Power and Actuation Autonomy for Flapping Wing Insect-scale Robots

Johannes James

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2024

Reading Committee:

Sawyer B. Fuller, Chair

Joshua Smith

Santosh Devasia

Program Authorized to Offer Degree:
Mechanical Engineering

University of Washington

Abstract

Power and Actuation Autonomy for Flapping Wing Insect-scale Robots

Johannes James

Chair of the Supervisory Committee:
Sawyer B. Fuller
Mechanical Engineering

This work aims to advance the field of Flying Insect-scale Robots (FIRs) by presenting solutions for key outstanding problems in wireless power, control, and actuation autonomy. My work in [30] was the first ever wireless lift-off of an FIR powered by external laser, and subsequent work [31] has shown an additional wireless power solution using magnetically coupled resonators to be more than adequate for FIRs and also human safe. In [29] I was the first to prove that the onboard PEU can linearly modulate the thrust required for roll and altitude control, which was an important prerequisite for flight control. Further, the work in [29] also demonstrates that an unaided repetitive control law can learn to regulate output waveforms to reference with zero phase within the span of a few wing strokes, thus providing a valuable feedforward term to waveform generator control in a computationally lightweight form.

The primary aim of this body of work is to advance the technology needed for wireless operation of the FIRs. No discrete or off-the-shelf solution yet existed to drive two independent piezoelectric actuators typically selected for FIR flight muscles within the size, weight, and power (SWaP) budget of a robot weighing less than 200 mg. To-date, FIRs have relied on a physical tether to large specialized benchtop systems for power and control signals to the flight muscles and any onboard systems for experiments, thus severely restricting their utility. Power and actuation autonomy represents a key bottleneck in the development of this technology, yet for decades there has been relatively little advancement toward true autonomy. This work presents an exploration of state of the art solutions, and presents new power electronics and waveform generation control to enable such advancement. Specifically, the waveform generation and power electronics work together to result in a doubling of the effective power capacity of our FIRs by improving end-to-end efficiency over the previous state of the art, from 28% to 54% with approximately 75% the weight. This advance in waveform generation and power electronics is possibly the first to demonstrate independent wing control in an onboard PEU. This unlocks new frontiers in onboard flight control, sensing, and payload.

TABLE OF CONTENTS

LIST OF FIGURES

Figure Number Page

GLOSSARY

FIR: Flying Insect-scale Robot.

FWMAV: Flapping Wing Micro Aerial Vehicle.

MAV: Micro Aerial Vehicle. This term has been used to describe relatively large aerial vehicles, up to approximately a meter. Due to the need to differentiate such large vehicles from insect-scale, a number of acronyms and terms have arisen to designate bumble-bee sized robots. Therefore “MAV” is not a preferred term for FWMAV’s or FIR’s in this work, but appears in many references and thus requires explanation.

SWAP: Size, Weight, and Power. This acronym refers to the quantifiable needs and allowable budget for light weight, small size, and low power demanded of an FIR’s subsystems. A preceding ‘S’, as in SSWaP, indicates computational and actuation speed in addition to SWaP.

MOCAP: Motion Capture. In usage here, also refers to the physical space, ”the flight arena”, covered by the four cameras (Prime 13 cameras, OptiTrak, Inc., Salem, OR) which are positioned at the upper corners of the rectangular solid flight arena.

PIEZO: Due to the numerous occurrences of piezoelectric actuator in this work, “piezo” will often serve as shorthand to indicate piezoelectric actuators.

V_S: Many references in circuit schematics and text will refer to an electrical node denoted V_S which is the low-voltage source on the robot, typically between 3-7 volts. V_S can be supplied from multiple different onboard power sources.

V_B, V_{BIAS}: Refers to the DC bias voltage $V_b \approx 200 - 300$ volts, for the bimorph piezo actuators configured for 'simultaneous drive'.

V_{SIG1}, V_{CH1}: Time-varying between $V_S - V_b$ volts. Refers to the unipolar sinusoidal flapping signal for one of the two actuators in the typical 2-winged FIR of this work.

V_{SIG2}, V_{CH2}: Time-varying between $V_S - V_b$ volts. Refers to the unipolar sinusoidal flapping signal for the alternate one of the two actuators in the typical 2-winged FIR of this work.

BVD: Butterworth-Van Dyke model, a representation of a piezoelectric actuator with n attached resonant modes. This work uses the BVD model augmented to account for dielectric losses, but not hysteresis losses.

WPT: Wireless Power Transmission. In this work refers to power transmitted from an source external to the robot to the robot, without means of attachment by physical conductors.

MCR: Magnetically Coupled Resonators. Two or more electrical resonators which are coupled via magnetic waves, providing a means of wireless power transmission.

MRC: Magnetic Resonance Coupling. A means of wireless power transmission using magnetic coupling of physically disconnected resonators.

PV: Photo-voltaic.

LPB: Laser power beaming. Wireless power transfer using high intensity focused light source providing input power to the robot and onboard PV cell.

XPC TARGET: This is a desktop PC which has been repurposed to exclusively run MATLAB Simulink controllers in real time which are compiled by the host PC. It has a few important periphers such as the DAQ cards for ADC and DAC needed to generate the low voltage control signals which are amplified by the piezo amplifiers.

HOST PC: The host PC running a Windows OS which performs a few important tasks: 1) Using MATLAB compiles control models and programs the xPC target with such 2) interfaces with the MoCap cameras to perform pose estimation and sends these updates at 250 Hz to the xPC target.

PEU: Power Electronics Unit. The PEU is the electronics that take as input the power source, and output the electrical signals needed to drive the actuator(s).

PWM: Pulse-Width-Modulation. A method of modulating a continuous signal into a digital output by adjusting the period of pulses in a constant-frequency digital output, thereby changing the duty-cycle.

PFM: Pulse-Frequency-Modulation. A method of modulating a continuous signal into a digital output by adjusting the frequency of pulses while holding the pulse-width constant, thereby changing the duty-cycle.

LUT: Lookup-table. An n-dimensional array of values which can be indexed. e.g. in discrete time during a sequence of steps in a flapping cycle which have different prescribed values, or by interpolating a variable of interest within some range. For instance the sin function can be precomputed in an LUT and subsequently indexed by interpolation so as to reduce subsequent computation effort.

FET: Field Effect Transistor, either P-channel or N-channel and used extensively in the power electronics for switching roles in DC-DC converters and piezo drive electronics. Ubiquitous silicon based technology used in this work includes Metal-Oxide-Semiconductor FET (MOSFET), and GaNFETs (below).

GANFET: Gallium Nitride FET. Similar to an enhancement mode N-channel MOSFET but with superior charge mobility, lower input capacitance, and superior reverse recovery, making it more efficient for applications in this work than Si FETs.

BJT: Bipolar Junction Transistor. A semiconductor switching element either P-doped (PNP) or N-doped (NPN). In this work both PNP and NPN have applications similar to FETs serving in high-side or low-side (described below) switching roles, respectively.

MUX: Multiplexer. As used in "MUXed PEU" refers to the multiplexing (and also the de-multiplexing not differentiated for simplicity) of the inputs and outputs of the converter elements.

LFBD: Leaky Flyback gate Driver. The vital component which activates and deactives the high-side N-ch transistor switching elements which perform the MUX role in the new PEU design.

HIGH-SIDE: Refers to a switching element which is on the higher voltage side of a downstream load such as an actuator. This is the most intuitive placement of a switch for most people, e.g. most envision switching a light bulb on-and-off by switching the 'hot' wire.

LOW-SIDE: Refers to a switching element which switches the 'downstream' lower voltage side of an upstream load. Due to transistor design, N-channel FETs are more conveniently used in low-side application without additional means to drive them from some lower voltage such as a microcontroller.

ACKNOWLEDGMENTS

The author wishes to express his appreciation for his advisor, Professor Fuller, and advising committee Prof. J. Smith, Prof. S. Devasia, and Prof. J. Garbini for their support, knowledge, educational materials, and valuable insights at key points along the way. The author also thanks collaborators, past and present, particularly Yogesh Chukewad, Daksh Dhingra, Vikram Iyer, and Xingyi Shi. Lastly, the author thanks his wife Rose Hendrix and his human and inhuman family for their unwavering support through the many challenges overcome during this endeavor.

Chapter 1 Introduction

1.1 Motivation

Robots the size of common insects like a honeybee ($\sim$100 mg) are an unexplored technology compared to established classes of heavier aerial robots. Flying Insect-scale Robots (FIRs) have enormous potential advantages stemming from physics scaling effects of reduced principal dimensions. FIRs can potentially outperform larger aerial robots in tasks that benefit from the small size, economies of scale, or swarm capabilities enabled by large deployment numbers. Potential applications include sensing networks comprised by mobile nodes cabable of individual positioning and locomotion, gas leak detection, agricultural assistance, safe operation around humans without impact hazard, and exploration of hazardous or remote space. Although manufacture and system design is challenging and complex, analogous technological developments in consumer electronics suggest that economy of scale may be a disproportionately large advantage for such small robots. It is also important to develop such technology responsibly, so this work also aims to elucidate limitations and realities of the technology.

Since they were originally proposed as “gnat robots” in 1989 by Brooks and Flynn [3] and attempted in earnest by the Berkeley Micro Robotic Fly project starting in the early 2000s [15], progress toward truly autonomous insect scale robots has seen important milestones. These include the first lift greater than weight of a 100 mg robot [63], subsequent controlled flight [46], sensor integration [19], and expanded capabilities such as landing [22]. However, in the decade that has followed first liftoff, not one of these 100 mg robots has been able to fly *continuously* in a *controlled* fashion without tiny wires to control it. Foundational engineering work is needed to implement onboard those crucial tasks of flight control, sensing, telemetry, power conversion, and actuation.

This work presents key advancements towards “cutting the cord” and enabling true autonomous wireless flight. The main components of this are advancements in onboard power and developing a capable but lightweight signal generator to generate the precise voltage signals needed using onboard resources alone.

One variant of the Robofly is shown in Fig. 1.1a, and a more recent design ‘Robofly Expanded’ is shown in Fig. 1.1b. While these designs, their differences, and design choices are not the subject of this work, both robots require the same fundamental advances in wireless power and control.

(a) [older Robofly design] (b) [current design: Robofly Expanded]

Figure 1.1: [left] older robofly design, used in [30]. [right] Robofly Expanded design [11] featuring larger actuators and improved body geometry.

Realizing wireless flight requires solving three key challenges that arise from the small scale:

1. Insect scale at < 200 mg discourages traditional forms of propulsion such as a propeller driven by electromagnetic coils because unfavorable physics scaling [58]. Instead, flapping wings driven by piezo-electric actuators are more efficient [64]. While piezo-driven robots have been successfully used for flight, they require high potentials over 200 V [63]. Generating the necessary voltage signals has so far required large electronic components with a prohibitive weight relative to insect-scale aerial payload capacity (e.g. [35] at 675 mg).

2. Wireless flight requires an energy source to power the electronic and mechanical components. To date, the smallest high-drain (>10 C) batteries available are too heavy

at 350 mg (GM300910, PowerStream Technology, Orem, Utah). The only currently viable alternative is a battery-free design.

The onboard power conversion to the relatively high-voltage from 3–8 V to ≈180–300 V imposes requirements for the onboard power source. Namely, the powersource must be capable of sourcing high peak currents as detailed in §3.1.1. This need can be addressed by 1) large storage capacitors with low equivalent series resistance (ESR), or 2) the power source itself being low impedance. The former is undesirable for FIRs due to weight of capacitor components, while the latter suggests a way to select from alternatives available.

3. Finally, all the required digital processing has to be performed on the aircraft. Onboard computation that operates within the size, weight and power (SWaP) requirements is not only necessary to control the electronics and piezo driver, but also a basic requirement for truly autonomous insect robots capable of sensing and more complex functionality.

Wireless operation challenges control as well as power. Control of the Robofly of course entails *flight control* as can be seen below in Fig. 1.2, however traditional regulation of the output voltage is done using an inner control loop needed; this is *the waveform generator*, whose input is the commanded signals from the controller and whose output is the electrical signals to the piezo. The motional control scheme I propose below possibly obviates this inner coltage control loop. In Fig. 1.2 this waveform generator is embodied by the signal gen + DAC and piezo amplifier blocks. This is very much analogous to valve positioners on industrial feedback controlled processes, which actively position the valve (in closed loop) in order to effectively servo the control variable to equal the control signal or in other words ‘make sure the actuator is doing its job’, despite load or valve dynamics. This inner control must typically have much higher bandwidth than the outer control loop.

Although the system of Fig. 1.2 featuring this inner control loop represented by the Piezo Amplifier block, has been the paradigm for laboratory development of FIRs performing controlled flight since their inception, the applications of the Robofly are limited due to

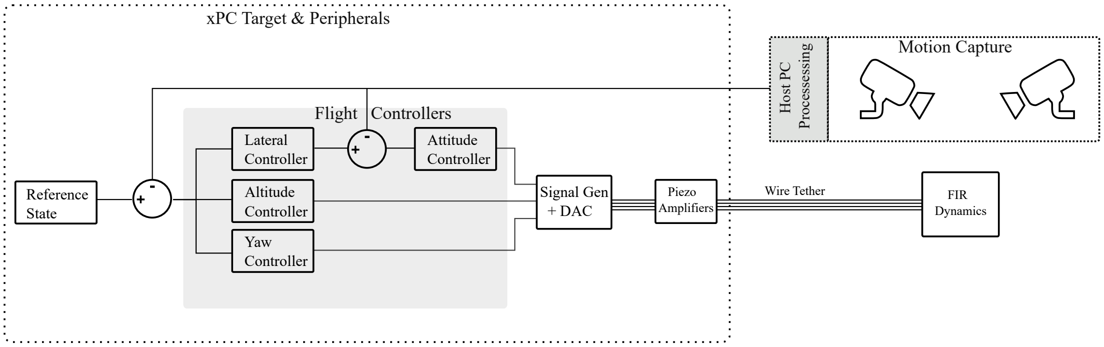

Figure 1.2: Flight control system for 2-wing Robofly in motion capture arena. Specifically, this inner loop is implemented in the non-inverting piezo amplifier element which multiplies the low voltage control signals from the DAC to the higher voltages needed by the piezos. When the weight of circuitry to implement these amplifiers is not a problem, this is easy. The challenge is that this system as it is cannot be simply miniaturized and put onboard.

the need for the physical wire tether to the robot and the large fixed desktop components including computer (xPC target), DAC, and piezo amplifiers (Trek 2205, Lockport, New York). The subject of §3 is to answer the question of how to provide *power* to the robot so that we could "cut the cord" and go wireless, but the other necessary step toward control autonomy is to develop a sufficiently lightweight onboard signal generator which generate the precise voltage signals needed using onboard resources alone - the subject of §5.

1.2 Main contributions and associated challenges

The main contribution of this work is the development of onboard systems fitting within the size, speed, weight, and power (SSWaP) budget of the Robofly in physically untethered flight - free of wire attachments to large benchtop systems used traditionally. The power electronics unit (PEU) and the control board are the two major components and meet three major needs: 1) to be capable of performing flight control including sensing, 2) to generate the relatively high-voltage electrical signals to precisely drive the piezoelectric flight muscles of the robot, given limited available wireless power sources, and 3) to collate and transmit telemetry and receive commands over wireless communication link.

The foremost challenge was in understanding the reasons behind existing PEU inadequacy, and then designing, fabricating, and integrating multiple novel DC-DC conversion piezo driver circuit topologies which overcome the limitations of existing technologies and most importantly can achieve independent wing control within the power and weight budget of the robot. Generation of precise actuator waveforms with minimal computational resources and in anticipation of higher-priority missions such as flight control, also required development of nontraditional control strategies for the waveform generator.

Untethered flights of FIRs to-date has only been shown with electronics sufficient to drive only one of two wings, and the robot systems used are therefore not capable of controlled flight, and worse, do not extend to FIRs with larger payloads for integrating all systems needed for true autonomy. There are no commercially available electronics sufficiently light or high voltage to serve as the PEU. In literature, existing power electronics topologies have also proven incapable of driving the larger piezoelectric flight muscles of the Robofly given input power below 300mW. A contribution of this work is to develop a sufficiently lightweight PEU capable of truly autonomous flight.

The ability to design simulations and coordinate simulation and analysis using simultaneous combinations of software such as Python/MATLAB, LTspice, electromagnetic FEA, and LabVIEW was a significant effort. Detailed simulations using exact component models of the components used are vital and the time spent on them crucial, since component particularities often make or break a design which might seem fine theoretically but often have some seemingly benign show-stopper. This skillset and the desire to dig into the heart of the problems in simulation where otherwise inaccessible signals can be probed, is what finally unlocked understanding of the real issues and creative design solutions for the power electronics. For all my years here I have been locked in stubborn battle (and at times utterly frustrated) with the same power electronics inefficiency that has stunted the field of FIR power and actuation autonomy for decades. Although this work presents early results, the contributions of multiple different power electronics topologies which are superior to state of art, and waveform generator control strategies to utilize them, have huge potential.

The solutions and work demonstrated to overcome these challenges is evidenced in several publications, especially the first ever wireless operation of an FIR at ICRA 2018, and detailed below. In summary, this document aims to

1. Summarize the robot systems.
2. Document the work done thus far toward developing wireless power capability in §3.
3. Document the advances toward wireless control in §5.
4. Explore potential future avenues in both power and control to unlock demonstration of the first wireless controlled flight of FIRs.

1.3 *Flying Insect Robot: The Robofly*

1.3.1 How the Robofly works

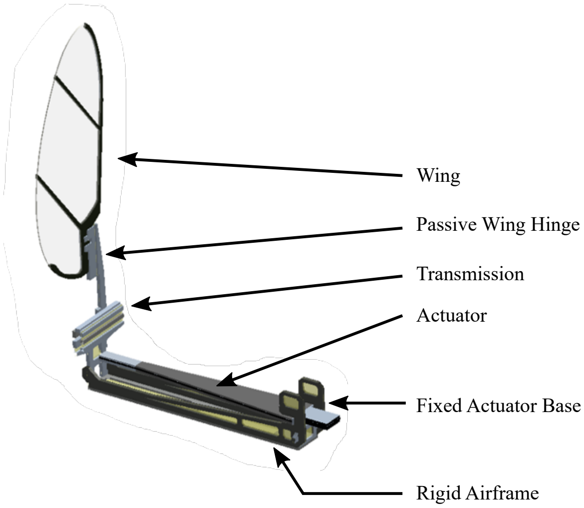

Figure 1.3: Half-fly assembly showing the rigid airframe, bending cantilevered beam bimorph piezo actuator, 4-bar transmission, passive wing hinge, and wing.

The Robofly in its various forms consists of two sub-units 'half-flies' which are adhered together for the whole robot at the time of this writing. Fig. 1.4 and Fig. 1.3 show the half-fly sub-unit. Each sub-unit has a rigid airframe which fixes the base of a 2-layer (bimorph) bending-beam piezoelectric actuator, and at the other end of the airframe is the 4-bar transmission mechanism. The transmission rotates the wing as the actuator tips moves, as shown in Fig. 1.4. Wing angle of attack results from the aerodynamic drag resisting the stroke of the wing and pushing the wing to a desirable angle-of-attack against the stiff wing hinge. The flapping stroke and angle of attack resembles a swimmer's arm and hand while treading water. However in order to reduce FIR actuators, the angle of attack is achieved passively instead of by actively rotating the wing.

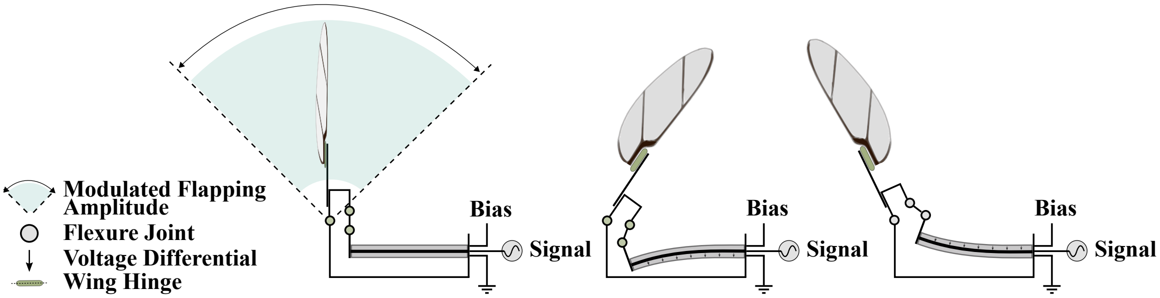

Figure 1.4: Diagram showing the principle of operation of a single piezo-actuated wing in the UW Robofly. (left) A bimorph cantilever-beam piezoelectric actuator is fastened rigidly to the airframe at one end. At the other, it is connected to a transmission with a short lever arm, providing a large angular amplification of the small tip motion of the actuator. Under typical conditions, the wing amplitude is around 90°. (middle and right) As the voltage signal oscillates, the top and bottom piezo actuators are alternately charged. The resulting electric field induces a small strain in the piezoelectric material, which results in motion at the tip of the actuator. This causes the flapping motion. The wing rotates passively around a wing hinge due to aerodynamic thrust, resulting in insect-like flapping kinematics and wing angle-of-attack that produces lift.

1.3.2 How control signals achieve desired wing flapping

As can be seen in Fig. 1.4 and an isometric view of the half-fly subassembly in Fig. 1.3, the motion of the actuator tip causes wing flapping. The actuator tip motion is caused by differential strain of the top layer relative to the bottom layer of the bimorph, which is cantilevered from the fixed base. The differential strain of these two layers is achieved by the inverse piezoelectric effect with the varying the electric field of the top piezo layer compared to the bottom. The varying electrical field results from applying the electrical signals V_{Bias}, V_{Signal}, and gnd as shown in Fig. 1.4. Because the piezos have extremely high bandwidth relative to the flapping frequency and dynamics of the mechanism, these applied voltages result in relatively immediate corresponding flapping stroke. The wing angle-of-attack lags somewhat behind the stroke angle, as studied in [8].

$$V_{signal}(t) \;=\; \frac{V_{Bias}}{2} \pm \left(\sin(\omega t) + \mu \sin(2\omega t)\right)\left(V_{amp} + V_{roll}\right) + V_{pitch} \qquad (1.1)$$

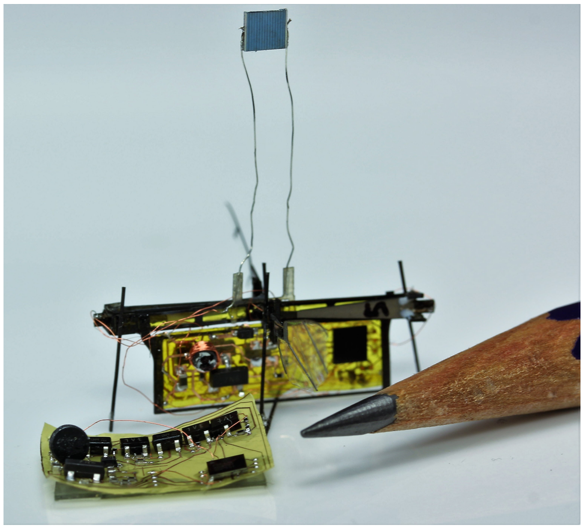

Figure 1.5: The UW Robofly shown with flight-weight power electronics, on-board PV cell power source, and microcontroller. A pencil tip is shown for scale. Shown at lower left is a flight-weight implementation of the updated power electronics unit (PEU) reported here that additionally features two complete half-bridge drivers for independent dual-wing control.

There are also normalization factors in Eq. 1.1 to correct the perturbed total amplitude as in [44], but these are not shown for simplicity.

The electrical signal to send to the actuators for the left and right wings is computed by the flight controller according to Eq. 1.1, where μ is the second harmonic content (yaw torque control knob), and $V_{amp}, V_{roll}, V_{pitch}$ are the control knobs for altitude, roll, and pitch respectively. These major control gestures *roll, pitch, yaw* are shown graphically as viewed from above the robot in Fig. 1.6. Examples of the signals for roll, pitch, and yaw control actions (for one of the two wings) are shown in Fig. 1.7.

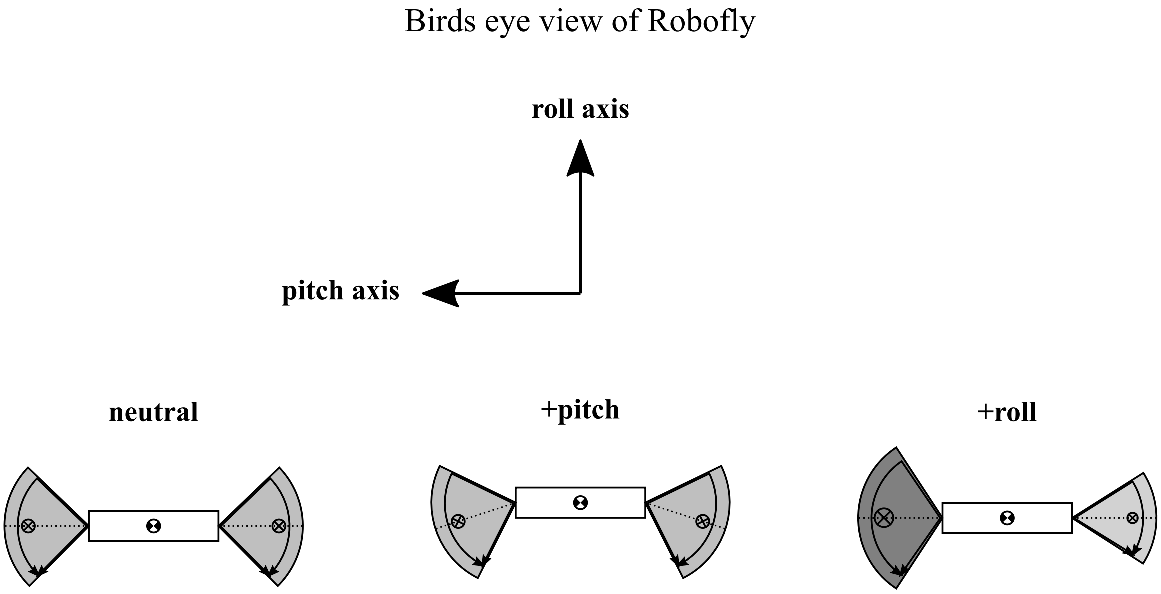

Figure 1.6: Cartoon view demonstrating the different wing flapping envelopes to generate flight control torques.

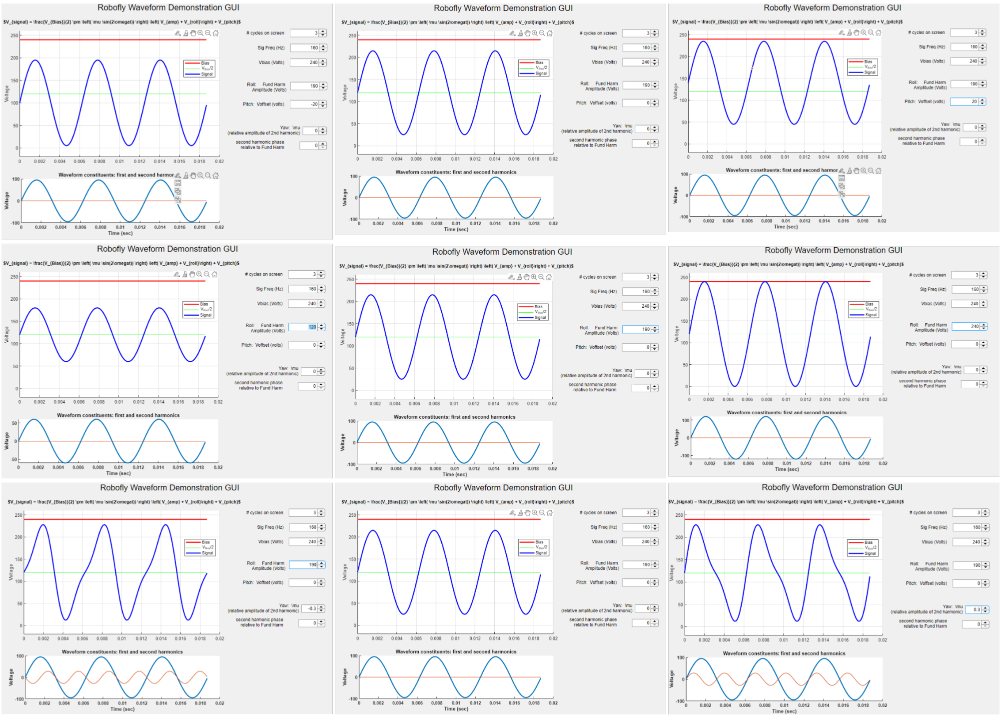

a) Varied amplitude
control roll, altitude

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): -20
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): 20
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

b) Varied DC
Offset
control pitch

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 120
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 240
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

c) Varied second
harmonic
control yaw

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): -0.3
second harmonic phase relative to Fund Harm: 0

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): 0
second harmonic phase relative to Fund Harm: 0

Robofly Waveform Demonstration GUI

`$V_{signal} = \frac{V_{Bias}}{2} \pm \left( \mu \sin(2\omegat)) \right) \left( V_{amp} + V_{roll}\right) + V_{pitch}$`

cycles on screen: 3
Sig Freq (Hz): 160
Vbias (Volts): 240
Roll: Fund Harm Amplitude (Volts): 190
Pitch: Voffset (volts): 0
Yaw: \mu (relative amplitude of 2nd harmonic): 0.3
second harmonic phase relative to Fund Harm: 0

Figure 1.7: Demonstration of the different piezo actuator signals in order to achieve different control actions. [top] 3 cases of varied amplitude, as for roll or altitude control. [center] 3 cases of varied DC offset, as for control of pitch. [bottom] 3 cases of varied second harmonic component, as for yaw control. Variation of the speed of the upstroke relative to the speed of the downstroke causes different drag forces on the wing, creating the yaw moment.

