Online]. Available: http://doi.acm.org/10.1145/581630.581668 pp. 238–246.

[15] K. Rajamani, H. Hanson, J. Rubio, S. Ghiasi, and F. Rawson, "Application-aware power management," in *Workload Characterization, 2006 IEEE International Symposium on*, 2006, pp. 39–48.

[16] Y. Zhu and V. J. Reddi, "High-performance and energy-efficient mobile web browsing on big/little systems," in *Proceedings of the 2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, ser. HPCA '13. Washington, DC, USA: IEEE Computer Society, 2013. [Online]. Available: http://dx.doi.org/10.1109/HPCA.2013.6522303 pp. 13–24.

[17] A. Carroll and G. Heiser, "An analysis of power consumption in a smartphone," in *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*, ser. USENIXATC'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 21–21.

[18] G. Contreras and M. Martonosi, "Power prediction for intel xscale processors using performance monitoring unit events," in *Proceedings of the 2005 international symposium on Low power electronics and design*, ser. ISLPED '05. New York, NY, USA: ACM, 2005. [Online]. Available: http://doi.acm.org/10.1145/1077603.1077657 pp. 221–226.

[19] M. S. S. Govindan, S. W. Keckler, and D. Burger, "End-to-end validation of architectural power models," in *Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design*, ser. ISLPED '09. New York, NY, USA: ACM, 2009, pp. 383–388.

[20] H. Esmaeilzadeh, T. Cao, Y. Xi, S. M. Blackburn, and K. S. McKinley, "Looking back on the language and hardware revolutions: measured power, performance, and scaling," in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems*, ser. ASPLOS XVI. New York, NY, USA: ACM, 2011, pp. 319–332.

[21] T. Do, S. Rawshdeh, and W. Shi, "pTop: A Process-level Power Profiling Tool," in *HotPower '09: Proceedings of the Workshop on Power Aware Computing and Systems.* New York, NY, USA: ACM, Oct 2009.

[22] F. Barany, "System Bandwidth vs. Resolution for Analog Video," Analog Devices, Tech. Rep. AN-945, 2007.

[23] J. Leng, T. Hetherington, A. ElTantawy, S. Gilani, N. S. Kim, T. M. Aamodt, and V. J. Reddi, "Gpuwattch: Enabling energy optimizations in gpgpus," in *Proceedings of the 40th Annual International Symposium on Computer Architecture*, ser. ISCA '13. New York, NY, USA: ACM, 2013. [Online]. Available: http://doi.acm.org/10.1145/2485922.2485964 pp. 487–498.

[24] W. Jung, C. Kang, C. Yoon, D. Kim, and H. Cha, "Devscope: A nonintrusive and online power analysis tool for smartphone hardware components," in *Proceedings of the Eighth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, ser. CODES+ISSS '12. New York, NY, USA: ACM, 2012. [Online]. Available: http://doi.acm.org/10.1145/2380445.2380502 pp. 353–362.

[25] Monsoon Solutions, Inc., "Mobile device power monitor manual, version 1.14," May 2014.

[26] "Using ARM Streamline," ARM, Tech. Rep. DUI 0482H, May 2012.

[27] M. S. Gupta, J. L. Oatley, R. Joseph, G.-Y. Wei, and D. M. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *Proceedings of the Conference on Design, Automation and Test in Europe*, ser. DATE '07. San Jose, CA, USA: EDA Consortium, 2007. [Online]. Available: http://dl.acm.org/citation.cfm?id=1266366.1266498 pp. 624–629.

[28] K. Shringarpure, S. Pan, J. Kim, B. Achkir, B. Archambeault, J. Fan, and J. Drewniak, "Innovative pdn design guidelines for practical high layer-count pcbs," in *Proceedings of DesignCon 2013*, ser. DesignCon '13, 2013.

[29] M. Jones, "Simulating power planes with LTSpice IV," April 2013. [Online]. Available: http://www.linear.com/solutions/1810

[30] K. R. Carver and J. Mink, "Microstrip antenna technology," *Antennas and Propagation, IEEE Transactions on*, vol. 29, no. 1, pp. 2–24, Jan 1981.

[31] "Printed circuit board (PCB) power delivery network (PDN) design methodology," Altera Corporation, Tech. Rep. AN 574, May 2009.