Chapter 2 Background and Related Work

2.1 Insect-scale robotics

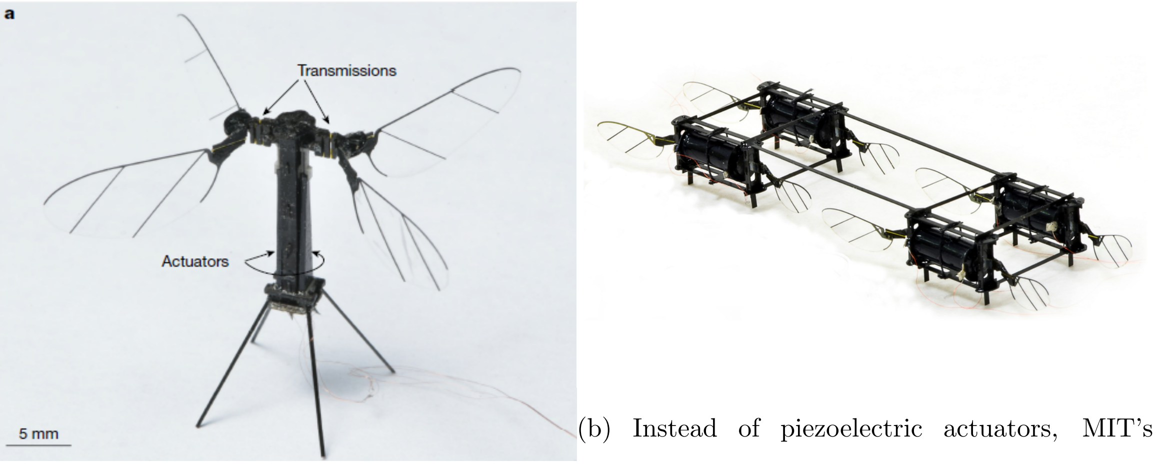

(a) Harvard RoboBee X-wing, the celebrated FIR featured in [27].

(b) Instead of piezoelectric actuators, MIT's FIR is powered by dielectric elastomer actuators (DEA) [42].

Figure 2.1: Related FIRs from the field. (Left) The Harvard Robobee "X-wing" [27], and (right) An rigidly coupled multi-body FIR featuring soft DEA actuators from the Chen lab at MIT [42]

Harvard's RoboBee [28], shown in Fig. 2.1a in its 4-wing configuration, shares many commonalities as well as limitations with the RoboFly. MIT's FIR (Fig. 2.1b) [6] uses dielectric elastomer actuators (DEA's) instead of piezoelectric actuators. These muscle-inspired flexible actuators are more resilient than piezos. However, they require even higher voltages and thus are not comparable for the scope of this work.

2.2 Flight-weight power

The oscillating motion of the bimorph piezo beam actuators that flap FIR wings must be driven by a sinusoidal, high-voltage signal. This is typically in the range of 200–300 V to maximize the power density produced by the actuator. The need for low weight and high efficiency strongly influence design choices. Possible approaches for FIRs include laser power beaming to a photovoltaic (PV) cell, solar power with a PV cell, onboard batteries, large capacitors, and magnetic resonant coupling.

2.3 Piezo-electric actuators, drivers, and power electronics

2.3.1 Piezoelectric flight muscles

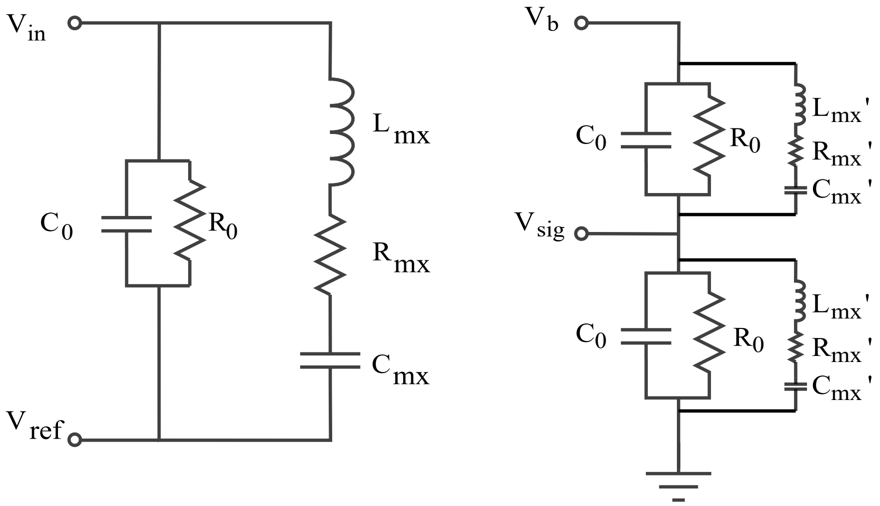

Figure 2.2: [Left] The Butterworth-Van Dyke (BVD) model of a monomorph piezoelectric actuator with dielectric loss and primary resonant mode of the robofly wing, represented by series resistor-capacitor-inductor (RLC) circuit equivalent. Subscript x just denotes the resonant mode, $x = 1$ for the flapping wing stroke mode being the most relevant to FIRs. [Right] Approximate model of the bimorph actuator based on the BVD model at left, with primary resonant mode. This diagram makes clear that the bimorph is in fact two monomorph piezo layers in parallel and is useful in circuit simulations if care is taken to divide the total power flow by two, and reflect the coupling of the two layers to the same resonant mode via their adhesion to the same passive layer of the bimorph - i.e. extra modeling as by using current sources or at minimum sanity-checks of the results. However the left version is sufficient for modeling with simplifying assumptions and if the Bias voltage V_b is held constant in a steady-state hovering flight as discussed in chap. 4.

This work is concerned with bimorph piezo bending actuators configured in parallel "simultaneous drive" as described in [65, 63]. Bimorph Piezoelectric bending beam actuators and attached loads are often modeled using the Butterworth-Van Dyke (BVD) model described in IEEE piezoelectricity standards and also detailed in [38]. This model, shown in Fig. 2.2, features the principal or 'undeformed' capacitance C_0 of the piezoelectric actuator, and resistor R_0 in parallel with C_0 which accounts for various dielectric losses such as charge leakage, and actuator load(s) are modeled as lumped parameter series RCL equivalent circuits attached in parallel as shown above. The piezo material capacitance varies when driven dynamically as plotted over driving voltages in [38], and dielectric loss is a complicated function of excitation frequency, principal capacitance, and a 'loss-tangent' term $\tan\delta$ corresponding to the piezoelectric material found in various works to be a function of applied voltage for a given piezoelectric.

This work is primarily concerned with just the first resonant mode, wing-transmission-actuator system, with frequency response depicted in Fig. A.1. Additional RLC branches can be added in parallel to represent additional vibration modes. Although largely neglected in [16], it was found in [8] that the wing flapping angle of attack is not negligible for control purposes, but from a power and efficiency perspective in this work it is not vital to include an additional resonance modes in the circuit representations of the loaded actuator.

An additional caveat to the BVD model is that aerodynamic wing loading is not well represented by a constant. It is a function of the square of wing stroke velocity $\dot{\Phi}^2$, and so a nonlinear dissipative element may be substituted into the first vibration mode for simulation purposes. This can be done in practice by defining a nonlinear relationship for R_{mx} in the circuit simulation software, details in appendix.

This lumped parameter model is convenient for analysis of linear load models, lends itself to circuit simulation, and is well suited to the Robofly, whose wing mechanism is adequately represented by a second order oscillator [16, 8] and equivalent vibration mode. This model fails to capture the rigid mechanical coupling between the PZT layers of the monomorph, so the right side of Fig. 2.2 is to be considered with a constant bias voltage and small perturbations of signal voltage. Attempts to improve this basic model include adding a series resistor to represent hysteresis and using nonlinear elements for attached loads [54].

2.3.2 Dielectric losses

The resistance R_0 in parallel with the capacitance C_0 is of particular importance in this work, since most of the energy in the system is in the piezo's capacitance, to the extent that the piezoelectric actuator operated below its resonant frequency (quasi-steady state) is considered to be principally capacitive.

Significantly, the bending beam piezoelectric actuator has been found to have a relatively low electromechanical coupling coefficient [59]. A piezo actuator with an attached resonant load can be thought to form a coupled resonator system across electrical-mechanical-fluid domains through the inverse piezoelectric effect, actuator tip motion, transmission, and wing aerodynamics, much like the wireless power transfer method using magnetically coupled resonators discussed in chapter 3 and [53]. With coupled resonators, it is generally possible to maintain high efficiency of power transfer *if energy not taken by the coupled load remains in the transmitter* - an insight explained in [53]. However since R_0 is in parallel with the principal capacitance for the loaded piezo, energy not taken by the load represented by the coupled resonant mode in Fig. 2.2, is largely dissipated by R_0.

Further, since the dielectric loss R_0 has been found to be represented by

$$R_0 = \frac{1}{2\pi f C_0 \tan\delta} \tag{2.1}$$

where $\tan\delta$ is the loss tangent, f is the excitation frequency, and C_0 is the principal capacitance. The loss tangent $\tan\delta \approx 0.115$ over the operating voltages of interest here $(150 - 300\text{V})$ is a function of input voltage, characterized for the 5H PZT material utilized here by [37]. For the piezo bimorphs utilized in this work as in [9, 11, 12], assuming the same basic loss tangent $\tan(\delta)$ for the 5H piezo material and operating at $C_0 = 9\text{nF}$ and $f = 180$ Hz, $R_0 = 0.85$ MΩ. Values for R_0 computed using Eq. 2.1 agree with the power requirements measured by benchtop piezo amplifiers driving $V_b \approx 200 - 250V, V_{sig} \approx 20 - 230V$ as used for robots in hovering flight.

This relatively low impedance of R_0 is a significant portion of the load especially when C_0 or ω_{flap} is increased, so Eq. 2.1 is coded into the BVD model used in circuit simulations throughout this work. Additionally, the computed R_0 is consistently added to the capacitive ‘dummy’ loads representing bimorph actuators in this work. It is unclear if PEU designs such as [38, 44, 43] instantiate a physical resistor to represent R_0 during development. Physical RLC branches for the BVD model such as measured by an impedance analyzer in [38] may be unrealizable for the purposes of breadboard development due to large inductances, but parallel R_0 is a convenient improvement to the accuracy of a bimorph RC ‘dummy load’ used in lieu of a robot for development purposes.

Additionally, the excitation frequency f in Eq. 2.1 should be considered as more than just the principal harmonic of the piezo’s input signal (the flapping frequency) but likely includes additional harmonic content such as the high frequency switching of DC-DC converter/drivers. High frequency signals as from the onboard PEU during individual pulses of the switches, may lose additional power through dielectric loss because of this frequency dependence of R_0.

2.3.3 Piezo actuator load characterization

Although conventional wisdom and actuator design for energy density as guided by [65] emphasizes operation near the max power point of the actuator's force-displacement curve, increases in both principal capacitance and excitation frequency in order to increase flapping wing lift also affect power requirements. Namely, increased dissipation of power through dielectric losses from Eq. 2.1 and significant increases in reactive power requirements due to $I_C = C^{dV}/_{dt}$ which are difficult to efficiently handle with onboard PEUs. Standard definitions of FIR efficiency in literature do not distinguish between useless real power dissipated by R_0 and useful real power dissipated by aerodynamic damping.

By inspecting the Bode magnitude plot of the wing-transmission-actuator system (Appendix Fig. A.1), it can be observed that flapping excitation frequencies outside of the resonant peak result in decreased stroke amplitude gain - which simply means that we want to flap at mechanical resonance which is well understood. The electrical implication of this basic feature is that, when flapping frequency excites the system at resonance, the impedance of the first resonant mode in Fig. 2.2 approaches the purely real resistance R_{mx} which represents the useful (mostly) wing aerodynamic damping. Meanwhile, the undesirable dielectric loss R_0 changes much slower for the same excursions away from ω_n.

Data from a hovering flight of the Robofly was collected for power analysis for [9]. The instantaneous power is simply given by $P = VI$, which includes the real and reactive components of power. The average power over integer multiples of the flapping period yields the stroke-averaged real component of the power, but it is important to note that a signifiant real power consumer is actually the dielectric losses represented by R_0 in Fig. 2.2, as well as the useful real power going to aerodynamic lift and drag represented by R_{m_x} in Fig. 2.2.

The voltage and current waveforms can also be used to determine the phase angle of the load. As part of an effort to compute this phase angle, I used autocorrelation to compute the phase of current to voltage but more typically datamarkers at peaks and timestamps. Visual confirmation of the measurements of typical phase is shown in Fig. 2.3. Using the time-stamps from data markers placed at peaks of the voltage and current waveforms, phase angle in degrees between voltage and current is obtained by

$$\phi = 360\frac{T_{Ipk} - T_{Vpk}}{T_{cycle}} = 360\frac{0.2525 - 0.2466}{0.2537 - 0.2466} = 299^\circ \tag{2.2}$$

with T_{Ipk}, T_{Vpk} being the timestamps (seconds) of the first voltage peak and the current peak data marker, and T_{cycle} being the time period between the first and second voltage peaks. The computed phase of applied voltage to current of 299° corresponds to the current leading the voltage by 61° as expected for a primarily capacitive load with some real power going to system's dissipative elements. Dissipative elements in this system are resistors (including the dielectric loss R_0) in the electrical domain, friction and structural damping in mechanical domain, and aerodynamic drag/lift in fluid domain. The lumped parameter model of the actuator-transmission-wing mechanism is considered in [37], [16], [26].

This phase measurement gives us important insight into the character of the load and actuator. With undeformed piezo capacitance C_0 measured easily during assembly, we can distinguish between real and reactive power components by looking at the phase angle and computing power from actuator voltage and current measurements. Exact system identification using only phase and power measurements is muddled by variation of the piezo material's capacitance during dynamic strain and electric field during flight, but useful approximations about the system can be made nonetheless. Phase angle is a measure that we can obtain empirically from tethered flight experiments using voltage and current monitors of benchtop amplifiers, and can be related to the repetitive control element designed in chap. 5.

Empirical load characterization example

For example, the impressive result of the 4-winged "Robobee X-wing" in [27], reported real power $P_{re} = 26 - 35$ mW, and reactive power $P_{im} = 50 - 59$ mW (mVAR). Given that $\Phi = \arctan\left(\frac{P_{im}}{P_{re}}\right)$, at the upper end of the reported power ranges the phase angle $\Phi = 59.3^\circ$ would be observable in plots of voltage and current from experiments, similar to Fig. 2.3. In [27], the reported flapping frequency was $\omega_{flap} = 170$ Hz, and the voltage amplitude of the drive signal was 200 V_{pp} with DC offset $V_{dc} = 100\ V$. There was no bias rail utilized in driving the bimorph actuators so the voltage input was just the unipolar sinusoid. As discussed above, since ω_{flap} was equal to the resonant frequency of the wing system which is represented by the RLC branch in left of Fig. 2.2, the reactances of the L_{m1} and C_{m1} cancel, leaving the net real part of the model's total impedance to be well represented by just the parallel combination of $R_0||R_{m1}$. Likewise, the reactive component of the overall load is predominantly the capacitance C_0.

The reactive component of 59 (mVAR) reported in [27] can be used to back out the capacitance of their actuators if such measurements are not readily available or reported. Monitoring actuator capacitance online can obviously provide indication of actuator damage as well and is done in several relatively bulky commercially available piezo haptic driver ICs, as well as considered for microrobotics in [32]. Considering the flapping frequency, voltages, and power numbers reported, we can estimate the capacitance of the piezo layers C by: $P_{im} = CV^2f \implies \frac{0.059}{(2\pi 170)({}^{200}/\sqrt{2})^2)} = C = 2.76$ nF. Extending from here, estimating this obtained value to approximate the principal capacitance C_0 (which in undeformed state would be less than or equal to), we can estimate dielectric loss using Eq. 2.1 and then the remainder of real power output.

While the actuator voltage signals in controlled flight will generally be variable, the real component of the load power given a nominal unipolar sinusoid at some DC offset for this linear system can be considered using Eq. 2.3 where $P = \frac{V^2}{R}$ is just separated for the DC and primary harmonic content of the input signal. The waveform of [27] had a third harmonic at a low amplitude which was dismissed in that work as not greatly distorting the sinusoidal flapping motion, so this is an adequate approximation for the reported values.

$$P_{re} = \frac{(V_{dc})^2}{Z}\Big|_{s=0j} + \frac{\left({}^{V_{pp}}/(2\sqrt{2})\right)^2}{Z}\Big|_{s=j\omega_{flap}} \tag{2.3}$$

where $V_{dc} = V_b/2$ in nominal hovering flight with zero pitch offset, per the standard flight control waveform function defined for FIR flight control in [10, 9] .

The generalized impedance of the BVD model consists of the impedance of the resonant mode branch(es) in parallel with the impedance of the actuator itself. Referring to the left side of Fig. 2.2, the generalized impedance of a single layer piezo actuator with just the primary resonant mode is given by Eq. 2.4.

$$Z_{eq} = \left(\frac{1}{(C_0 s + \frac{1}{R_0})^{-1}} + \frac{1}{(R_{m1} + L_{m1} s + \frac{1}{C_{m1} s})} \right)^{-1} \tag{2.4}$$

.

The power flow from the DC bias rail and the power flow from the input signal can both be used to estimate the dissipative parameters of the actuator's BVD model. Since power flow from bias rail should be much less than the power from the input signal V_{Sig} as noted in chap. 4, this work instead considers a rough approximation of system parameters using phase and total power requirements. When the input ω_{flap} is at the resonant frequency of the series RLC branch, the reactance of the C_{m1}, L_{m1} are equal and opposite about the real axis, and we can simplify to see the real part of the Z_{eq} in Eq. 2.4 consists of the parallel R_0, R_{m1}. This is for the time-varying flapping signals, at the V_{sig} terminal, not for static DC bias where the impedance offered by the resonant branches is of course high at $\omega << \omega_{flap}$. When driven by V_{sig} at resonance, the equivalent real-valued impedance is Eq. 2.5.

$$Z_{re} = 2(\frac{1}{R_0} + \frac{1}{R_{m1}}')^{-1} \tag{2.5}$$

for the bimorph model of Fig. 2.2 at mechanical resonance.

Using Eq. 2.1 and the reported peak operating voltage $V_{pp} = 200$ V from [27], estimating the reported actuator capacitance $C \approx 2.7$ nF as above, and then assuming $\tan(\delta) \approx 0.115$ described above, the resistor representing dielectric loss can be estimated. $R_0 = (2\pi f C_0 \tan(\delta))^{-1} \implies R_0 \approx 3$ MΩ.

Finally, given the reported real power measurement $P_{re} = 35$ mW from [27], and using the results found above, Eq. 2.5 and Eq. 2.3 can be solved as in Eq. 2.6, yielding an estimate $R_{m1} \approx 2.4$ MΩ. This relatively low value makes sense given the extensive effort in [27] to increase the lift of the FIR at all costs, most obviously by increased flapping frequency and doubling the wing area with a second set of wings attached to the original transmission output similar that shown in Fig. 1.4 which typically has just one wing. This 4-wing design (where the wings are *not* independent) is shown in Fig. 2.1a. Another significant improvement which is harder to immediately see, is in better piezo utilization. The piezo utilization in [27] was improved substantially by implementing the drive scheme described further in §5.2.4, which "achieves greater deflection and force [for a given applied waveform] by using the full range of available positive and negative electric field" of the bimorph, as proposed in [52].

Empirical characterization from apparent power measurements in tethered flights

Since voltage and current measurements are convenient during tethered flights such as [9], Eq. 2.3 is not necessary to compute at all given direct measurements of V, I which give the power $P_{avg} = \overline{I(t)V(t)}$, as represented by the dashed line in the bottom plot of Fig. 2.4. A purely reactive load (neglecting wire tether resistances etc) would have an average $P_{avg} = 0$, where the reverse power returning from the load back to the amplifiers $(P = I(t)V(t) < 0)$ equals the power leaving the amplifiers. At the low levels of power involved in FIR flight, measurement noise such as from common mode or ground loops in analog inputs is not necessarily negligible so care must be taken. In the flight experiments such as in Fig. 2.4, the remainder of the instantaneous power above the x-axis in the lower plot of Fig. 2.4 is what has gone to the load and not returned. This real power includes both useful power out as well as dielectric losses. Just as done above, the dielectric loss parameter R_0 can be estimated using the relationship of Eq. 2.1.

Extending this, a simple closed-form estimate for the 'useful' power term R_{m1}, can be found in Eq. 2.6 given the known sinusoidal input to the two bimorph piezo actuators as in [9, 18], experimental measurement of average power P_{re}, and an estimate of R_0 using Eq. 2.1. Under these simplifying conditions, the real power is equally distributed to all four layers of the two bimorph actuators. Eq. 2.6 is obtained by substituting Eq. 2.5 into Eq. 2.3 and solving for R_{m1}, with derivation symbolic code in the github linked in appendix A.3.1. This estimate is approximate for a bimorph actuator assuming flapping in nominal hovering flight at resonance of the wing system, and the capacitance C_0 is roughly representative of the piezo capacitance during flight conditions; as noted there is some change in capacitance when piezo is driven dynamically as found in [37] which muddles the ability to identify R_{m1}

this way. The generalized impedance of the BVD model from Eq. 2.4 can be used instead of Eq. 2.5 when the input signal is not at the resonant frequency, however this requires more exact modeling of L_{m1}, C_{m1}.

$$R_{m1} = 2\frac{R_0 \cdot \left(8V_{dc}^2 + V_{pp}^2\right)}{4P_{re}R_0 - 8V_{dc}^2 - V_{pp}^2} \tag{2.6}$$

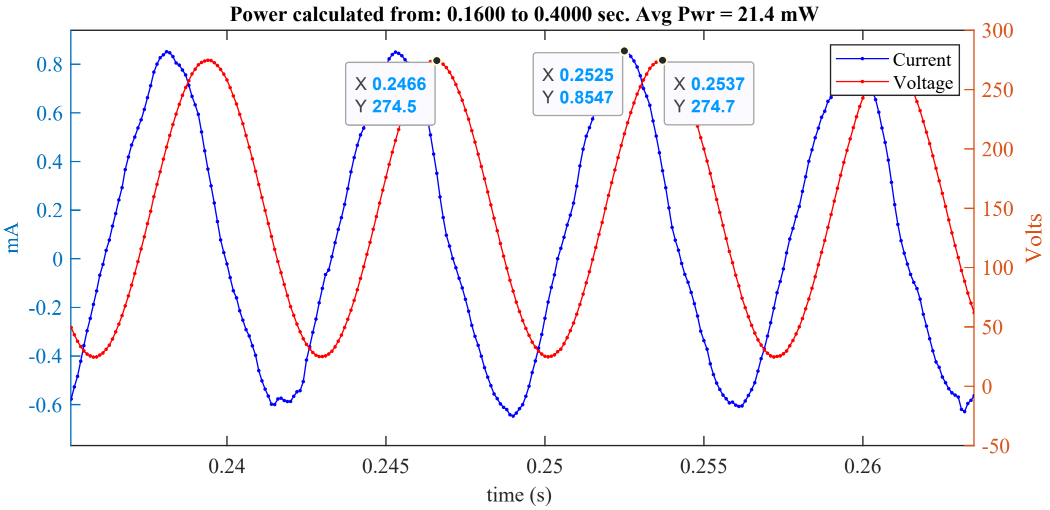

Figure 2.3: Voltage and current waveforms from a flight of the single-wide actuator Robofly (earlier version) whose actuators have principal capacitance of 2.7 nF from each layer of the bimorph. The waveforms were recorded and studied for power analysis of the hovering Robofly and so that the phase of current to voltage can be visually confirmed using the placed datamarkers.

Similar analysis is performed during experiments of this work, the author has written actuator test code in common usage for lab actuator test procedures which presents the voltage and current measurements along with thrust force measurements if used, and some optional power analysis. An example of voltage, current, and power measurements obtained from Trek amplifiers during flight of an improved design which was measured to provide approximately 225 mgf thrust in addition to the weight of the robot is shown in Fig. 2.4. This flight shows one of the better robots in this work which featured increased lift along with

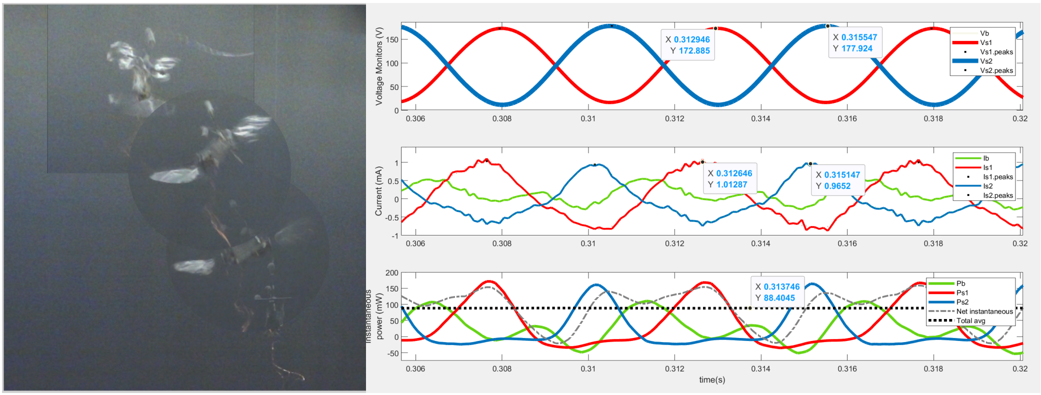

Figure 2.4: Voltage and current waveforms from a flight of the more recent triple-wide actuator Robofly whose actuators had baseline capacitance of 8.1 nF for each layer of the bimorph, measured by LCR meter. The waveforms were recorded and studied for power analysis of the Robofly during this vertical flight, as shown in the still frame sequence. $V_b = 200, A = 160\ V_{pp}$ with no DC offset. Average apparent power to the robot from the amplifiers for this experiment were 90 mW for both wings. $V_b, V_{s1}, V_{s2}, I_b, I_{s1}, I_{s2}, P_b, P_{s1}, P_{s2}$ are the instantaneous voltage, current, and power to the bias, channel 1, and channel 2 respectively. The instantaneous power peaks slightly before the actual peak voltage, due to the capacitors' current $\propto C\frac{dV}{dt}$. Corresponding input power requirements can be projected for an onboard PEU by dividing the instantaneous and average power measurements by typical PEU efficiency, for the typical boost converter + linear driver $\eta = 15{-}25\%$ depending on power source voltage and impedance, actuator size, and output voltage and frequency.

the increased power requirements of larger actuators, indicating the need for increases to PEU efficiency sufficient to handle the increased power requirements, potentially offsetting added weight of electronics, larger airframe, and larger actuators.

Operation at resonance

As is made clear above and mechanically intuitive, it is vital to flap at the mechanical resonance of the wing system. While it is convenient to verify the resonant frequency of an FIR when first fabricating and during initial tests, this can also be done online from power analysis principles as above and demonstrated in piezo microrobots in [32]. Any difference in loaded vs unloaded phase angle can provide another indication of whether or not the robot is flapping at the natural frequency of the wing-transmission-actuator system. If so, BVD resonant mode would present a predominantly real load whose reactances of L_{m1}, C_{m1} are equal and opposite about the real axis, leaving the real component to be the predominant load seen by the actuator. When flapping frequency deviates from resonance, there is increasing magnitude of phase shift between the voltage input and the through variable in the RLC branch representing the first resonant mode of the BVD model of the FIR actuator-wing system. This in turn shifts the phase of the actuator current relative to input voltage, away from the value observed at nominal flapping at resonance.

Unfortunately, much of any apparent change in phase angle of a given robot's V, I measurements during flight is hard to differentiate between the consequences of variation in capacitance of the piezo and associated changes in R_0 which are not *useful* real power, and actually useful aerodynamic damping may not have increased to the same degree.

Actuator size consideration given dielectric losses

Increasing the actuator size and flapping frequency should be done very carefully to obtain the minimum actuator force/displacement and lift/drag required while considering increased dielectric losses and reactive power requirements which onboard PEUs are ill-suited to handle. This is made very clear by considering the ratio P_{useful} of estimated useful real power output through admittance $\frac{1}{R_{m1}}$ to the admittance of the dielectric loss $\frac{1}{R_0}$. This ratio is given in Eq. 2.7.

$$P_{useful} = \frac{1/R_{m1}}{1/R_0} \tag{2.7}$$

As seen above, the successful flight of the Robobee X-wing had a favorable ratio $P_{useful} \approx 1.25 = \frac{1/2.4\text{M}\Omega}{1/3\ \text{M}\Omega}$ but is frankly maxed out in its ability to carry additional payload for independent wing control, any sensors, superior microcontroller, or RF communications. On the other hand, the FIR featuring 'triple-wide' 9nF bimorhph actuators used in Fig. 2.4, had greatly increased power requirements in exchange for some more lift and payload. For this FIR, we face a worse ratio $P_{useful} \approx 0.95 = \frac{0.85\ \text{M}\Omega}{0.893\text{M}\Omega}$ using Eq. 2.1 and Eq. 2.6 for the

45 mW per wing measured in Fig. 2.4.

2.3.4 Waveforms for piezo-electric drive

It was demonstrated in [46] that bimorph (two-layer) bending beam piezoelectric actuators can be used to flap the wings of FIRs and achieve stable flight control, but these ‘flight muscles’ require precise high voltage electrical signals. To date, flight control demonstrations of piezo-actuated robots have required benchtop piezo amplifiers (e.g. Trek model 2205 in [11]). These serve to amplify the low voltages (e.g. 0–5 V) produced by the flight control system to 200–300 V required by the piezoelectric actuators. Alternatives to piezoelectric actuation that do not have a high voltage requirement, such as electromagnetic coils [66, 2] have not yet demonstrated sufficient power density or independent wing control at the size scale we are concerned with here (less than 500 mg).

This work utilizes bimorph piezoelectric actuators driven in ‘simultaneous drive’ mode to provide a balance between low weight and efficiency [38, 52, 65]. In simultaneous drive, the two outer surfaces of the piezo cantilevers are held at a fixed potential difference, V_{Bias}, while the middle layer between the two piezos is excited by a unipolar sinusoidal signal V_{Sig} (Fig. 1.4). As in Eq. 1.1, modifications to the baseline sinusoid V_{Sig} including changes in its amplitude V_{pp}, DC offset, and adding a second harmonic, can produce all of the torques and forces necessary to control a two-winged FIR in flight [47, 46]. Furthermore, it has been recently shown that these same signal perturbations are sufficient to actuate FIRs underwater [5] and on the ground and the surface of the water [10].

Motivated by a plan to eventually bring the PEU on-board, the flight control demonstrations of [18, 46, 9] all used piezo actuators in this simultaneous drive configuration and this configuration is the focus of this work. Subsequent works by the same researchers [38, 65] have suggested improvement by switching to ‘alternating drive’ configuration in which the bimorph pzt layers poling directions are reversed and the middle layer is grounded, requiring two time-varying signals for the same actuator motion. However, this is counter-indicated by the power electronics design performed in this work.

2.3.5 Onboard Waveform Generation

Chap. 5 further discusses the usage of the onboard power electronics of chap. 4 to achieve the voltage signals for the piezos. This task is foreshadowed by the difficult of simply miniaturizing the traditional benchtop systems which perform this task with relative ease. Using benchtop piezo amplifiers, connected to the FIR through a thin wire tether, is convenient to demonstrate functionality of the flight actuation apparatus [63, 46]. This is because high voltage, low-impedance amplifiers which have little to no resitrction on weight, complexity, and power consumption and are commercially available in form factors suitable for the benchtop. However, use of a wire tether is clearly undesirable because of how it limits the useful range of the robots and confounds many aspects of flight control [18].

There are smaller, commercially available piezo drivers, but they are either too heavy for FIRs which weigh < 200 mg or not suitable for driving bimorph actuators at sufficiently high voltage. Boréas Technologies, (Bromont, QC, Canada) offers piezo driver ICs such as BOS1901 which stem from [4] and are impressively efficient and within SWaP constraints of FIRs. But these can generate a maximum 190 V_{pp} differential drive signal (*not* ground referenced) , but only half that when ground referenced and still and don't supply a bias rail in this higher voltage operation mode. Alternatively, custom power electronics using discrete components provide design flexibility and can meet design needs within size, weight, and power (SWaP) constraints. Such custom flight-weight power electronics have been presented [38] and utilized [30, 27]. Recently, progress has been made in on-board power systems on FIRs such as the UW Robofly [30] and the Harvard Robobee [27]. These robots were *powered* wirelessly. However, their flights were in open loop, that is, with *uncontrolled* flapping, resulting in very short flights, all of which were less than a second.

This work favors discrete component PEU designs to leverage the abundance of miniaturized consumer electronics and enable rapid prototyping for design innovation.

A brief introduction to the onboard power electronics and their usage for typical waveform generation follows.

2.3.6 Drive Signals from the traditional waveform generator

Onboard insect sized robots as opposed to with wire tether attachment to benchtop systems, the needed waveforms have traditionally been obtained in literature by two primary approaches: a boost converter + halfbridge drive stage which provides a DC bias $V_b \approx 200-250$ while the drive stage takes from the bias rail and provides a unipolar sinusoid $V_{sig} \approx 5 - V_b$ V; and a direct conversion approach which only produces a single signal V_s. In this work and in literature, the first option in several variations aimed at increasing poor total efficiency is the most demonstrated. This is treated in greater detail in chap. 4, but introduced in simple form here.

The nominal voltage over time waveform for a balanced Robofly using bimorph piezo actuators in simultaneous drive configuration [65] in hover would be a unipolar sinusoid exactly centered between ground and V_{Bias} (zero pitch) and having no second harmonic content (zero yaw). An example of this nominal waveform is depicted in Fig. 2.6 Onboard, this waveform can be obtained with digital control signals from a microcontroller consisting of a pulse train of varying duty-cycle to control the power electronics which charge or discharge the actuator in order to change its voltage. The purpose of the pulse trains is to operate the power electronics so as to drive the output up (red) or down (blue) in voltage toward reference voltage as shown in the middle subplot, and the simulated output of an onboard PEU in the bottom subplot. Simply put, increasing or decreasing the actuator voltage is generally done by charging or discharging the charge on the actuator capacitance according to $V = \frac{Q}{C}$. Dynamic load, capacitance variation during operation, and PEU dynamics complicate the control of actuator voltage.

The basic idea of the boost converter plus driver is to use the ‘high-side’ of the driver to increase the signal voltage by allowing current to conduct to the actuator from the bias rail, and conversely to use the ‘low-side’ to sink current from the actuator and cause the signal voltage to decrease. A simplification is shown in Fig. 2.5, and a corresponding circuit schematic including voltage feedback is shown in chap. 4 in Fig. 4.1. .

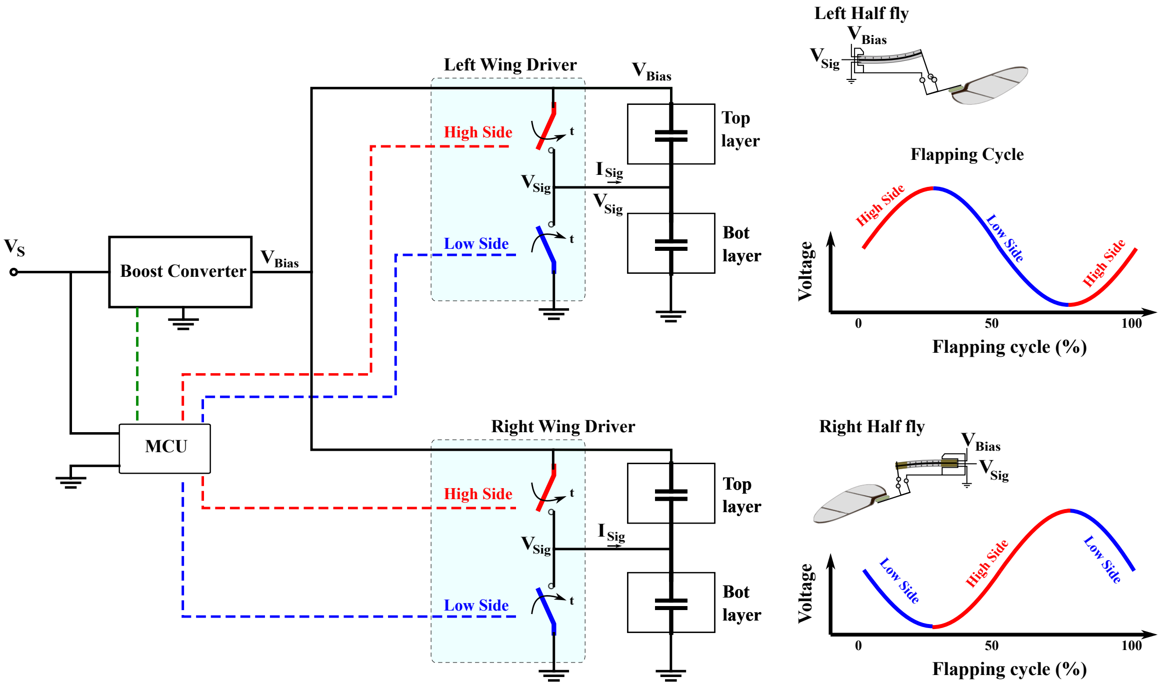

Figure 2.5: A simplification of the simple and light but inefficient FIR power electronics topology, showing the basic operational idea of driving the piezo electric actuators. In hydraulic analogy, the boost converter is similar to a pump providing a high pressure supply (the bias rail V_b) which is then used by a set of “highside” valves for two separate wing channels to increase the pressure (voltage) in the piezo during one part of the cycle. In the other part of the cycle when we desire to decrease the voltage, the “lowside” valves (switches) drain the actuator by connecting the actuator back to the reservoir (ground).

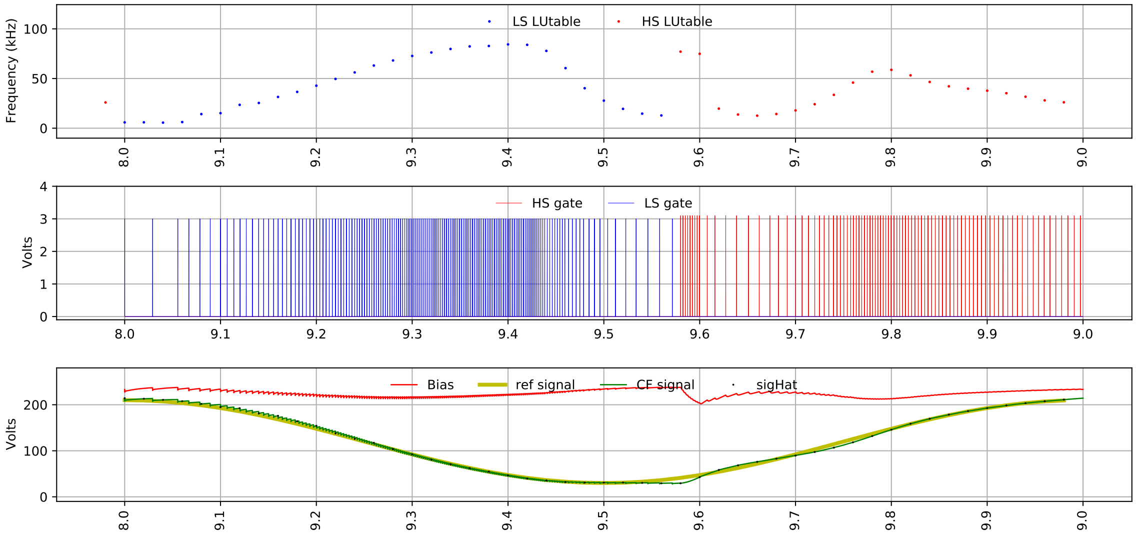

Figure 2.6: [Top] An example pulse train for the linear driver highside and lowside outputs. [Middle] The pulse trains corresponding to the commanded frequency set at each sample of the lookup table (50 controller samples per cycle). There are many pulses possible within a single control sample period, since the digital timers are running at *much* higher frequency than is the digital controller. Each pulse period is an integer multiple of the system clock period, typically 0.3–5 μsec. The pulse period is adjustable to change driver authority. [Bottom] The output waveform from the simulated circuit with a high-impedance bias voltage source representing the boost converter for improved simulation speed. Significant periodic loading of the bias rail is somewhat compensated by the repetitive controller prior to convergence in this case.

Chapter 3 Power autonomy: flight weight wireless power sources

Onboard power is ultimately necessary for untethered flight. The main alternatives for onboard power considered in this work are solar PV, wireless power transmission (WPT) by magnetically coupled resonators, and laser power beaming. Efforts at these possibilities are discussed below, and MCR is preferred as the best WPT option due to greater safety and reliability.

3.1 Laser power beaming solution for first wireless takeoff

Figure 3.1: Robofly used for [30] equipped with PEU and PV cell, with visible laser showing an example of how the near-infrared beam is aimed at the PV cell for wireless power delivery.

Laser power beaming was the first choice in wireless power options. The first wireless takeoff of an FIR in [30] is a major result for the field achieved during the course of this doctoral work. This section will describe that effort, including the power conversion and drive circuitry, laser power beaming method, and takeoff result.

One potential vision of an intermediate step toward autonomy using the laser power-beaming is depicted in Fig. 3.2

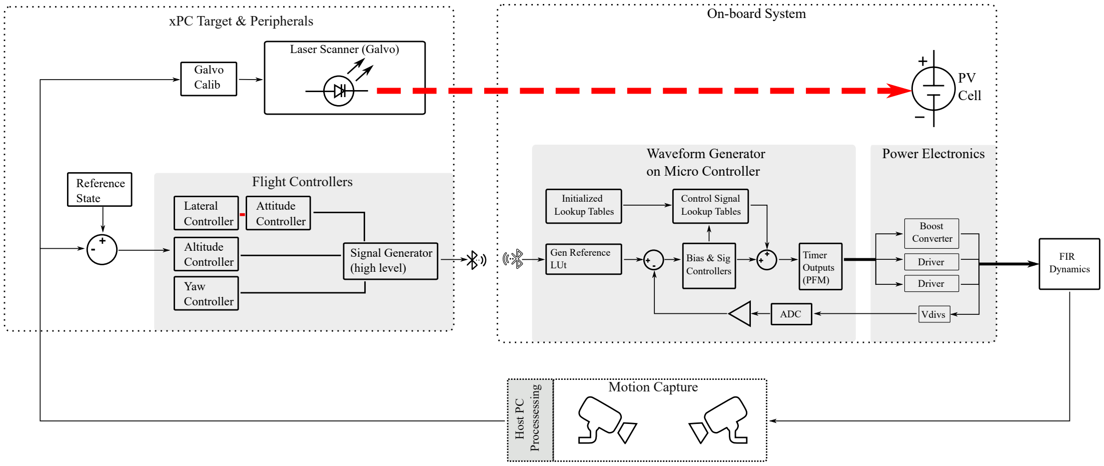

Figure 3.2: Flight control system for fully wireless 2-wing Robofly in motion capture arena, powered by optical power transmitted by laser. Notably, the wire tether in Fig. 1.2 has been replaced by the wireless communication link from the flight controllers running on the xPC target to the onboard signal generator. The signal transmitted wirelessly is envisioned to be high-level commands such as the amplitude, DC offset, and second harmonic content of the sinusoidal voltage signals. Wireless power transmission to the robot is done by steering a laser to the onboard PV cell array.

3.1.1 Power conversion and piezo drive

The oscillating motion of the bimorph piezo beam actuators that flap the wings must be driven by a sinusoidal, high-voltage signal. This must be in the range of 200–300 V to maximize the power density produced by the actuator. The need for low weight and high efficiency strongly influence our design.

Efficiency would be improved if the actuator and wing assembly operated at both electrical and mechanical resonance. However, because the capacitance of the actuators is approximately 5 nF, an inductance needed to achieve electrical resonance at the flapping frequencies would be prohibitively heavy. Therefore, a design goal is that the sine wave should operate at a user-programmable frequency near the mechanical natural frequency of the actuator-wing system, 140–170 Hz [26].

We employ a design geared to bimorph actuators consisting of a constant, high-voltage bias signal and a sinusoidally varying signal channel, following the approach of [38]. An example signal appropriate for this configuration is shown at the top of Fig. 3.3.

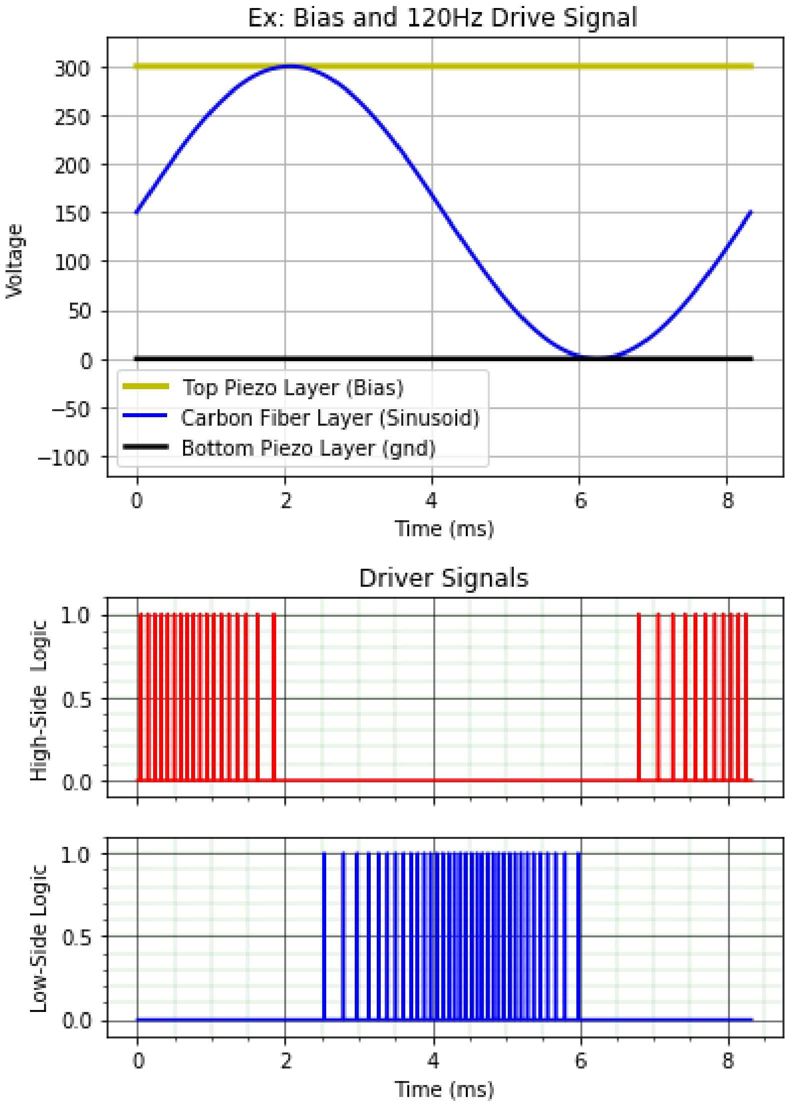

Figure 3.3: (Top) Example target waveforms for boost converter output (yellow) and sinusoidal driver output signal (blue); (Bottom) Example of driver pulse train varying by pulse frequency (PFM)

3.1.2 Circuit design

Commercially available piezo driver ICs (e.g. Texas Instruments DRV2700) cannot produce the required voltage for FIR waveforms. Commercially-available integrated solutions such as the PD100 (Piezo Drive, Callaghan, Australia) are too heavy at 500 mg. While a monolithic SoC design is the typical solution for reducing the size and weight of electronics, we instead focus on a simple switched mode design built with off-the-shelf components.

This approach has a number of advantages. First, it allows for greater design flexibility and rapid prototyping which is important considering that insect robots are still an active area of research with frequent design changes. For example, optimizing an IC for a particular actuator design precludes further improvement in that domain or later incorporation of additional features such as energy recovery mechanisms. Second, designing a single SoC solution that integrates the high voltage actuator drive circuitry presents a design tradeoff. Digital circuits for processing can take advantage of device scaling to operate at low power and reduce size. However, these devices cannot tolerate the high voltages needed for the drive electronics. We instead choose to use a commercially available ARM microcontroller (STM32F051) to implement the timing and control which allows us to leverage the plethora of commercial products which are thoroughly tested and highly optimized for low power and size.

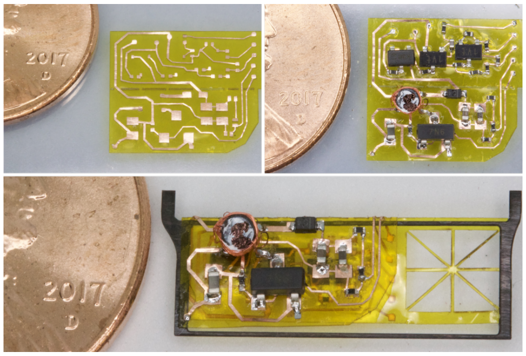

Figure 3.4: PEU at several stages in the circuit fabrication process. Top left: bare unpopulated circuit. Top right: circuit populated with components including coupled inductor. Bottom: Assembled PEU with boost converter and driver, ready to mount to robot.

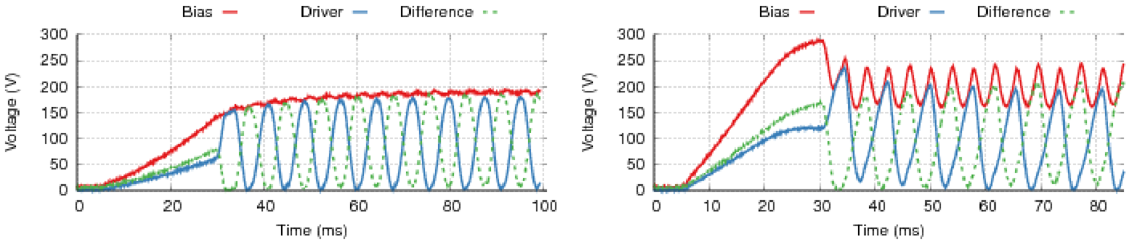

(a) Driver with 100 nF capacitor

(b) Driver with 660 pF capacitor

Figure 3.5: Waveform output by onboard driver effecting modest amplitude sinusoid of controllable frequency by high-side and low-side control signals with a 100 nF capacitor on the bias rail in 3.5a and a 660 pF capacitor in 3.5b.

A schematic of our boost converter is shown in Fig. 4.1. The switching mode boost converter switches electrical current through a coupled inductor with a high turns ratio at frequencies above 100 kHz [38]. The switch control signal is generated by the pulse width modulation (PWM) output of a microcontroller and connected to the gate of the MOSFET M1 in Fig. 4.1. Current through the primary winding of the coupled inductors stores energy in the magnetic field which is transferred to the secondary winding. Brief high voltage pulses on the output of the secondary winding after the MOSFET switches rapidly to nonconducting state are rectified through a fast diode. The diode's output charges a high voltage capacitor for storage and this output is connected to the load. The load in this case is the driver circuit, which linearly regulates the center node of the bimorph actuator in order to drive sinusoidal displacement. The bias and driver waveforms seen in Fig. 3.5 are connected directly to the top piezo surface and the carbon fiber layer of both the bimorph actuators respectively.

The driver circuit is designed to source or sink current at a sinusoidal rate to the center node of the bimorph actuator in simultaneous drive configuration. Transistors Q1 and Q2 are configured as a two stage amplifier designed to source current from the bias rail, implementing the "high side" to increase the sinusoid to its maximum voltage. Transistor Q3 generates the "low side" of the waveform by sinking current from the center node of each actuator to ground.

We use bipolar junction transistors (BJTs) as opposed to FETS standard in current aerial microrobotic research [43] simply due to the ability to tolerate higher voltages than FETs in somewhat smaller commercially available SMT packages, and simple gate biasing design for both linear operation and the rapid pulsing required for potential use of inductive energy recovery schemes [38]. Since the required sinusoid is at a low frequency compared to the clocks of microcontrollers, we generated the sinusoid using pulse width modulation, as depicted in Fig. 3.3. Since the microcontroller clock frequency is a significant factor in its power consumption, the frequency of the pulse width modulation was selected to generate a sufficiently smooth sinusoid with adequate PWM resolution, without excessively high internal oscillator frequency. Future work will investigate whether the DACs built into microcontrollers, PSoCs, FPGAs, or even passive oscillator circuits could generate the waveform in a way that improved efficiency.

Because actuator displacement depends on the voltage difference between the piezo layers and the center carbon fiber layer, our goal is to maintain a constant high voltage bias while the sinusoid varies over time. Dynamic common mode control of the bias rail as in [43] is desirable but must be conducted carefully so as to achieve the correct sinusoidal driver output which is effectively equal to $V_{bias} - V_{signal}$. The complications of common mode control are evident in Fig. 3.5b Because the actuator load varies dynamically during flight and the input power source may be unstable, we design a feedback controller to help regulate a constant bias voltage. We use a simple resistive voltage divider to reduce the bias voltage to within the 3 V operating range of the microcontroller and use its ADC to digitize the value. Based on the ADC reading we adjust the duty cycle using a basic proportional controller in addition to feedforward terms anticipating dynamic load increases of the driver throughout the low frequency actuator cycle.

3.1.3 Circuit Performance

Fig. 3.5 shows the high voltage output waveforms generated by open-loop boost converter and driver for different storage capacitor values. As expected, a large 100 nF capacitor as shown in Fig. 3.5a produces a very consistent bias output thereby reducing the need for feedback control to maximize wing displacement. In contrast, the waveform with a smaller 660 pF capacitor varies noticeably, challenging feedback control and complicating the driving of consistent waveforms. While a large capacitor reduces bias variation it comes at a cost of 16 mg of weight. These waveforms also demonstrate two other important results. First, in both cases the circuit produces an amplitude greater than 170 V which we verified experimentally is the minimum required to lift our MAV. Second, the sinusoid waveform is smooth and does not have abrupt discontinuities that could potentially stress and damage the actuator.

3.1.4 Laser power transfer

Achieving wireless liftoff requires powering all the above components. Our robot requires 200-300 mW of power for liftoff and requires a total of 25 mA of current.

The required energy density and peak current draw are beyond the capabilities of existing battery technologies, but a potential alternative is to use super-capacitors. A 7.5 mF capacitor for example has a maximum voltage of 2.6 V [24]. This voltage however is insufficient to run the boost converter even in simulation. A series parallel combination of 4 supercapacitors could theoretically provide power for 250 ms of flight at which point the capacitors would discharge from the total 5.2 V to below 4 V at which the boost converter stops functioning in empirical evaluation. Perhaps even more important than their inability to support sustained flight though is their combined weight of 96 mg which is greater than the weight of the entire boost converter.

3.1.5 Optical wireless power transfer

Since on-board energy storage cannot meet the requirements for flight, we look to wireless power technologies instead. A practical wireless power solution for an insect scale robot must meet two criteria: 1) it must be able to deliver the 250 mW of power required for flight, and 2) it should have an operating range that allows for flight. Near field magnetic induction can provide efficient power delivery and have been demonstrated for walking robots [35], however its range is fundamentally constrained to tens of centimeters. Far-field microwave approaches (e.g., Wi-Fi) can operate at longer ranges but suffers from efficiencies less than 1% due to RF path loss [56]. We instead select an optical approach as lasers provide a collimated beam with high power density that can be harvested by photovoltaic (PV) cells with conversion efficiencies of over 20%.

Our laser power delivery system consists of a 976 nm laser source and a photovoltaic (PV) cell. For our laser source we use the MHGoPower LSM-010 976 nm laser source capable of providing 10 W of optical power. We connect the fiber output to a collimator (Thor Labs F220FC-980) to produce a beam in free space. An ideal laser should produce a perfect collimated beam that does not diverge in space, however the internal focusing optics of this laser and the use of multi-mode fiber causes measurable beam divergence in space.

Unlike typical PV cells designed to harvest broad spectrum solar energy, our system should be optimized for a single wavelength and high power densities. We therefore select a vertical multi-junction PV cell (MH GoPower 5S0303.4) [48] which consists of serially interconnected p-n junctions bonded together to form a small PV array with low series resistance that performs well under high intensity light [48]. The PV cell measures 2.88 mm x 2.95 mm and weighs 8 mg with an additional 5 mg of wires. This is well within the size and weight constraints of our Robofly. We find experimentally that at intensities up to 20 W/cm^2 the cells achieve maximum power output when operating at 8.8 V with efficiencies of up to 40% for short pulses.

While a power source for liftoff only requires a limited range, we note that laser power beaming can be extended to longer ranges. For example, our laser can deliver sufficient power to the Robofly up to ranges of 1.23 m indoors. This range is not fundamentally limited, but rather determined by the beam divergence and output power of our specific laser source. At ranges beyond 1.23 m, our beam expands to a point where insufficient power is available over the small area of cell. Thus, we can in principle achieve tens of meters of range using commercially available lasers with higher output power or a more collimated beam.