[32] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '11. New York, NY, USA: ACM, 2011. [Online]. Available: http://doi.acm.org/10.1145/2063384.2063454 pp. 52:1–52:12.

[33] H. Patil and T. E. Carlson, "Pinballs: Portable and shareable user-level checkpoints for reproducible analysis and simulation," in *2014 Workshop on Reproducible Research Methodologies (REPRODUCE 2014)*, February 2014.

[34] E. Perelman, G. Hamerly, M. Van Biesbrouck, T. Sherwood, and B. Calder, "Using simpoint for accurate and efficient simulation," in *Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '03. New York, NY, USA: ACM, 2003. [Online]. Available: http://doi.acm.org/10.1145/781027.781076 pp. 318–319.

[35] S. Li, J. Ahn, J. B. Brockman, and N. P. Jouppi, "McPAT 1.0: An integrated power, area, and timing modeling framework for multicore architectures," HP Labs, Tech. Rep., 2009.

[36] Analog Modules Inc., "The quest for accurate current sensing," http://www.omnipulsetechnology.com/products/pdfs/The%20Quest%20for%20Accurate%20Current%20Sensing.pdf, Tech. Rep., January 2010.

[37] F. Zandman and J. Szwarc, "Non-Linearity of Resistance/Temperature Characteristic: Its Influence on Performance of Precision Resistors," Vishay Precision Group, Tech. Rep. 108, February 2013.

[38] K. Blake, "Op Amp Precision Design: PCB Layout Techniques," Microchip Technology Inc., Tech. Rep. AN1258, 2009.

[39] E. Balestrieri, P. Daponte, and S. Rapuano, "A state of the art on adc error compensation methods," *Instrumentation and Measurement, IEEE Transactions on*, vol. 54, no. 4, pp. 1388–1394, Aug 2005.

[40] M. Mishali and Y. Eldar, "Blind multiband signal reconstruction: Compressed sensing for analog signals," *Signal Processing, IEEE Transactions on*, vol. 57, no. 3, pp. 993–1009, March 2009.

[41] G. P. Lepage, "A New Algorithm for Adaptive Multidimensional Integration," *Journal of Computational Physics*, vol. 27, p. 192, May 1978.

[42] F. Papenfuss, Y. Artyukh, E. Boole, and D. Timmermann, "Nonuniform sampling driver design for optimal adc utilization," in *Circuits and Systems, 2003. ISCAS '03. Proceedings of the 2003 International Symposium on*, vol. 4, May 2003, pp. IV–516–IV–519 vol.4.

[43] Analog Devices Inc., *Op Amp Bandwidth and Bandwidth Flatness*, ser. Tutorial MT-045, October 2008.

[44] "Noise analysis in operational amplifier circuits," Texas Instruments, Tech. Rep. SLVA043B, 2007.

[45] L. Schuchman, "Dither signals and their effect on quantization noise," *Communication Technology, IEEE Transactions on*, vol. 12, no. 4, pp. 162–165, December 1964.

[46] L. Melkonian, "Improving A/D Converter Performance using Dither," National Semiconductor Corporation, Tech. Rep. 804, February 1992.

[47] T. Laakso, V. Valimaki, M. Karjalainen, and U. Laine, "Splitting the unit delay [fir/all pass filters design]," *Signal Processing Magazine, IEEE*, vol. 13, no. 1, pp. 30–60, Jan 1996.

[48] Z. Nakutis, "A current consumption measurement approach for fpga-based embedded systems," *Instrumentation and Measurement, IEEE Transactions on*, vol. 62, no. 5, pp. 1130–1137, May 2013.

[49] W. C. Losinger, "A review of the GUM workbench," *The American Statistician*, vol. 58, no. 2, pp. 165–167, 2004.

[50] *Part Number Data Sheet*, Ecliptek Corporation, 6 2012, revision V.

[51] *Mixed-Signal Front-End (MxFE^{TM}) for Broadband Communications*, Analog Devices, 2002, revision 0.

[52] *CMOS 7x5x2.5mm SMD, 'V' Group*, Euroquartz Limited.

[53] "Clock (CLK) Jitter and Phase Noise Conversion," Maxim Integrated, Tech. Rep. APP 3359, December 2004.

[54] "Aperture Jitter Calculator for ADCs," Maxim Integrated, Tech. Rep. APP 4466, September 2009.