While lasers are capable of powering the Robofly, their use raises other practical questions as well. First, to maintain flight, the laser must track the Robofly. Although tracking is beyond the scope of this work, potential solutions include using motion capture systems as demonstrated in [46] for robot control to track the position of the Robofly and direct the laser appropriately using a device like a galvo mirror. Additionally, we can simplify this problem by placing the laser on a ceiling or floor requiring it to move along only 2 axes. Another alternative to vision based approaches is to use optical feedback from a device like a retroreflector. By placing a light weight retroreflector on the Robofly, we can use an additional laser to verify alignment to the robot. Such tracking systems could be attached to fixed or moving chase vehicles acting as laser power base stations.

Second, 976 nm laser radiation at the levels required for flight are above safe exposure limits. While the area within the laser may not be safe, we can exploit the fact that laser's power is highly focused, therefore guaranteeing that the unsafe areas are limited to the beam itself [25]. By using one of the tracking methods proposed above, we could recognize humans before they enter the beam to immediately turn off the laser source, thus complying with the exposure limits.

3.2 Implementation and evaluation

We begin by describing the mechanical structure of our Robofly MAV followed by a detailed description of our setup for flight experiments and discussion of results.

3.2.1 Robofly design and fabrication

The basic principle of the University of Washington (UW) Robofly design [10, 9] inherits from [63]: a bimorph piezo cantilever actuator amplifies the small field-induced strains into relatively larger motions at the tip. This is further geometrically amplified by a transmission structure consisting of flexure joints to attain a $\approx 90^\circ$ stroke amplitude. The wing's angle of attack is allowed to rotate passively around a torsional spring consisting of a flexure at the base of the wing, resulting in a simple mechanism that produces insect-like wing kinematics. The airframe consists of a single folded structure made from laser micro-machined unidirectional carbon fiber composite bonded to polyimide flexural material. Slight changes to transmission/flexure geometry have greatly increased lift [26].

In addition to simplifying fabrication relative to [46], by employing only a single part for the airframe, our Robofly has a different arrangement of piezo actuators, that are oriented

Component	**Weight (mg)**
DC-DC Converter & Driver Subtotal	**73.7**
Coupled Inductor	21
MOSFET	9.2
V_sCapacitor	2.6
Diode	1.5
Driver Transistors	17.5
Cu traces	6.3
Circuit Substrate	10.0
Assorted Resistors	0.4
Solder & Conductors	3.0
Carbon Fiber Frame	2.2
MCU Assembly	**17.5**
PV Cell & Leads	**13**
Robofly without PEU	**73**
Misc. Glue & Wiring	**13**
Total Robot Weight	**190**

Table 3.1: Total weight of all robot components for RoboFly version in [30] including body, power electronics, microcontroller, and PV cell.

horizontally. This facilitates easy integration of electronics directly below, as shown in Fig. 3.7. The position of the electronics package was chosen to facilitate assembly and rework while avoiding adverse impacts on thrust and stability. The PV cell is positioned above the robot to achieve a direct line-of-sight path to the laser source. Without position tracking, liftoff will move the cell out of the laser beam and cut off power to the fly. We therefore assume a small flight altitude and position the cell at a height of 20 mm above the fly body. To drive the piezo bimorph to produce wing-flapping oscillations requires a roughly constant high-voltage bias signal, a ground signal, and an oscillating drive signal that is

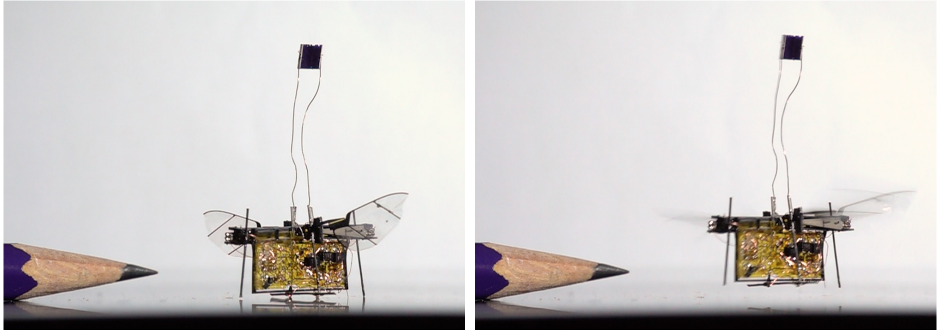

Figure 3.6: (Top) The 190 mg RoboFly and power system before liftoff. (Bottom) After the laser is powered on, power reaches the robot through photovoltaic cell at top. Onboard electronics generate the waveform to drive the wings, causing the robot to lift off. After liftoff the robot is no longer in contact with its reflection on the surface below.

roughly sinusoidal, as shown in the top half of Fig. 3.3.

3.2.2 Setup and takeoff results

To demonstrate wireless liftoff capability, we position the fly at a distance of 1 m from the collimator output. With the beam divergence of the laser, this results in a 13 mm spot size at the PV cell which is more than sufficient to cover it. We design a $0.6 \times 0.75 \times 0.6$ m enclosure and use a series of 2 mirrors to achieve the 1 m distance and to align the beam on the cell. We program the microcontroller to flap both wings continuously at maximum possible amplitude at a frequency of 170 Hz using a single driver circuit to maximize lift. Because the fly is dynamically unstable and our goal is simply to demonstrate liftoff, we attach a carbon fiber rod across the base of the fly in order to minimize risk of structural damage during repeated experiments. We perform experiments by placing a digital camera inside the laser enclosure recording at a 240 Hz and apply short pulses from our laser power source.

Prior to performing flight experiments with the electronics attached, we verify that the fly is capable of liftoff when driven by a 190 Vpp sinusoid at 170 Hz with a 130 mg toothpick attached as a dummy load. Table. 3.1 shows the final weight of the power electronics amounts to 104 mg, which is well within this weight budget. We also verify the functionality of the full system by measuring the output of the electronics prior to final attachment on the fly while powered by the PV cell and driving the actual fly. We measure that the output of the PV cell is capable of providing over 250 mW to supply the power demands of the boost converter, drive circuit and microcontroller. Operating the boost converter at 150 kHz with a 6% duty cycle yields unloaded bias voltages of over 250 V as shown in Fig. 3.5. In this configuration, despite variation in the bias rail due to use of a smaller storage capacitor, the voltage difference between the bias and sinusoid is more than 170 V_{PP} at 170 Hz.

As seen in Fig. 3.6 and the supplementary video [51], we demonstrate a completely wireless RoboFly liftoff using only onboard electronics and wireless power transfer. We note that the altitude of the flight could be easily improved in future experiments. Specifically, the prototype fly shown in Fig. 3.7 includes a variety of fabrication errors and repairs which may have made flight even more difficult. Additionally, lighter components for the boost converter such as a sub milligram single chip voltage regulator to replace the multicomponent shunt regulator used for the microcontroller, a lighter 5 mg MOSFET as well as laser micromachining to remove unpopulated areas of the circuit substrate would easily improve the payload margin. Performance could further be improved by reducing the power consumption of the micro-controller using low power optimizations on the chip or using alternative chips available in the same size or smaller thereby allowing a greater fraction of the total laser power to be delivered to the wings. These weight reductions could allow the use of a larger storage capacitor which would improve the boost converter and driver performance thereby increasing wing stroke amplitude and lift.

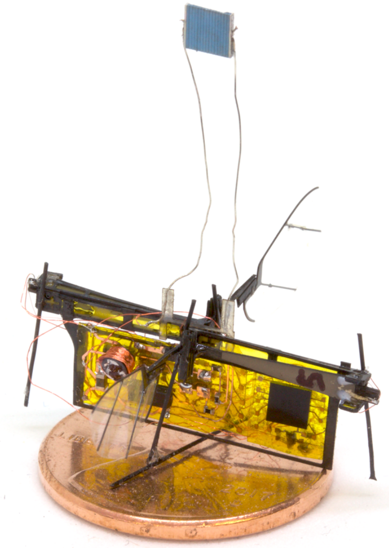

Figure 3.7: Full insect scale robotic fly placed on a US penny for scale. The power electronics and microcontroller are below the robot and the PV cell is 20 mm above it.

3.3 Solar photo-voltaic direct power

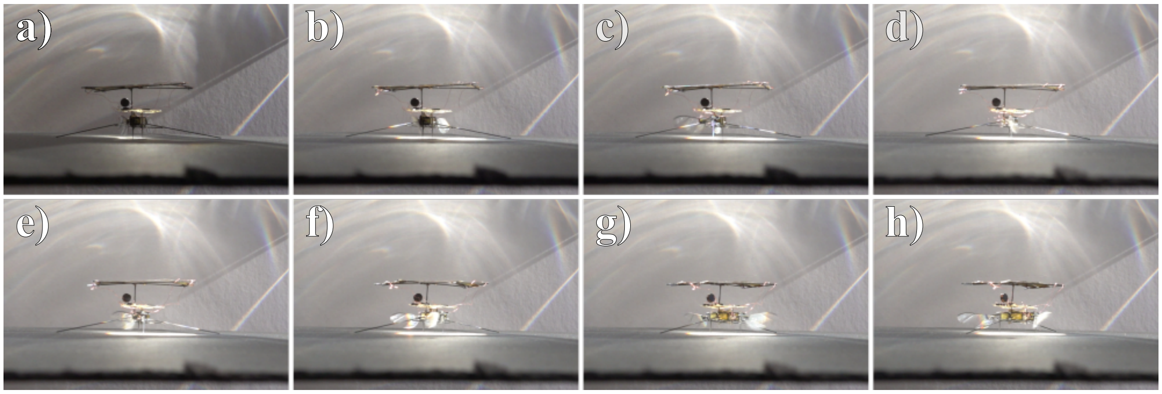

Figure 3.8: Progressing from upper left (a) to lower right (h) time sequence over 0.15 seconds during a flight attempt under PV power. The Robofly Expanded is illuminated under concentrated natural light in February. The light was first reflected to near perpendicular, then concentrated by acryllic Fresnel lens.

One experiment attempting solar-powered wireless takeoff is shown in Fig. 3.8. Improvements could be made to achieve lift-to-weight > 1, in particular in the solar spectrum applied to the PV cells in indoor conditions which was dominated by IR due to the warm-up period of the halogen bulb, causing PV cell output power to be insufficient. Pneumatic and solenoid powered shutter systems were made to allow the light to warm up, but due to heating of components could not safely allow the light to fully warm up. An optoisolated relay and triac circuit was used to control the AC lights from low voltage datacollection system. A Fresnel lens was used to concentrate the light, but this unfortunately prevented complimenting the light spectrum with LED component. Other factors were that the light was not well positioned in outdoor apparatus. The largest problem areas were in the fundamental inefficiency of the boost converter + half-bridge driver PEU, small wing size, and predominantly that the onboard PEU controller did a poor job of controlling the waveforms to flap the wings effectively.

In particular, the waveforms produced by the PEU lacked good control over the bias voltage and the signal voltage (delivered to both wings) which is especially needed given the loading of the bias rail by the driver and varying input power. Those deficiencies and others caused the robot to yaw substantially, as can be seen in the rotation as the series progresses is Fig. 3.8.

3.4 Wireless power via magnetically coupled resonators

Wireless power transfer (WPT) via magnetically coupled resonators was used to power an untethered Robot complete with receiver and power conditioning circuitry, single-wing PEU, and microcontroller. A transmitter and controller loaned by WiBotic, Inc. was used. MCRs were successfully used to power the robot despite unrelated failure to fully lift-off, and represent a milestone toward controlled wireless flight of a sub-gram flying insect robot (FIR) powered through a human-safe wireless power transfer method. Previous wireless flights have been powered by light [30, 27] and were not under feedback control. Controlled flight to-date has required a wire tether for actuator drive signals, but a wire greatly diminishes the utility of the robots for envisioned applications such as insect flight research or agricultural support. Obtaining power onboard presents a challenge as available battery technologies do not meet the stringent size, weight, and power (SWaP) constraints of a flying robot weighing slightly more than a toothpick. As a result, obtaining such power autonomy for insect sized flying robots remains an open challenge. Our interest in magnetically coupled resonators (MCR) for wireless power transfer is motivated by two significant advantages over laser power beaming: inherent human safety, and the ability to bypass visual obstructions. I report scale-compatible MCR wireless power transfer and a voltage power conversion for piezoelectric actuators that is nearly sufficient to attain liftoff.

Insect sized flying robots represent a frontier of both opportunities and engineering challenges arising from physics scaling effects of shrinking down the characteristic dimensions flying robots. Large numbers of cheap, agile, and low-power robots are envisioned as a promising alternative to few, costly, and power-hungry robotic platforms [3]. Amongst advantageous scaling effects, the large strength-to-weight ratio observed in ants is an important example. But there are other equally advantageous, if less well known, scaling effects to leverage. These include low total power requirements and the concomitant potential to tap into ambient low-power energy sources such as WiFi [57], high maneuverability, and low materials cost.

While this dream has inspired decades of work since early realizations such as [14, 63], and has separately demonstrated key requirements such as fabrication [62], control [46],

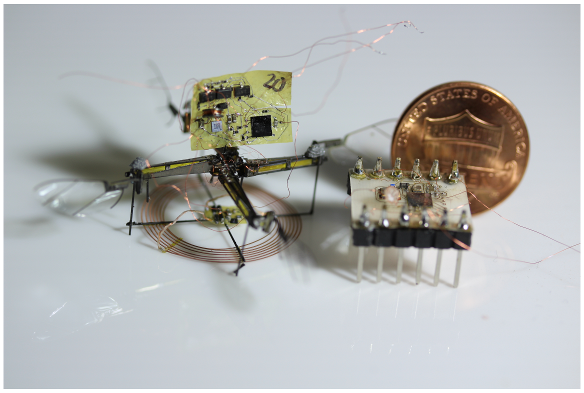

Figure 3.9: At left is a 4-winged Flying Insect Robot (FIR) similar in design to the 2-winged UW Robofly FIR, with receive coil and fully operational electronics suite for untethered flight. To the right is a testing breakout of a Bluetooth capable MCU planned for future implementation, and a US penny for scale.

sensing [20], controlled flight demonstrations involving piezo-actuated insect-sized robots have all relied on a physical wire tether for control and power signals. This wire tether confounds control and measurements of flight dynamics and nullifies the utility of the robots for envisioned applications such as insect flight research or agricultural support.

Obtaining power onboard presents a challenge as available battery technologies do not meet the stringent size, weight, and power (SWaP) constraints of a flying robot weighing slightly more than a toothpick. The key requirement of power-delivery to an untethered FIR has only been shown utilizing onboard photo-voltaic cells and light sources such as lasers or lamps equivalent to multiple standard suns, posing safety risks or unrealistic levels of illumination [30, 27]. In this work I demonstrate usage of magnetically coupled resonators to deliver power to the FIR safely [7], efficiently, and without needing line of sight. The 4-winged FIR platform used in this work is shown in Fig. 3.9. This FIR was developed in [18], and increases payload relative to a 2-winged FIR [9].

Wireless Power Transmission

Given that onboard batteries are not yet feasible, several different methods of wireless power transmission (WPT) were considered to power the Robofly over a range of distance, power capacity, safety, and efficiency. At short range but highly efficient, traditional inductive coupling of coils, as in common-place AC transformers used in AC power distribution worldwide, is limited to very short-range due to the need for a high degree of electromagnetic coupling between non-resonant windings which share a highly permeable magnetic core. While capable of high efficiency, the short operational range and need for a typically heavy magnetic core between coils makes this example of near-field WPT impractical for our FIR. Electromagnetic inductive coupling is also commonly used in wireless charging circuits typical in smartphones or electric toothbrushes. Non-resonant inductive WPT is easy to implement and has adequate power capacity and efficiency for the FIR, but it also requires the receive (Rx) and transmit (Tx) coil to be in close proximity. At higher transmit power inductive coupling also poses safety concerns for humans and equipment due to large electric fields.

On the other hand, far-field radiative power transmission plummets with distance proportional to $\frac{1}{r^2}$ in omni-directional transmission and is unsuitable for powering our FIR during flight. Focused and directional transmission, such as microwave power beaming or phased-array transmission, or laser power beaming [30] targets a receiver and can transmit power wirelessly over great distances and respectable efficiency such as the 40% efficient PV cells utilized in [30]. However, this mode of WPT is highly inefficient if the transmission is not precisely targeted, since electromagnetically radiated power is entirely lost once transmitted. Successful WPT requires precise tracking, nearly perpendicular orientation of the receiver, and clear line-of-sight (LOS). Safety around such radiative WPT systems is difficult to achieve.

3.4.1 Magnetically Coupled Resonators

At an intermediate range, WPT may be achieved using magnetically coupled resonators (MCR) and presents an elegant alternative to the deficiencies of the other WPT options for

FIR application. In WPT by MCR, the Tx, Rx, and any intermediate coils, are coupled magnetically and resonate with each other with minimal interaction with other objects. In contrast, inductive coupling can easily induce large electric potentials on nearby objects and human tissue. WPT efficiency via MCR does not depend on extremely high coupling as does inductive WPT. Furthermore, it can operate efficiently even if the coils are far apart and have lower coupling, since energy not taken by the Rx coil largely remains in the high-Q transmitter coil as opposed to be dissipated in space [53]. It is this fundamental difference in operating mechanism that makes MCR WPT safer and more efficient, within transmission range, than radiative WPT. Additionally, the Tx and Rx coils do not require line of sight and the Rx can change position and orientation without the need for the physical tracking and aiming the transmitter required for laser power beaming. Like all WPT methods considered in this work, the Rx must be oriented approximately parallel to the Tx. However, in MCR, misalignment of Tx and Rx coils during flight can be addressed by active frequency tracking and impedance matching handled by the offboard transmitter described in [53, 61] , thus the power management of the FIR is passive.

Magnetically coupled resonators for wireless power transmission have demonstrated the capability to power a piezo-actuated cockroach inspired ground microrobot in [36]. In that work the Rx coil weighed 180 mg, roughly three times our weight budget. Nonetheless, this demonstration on a ground robot serves as an intriguing and useful benchmark. As in [55], range of a single Tx, Rx can be efficiently increased for a mobile robot such as the FIR without significant loss of efficiency, by deployment of intermediate passive relay coils.

In this work I choose MCR to obtain power onboard the untethered FIR, using a transmitter similar to [36], and designing our custom Rx coil to fit within the stringent size, weight, and power (SWaP) constraints of the FIR.

Robotic platform

I utilize the 2-winged FIR described in [9] and the similar 4-winged variant developed in [18], shown in Fig 3.9. These robots use 3-wire bimorph bending beam piezoelectric actuators to actuate wing stroke, while the angle of attack desired for lift results from drag forces against a passive wing hinge during the flapping wing stroke. The wing is flapped at mechanical resonance near 170 Hz, generating enough lift for the robot and approximately 150 mg of additional payload at the achievable voltage from the onboard boost converter and waveform generator. Precise waveform measurements were not performed as wireless voltage and high-side current measurement apparatus is nontrivial and does not fit within the SWaP budget of the untethered robot. However, during benchtop testing the piezoelectric actuator bias and unipolar sinusoid signals were approximately 200 Vdc and 170 Vpp, and the system including MCU required approximately 250—300 mW input power from the Rx.

The power electronics unit (PEU) and piezo-electric actuators of the FIR are somewhat analogous to the brain, peripheral nervous system, and musculature of biological insects. The MCU must generate the control signals to operate the flight muscles. In our robot, lightweight high voltage amplifier circuitry on the PEU converts low-voltage power delivered by the Rx coil and power conditioning circuit into the high-voltage signals which cause the piezo-electric actuators to strain and produce mechanical power output by wings. The power conditioning unit attached to the Rx coil has the critical task of converting the power from the Rx coil into forms suitable for both the microcontroller (MCU) and the piezo actuator waveform generator. It does this in lightweight form by a full-wave rectifier (4 diodes since not center-tapped) and voltage regulator. In this work both MCU and high-voltage waveform generator comprise the power electronics unit (PEU) shown atop the FIR in Fig. 3.9. The power conditioning circuit was attached directly to the Rx coil and considered separate from the PEU, in keeping with the objective for an interchangeable power source for the FIR. The complete system diagram is shown in Fig. 3.10.

The PEU topology utilized for this experiment is shown again in Fig. 3.10, and is identical to that used in the laser power beaming experiment 3.1.

Receive coil self-resonant design

I make use of the inherent parallel-plate capacitance of a double-layered flex PCB material to produce a light-weight inductor with distributed capacitance so as to have a self-resonant frequency near 13.56 MHz matched to the MCR transmitter. Resonant tanks are feasible

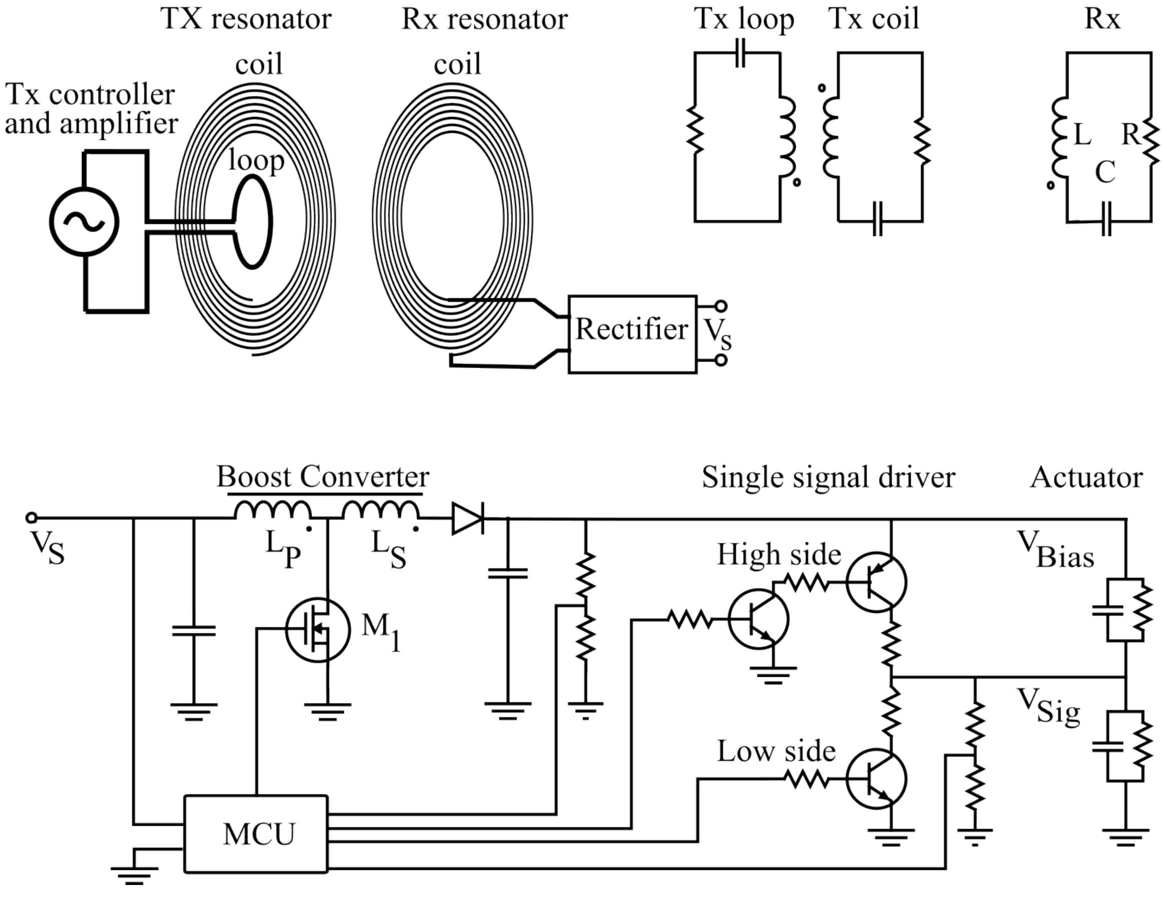

Figure 3.10: : (top left) 3-coil MCR wireless power system with power conditioning circuit at load side. (top right): equivalent circuit of the coupled Tx loop, Tx coil, and series compensated Rx coil. (bottom) Robofly FIR's power electronics unit with one of two signal drivers and piezoelectric bimorph actuator

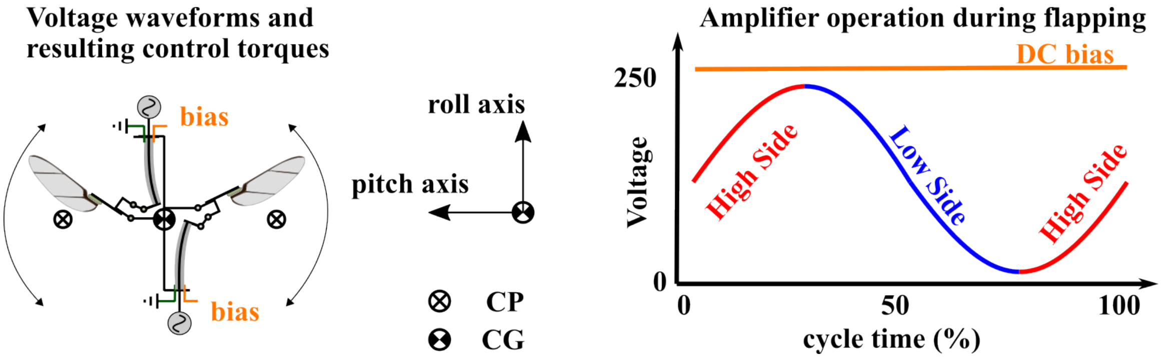

Figure 3.11: : (left) Top view of 2-winged robot showing two bending beam actuators deforming to actuate wing stroke via the 4-bar transmission mechanism. (right) A single unipolar sinusoid and DC bias signal are plotted for a single wing stroke, the red and blue arcs of the unipolar sinusoid indicate the approximate timing of high-side vs low-side, respectively, of the half-bridge amplifier switches.

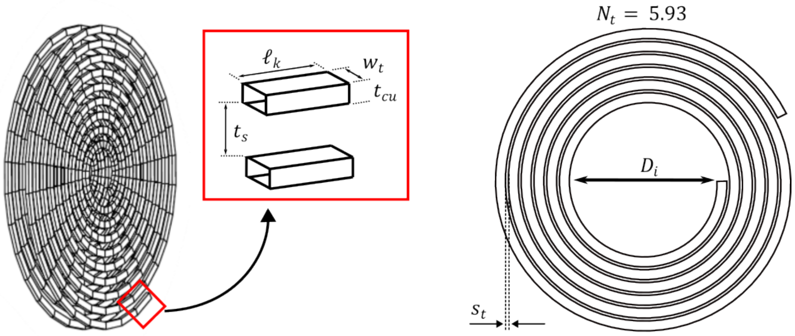

Figure 3.12: : (left) Finite element rendering of a double-sided embedded capacitance Rx coil. Image is adapted from software documentation to illustrate exaggerated substrate thickness. Two conductive trace elements on top and bottom with their thickness t_{cu}, length l_k, and width w_t. The traces are separated by substrate of thickness t_s. (right) 2D vector graphic of a coil used for visualization and manufacture, illustrating the trace separation distance s_t, inner diameter D_i, and number of spiral turns N_t

to implement using a single embedded capacitance substrate with traces on both sides, with top and bottom conductive traces connected for series or parallel arrangements of the compensation capacitor C. The system diagram shown in Fig. 3.10 depicts one possible arrangement of the coil capacitance C and inductance L. The Rx coil should weigh as little as possible while still having a large diameter for increased range, and the coil's distributed capacitance should not require any additional discrete compensation capacitor to achieve designed resonance – thus minimizing weight and additional power dissipation of the compensation capacitor. Further, this coil should have as high a quality factor Q as possible, so as to extend the range of magnetic coupling [53, 41] sufficiently to achieve minimum flying power requirements. Thus the resistance of the coil and skin and proximity effects should be minimized. A double-layered spiral design is selected for ease of parametric spiral generation and to achieve trace separation. For the finite element analysis described in further detail below, each spiral was segmented into linear elements with length and width l_k,w_k, as illustrated in Fig. 3.12.

Equation 3.1 gives the self-resonant frequency ω_{SR} in rad/sec of the Rx coil as a function of its capacitance C and inductance L. Equation 3.2 defines the Rx coil's quality factor Q for a series resonator, with R_{tot} representing the coil's dissipative losses including resistance, skin, and proximity effects. Equation 3.3 gives the coil's sum total capacitance C, simply equal to the summation over k of each element's conductor area $A_k = w_t \cdot l_k$, separated by dielectric of thickness t_s and having dielectric constant ϵ. Equation 3.4 likewise yields the summed weight of the k elements' substrate and copper traces where $A_k, t_s, \rho_s, t_{cu}, \rho_{cu}$ are the cross-sectional area, substrate thickness, substrate density, copper conductor layer thickness, copper density, and substrate dielectric coefficient of each element, respectively.

$$\omega_{SR} = \frac{1}{\sqrt{LC}} \tag{3.1}$$

$$Q = \frac{1}{R_{\text{tot}}}\sqrt{\frac{L}{C}} \tag{3.2}$$

$$C = \sum_k \frac{\epsilon A_k}{t_s} \tag{3.3}$$

$$m_c = \sum_k \left(\rho_s A_k t_s + 2\rho_{cu} A_k t_{cu}\right) \tag{3.4}$$

For modeling fidelity of the parametric spiral coil designs, I obtain the capacitance and weight numerically instead of analytically using the parametrized continuous spiral, in part to stand-in for envisioned finite element solutions to the capacitance along similar lines as the inductance, but also as an inbuilt checkpoint to compare the numerical results to the measured actual capacitance of the continuous analytically generated coil which is ultimately manufactured. Short radius curves and discontinuous turns are undesirable in the traces due to large local electric fields and impedance effects on the Rx coil, and thus I utilized a large number of segments in the finite element representation. I provide no analytical equation for the coil's inductance L or conduction losses R_{tot} coming from resistance and skin and proximity effects, because these are a much more complicated function of the coil's internal electromagnetic interactions, geometry, and material. Thus, while equation 3.1 can be readily solved for the compensation capacitance required for the desired ω_{SR}, the inductance and lumped resistances including skin effect varies simultaneously and are less trivial to predict for arbitrary Rx morphologies. While some analytical formulae for inductance and conduction losses are stated in [60] and provide high accuracy starting points to coil parameters for optimal analytic design, any modeling errors disproportionately impact the down-scaled coil needed for the FIR. Therefore in this work I explore automated finite element analysis for this multivariate optimization problem. The subsection below describes how I numerically solve the inductance L and high-frequency impedance characteristics as part of the Rx coil design.

I performed an optimization of the coil parameters P, which include:.

- D_i: Spiral inner diameter
- N_t: Number of turns
- w_t: Coil trace width
- s_t: Trace spacing

- ϵ: Substrate dielectric constant
- t_{cu}: Conductor thickness
- t_s: Substrate thickness

Inner diameter is a less-typical coil measure than nominal diameter but D_i is varied in this work as a convenient starting point for coil segmentation and since nominal coil diameter and outer diameter are constrained according to N_t, s_t. The Rx coil in a MCR WPT system must have high quality factor Q and resonate at the designed frequency in order to magnetically couple over meaningful distance to the Tx. The Rx coil on this robot in particular must be both light-weight and provide sufficient power output to the PEU at range. Since range, efficiency, coupling, and power transfer capacity are also dependent on the ratio of Rx and Tx coil diameters, the design must strike a balance between competing goals of high Q and power capacity vs. low weight, despite the resonance exhibiting high degree of sensitivity to all of the parameters enumerated above. In this work I utilize a gradient descent optimization routine to design the set of parameters P, given the complexity of the design space.

Guided by analytical Rx coil design principals stated in [60] and basic equations above, I develop and minimize a simple cost function J; a weighting vector W which assigns relative importance of the metrics for coil weight, natural frequency, total inductance, and quality factor Q of a coil; vector X quantifying the coils' computed mass and electrical properties, and I determine the set of parameters to vary. The construction of vectors X and penalty weights W represent an engineering judgment call between the tradeoffs of total Rx coil weight and quality factor, as a higher Q with higher power transfer capacity corresponding to inductance L, competes with the interest of low weight. The cost function J minimized in equation 7 provides us a scalar metric for each candidate coil design. This cost is then minimized by gradient descent as in equation 7 to yield locally optimal candidate coil C_{Rx} (P).

$$W = \begin{bmatrix} j_m & 0 & 0 & 0 \\ 0 & j_\omega & 0 & 0 \\ 0 & 0 & j_L & 0 \\ 0 & 0 & 0 & j_Q \end{bmatrix} \tag{3.5}$$

$$X = \left[m_c, |\omega_{Tx} - \omega_{SR}|, \frac{1}{L}, \frac{1}{Q} \right]^T \tag{3.6}$$

$$C_{Rx} = \min\{J\} = \min\{X^T W X\} \tag{3.7}$$

Where j_m, j_ω, j_L, j_Q are the weighting terms for total coil weight, natural frequency, inductance, and quality factor, respectively. ω_{Tx} is the operating frequency for transmission which is selected similarly as in [53] for coupling range, minimizing undesired magnetic interactions with surroundings, and Tx control via active frequency tracking and impedance matching. To provide an educated starting guess for appropriate weights, the weights were normalized to characteristics of coils from [60, 36]. This is assumed to be approximately optimal and the algorithm will converge to more optimal values for the FIR if possible.

Since the cost function of the multi-variate coil design is unknown but assumed nonconvex, multiple starting points for coils across the design space, e.g. large coils, skinny coils, range of spiral arc, trace spacing, etc, are chosen as starting points; ensuring that the optimization does not get stuck in sub-optimal local minima. In particular, I found that Rx coil Q and m_c were particularly sensitive to coil diameter and trace-width, and thus I iterated starting conditions across the full range of these variables. Other important design variables such as substrate thickness, different options of substrate with different dielectric constants, series vs. parallel connection and multiple coil layers, and copper cladding thickness are not suitable for continuous variation by the gradient descent, so gradient descent was performed within each of these sets of discrete options in order to formulate a comprehensive set of candidate designs.

Implementation

I implemented gradient descent of equation 3.7 using a custom Python script. During each candidate coil descent, the program varied coil design parameters from an initial guess P_0, generated a design file detailing the finite element model, called FastHenry 3D inductance solver [34], evaluated the resulting inductance analysis file, and chose the optimal parameter variation for subsequent descent iterations while saving locally optimal design candidates.

The FastHenry solver utilizes a 3D specification of geometry and properties of the coil whose finite elements (segments) of the coil are defined by the Python script and encoded into a template inductor file. A FastHenry rendering of a segmented coil design for the finite element analysis is shown in Fig. 3.12. The results of the inductance analyses were parsed and used to compute the cost gradient, and the algorithm chose the best available variable change and began a subsequent iteration. If the cost gradient converged to a low value then the optimization ceased and the coil design was chosen as a locally optimal candidate. Amongst locally-optimal candidate Rx coil designs, several designs featuring lowest cost were used by the same program to generate a scalable vector graphics file for visualization and a CAM file for manufacture.

Manufacture was done using 355nm frequency-tripled Nd:YAG Diode Pumped Solid State Laser (DPSS) micro-machining system, the same equipment used to fabricate the PEU flex PCB and the robot chassis. The Rx coils were realized by two methods: Firstly, spray polyurethane PCB etch resist was rastered off the double-clad flex PCB material by DPSS laser, then unwanted copper removed by chemical etchant. This raster-etch (RETCH) process is used for in-house fabrication of the PEU flex circuits as well, but a coil fabricated this way still has a large area of the material's bare substrate which serves to maintain the coil shape, but does not contribute electrically and adds 20-50% additional weight. Additionally, over-etching the RETCH'ed coil potentially caused deviation from design parameters. In the second method, I directly through-cut the coil design from the embedded capacitance flex PCB material. The second method was preferred for lighter weight but required insulated carbon fiber support rods to span the coil and maintain the precise designed shape as well as provide mounting points for robot assembly. In both methods,

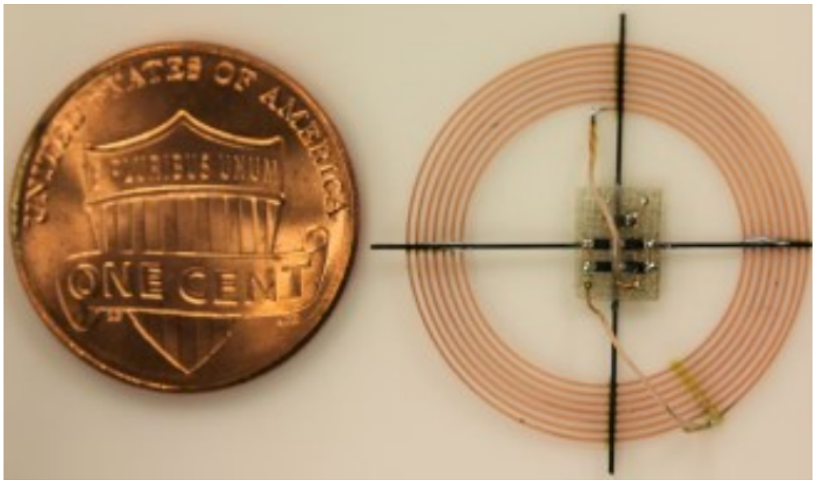

Figure 3.13: Rx coil and power conditioning circuit mounted to carbon fiber rods, shown next to a US penny for scale.

fabrication inaccuracy of the photo machining due to DPSS laser spot diameter of 20 μm, imprecision of the photo machining vector path planning, and xy positioning error of laser scanning was assumed to be negligible.

While several candidate designs were implemented using the 3M C2006 embedded capacitance material with 6 μm substrate and 35 μm copper thickness, it was found that the copper adherence to substrate made the material delicate to handle and susceptible to tearing and delamination, ruining the coil. Polyimide substrates from the DuPont Pyralux series with 12 μm – 25 μm substrate thickness and 12 μm - 25 μm copper trace thickness were also found to have suitable dielectric constant with similar degrees of repeatability but superior copper adhesion. Thus in this work I favored the Pyralux flex PCB material used in our PEU fabrication.

Manufactured candidate coils were tested using an LCR meter to confirm designed inductance, capacitance, and resistance. The coils were then adhered to insulated support structure, and attached to the power conditioning circuit by Litz wire approaching at perpendicular to the coil, so as to minimize unintended magnetic interactions with the WPT. An assembled coil and power conditioning circuit is shown in Fig. 3.13.

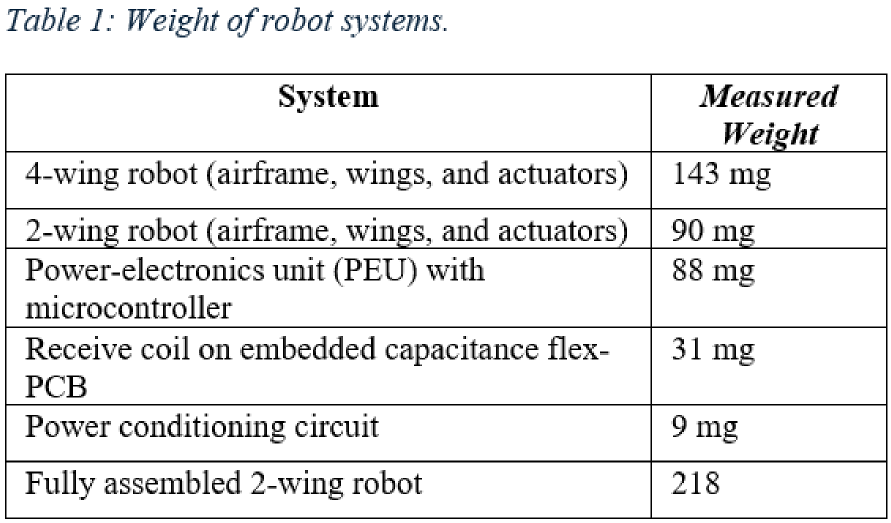

Table 1: Weight of robot systems.

System	***Measured Weight***
4-wing robot (airframe, wings, and actuators)	143 mg
2-wing robot (airframe, wings, and actuators)	90 mg
Power-electronics unit (PEU) with microcontroller	88 mg
Receive coil on embedded capacitance flex-PCB	31 mg
Power conditioning circuit	9 mg
Fully assembled 2-wing robot	218

Figure 3.14: Weight of robot systems.

Experimental results

Assembled robots for untethered experiments were constructed in 2-wing and 4-wing variants, with several selections of optimized candidate coils. Table 3.14 shows the weights of the systems.

A fully untethered FIR powered by the Rx coil is shown in Fig. 3.16, with transmitter placed approximately 25 mm underneath the acrylic surface. Lift was enough to move across the surface but not to take off. Our design and simulation predicted slightly higher $|s_{21}|$ than achieved in the best performing coils. The primary hindrance to sustained flight in this case was limitations in the PEU efficiency and FIR wear. More generally, the Rx coil design is difficult to achieve in practice since the capacitance and self-resonance of the ultra-light Rx coil are extraordinarily sensitive to the parameters of coil design such as dielectric properties which can be impacted by manufacturing variability, soldering heat or pinching during assembly, excessive solder, heat during operation, or flexing and vibration. For example, assuming uncorrelated changes to L and R, even small δC variations in capacitance, $C = C_{nominal} \pm \delta C$, significantly impact Q , especially given low nominal values of L,C obtainable with SWaP budget of the FIR Rx coil. It is possible that smearing of the traces across the

Figure 3.15: The measured scattering parameter $|S_{21}|$ diminishes with distance. Decreased efficiency at lower distance is compensated by active frequency tracking and impedance matching.

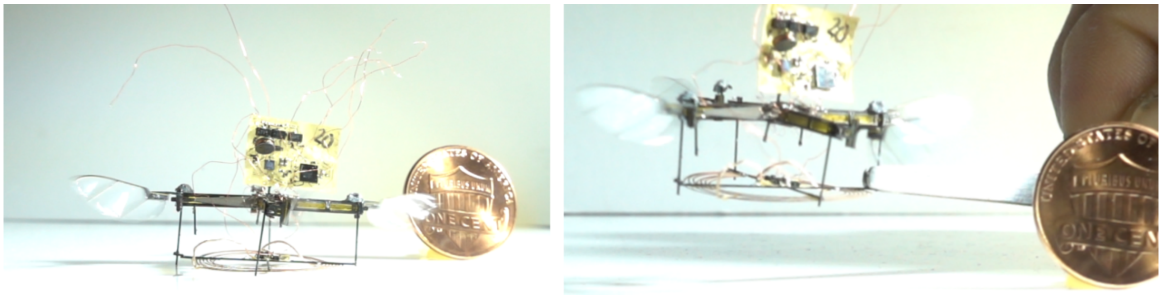

Figure 3.16: Still-frames of a fully assembled 4-wing robot recording flapping its wings powered wirelessly at up to 1.25 inch distance from Tx coil located underneath the acrylic platform. (right) the robot is gripped by tweezers and lifted above the platform, still receiving enough power to wake the PEU and flap the wings. Here, the 51-AWG wire tether I used for measurement and programming has been cut.

dielectric impacted the coils as well. This can be caused by laser through-cutting of the coils leaving microdeposits of conductor spanning the dielectric where cut, a phenomenon similarly impacting piezoelectric actuator fabrication during laser micro-machining.

Regarding heat dissipation, several hundred mW of power are received by the Rx coil during the WPT flight experiments, which represents a large power density. Since the coil must be light, there is both minimal conductor surface area available to cool the coil and little thermal capacitance to begin with. Resulting temperature fluctuations may change the electrical properties of the coil, adversely impacting the high quality factor Q needed for MCR WPT.

Conclusions for MCR feasibility

The coil design process can be meaningfully improved in several ways. Firstly, simulating coupling to Tx and relay coils in addition to the Rx as well as utilizing finite element analysis of the capacitance of the Rx coil to supplement the analytical calculation would improve accuracy. Incorporating finite element analysis of parasitic capacitance was not done in this work but would be particularly useful to optimize around skin and proximity effects. Lastly, the cost function can readily be augmented with additional or alternative metrics for particular designs, e.g. including expressions in X along with corresponding weights in W which would penalize design sensitivity to uncertainty in parameters, rewarding desired inductance so as to specify Rx output voltage, or including terms which are more directly related to impedance and power capacity of the Rx. It's difficult to quantify the effects of physical deformation of the Rx coil during operation and from handling during manual assembly. Modifications to the fabrication process to improve this aspect could include adhering the spiral coil immediately to a support before removal and handling, and switching to a support with more radial attachment points than the four used here. Nonetheless, power electronics improvements to better utilize the MCR WPT are the strongest indication.

WPT using MCR is shown to work to power the robot sufficiently to flap the wings and nearly take off despite inefficiency in the power electronics, actuator breakage, and undesired torque biases in the robot during flapping which reduced lift. A more power efficient power electronics within the SWaP budget of FIRs is an ongoing design effort which will unlock successful demonstration of FIR powered by MCR. Regardless, MCR power worked impressively on a robot too small to carry batteries and represents perhaps the lightest demonstration of this human-safe approach yet. Additionally, the Rx coil itself provides incidental benefits such as shielding the wings from collision and easing the burden of flight control somewhat by adding lightweight angular inertia to the otherwise very fast dynamics of the FIR. I anticipate that only a few refinements will be necessary to attain free-flight under this form of WPT.

3.5 Conclusion

The laser powered liftoff result presents a significant milestone towards the achievement of flight autonomy for honeybee-sized insect-scale robots, demonstrating a wireless takeoff. Specifically, we present the lightest wireless robotic flight to date by showing liftoff of a 190 mg robot. We demonstrate the first power electronics package fully integrated into a functional aerial robot and fabricated by a unique application of laser micro-machining techniques for fast-turnaround circuit fabrication. We also successfully demonstrate optical wireless power transfer sufficient to run the power electronics and onboard microcontroller.

Despite not achieving liftoff due to unrelated hardware issues, the work with magnetically coupled resonators also proves functionality of an additional light-weight power source which has distinct advantages in human safety and invulnerability to line of sight obstruction to complement the drawbacks of laser power beaming. Solar power under ambient lighting conditions, remains slightly too costly for the SSWaP budget, requiring at least $\frac{1mg}{mW}$ weight of PV cells for given power requirements, thus putting it slightly out of reach for the Robofly with onboard systems, projected from the weight (including additional wing driver) of table 3.1.

This work serves as a platform for enabling a multitude of new research directions advancing MAVs closer to the vision of autonomous flight. The integration of an onboard microcontroller presents the opportunity to add sensing and communication capabilities which is a necessity for enabling onboard control and achieving extended stable flight. Laser based power for insect scale aerial robots also opens multiple research directions for extended flight including tracking, extending range, and laser-based communication.

Chapter 4 Power Electronics - DC-DC Conversion to Drive Piezoelectric actuators

4.1 Introduction

This chapter discusses discusses the power electronics used in the state-of-art, presents analysis of performance of existing power electronics units for piezo actuation FIRs with an emphasis on considerations for electronics efficiency improvements required to advance the field, and finally presents my contributions to the field of power electronics for FIRs. The new PEU designs presented in §4.3 are of high value to the field, even in an early stage of development, because of several key advantages over state of art which are prerequisite to enable the full autonomy of the Robofly. Again, FIRs such as the Robofly, Robobee, all the way back in time to the Berkeley MFI project [15], have been almost entirely tethered to the lab benchtop equipment by a wire tether. The blocking issue is not the power source as solutions have already been presented in this work, but it is predominantly in the power electronics efficient enough to drive the large actuators needed with given power for the autonomous payload and and independent wing control.

I begin by considering several fundamental equations below which characterize the design choices for switching mode energy conversion using inductors.

In order to select coupled inductors for the conversion of low voltage power sources onboard, to the high voltages needed for piezo actuator, it is important to choose a proper winding turns ratio to transform power on one side to power with a different ratio of through and across variables on the other. Since this linear transformation takes place in the magnetic domain and we wish to minimize switching losses, the core permeability and air-gap of the power magnetic material should ensure large magnetic energy capacity so as to minimize required numbers of conversions - i.e. it should have a large possible current saturation before saturation, but this is a tradeoff for weight. The design of the power magnetic material is the focus of a great many electrical engineering and material science works, but for the purposes of this work is has been sufficient to select for the weight of the core, design the turns ratio, ensure that custom windings are tightly grouped and adequately insulated from adjacent windings, and to use power magnetic cores from trusted vendors such as Coilcraft.

To choose turns ratio for the switch tapped inductor boost converter heavily utilized in this work. I anticipate operation at up to 500 mW source requirement with $I_{avg} \approx 100$ mA. A voltage drop is expected across equivalent impedance typically 25 Ω of an onboard 5V nominal voltage source. Therefore, $V_{in} \approx 5.0 - 25\frac{500 \text{ mW}}{5 \text{ V}}$ is expected during operation. Targeting, e.g. 300 V as the maximum output voltage, Eq. 4.1 is solved for $N \geq 7.7$ for a desirable turns ratio for the coupled inductors. When voltage droop under load is less severe, care must be taken to reduce converter switch period so as to avoid saturating the transformer core.

$$V_{tap} = \frac{V_{in} + V_{HV}}{N + 1} \tag{4.1}$$

Where V_{tap} is the voltage at the switch tap in the tapped-inductor boost converter, here confirmed in simulation to be ≈ 35 V during operation which is safely below the 100 V repetitive voltage rating of the GaNFET boost converter switch used primarily here. Difficult to model parasitics capacitance in the coupled inductors can easily increase the voltage stress on the switch as seen during simulation and oscilloscope measurement, therefore a higher turns ratio of 8:1 – 10:1 is is typically chosen in this work, with 10:1 being a commercially available coupled inductors from Coilcraft. For custom wound coils, Coilcraft 30xx and LPR4012xx cores have been chosen for size, energy storage, wide flanges to limit leakage flux for custom wild-wound windings, and easily removed magnetic shielding to reduce weight.

To select a core size, inductor saturation current is obtained either by direct measurement of inductor current over time in a step voltage input, or referring to the saturation current I_{sat} rating if such is available. Then, using Eq. 4.2, plugging in the saturation current and the inductance, an upper limit can be obtained for how much energy the inductor can store during one conversion.

$$E = \frac{1}{2} I^2 L \tag{4.2}$$

Since inductor current I_L can be assumed to be a simple function only of the inductance L neglecting small winding resistance, the current in the inductor during conversion ramps up linearly with the conduction time, according to Eq. 4.3.

$$I_L = \frac{VT_{on}}{L} \tag{4.3}$$

This value for $I_L(V, T_{ON}, L$ is used to check that the microcontroller is able to complete a switching cycle given its hardware timer speed (for PFM or PWM), *before* the inductor current exceeds the value found in Eq. 4.2. Likewise, to find the allowable on-time for the converter's switch, T_{ON}. Eq. 4.5 is solved for T_{ON} rounded down to the nearest integer number of digital clock ticks, in order to ensure that the switch does not conduct for too long, else the inductor offers less impedance to the changing current and approaches a short circuit without storing any energy. Equivalently, Eq. 4.3 can be solved for T_{on} for a given saturation current of an inductor L, setting $I_L = I_{sat}$, T_{on} is given in Eq. 4.4.

This value is used in simulations for setting the pulse duration just below the saturation limit for the inductor, so as to minimize switching losses in conversions of DCDC converters.

$$T_{on} = \frac{LI_{sat}}{V} \tag{4.4}$$

$$E_L = \frac{1}{2} \left(\frac{VT_{on}}{L} \right)^2 L \tag{4.5}$$

Similarly, the energy stored in a capacitor is given by Eq. 4.6 and is used extensively below to analyze the PEU and system, make reasoned design choices, and determine one way to obtain approximate (allowing for model errors) switch timing for PEU control signals.

$$E_C = \frac{1}{2} CV^2 \tag{4.6}$$

4.1.1 Loss mechanisms

Ideal energy conversion would pass all of the energy stored in the conversion inductor to the output, changing only the ratio of across variable V to through variable I. In stark contrast, a linear regulator such as the linear half bridge driver used in [38], [30], [29], does not preserve the total energy and power from one side to the other. Instead, an output voltage less than the input voltage (e.g. bias rail being regulated to the actuator's voltage at some lower potential) is obtained by allowing current to conduct through a linear dissipative element causing a voltage drop from the higher voltage to the lower voltage. In the linear half-bridge driver, in addition to the series elements deliberately utilized to smooth the output, the transistors and series impedances such as the dissipative trace resistances along the way are what create this voltage drop when the switch conducts to the capacitor. Since the capacitor's impedance $Z_C = \frac{1}{Cs}$ is frequency dependent, the near instantaneous closing of the switch would result in current which is limited in reality the by the dissipative elements and (to a less important degree the other reactive elements) in series. Thus the transistors of the linear half-bridge driver behave as do the resistors in a potentiometer which is swept dynamically so as to create a sinusoidal signal in the manner of the FIR wing driver - the output voltage regulation is linear and the potentiometer is significantly loaded by the actuator system.

The dissipative element can only waste power, not temporarily store it and return it back or to somewhere else, as do the switching converters. That said, a linear regulation is capable of smoothly continuous output with less harmonic distortion (and risk of damaging piezo actuators) due to its ability to continuously modulate current, where a switching converter using inductive elements as here takes relatively large amounts of energy from the actuator with every conversion, resulting in 'jagged' output especially at higher voltage and limited to a minimum discrete pulse periods of a digital output. This is a tradeoff for its improved efficiency.

However switching mode converters suffer realistic losses too, which are considered briefly here in order for analysis and design, but better understood referring to [50]. Loss mechanisms of particular interest here include conduction losses I^2R of traces and winding wires, conductive losses of transistors due to high R_{DSON} or linear operation while turning on/off, diode drops across BJTs and diodes which are conducting current, and switching losses such as charging and discharging of capacitances at diodes or transistor terminals. Magnetic losses such as hysteresis and leakage inductance are out of scope for this work but avoided by practices laid out in references such as [50, 23]. Sec.4.2 discusses these loss mechanisms in more detail when analyzing PEU performance.

4.2 Existing technologies, results on the Robofly, and motivation for new PEU design

Although many topologies for DCDC conversion exist, few are realizable on an FIR within SWaP budget. The state of art power electronics topologies used in literature for FIRs generally include variations on the tapped-inductor boost converter + half bridge driver with and without energy recovery, and the bidirectional flyback converter. These topologies have proven inadequate to drive larger actuators and carry additional payload of onboard PEU and power source in this work. These topologies are shown in Figs. 4.5 and 4.6, and will be discussed further below. In summary, despite demonstrations of ‘flights’ having only one set of wings driven independently, it is inadvisable to implement these topologies for larger FIRs by simply adding additional driver channels for independent wing flapping. This is to avoid the run-away problem of requiring additional heavy components which in turn require disproportionately more input power due to their fundamental inefficiency, which likewise require a heavier power source, needing again larger payload and actuators, etc.

Boost converter plus linear drive stage

In case of the two-stage tapped inductor plus linear halfbridge driver shown in Fig. 4.1, a relatively efficient tapped-inductor boost converter supplies a high voltage DC bias rail V_B with boost converter efficiency $\eta \approx 40 - 60\%$ over expected bias voltages $V_b = 200 - 240$ V. However in this topology, the bias rail has to in turn supply each wing driver channels with unipolar sinusoids $20 - 240$ $V_{pp} \approx 1 - 2$ mA depending on the actuator capacitance for typical wing loading during hovering flight as observed in my work [29]. Additional driver outputs (e.g. for independently controlled wings) can be driven by the addition

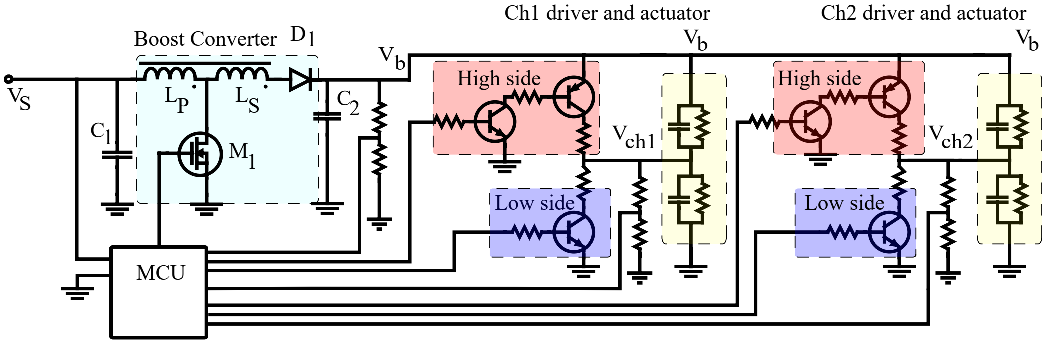

Figure 4.1: Standard dual-stage boost converter plus driver topology extensively utilized to-date in FIRs and in previous work [30], [29].

of a mere 3 transistors each at a weight cost of ≈ 20 mg per channel. However, when flapping frequency is increased from 160 Hz for the 2-winged Robofly using piezo bimorph actuators with PZT layers having principle capacitance 9 nF shown in Fig. 2.4, additional power is required. Considering the net efficiency of this two-stage conversion process of bst converter + driver(s) being $\approx 15 - -30\%$ (where efficiency decreases with higher voltage and current expected for flight with payload), the PEU would require instantaneous power input exceeding 500 mW, which is not within SSWaP budget in this consideration.

The heart of the boost conversion stage is the tapped inductor. The tapped inductor is member of the class of coupled inductors. The tapped inductor in this usage consists of two windings L_p, L_s which are coupled together by the same magnetic core and electrically connected at a single node with respect to the winding direction indicated by the schematic dot symbol at the inductors.

I'll note here that the tapped inductor here such as drawn in Fig. 4.1 is identical in structure and schematic symbol to an auto-transformer, but the switching mode conversion is an entirely different operating mechanism than the operation of an electrical transformer. The converter first builds energy in the electromagnetic core according to Eq. 4.2, then discharges that energy into the secondary, whereas a transformer electromagnetically couples the two windings together, inducing time-varying current in the secondary simultaneously with the time-varying current in the primary - there is no build up of energy, a collapse of magnetic field, and then a subsequent delivery of energy to the other side as in the converter. Thus, despite having similar symbols and structure, the operating mechanism is entirely different and the core material used in the converter is designed for this temporary energy storage in the magnetic 'air-gap' of the material. These 'air-gaps' are generalized across-variable elements in the magnetic domain, like magnetic capacitors.

These windings are formed with different wire guage for purposes of resistance and geometry for inductance. Generally, a single winding can be broken out into a multitude of windings at any 'tap' point along the length of conductor, hence the name of 'tapped inductor'. In the case of the PEU's in this work, this tap point was physically broken out at a solder pad on the pcb. If the windings are electrically isolated from each other then the coupled inductors can be used in a flyback converter or otherwise, and have a schematic symbol similar to a typical two-winding transformer. There are several varieties of tapped inductor DCDC converters depending on how the tap point is utilized: rail-tapped, diode-tapped, and switch-tapped are common configurations. In this case, we utilize a switch at the tap point and operate it as a boost converter. The tapped inductor boost converter in this usage benefits from the high turns ratio to help with the large voltage conversion range, and similar to a boost converter the output is directly in series with the low voltage source V_s. For the tapped inductor boost convert, at the moment of the magnetic field collapsing to deliver the pulse of energy to the output, the magneto-motive-force (MMF) acts on the combined inductance of the primary *and* the secondary winding. Thus the tapped inductor boost converter is an advantageous combination of the classic boost converter and flyback converter. A flyback converter (discussed below) can likewise have a high turns ratio but does not benefit from the series boost of V_S through the L_p winding.

While the boost conversion stage can be efficient in best realistic FIR PEU conditions at nearly 70% as reported in [44] and [38], the linear half bridge driver is fundamentally inefficient with capacitive loads for two primary reasons: 1) The linear driver regulates the output voltage by virtue of linear voltage drop across dissipative elements, much like a simple resistive voltage divider. This is very inefficient when the output is at low voltage compared to the high voltage source, and even worse when the output demands more current. 2) When output voltage is decreasing as during negative $\frac{dVout}{dt}$, the lowside switch of the linear driver simply dumps charge to ground, dissipating power as heat in order to regulate the output voltage low. Generally, all reactive power to the load is wholly lost in the linear driver. Therefore the larger the capacitive load of a piezoelectric actuator, the worse the power efficiency becomes. Fig. 2.3 shows actual measurements taken during controlled flight, from which the power requirements for a linear driver can be inferred multiple ways but perhaps simplest by just dividing the measured apparent power average by the efficiency of the boost converter + linear driver topology demonstrated in practice, $\eta \approx 20\%$.

The instantaneous power requirement throughout the flapping cycle varies substantially, especially under the different waveforms required for the roll, pitch, and yawing flight control as I show below and was also found in [44]. Instantaneous power required for the PEU is generally maximum during rapidly changing voltage, as expected due to the characteristic equation of the capacitor, $I_C = C{}^{dV}\!/_{dt}$, and the increased peak voltages required during wing stroke for flight control commands in roll, pitch, and yaw.

Boost converter plus half bridge driver with energy recovery

The addition of an energy recovery inductor and two diodes per signal, L_H and D_L, D_H respectively in Fig. 4.2, regains some lost efficiency. However, this requires faster driver switching elements than the version without energy recovery, especially difficult in the high-side switch M_{HS} which is typically P-ch FET or PNP BJT. Both of these are slower than their low-side counterparts, thus resulting in large switching losses and limited efficiency gains. Nonetheless, this topology was implemented with discrete components in [37] as well as in monolithic forms in [44]. While the drive stage showed significant expected improvements over a linear drive stage such as Fig. 4.1 and the monolithic implementation efforts also devised additional efficiency improvements such as ‘charge sharing’ between channels [43], the total efficiency of the topology was not improved sufficiently to ‘pay’ for the weight of the extra inductors. The impressive work of the monolithic implementations of this two-stage PEU with energy recovery was not utilized in subsequent FIR demonstrations by the same research groups. This is because the two-stage conversion has decreased overall efficiency due to series stacked of conversions - so improvements to a single stage on their

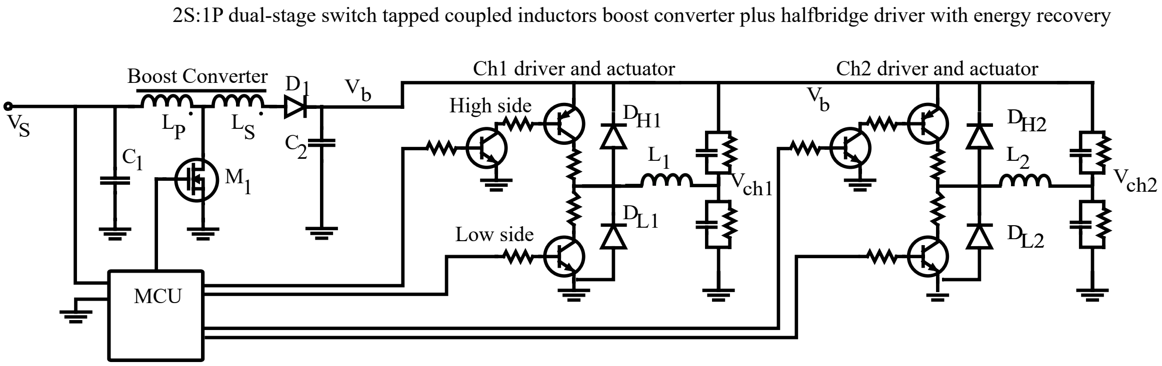

Figure 4.2: The two-stage conversion PEU with two additional heavy inductors L_1, L_2 and diodeds $D_{L1}, D_{L2}, D_{H1}, D_{H2}$ which enable 'energy recovery' as proposed in [37, 38]. These additional components convert the highside and lowside drivers into buck and boost converters respectively. In this representation, the highside and lowside switches are left as BJTs just for agreement with Fig. 4.1, but may in practice be implemented by FETs and series resistances minimized for better performance of the buck and boost operations.

own are doomed to fundamentally limited benefits to overall efficiency, and input power requirements are *still* excessive for autonomous FIRs.

The basic operation of this PEU is quite similar to the PEU of Fig. 4.1, in that the highside and lowside switches of the half-bridge transistor stages are used alternately to charge and discharge the attached load through the added inductors. The energy recovery inductors simply change the driver into a buck-boost driver, similar to class D-amplifiers. Specifically, during discharge when the lowside switch conducts, energy in the magnetic field of the recovery inductor rapidly builds until the lowside transistor is turned off, at which time a large back EMF is induced against the $\frac{dI_L}{dt}$, forward biasing the diodes and returning some of the energy taken from the actuator to the bias rail. This conversion is not efficient at large conversion ratios when the signal is low. Conversely, to charge the actuator, the highside complementary pair of transistors as drawn (or more typically P-ch FET) is turned on, allowing current from the bias rail to conduct through L_H to charge the actuator. When M_H is turned off, the inductor L_H opposes the change in the current, causing the diode D_L to be forward biased and conduct current to the actuator via the inductor in order to balance the stored inductor energy minus losses. In a sense, this is a simple buck conversion in which current from the high energy bias rail is mixed with low-energy ground current to produce current at the voltage of the signal layer - akin to a mixing valve in the fluid domain which mixes high temperature water with colder water to obtain an intermediate across variable without the thermal energy loss that would be incurred by instead (inefficiently) dissipating heat in order to obtain the same lower temperature. Thus, the driver functions alternately as a boost converter when discharging the signal layer up to the bias rail, and as a buck converter when charging the actuator.

A tricky aspect of this design is the driving of the highside switch for the important buck conversion task when charging the actuator. The complementary pair of BJTs depicted in Fig. 4.2 features a PNP which is generally too slow to perform this role efficiently avoiding conduction and switching losses. Meanwhile, P-ch MOSFETs are also relatively slow, with limited commercially available options, are difficult to bias while keeping the gate voltage within a maximum allowance to avoid damage, and have similarly large losses in this role.

[4] presents a highly integrated converter suitable for generating a high-voltage bias V_B featuring a capacitive element in the power conversion similar to a SEPIC converter, which attained similar efficiency to the tapped inductor boost converter documented in [38]. However DCDC converters utilizing capacitive energy conversion are difficult to implement while avoiding the weight of suitable discrete capacitors in addition to inductors. Moreover, the conversion stage is only half the story since a downstream drive stage powered by the conversion stage still results in series stacks of efficiency losses.

The common theme in two-stage converter + driver topologies discussed in greater detail in ch. 4 is that nominal efficiency improvements gained in the driver stage of the boost converter + driver topology just don't pay for their weight, since the upper limit of total efficiency is still capped by the series stack of losses. This is the result of how total power requirements relate to end-to-end efficiency of a series stack of consecutive conversion stages, as expressed in eqs.4.7, 4.8.

$$\eta_{net} = \eta_{s1}\eta_{s2}...\eta_{sn-1}\eta_{sn} \tag{4.7}$$

$$P_{in} = \frac{P_{out}}{\eta_{net}} + P_{aux} \tag{4.8}$$

where η_{net}, $\eta_1 - \eta_n$ are the net PEU efficiency and the efficiency of the consecutive stages of the individual converter subsystems in a series stack, P_{out} is the real and reactive power to the actuator/wing system, and P_{in} is the total input power required for the PEU supplied by the low voltage onboard power source. It is important to note that power efficiency for FIRs reported by [37, 39, 38] is defined as $\eta = {}^{P_{re}}/_{P_{tot}}$, and this work uses this definition for consistency. However, this is a poor metric since the measurable real power includes dielectric losses R_0. These dielectric losses are difficult to measure but have relationships found in [38], and they can be larger than the real power actually going to aerodynamic lift and drag which are represented by lumped dissipative element R_{m1} of the parallel resonance mode in the Butterworth-Van Dyke model illustrated in Fig. 2.2.

The voltage and current measurements of the flight experiment in Fig. 2.3 were used to compute the power consumption of a linear driver which does not recover the reverse power flow, which was measured over integer number of flapping cycles during a period of near neutral hovering, as actually slightly more than 300 mW in that experiment, confirming estimations that average power requirements were below budgeted 500 mW possible for onboard power sources explored in chapter 3. However the instantaneous power requirements peak higher than this alotment and more so with increasing reactive power requirements. The weight of all onboard systems requires an increase in actuator size and flapping frequency, increaasing reactive power requirements. The 9 nF actuators“triple-wides” utilized at the time of this writing and used in [11, 9, 12] require substantially more power depending greatly on the power factor of the loaded actuator system.

An intermediate increase in actuator size, to bimorph actuators whose PZT layers were measured to be 5.4 nF each, had experiments conducted using custom designed highside current measurement so that the power consumption could be measured both to the boost converter and then to the linear halfbridge wing driver. This data is supplementary to the publication in [29] and is shown as on the experimental datacollection software in Fig. 4.3. This experimental observations show that the boost converter plus linear driver is at the extreme range of power capacity, and struggles to drive the peak instantaneous power required for even double-wide actuators.

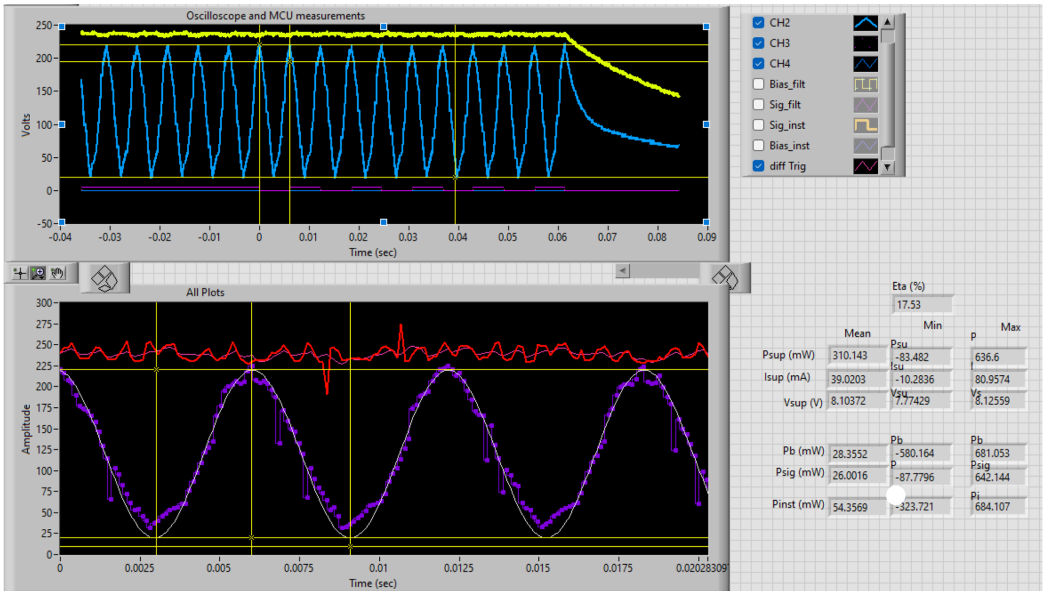

☑ CH2
☑ CH3
☑ CH4
☐ Bias_filt
☐ Sig_filt
☐ Sig_inst
☐ Bias_inst
☑ diff Trig

Eta (%) 17.53

	Mean	Min	Max
Psup (mW)	310.143	Psu -83.482	p 636.6
Isup (mA)	39.0203	Isu -10.2836	I 80.9574
Vsup (V)	8.10372	Vsu 7.77429	Vs 8.12559
Pb (mW)	28.3552	Pb -580.164	Pb 681.053
Psig (mW)	26.0016	P -87.7796	Psig 642.144
Pinst (mW)	54.3569	-323.721	Pi 684.107

Figure 4.3: Experimentally measured power flow during operationas seen directly on the experimental data collection apparatus. Waveforms displayed are voltage and power measurements to a full Robofly with 5.4nF bimorph actuators and wings flapping at 160 Hz with 200 V amplitude and 240 V bias. The average power required from the 8 V supply was 310 mW, with 40 mA of current, and the instantaneous power ranged throughout the stroke ranged from 0–680 mW, agreeing exactly exactly with circuit simulations. Each of the two wings consumed 26 mW real power.

4.2.1 Bidirectional Flyback converter PEU topology options

A flyback converter with high turns ratio is suitable to generate the sinusoidal signals needed for piezo actuated FIRs of this work, $V_{sig} = V_S - 250$ V $\omega_{flap} = 120 - 200$ Hz [37, 38]. A birdirectional flyback converter can perform a single-stage conversion from the low voltage source to charge the actuator and then return energy from the actuator back to the source when discharging it. The power flow is controllable in both directions, hence the name. The basic bidirectional flyback converter topology was proposed for FIRs as a single-stage alternative to the two-stage PEU in the same seminal works of [37, 38]. This topology was reported initially as being up to 60% efficient, but was not utilized for FIRs going forward in the same research group due to the requirement for four separate flyback converters in order to drive the 4 layers of piezo for independent wing control of a bimorph FIR. Fig. 4.4 shows the basic topology extended for four independent channels.

In the case of the bidirectional flyback converter, the work done in [27] represents a superlative implementation with a slightly different circuit and scheme as shown in Fig. 5.1b. To better utilize the full range of electric field for the bimorph for substantially increased deflection (an idea presented in [52]), a switch Q_S was used to clamp the lower layer of the actuator to ground. While this improved performance, the PEU topology is still fundamentally the same bidirectional flyback converters. Total total efficiency reported when actually implemented on an FIR in this work is only $\eta = 28\%$ in a 91 mg single-bimorph PEU. This value agrees with simulations performed in this work as in Fig. 4.9. This single-wing PEU would have to approximately double in weight in order to drive two wings independently and may require a different heavier microcontroller to provide the additional IO's needed. Though this configuration is only a single stage conversion and thereby avoids series stacked inefficiency, duplicating the single channel implementation remains a poor utilization of heavy power magnetics resulting in low power density of the PEU. The weight of four inductors required for independent flapping is a non-starter at the time of this writing. Further discussion of the cause of the surprising inefficiency of the bidirectional flyback driver for high voltage piezoelectric actuators is in a new proposed topology presented in sec. 4.3.

Thus, a new power electronics topology is indicated for larger actuators.

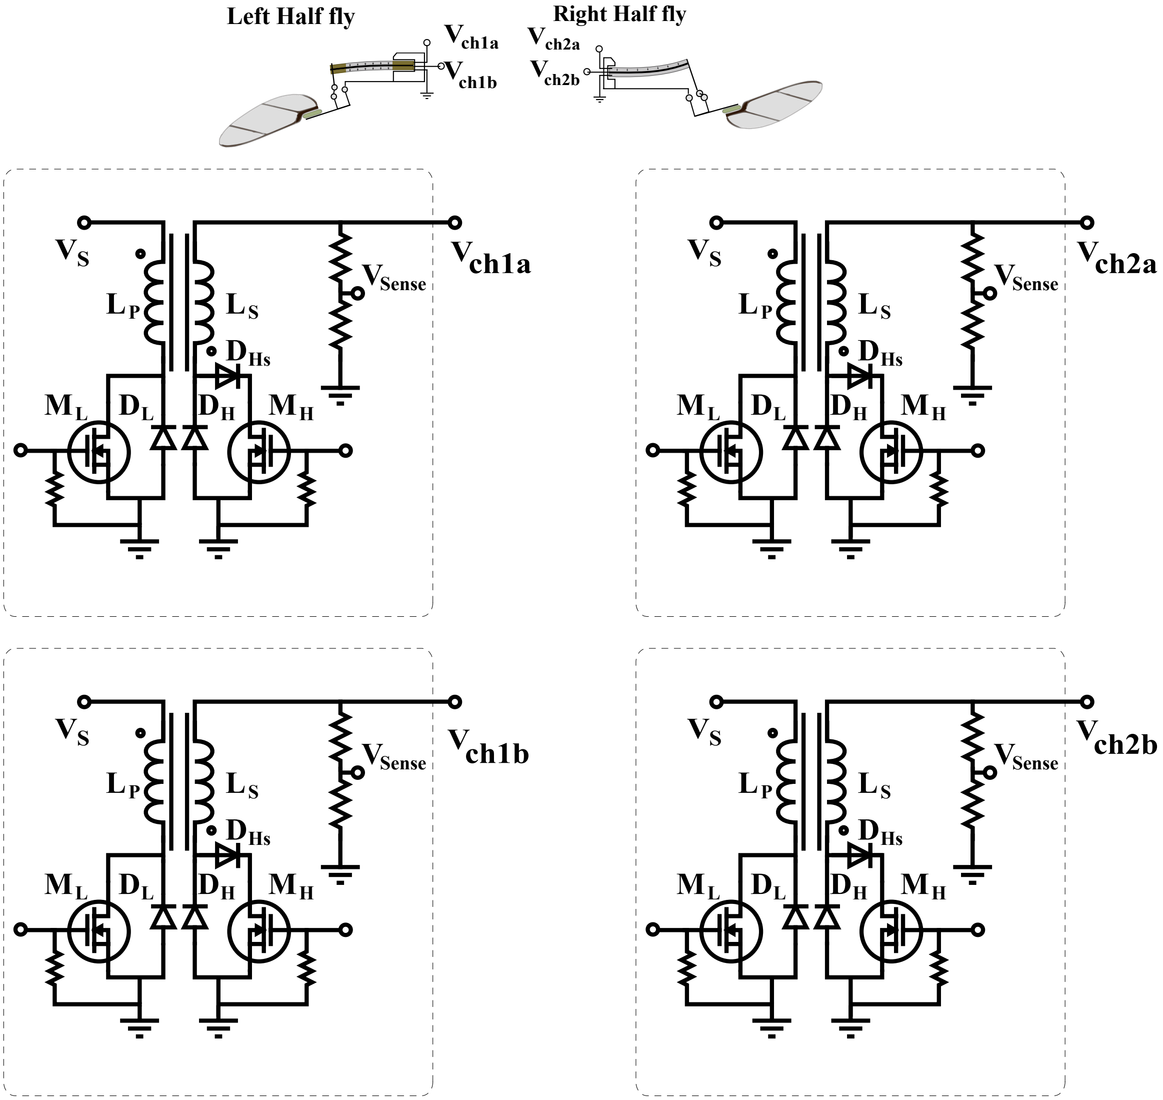

Figure 4.4: The basic bidirectional flyback converter for driving piezo actuator signals. For an FIR using bimorph piezos, 4 independent signals would be needed for the piezo signals indicated in the FIR cartoon at top. A slightly different variation on getting these independent signals was done in [27] and a single-wing version is shown in Fig. 5.1b, but regardless of specifics the fundamental issues remain: the weight of 4 coupled inductors needed for this PEU topology presented in literature, and the inefficiency of the high-ratio bidirectional conversion as discussed further below.

4.3 *New power electronics*

4.3.1 3-inductor boost converter and dual bi-directional flyback drivers

The bidirectional flyback converter topology in [27] and shown in Fig. 5.1b, can only drive a single channel, so a bimorph actuator requiring a unipolar signal across the top PZT layer must be supplied by a bias voltage using another converter channel, or can be driven alternately in one half-cycle and allowed to float during the other half-cycle as done in [27].

The topology shown in Fig. 4.5, is an adaptation combining topology ideas from [38] and [27] to workaround the limitation of the result in [27], which was that the PEU could only drive a single independent bimorph actuator. Here, a switch tapped-inductor boost converter to supply the bias rail necessary for a bimorph actuator and only two bidirectional flyback converters for the unipolar sinusoids V_{sig1}, V_{sig2}.

This topology was simulated in Spice software and instantiated for tests with RC dummy load, and found to have efficiency of 22% when driving 6nF actuators at 200Vpkpk with bias voltage $V_b = 240$ under proportional control, which is less efficiency than that reported when driving piezos with less nominal capacitance in [38, 27]. Studying the circuit in detailed simulation yielded two primary causes of this inefficiency, both originating with the large conversion ratio between V_{sig} and the low voltage V_S. Decreasing efficiency as the conversion ratio increases is particularly bad because the actuators have maximum energy at higher voltage as per Eq. 4.6 as discussed in §4.3.2. In short the flyback converter returning energy from high voltage signal back to low voltage source is efficient when the actuator has little stored energy, but inefficient when the actuator has maximum stored energy - exactly the opposite of what we want.

Observed losses in simulation were 1) Conduction losses of the high-voltage switch are significant when the transistor is turning on and turning off, not only due to the charge and discharge of switch capacitance but mainly due to the transistor operating in its linear region, causing it to develop a voltage drop and associated I^2R losses instead of conducting with minimal resistance - although I is small, the voltage drop during this turn-on and turn-off periods is significant. 2) Due to the large turns ratio of the flyback transformer needed for the large conversion ratio, the current “recovered” back to V_s through the discharging

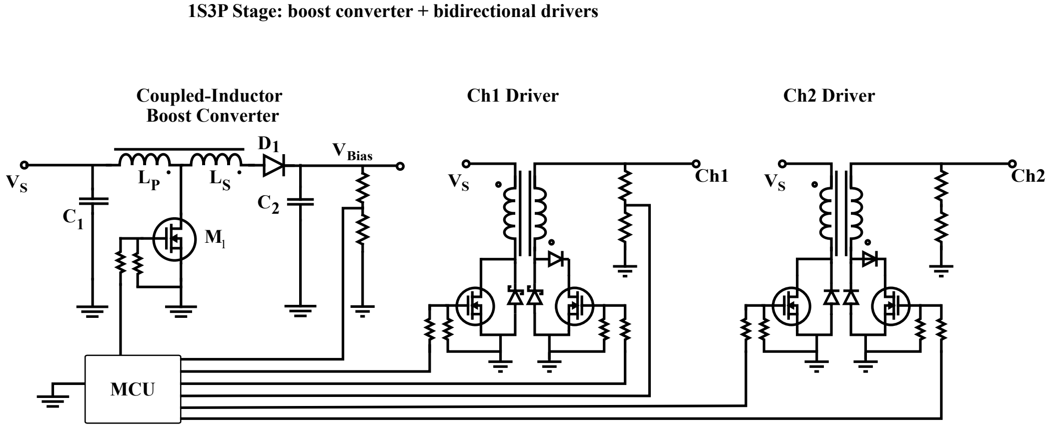

Figure 4.5: An extension of the single-stage, bi-directional flyback converter, originally proposed in [38] and utilized in [27], with the addition of a conventional switch tapped inductor boost converter to supply the bias rail for driving two independent bimorph actuators.

conversion phase, shown at left in Fig. 4.6 is a large value. Generally, to make full use of the heavy power magnetic and minimize the number of switching events, the flyback converter should store near its maximum amount of energy per pulse; for the Coilcraft 3015 and Coilcraft 4012 series cores, using Eq. 4.2 this corresponds to a primary side current of 1.3-1.7 amps (the saturation current rating of the inductor). Given the forward voltage drop $V_d \approx 1.1$V of the diode D_L through which the energy is returned to the low voltage source as indicated in the red arrow in left side of Fig. 4.6, this results in $P_{diode} = IV_d$ of nearly 1000 mW dissipation during a switching event, resulting in huge losses. Unfortunately, the current I_p cannot be arbitrarily reduced due to the minimum pulse duration of the microcontroller (in [27] this was 8 MHz) and generally increasing proportion of switching losses for given transistors operated with shorter pulses.

Thus, with the unsatisfactory efficiency very nearly the same as the $\eta = 28\%$ reported in [27], and *still* requiring 3 heavy inductors, this topology was not selected for development. Neither was the four-flyback converter topology shown in Fig. 4.4.

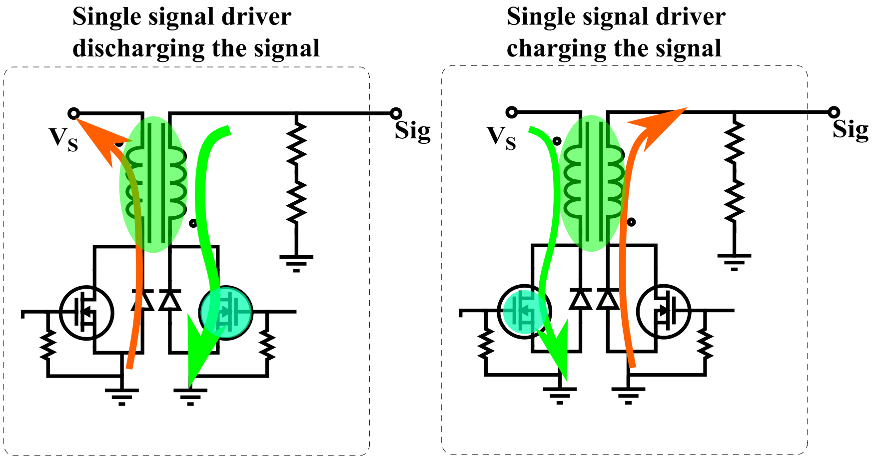
Single signal driver discharging the signal

Single signal driver charging the signal

Figure 4.6: The energy and current flow during charging (left) and discharging (right) modes of operation is shown. In charging mode, the switch M_1 conducts allowing primary current I_p (in green) to ramp up. When M_1 is ‘turned off’ the energy stored collapses and large MMF is induced on both windings according to the winding polarity indicated by the dot, since the electrical circuit on the primary side presents a nearly open circuit due to the diode D_L being reverse biased and M_1 is not conducting, this energy goes through the secondary winding, adding energy to current from the ground through forward biased diode D_H.

4.3.2 Design principles

In examining the deficiencies of existing FIR PEU topologies, the unsatisfactory efficiency and high power requirement of the simple and lightweight boost converter + linear driver stood in stark contrast to the power requirements of bumblebees in forward flight. Work in [13] measured specific power for bumblebees in forward flight at $P_d = 25\ \frac{\mathrm{W}}{\mathrm{kg}}$; a bumblebee comparable to a 200 mg FIR would need substantially less power than the hundreds of mW measured in [30, 29] and found repeatedly in circuit simulations using the load model and high-fidelity PEU simulation schematics with manufacturer provided component models.

A return to the bimorph actuator model as shown in Fig. 4.7 led to several insights into the energy flow during PEU operation to drive piezos. For non-static inputs V_{sig} and static bias rail V_b, a large bias capacitance $C_b >> C_{top}$ effectively anchors the top layer of the bimorph to an unmoving reference, which we can equate to a ground. Electrically this is seen as shown above; C_b is simply in series with C_{top} with respect to V_{sig}. In the mechanical domain analogy, this is similar to placing an infinitely large mass M_{bias} in series with a smaller mass M_{top}. From the perspective of a time-varying (non-static) input through-variable to the V_{sig} node (electrical current or analogously mechanical force), the large bias mass would not 'care' - that is to say it would not exhibit significant change in across variable for the applied through-variable because its generalized impedance (defined as across variable out per through variable in) is $Z_{M_b} = {}^{1}\!/\!(M_b s)$. Similarly, the generalized impedance of the capacitance C_{top} is $Z_{C_{top}} = {}^{1}\!/\!(C_{top} s)$. Therefore the C_{top} in series with it effectively has *an unmoving reference*; it is effectively grounded for the purposes of simplifying the loading of the input signal and when V_b is static and $C_b >> C_{top}$.

This simple but somewhat non-intuitive simplification leads to useful insights: Firstly, the sufficiently large bias capacitance means that from the perspective of the input signal, the PEU design shouldn't fixate on charging the bias rail first; the 'job' of the bias rail is just to provide an unmoving reference for the top layer. When one bimorph charges up, charge Q_{in} with energy E_{in} will be injected into the middle node V_{sig} of the bimorph, pushing charge from $Ctop$ into the bias rail (minus some energy going to real work from that piezo layer). This can be seen in the cartoon of Fig. 4.10. Charging the middle

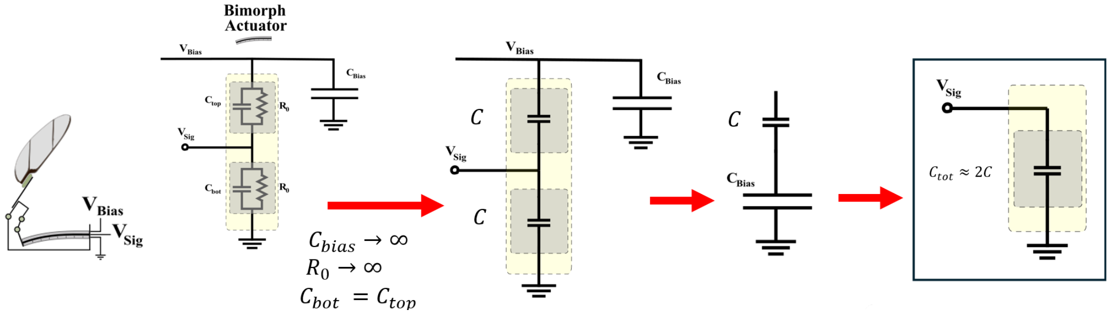

Figure 4.7: A bimorph actuator in simultaneous drive [65], such as used in the FIR 'half-flies' of this work depicted at left can be simplified as shown to a single equivalent capacitance comprised by the two bimorph layers in parallel to ground when the bias rail is static. In this simplification, the bias capacitor and the capacitance C_{top} of the top bimorph layer are in series to ground, with respect to the input signal V_{sig}.

node of a bimorph also charges the bias rail. Likewise when the bimorph discharges, the same charge will be drawn from the bias rail into the C_{top} as shown in Fig. 4.10. This discharge can return a portion of the energy to the bias as through a boost converter. This reality is reflected in the power measurements of Fig. 2.4, where the bias rail power is much lesser in amplitude than the power transferred in the signal wires to V_{sig1}, V_{sig2}. Therefore without directly charging the bias rail, we can still charge it just during charge/discharge of the actuators without risk of depoling (inducing a damaging negative electric field across a piezo layer as by $V_{sig} > V_b$). Again, the 'job' of the bias rail is just to provide an unmoving reference to the top layer. The take-away is that we just have to keep the bias 'topped off' to compensate for the hopefully minimal dielectric losses and capacitive coupling 'draining' the bias due to high-frequency content of the input signals arising from the PEU pulsetrains on the charge/discharge stages. As already noted it is important to keep in mind that we cannot completely neglect the dielectric losses through R_0 in FIR application. Including the R_0 in the bimorph model, makes clear that there is a real power dissipation from bias rail through the two series R_0 of the bimorph as well as the useful power.

Secondly, it is the input signal $V_{sig}(t)$ that really matters most here. It may take only a fixed amount of charge Q to result in a change in voltage δV, but the energy involved is a function of the total charge Q *and* the voltage! This is represented below in Eq. 4.9 Increasing efficiency of the PEU should be primarily concerned with being efficient at the voltages where the most energy is transferred for a given reference signal and actuator. This is not necessarily at the highest voltage present in the signal, e.g. in the sinusoidal waveforms driven in Fig. 2.4, the instantaneous power peaks near the peak of $I \propto C\frac{dV}{dT}$, not at the peak in voltage V_{sig1}, V_{sig2}. Simply put, the PEU should be efficient where we care about it.

$$E(V) = \frac{C_{tot}V^2}{2} \rightarrow \delta E(k) = C_{tot}\left((V_{sig}(k))^2 - (V_{sig}(k-1))^2\right) \tag{4.9}$$

- Most of the energy in the system is at high voltages. Prioritize efficiency there by choosing the charge and discharge converters according to efficieny at the operating point of maximum energy transfer for a given reference signal.

- Avoid returning energy back to the source voltage V_S. It is desired to add energy only as needed to maintain oscillation, keeping the energy in the system instead of returning it to the source in order to discharge the actuator. It makes little sense to inject energy into the actuator at high voltage only to pull it back to the low voltage source at a loss. It is better, if possible, to only feed input energy to maintain the oscillation. Leaving energy in the actuator faces dissipation by dielectric losses but this is better than attempting an inefficient recovery back to the source.

- Avoid series stages of conversion. $\eta_{tot} = \eta_1\eta_2 ... \eta_{n-1}\eta_n$

- Maximize inductor utilization. Deeply discontinuous operated means that there is considerable 'idle' time between conversion pulses such as for the tapped inductor boost converter utilized. At the extreme worst end of power-magnetic utilization, the implementation of [27] had each flyback converter sitting idle during a full half-cycle, resulting in very poor power density of the heavy power magnetics.

In terms of driving the actuator voltage back down during discharge phase, the logical way to keep the high voltage energy in the system and avoid energy losses from inefficient high conversion ratio transformations is to put the charge from the signal layer back up to the top layer of the bimorph or to the other bimorph if appropriate to do so, instead of converting back to the low voltage source V_S. This is the reasoning behind the 'energy recovery' inductor added to the two-stage boost converter + driver w/ energy recovery presented in [38], which allows the low-side of the half-bridge driver to function as a signal-to-bias boost converter.

4.3.3 comparison of conversion topologies

Given a set of principles by which to design a PEU, the next step was to envision possible energy pathways for driving the bimorph actuators. Some possible conversion pathways are shown in Fig. 4.8.

It was desired to plot the conversion efficiency over voltage conversion for different basic forms of DC-DC converter circuits which were possible to implement with flight weight components. DC-DC conversion options were envisioned to perform the possible energy conversion pathways indicated in Fig. 4.8. Because the options for discrete components suitable for the high voltages and extreme light weight are relatively few and component particularities often make or break a design, it was important to use the most accurate simulations possible with exact component models of the components available to guide the design work. Basic prototypes were constructed and tested in order to validate more extensive simulation results.

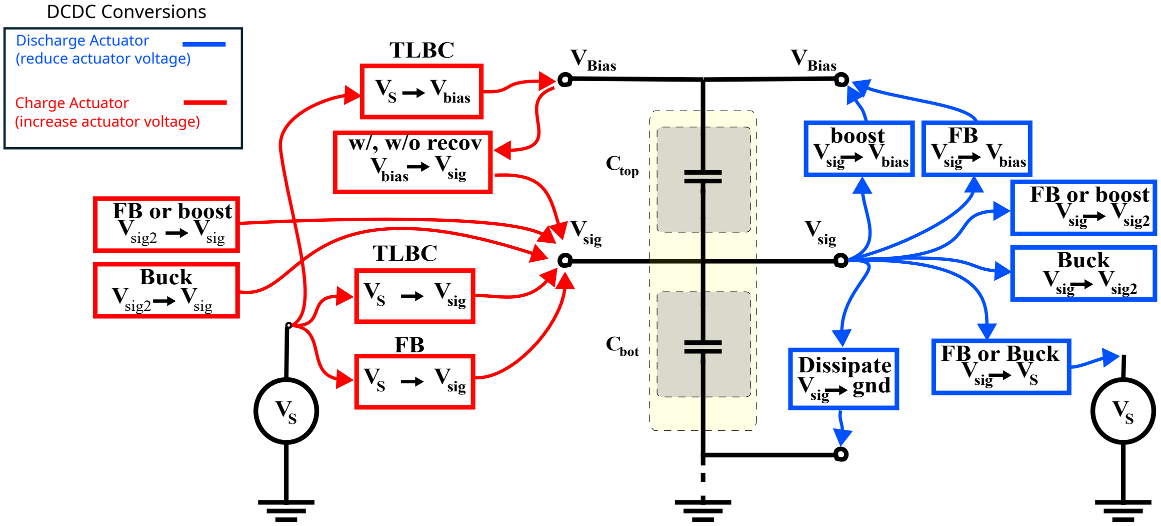

Figure 4.8: This graphic shows some of the possible energy conversion pathways for two independent bimorph piezo actuators (only one shown for simplicity). Charging stages (increasing the actuator voltage) are shown in red, and discharging conversion stages (decreasing actuator voltage) are shown in blue. *FB* denotes a ‘flyback’ converter, and *TLBC* denotes a tapped inductor boost converter. Charging vs discharging does not necessarily correspond to whether the output of the conversion stage is greater or less than the input voltage; for example, a boost converter can discharge the actuator by boosting up to the bias voltage V_b. Some conversion pathways in this graphic are documented and used in literature, while others are novel. Combinations of these conversion pathways were drawn in pen-and-paper schematics for initial design and estimation of complexity and needed component weight, and desirable options were simulated in Spice circuit simulations. Simulations are available in the author’s github.

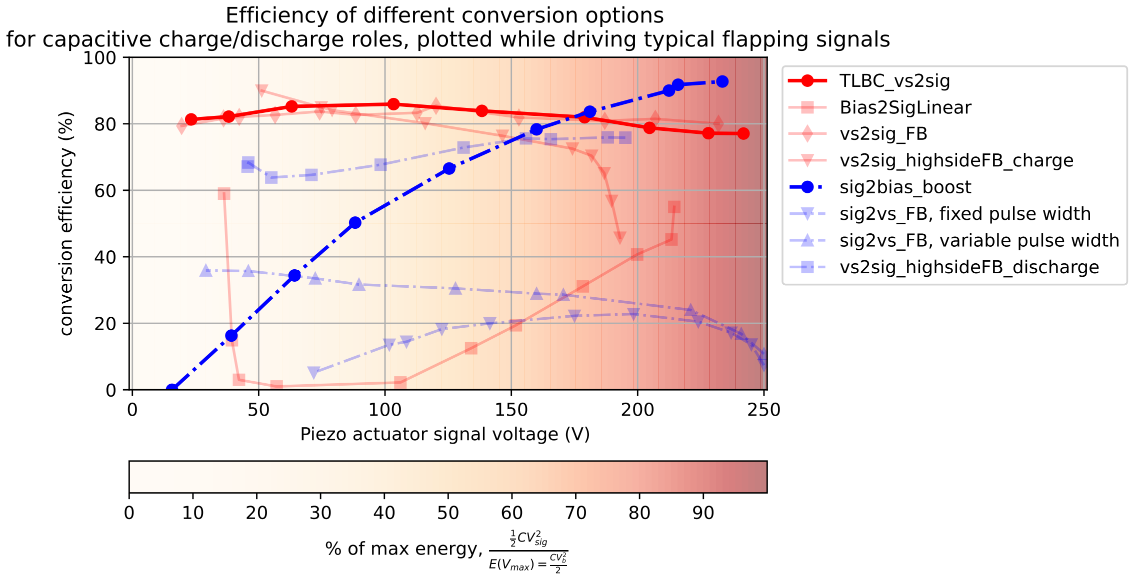

Figure 4.9: This plot aids selection of PEU DC-DC converter stages, by indicating best performing instantaneous conversion efficiency over voltage for the typical sinusoidal flapping signals. The color map indicates the energy concentration in the capacitor over the flapping signal, but in general the color map should be scaled to the power density of an arbitrary reference signal. Some conversion pathways in this graphic are documented and used in literature, while others are novel.

Instantaneous efficiency is considered here instead of conventional static efficiency metrics used in many FIR PEUs in literature, because the flapping signals have a wide voltage range and it is important to pick an operating point for the design given where along the voltage axis is the most power transfer happening for a given flapping signal. For the plot of Fig. 4.9, efficiency of these conversion options was obtained from simulations using high-fidelity models of the actual PEU components, examining the efficiency of just the energy conversion from input to output during individual pulses, not including other losses in a fully realized circuit; thus efficiencies appear higher than in practice such as for TLBC without accounting for blocking diode reverse current, etc. There is some noise due to some simulation parameters such as voltage source and parasistic modeling being slightly different in some simulations. Nonetheless this is informative about which design options perform best at the most power dense parts of the operating flapping waveforms.

In general, the distribution of a reference voltage signal V_{ref} can be weighted according to the instantaneous power requirements and projected onto the colormap in Fig. 4.9, and the density of this projection indicates where the PEU design efficiency should be concentrated. Equivalently, the power flow of a given reference signal for a given actuator can be factored with the instantaneous efficiency to yield a metric relating to the total PEU input power requirements, providing the same information we need to make this design choice. For the capacitive load $I \propto C\frac{dV_{ref}}{dt}$ informs this. Graphically, it is easy to see that most of the energy in the capacitive actuator over the expected range of voltages is simply at high voltages, making our choice clear from this standpoint. Therefore, since the piezo actuated FIR benefits from having a high possible voltage for both payload and control authority, we wish to choose charge and discharge stages which are maximally efficient at large voltages and still perform moderately well at intermediate voltages.

Of note, the bidirectional flyback converter topology when used in discharging role is plotted for two modes of operation, one is for a fixed pulse width and another is when the pulse width is varied according to Eq. 4.4 in order to minimize the number of pulses and reduce the proportion of switching losses to energy transferred. Regardless, it is clear that this topology which converts from $V_{sig} \rightarrow V_S$ is not a good choice due to the high conversion ratio and for efficiency at the higher voltages where we want it.

The plot of the selected boost converter was not generated using variable pulse width using Eq. 4.4, which improves the efficiency at the less important lower voltages when done in practice.

Referring to Fig. 4.9, The conversion options chosen are the switch tapped inductor boost converter for the charge stage, and for discharge stage the Sig-to-Bias boost converter topology + a hybrid option for less frequent usage to flyback to V_s when the conversion is more efficient than boosting to bias or to avoid charging the bias if desired. In some cases considered here, multiple converter types can utilize the same inductor, and these topologies are in development at the time of writing. The muxed peu comprised by the TLBC and boost converter hybrid will be considered below.

Major differences of the two-stage boost converter + driver w/ energy recovery compared to the topology proposed below include 1) avoids series stages 2) use of time-multiplexing and hybrid usage of the converter inductors to reduce the total inductor count significantly 3) more efficient operation during charging comparing to the two-stage conversion by using the TLBC directly instead of series boost conversion from V_S to V_b + buck conversion from V_b to V_{sig}. In other words, the energy recovery addition to the two stage converter in literature was 'on the right track' but was doomed because of the series stack of efficiency.

Further evaluation must consider the component weight and a measure of the complexity of the implementation.

4.3.4 Muxed PEU

This design thrust is one of a novel contribution to implement one of the new combinations of converter topologies found above; the switch tapped inductor boost converter to charge, and the signal-to-bias boost converter to discharge, with hybrid option to flyback to V_S instead of boosting to V_{bias}. The robots are designed so that actuators are out of phase and the wings stroke together, as do a humming bird's wings in flight, reducing vibration of the robot and spreading out the peak power requirements of the actuators, compared to if the actuators were in phase. The muxed PEU takes advantage of this typical actuator phase configuration by performing time multiplexing of the converters, thereby getting 'double duty' out of the heavy power magnetics.

MUX switches connect these two individual charge/discharge converter blocks to the appropriate wing channel to perform this time multiplexing, thus I require only two inductors to drive the independent wings in contrast to the four inductors needed to extend the state of art topology in [27]. The cartoon representation in Figs. 4.10, 4.10, 4.12 demonstrate the major pieces of the topology in normal operations. The full schematic in Fig. 4.13 is used to explain the operation of the hybrid discharge converter.

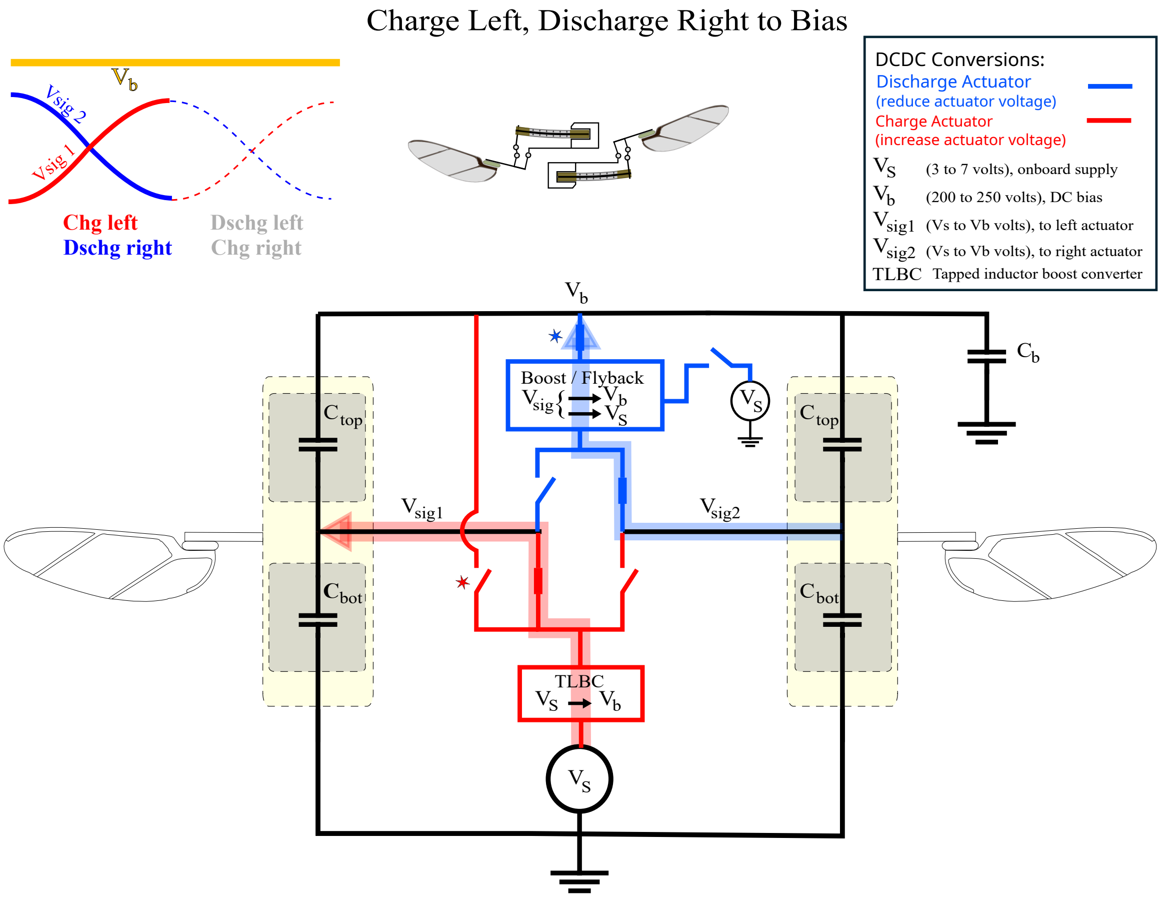

Figure 4.10: Cartoon simplification of the muxed PEU showing operational states for charging the left wing actuator V_{sig1} and discharging the right wing actuator V_{sig2}. The discharge converter operates as a boost converter to the bias rail V_b.

$V_{sig1}, V_{sig2}, V_{bias}$ are the left actuator, right actuator, and DC bias voltage signals, and $V_S \approx 3 - 7$ V is the low voltage onboard supply. $TLBC$ is the (switch) tapped inductor boost converter, chosen as the charge element, and the discharging element of the PEU is a hybrid boost/flyback converter which normally operates to boost V_{sig1}/V_{sig2} to bias rail, but can be converted into a flyback converter to return energy to V_S as shown below.

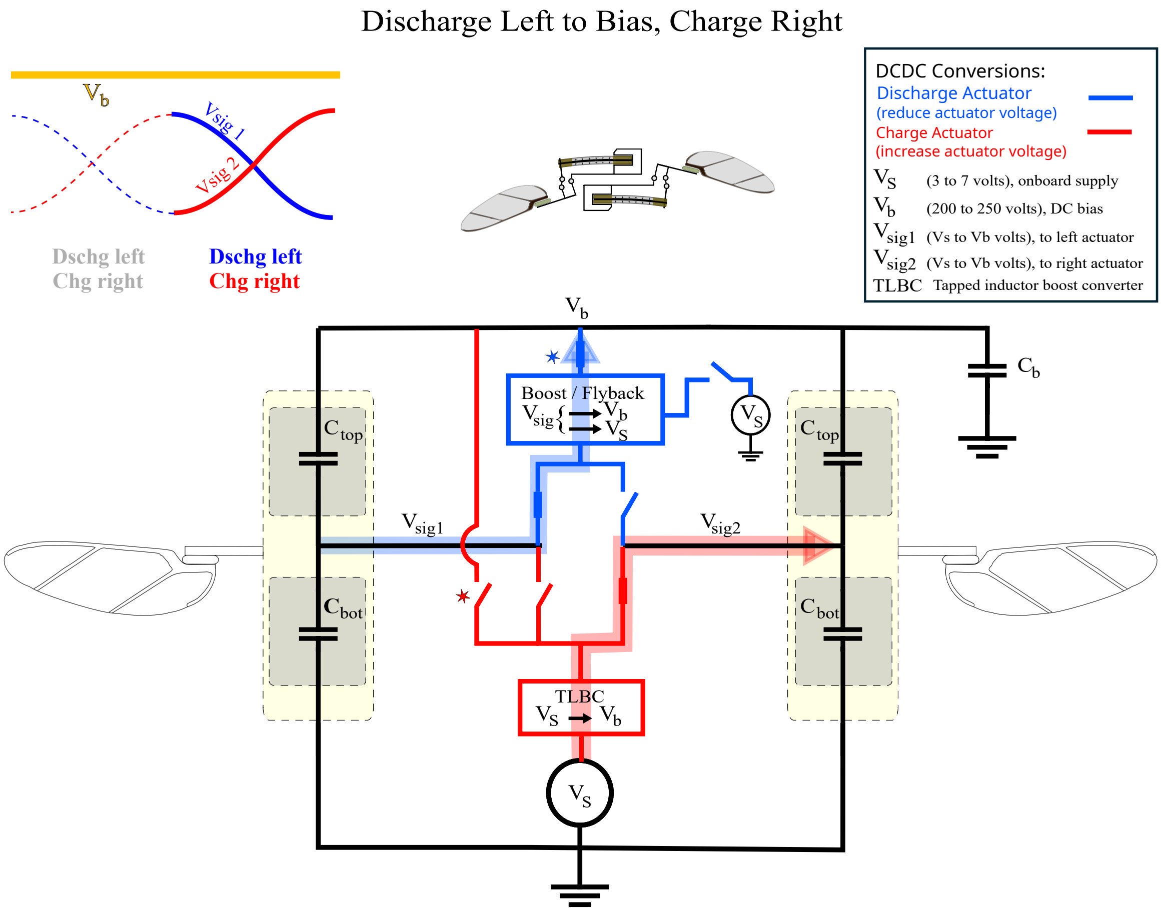

Figure 4.11: Cartoon simplification of the muxed PEU showing operational states for charging the right wing actuator V_{sig2} and discharging the left wing actuator V_{sig1} in normal operation. The discharge converter operates as a boost converter to the bias rail V_{bias}. $V_{sig1}, V_{sig2}, V_{bias}$ are the left actuator, right actuator, and DC bias voltage signals, and $V_S \approx 3 - 7$ V is the low voltage onboard supply. $TLBC$ is the (switch) tapped inductor boost converter, chosen as the charge element, and the discharging element of the PEU is a hybrid boost/flyback converter which normally operates to boost V_{sig1}/V_{sig2} to bias rail, but can be converted into a flyback converter to return energy to V_S as shown below.

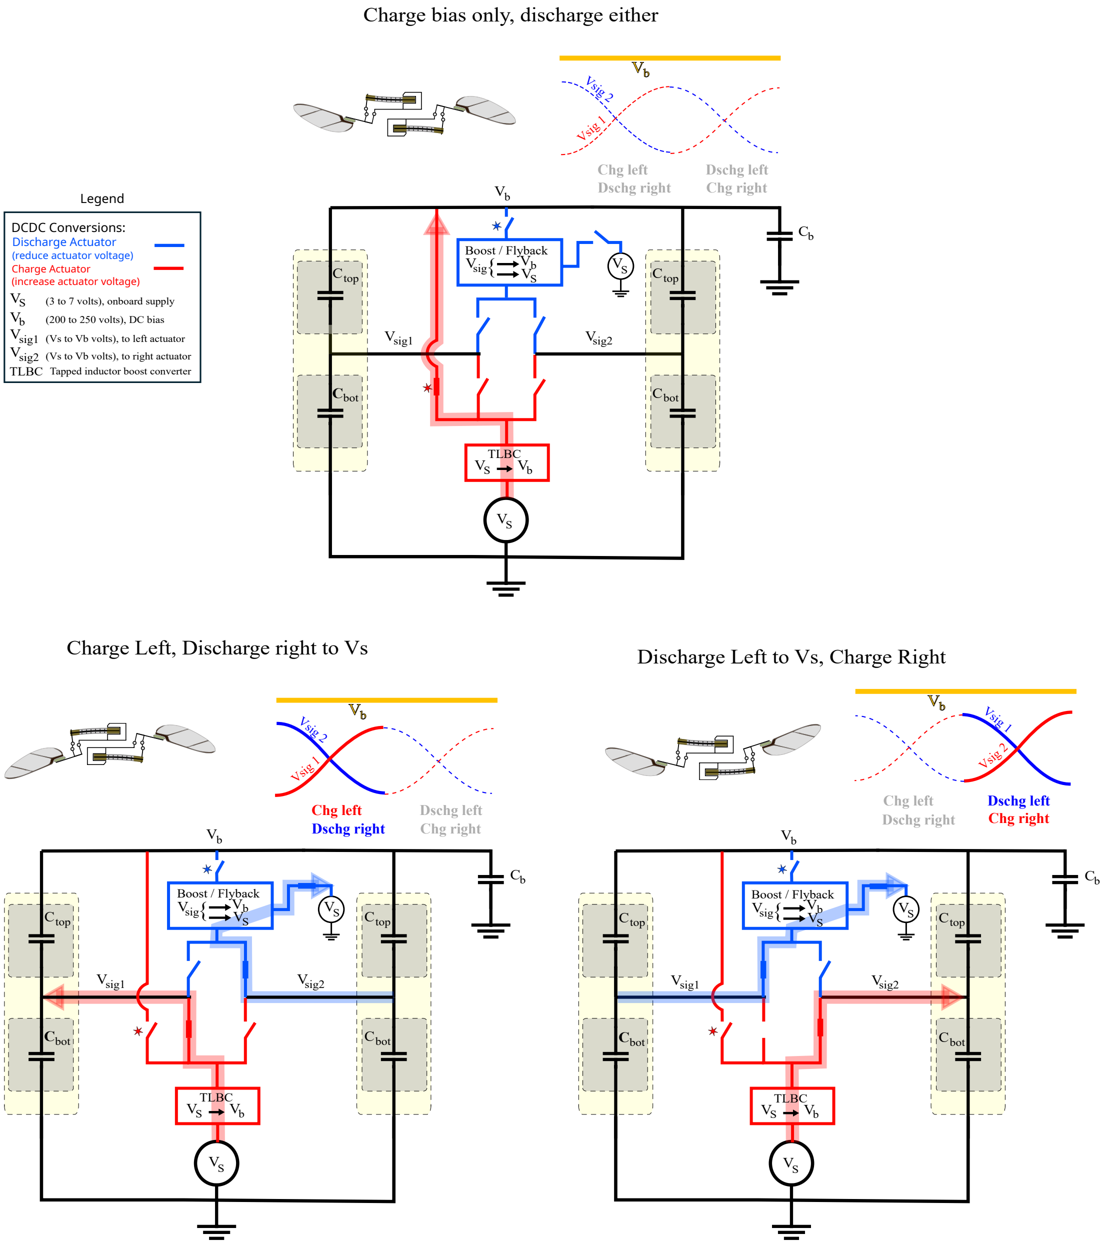

Figure 4.12: Less typical operational modes of the Muxed PEU switches, showing an option to charge the bias only, or discharge either signal to V_s instead of boosting to bias. (Center top) Muxed PEU showing operational states for charging the bias rail V_b only. (Left lower) The operational states for charging the left wing actuator V_{sig1} and discharging the right wing actuator V_{sig2} transferring power back to V_S. (Right lower) Charging the right actuator V_{sig2} and discharging the left actuator V_{sig1}.

To charge bias only, as in the top subfigure of Fig. 4.12, the switch states are achieved by leaving the charge mux switches open, and the diode (the red starred switch) becomes forward biased. The discharge switch states can be in any configuration.

To enable the flyback option to transfer power back to V_S in order to prevent charging bias or to be more efficient than the boost converter when $V_{sig2} << V_bias$, the selector switch (M_r from Fig. 4.13) is closed. This causes the diode indicated by the blue starred switch to become reverse biased as indicated by the blue-starred switch opening, and the discharge converter then operates as a flyback converter returning energy back to the low voltage supply V_S.

Likewise for the left wing (A.K.A. channel 1, or V_{sig1}), the discharge element can be used as a flyback converter back to the low voltage source V_S instead of boosting to bias by operating the same M_r switch, but with V_{sig2} as the selected source of the discharge stage.

Specifically, when M_r is closed, the flyback option preferentially takes the energy from the discharge stage's output because the recovery pathway through shottkey diode D_r, extra winding L_r, and through M_r back to V_S presents an easier path for the pulse of energy to discharge compared to rising higher in voltage all the way up to V_b, V_{bias} in order to forward bias the diode D_{b2}.

Muxed PEU schematic

The schematic diagram of the proposed PEU is shown in Fig. 4.13, and in cartoon representation in Figs. 4.10, 4.10, 4.12. The physical realization of this design is in Fig. 5.21.

The extra winding L_r is coupled to L_{sd} and consists of either a few mg of 51-AWG wire wrapped around an existing Coilcraft inductor (L_{sd} in schematic 4.13) which in prototyping had inductance $L = 220 - 440$ or as an existing 1:10 turns ratio coupled inductor from the LPD3015 or LPR4012 product line. D_r, M_r, Q_r and associated passives in the bootstrap network are sub-mg surface mount components.

The design and implementation of the high-side mux switches is detailed in §4.3.5.

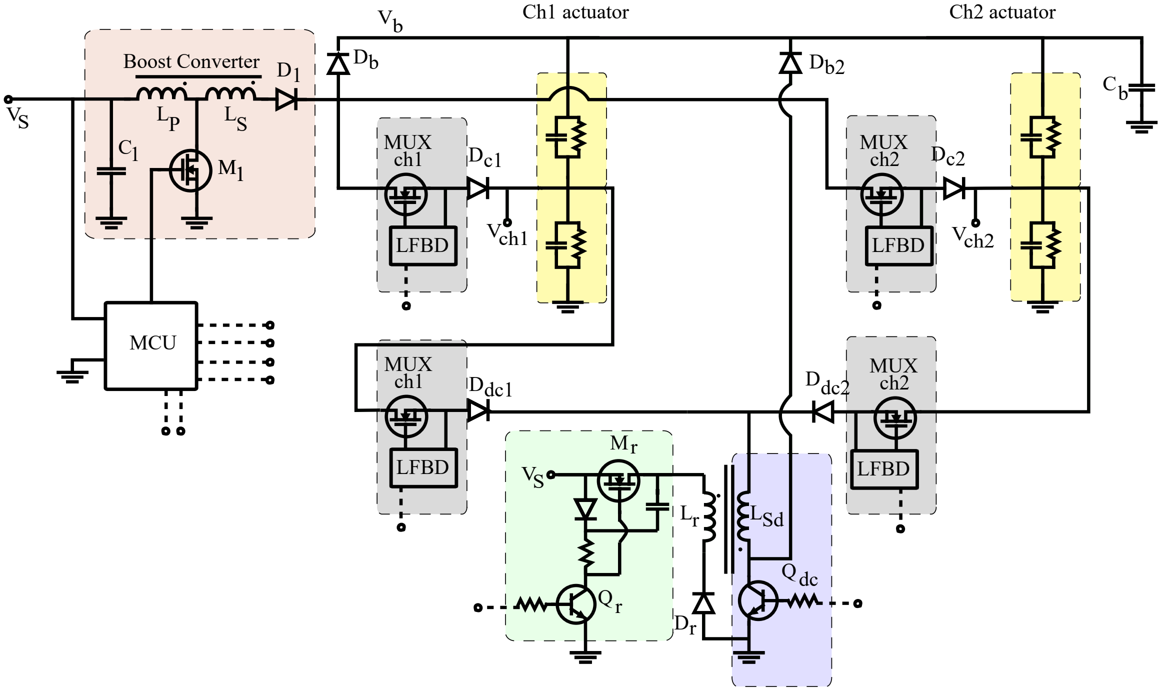

Figure 4.13: Muxed PEU schematic showing the bimorph actuator representation in yellow shaded regions and featuring switch tapped coupled inductor boost converter to charge and sig-to-bias boost converter to discharge with selectable energy recovery shunt M_R for bias down regulation. This PEU topology represents a near doubling of PEU efficiency, measured at 48-56% net efficiency (vs the 28% reported in [27] using benchtop voltage and current measurements and close to the 48-61 % efficiency found in simulation. The green shaded region is a flyback converter to V_S which can be toggled on or off by the bootstrapped high-side switch, in order to shunt excess energy away from the bias rail when desired and to allow for more efficient discharge of the actuator when the input to the discharge stage (ch1 or ch2) is at low voltage. The selectable recovery stage 'recovery shunt' transforms the discharge converter from a $V_{sig} \to V_b$ 'boost' converter into a $V_{sig} \to V_S$ flyback converter.

Pros and cons, and contributions to state-of-art by the muxed PEU

The nearly double efficiency over state-of-art; effective doubling of the power density of the heavy inductors vs the best efficient competitor; weighing ≈ 74% the overall projected weight of the closest state of art [27], with the strong likelihood of further weight reduction over the prototype shown below; and being the first demonstration of independent wing control onboard a flight-weight PEU; and possibly the lightest implementation of a 350 V isolated highside gate driver, are significant milestones on the way to power autonomy and for insect scale robots in general. The possibility of adding additional channels, given faster muxing, is of particular appeal for a multitude of actuators.

However, the bias control authority is limited since charging of the bias rail is only done passively via the PZT top layers while charging, and also when both charging muxes are 'off' at which point the charging converter can be activated to charge bias rail, or when one of the signals becomes greater to the bias voltage, causing the depoling diodes to be forward biased. To compensate for this, a heavy 17 mg capacitor C_{bias} is employed here, approximately 5× the bimorph's capacitance $2C_0$. Additionally, the MUX switches are complicated.

To our knowledge, several primary innovations underpin the worth of this contribution to the field:

- The first demonstration of time multiplexing in an FIR DC-DC converter stage, which *halves* the required weight of the heaviest component on the PEU: the power magnetic. This was enabled by the insight that the boost converter used is operated in deeply discontinuous mode, i.e. we should be making full use of that heavy inductor all the time, instead of *fully deactivating* each inductor half of the flapping cycle as done in [27]. A possible extension to multi-actuator robots is that additional channels can be added with faster switches, without requiring additional inductors up to the limit of full-time utilization of the inductor. In recent literature such as concerning DEA actuators instead of piezoelectric bending beam actuators, as in [42], the need for piezo bending beam actuators to have heavy inductors just to drive them is used to support the decision to abandon PZT actuators. This doubling of the power density of the inductor is particularly salient when considering the recent solar powered liftoff of the Robobee [27], which used a PEU topology that was only able to drive a single wing with the same number of inductors as this muxed PEU used to drive *two* independent wings.

- The first demonstration of hybridization of conversion topologies within SWaP of an FIR, as in the hybrid boost/flyback converter of the muxed PEU discharge stage.

- The insight that end-to-end efficiency is dramatically improved by avoiding series stacks of conversion inefficiency.

- The insight that the bias rail does not need to have an entire conversion stage just to charge it, since the top layer of the actuator is series with the bias capacitor referenced to ground (in simplification) and they they charge it when being charged themselves, and then the discharging phase charges the bias rail again. Bias ripple can merely be offset by increasing bias capacitance at a weight penalty. Additional notes on bias regulation are below, since this is more complicated in this topology than others.

- the development of custom, ultra-light and high-voltage isolated highside N-channel switches.

- Possibly the lightest demonstration of independent wing control on an FIR using an onboard PEU, driving very large capacitive loads and with relatively little power requirements due to the efficiency.

4.3.5 Custom high-side switch design

The new topologies require at least one type of high-side switch, this work will discuss the relatively low frequency selector ‘mux’ switch used to mux the input of the discharging converter and demux the output of the charging converter. This mux switch which should be able to switch from saturated conducting to saturated nonconducting states settling within 200 μs so as to be able to switch the input to the charge and discharge conversion stages from one channel to the other during the flapping cycle, within the time of one waveform generator control period T_{wfg}.

Unfortunately sufficiently high voltage switches which can function with SSWaP budget are not commercially available at the time of this writing. Conventional solutions to driving highside switches include optical gate drivers, discrete bootstrap capacitor circuits, charge-pump gate driver IC's or pulse-transformers for N-channel transistors, or simply accepting the vastly slower and higher resistance of P-channel transistors driven by an additional transistor (as in the complementary pair, or other dual transistor arrangements) of nearly the same voltage rating and size. The challenge of enhancement mode N-channel transistors on the highside is that the gate must be driven to a larger voltage than the source terminal.

P-channel transistors are not desired switches for this application due their higher resistance R_{DSON} which causes conduction losses, slower switching, and the greater weight relative to lighter possible weight of other solutions. The slow switching is more a problem when used for switching inductor current for DCDC conversion e.g when discharging the actuator. For example during a 'turn-off' of the switch highlighted green at the left side of Fig. 4.5, if the transistor transitions too slowly from saturation through its linear regime to a non-conducting state while the secondary current I_s represented by the green current arrow in the left schematic of Fig. 4.6 has already ramped to a large value, then the change in magnetic flux is decreased and the transistor presents large impedance during this transition causing power dissipation. This power dissipation $P_{diss} = I_S^2 R_{MH}$ is the mechanism by which the linear half-bridge driver's transistors dissipate power in Eq. B.1 in order to regulate the output.

Considering instead highside N-ch switches, optical gate drivers can drive the gate of the N-channel switch by using a low-voltage light source typical an LED directly driven by the microcontroller, to activate a photo-diode whose output is connected to the gate of the highside MOSFET. This is usually implemented inside an encapsulated IC but could potentially be done with discrete components and significant design effort, but would likely require several mA to drive in discrete form and is only available in relatively large and heavy package ICs. The current requirements to drive the photo-emitter exceed the power budget for the gate driver on the Robofly. Furthermore, the low-current on the order of μA output from the internal photocells cannot charge the gate capacitance of an N-channel FET quickly enough or saturate an NPN BJT. Thus, the optical gate driving solution is not selected.

Bootstrap capacitors are often used to drive highside N-channel switches. The bootstrap capacitor is implemented on the MUXed PEU in sec.4.3.4, where the source terminal of the recovery shunt switch M_{recov} in Fig. 4.13 is connected to the relatively constant low voltage supply V_S. The challenge of implementing this on the PEU's is that the discrete components used in the bootstrap network must be chosen by design for a specific range of operating voltages. Therefore the highside N-ch switch connected the to unipolar sinusoidal actuator signal varying from 0-V_b, makes this design impracticable.

Pulse transformers [45] are the last option mentioned above, which are two coupled inductors here used in a flyback operation giving isolation of the highside from the lowside. Discrete resistors are usually used to damp the pulse and drain off the output for a positive switching action both for BJT's and FET's. These are commercially available but too heavy for FIR SWaP.

For the role of the mux, an EPC2050 N-CH enhancement mode GaNFET is selected due to light weight at 5 mg, very high voltage rating of 350 V, and extremely fast switching relative to silicon MOSFETs.

Custom Pulse Transformers

The pulse transformers had to be extremely light weight for the gate driver to be worthwhile. They must also meet the isolation voltage need to protect the microcontroller from the full output voltage of the PEU. Suitably light commercially available options do not exist, so a custom design and fabrication was performed.

To obtain a small and light power magnetic for any converter, the compromise is generally that the switching must be faster because it must be switched fast enough that the tiny magnetic core does not instantly saturate. Based on manufacturer datasheet specification for the RMS current ($\propto$ saturation current I_{sat}), a surface mount choke suitable for the addition of a secondary winding was estimated using Eq. 4.2 to be able to store 0.037 μJ. Using Eq. 4.3 to determine the maximum allowable on-time for the low voltage switch, 2 clock cycles of the 16 mHz pwm timer of the NRF52840 were allowable before core

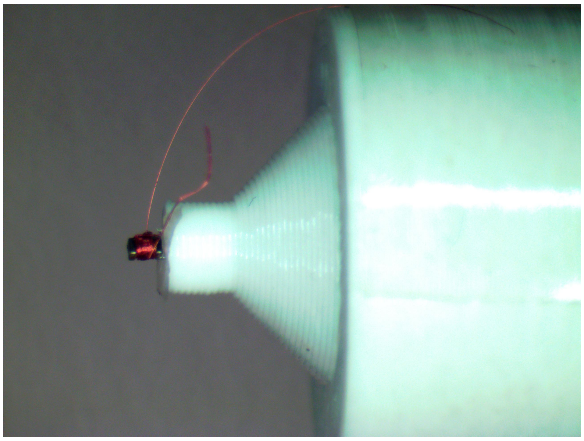

Figure 4.14: The winding of the custom pulse transformer on a 1.1 mg choke having inductance 3.3 μH. 30 turns were wound with attention to winding direction, obtaining a measured secondary inductance of 2 μH and the relatively high winding resistance of 4.5 Ω desirable to damp the pulse. The secondary resistance is desired high (counterintuitively for power conversion in general) because the output of the pulse transformer should be damped so as to avoid ringing of the output, and to obtain a designed pulse duration.

saturation current is exceeded given $V_S = 5.0$V.

LFBD

The leaky-flyback-driver (LFBD) depicted below fills the slow-mux role for the muxed PEU design. Here, a custom wound transformer was designed and fabricated by wrapping a secondary winding (adding 0.6 mg to the weight of 1.1 mg) around a Coilcraft SMT choke. By switching the isolated primary with low voltage supply V_S using the MCU digital output V_{DO} and ultra-small transistor Q_{mux}, energy is built up in the core of the 1.7mg transformer by rapidly rising current, similar to the boost converter discussed above. When the primary switch Q_{mux} opens, the collapsing magnetic field delivers a pulse of energy to the secondary side via L_S, series damping resistor R_S, and the diode D_S. This results in an increase in voltage $V_{gs} = V_g - V_{src}$ proportional to the total capacitance between gate V_g and source V_{src}, bringing the EPC2050 GaNFET gate above fully on threshold within the time of two

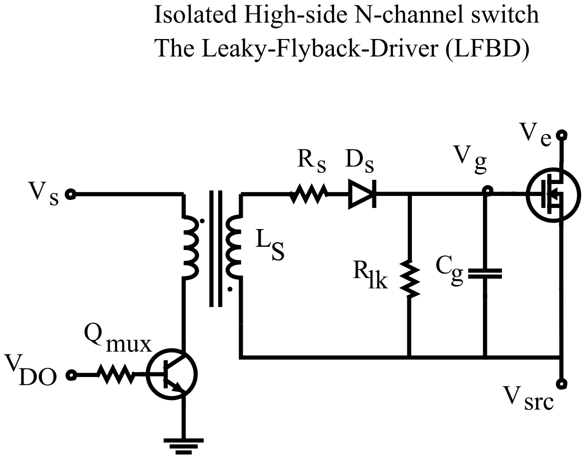

Figure 4.15: Schematic of the custom highside MUX switch, a 350 V GaNFET transistor driven by a fully isolated custom wound 1.7 mg pulse transformer and designed RC turnoff circuit. The assembled module enables switching *highside* loads up to 350 V above the microcontroller logic level, without intermediate gate driving or intermediate equally heavy transistors just to drive the highside switch. This switch activates and deactivates in microseconds, has extremely low R_{DSON}, and the power requirement to drive the pulse transformer is approximately 1 mW, which is much lower than the lighted commercially available opto-coupler or optical gate driver.

pulses, as shown below. A resistor is placed in parallel to the capacitor in order to result in a charge leakage which will turn off the GaNFET after a controlled first order decay. This resistor is selected to ensure that $\tau = R_L C_g$ is as fast as possible but does not risk draining the gate charge significantly within the time of two consecutive cycles of the low duty cycle 500kHz PWM activation signal.

To minimize the power consumption of driving the LFBD, the T_{ON} was minimized at 1 pwm clock period, 62.5 ns.

Challenges of this design with a microcontroller such as the NRF52840 which is not designed for gate driving tasks, were that the first design of the ILBD muxes worked in isolation but caused cross activation of additional switches when multiple MUXes were tested. After investigation, it was found that the ultra small mosfets to switch the low-

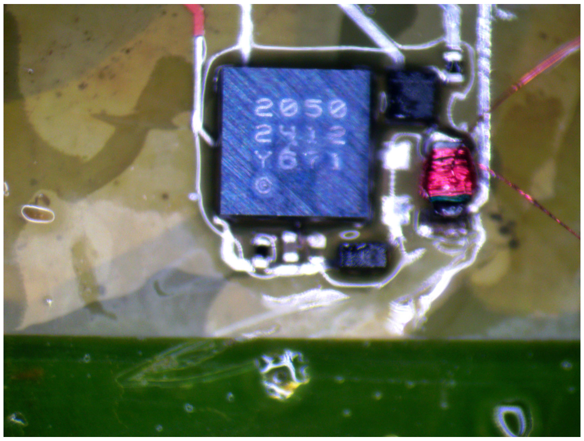

Figure 4.16: Fabricated and tested isolated high-side N-ch switch on 12 μm polimide flex pcb substrate, driven by the leaky-flyback-driver (LFBD) The largest component is the EPC2050 GaNFET weighing 5 mg, and the custom wound 1:1 30-turn pulse transformer weighs 1.7 mg. The other discrete components include $D_s, R_s, R_{lk}, C_g, Q_{MUX}, andR_{muxg}$, which weigh approximately 0.2-0.5 mg.

voltage primary side formed a second order system with their gate capacitance and realistic trace inductances. When switched at 500 kHz with a minimal pulse period of 62.5 ns, the test measurements of the LFBD modules showed excessive ringing on the DIO lines affecting all the microcontroller's outputs associated with the cross-activation. This can be seen in Fig. 4.17.

The ringing in the DIO's was caused by trace inductance, deliberately low series gate resistance, and the low voltage mosfet's gate capacitances causing large currents at the commencement of the pulse. This underdamped second order system formed resulted in each 62.5 ns pulse getting an impulse response with large overshoot, affecting the other microcontroller outputs. Simulation with trace inductance and internal coupling capacitance between the DIOs indicate this diagnosis.

The fix was actually quite counter-intuitive. A series resistance could not be added between the DIO and the low voltage pulse transformer MOSFET because the operating frequency would reduce below the limit imposed by the transformers's small size. Therefore,

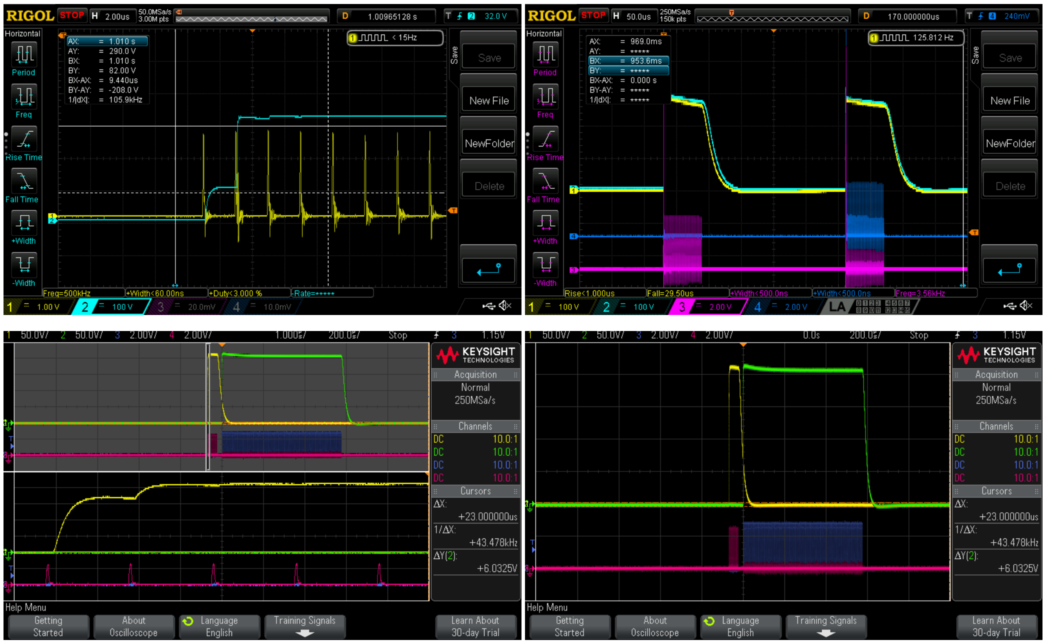

Figure 4.17: (Top left) A single highside MUX (LFBD) module, tested with an external power supply, the LFBD is shown switching the 250 V blue trace to a load resistor. The yellow trace is the LFBD drive signal from the DIO, which (exactly as in simulation) resulted in full saturation of the EPC2050 GaNFET after only two 62.5 ns pulses at 500 kHz. However, this yellow trace already shows signs of trouble in the excessive ringing on the DIO.

(Top right) Tests of two highside mux (LFBD's) are being driven by the magenta and blue traces at *separate* times from the MCU, in order to test reliable independent activation of the MUXes. The green and yellow traces are the switch high voltage signals, which are both activating together which could destroy PEU. Ringing on the DIO caused by excessive instantaneous gate charge currents is the culprit, confirmed in simulation.

(Bot left) After changing the design to BJT for the lowside FB driver, and reducing drive current to bare minimum of 30 mA, the LFBD switch modules clearly demonstrate less ringing on the DIO, as shown in the scope zoom of the initial pulses.

(Bot right) A full demonstration of independent highside switching after the redesign, consuming a mere 1 mW per channel while on!

to actually reduce the instantaneous current required of the microcontroller DIOs, the unusual choice was made to substitute the pulse transformer's low-voltage switching element from the typical mosfet, to a BJT, at which point the base current to the BJT could be deliberately reduced so as to protect the microcontroller. In fact, the BJT was *deliberately* operated in its linear region, so as to only allow the bare minimum inductor current to deliver an adequate pulse to the HV GaNFET gate - thus minimizing the power required to drive the mux.

4.3.6 Conclusion - Comparison of topologies

The design choices available for PEU topologies for FIRs has been been increased. While the efficiency over range of operation as plotted, measured, and computed, is one consideration, other considerations are total weight, a measure of relative complexity such as the number of discrete components and number of required control signals, and the ability of the topology to generate distinct and repeatable output voltages for separate piezoelectric actuators.

The tabulated weight of the muxed PEU is approximately 135 mg, not including microcontroller which is typically part of the separate control board. It requires only two PFM outputs, which are software muxed to the appropriate channel's control signal, four timer outputs for the muxes which do not have to change at high frequency, and an optional digital output to control the recovery shunt in order to reduce the charging of the bias rail if desired. The extension of the single-stage flyback converter/driver topology, shown in Fig. 4.5, weighs approximately 135 mg and is less efficient than the muxed PEU. The simple two-stage boost converter + dual linear drivers shown in Fig. 4.1 is light at $\approx 85 - 92$ mg and is the simplest to operate and fabricate, but has proven it is too inefficient at approximately 18% net efficiency to support flight with payload using anticipated power sources. Given the high efficiency and much more valuable contribution to the field, the muxed PEU is chosen for this work.

A comparison of salient aspects of discrete PEU implementations demonstrated to date on FIRs is shown in table 4.18

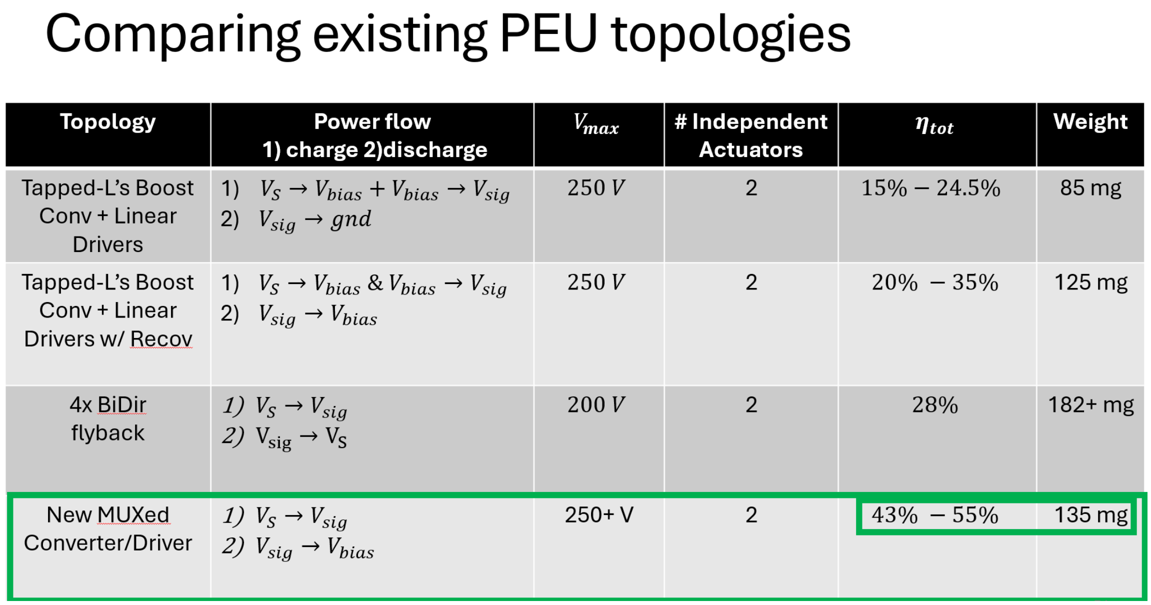

Comparing existing PEU topologies

Topology	**Power flow 1) charge 2)discharge**	V_{max}	**# Independent Actuators**	η_{tot}	**Weight**
Tapped-L's Boost Conv + Linear Drivers	1) $V_S \rightarrow V_{bias} + V_{bias} \rightarrow V_{sig}$ 2) $V_{sig} \rightarrow gnd$	250 V	2	15% – 24.5%	85 mg
Tapped-L's Boost Conv + Linear Drivers w/ Recov	1) $V_S \rightarrow V_{bias}$ & $V_{bias} \rightarrow V_{sig}$ 2) $V_{sig} \rightarrow V_{bias}$	250 V	2	20% – 35%	125 mg
4x BiDir flyback	1) $V_S \rightarrow V_{sig}$ 2) $V_{sig} \rightarrow V_S$	200 V	2	28%	182+ mg
New MUXed Converter/Driver	1) $V_S \rightarrow V_{sig}$ 2) $V_{sig} \rightarrow V_{bias}$	250+ V	2	43% – 55%	135 mg

Figure 4.18: A comparison of the new PEU design to salient aspects of discrete PEU implementations demonstrated on FIRs.

Chapter 5 Actuation autonomy and steps toward attitude stabilization – onboard waveform generation

This chapter discusses computationally light-weight control methods for precise and repeatable waveform generation, and the first steps toward implementing attitude stabilization by controlling angular rate using onboard measurement and waveform generation. Given the precise requirements on the output voltages' DC offset, amplitude, harmonic content, and the differential voltage across piezo bimorph layers, this task is nontrivial. A feedforward control term was utilized to reduce the 'work' required of feedback controllers. Obtaining these feedforward terms by model inversion and repetitive control is discussed in sections 5.3 and 5.4. Lastly, a feedforward waveform generator using motional feedback is proposed with early results, which does not strictly require output voltage feedback but is instead modulated by rate-gyro measurements.

5.1 Introduction

A critical aspect of any control system is the actuation. This is especially true for the fast and unstable dynamics of lightweight flying insect-sized robots (FIRs) such as [9] (Fig. 1.5). Considering the feedback control schematic of Fig. 1.2, the waveform generator block is what actuates the waveform references produced by the flight controller. This is basically a fast inner control loop utilizing feedback control of measured output voltage. Each PEU topology considered in this work actually has two separate 'actuators', one to charge the actuator ($\frac{dV}{dt}) > 0$), and another separate actuator with different dynamics to discharge the actuator ($\frac{dV}{dt}) > 0$).

Miniaturizing benchtop amplifiers or high voltage operational amplifiers is intractable on the robofly due to high parts count of discrete implementations or weight of the very few available monolithic high-voltage amplifier ICs. Considering digital feedback control using voltage measurements instead, simple proportional control or hysteretic (a.k.a "bang-

bang") control result in waveforms with significant harmonic distortion and large hysteretic offset, poor repeatability in flight control parameters of DC offset (pitch) and peak-to-peak amplitude (roll and thrust), and in this work were often unstable evincing low gain and phase margins with digital control at 4–10 kHz. To reduce distortion and waveform error, the gains had to be large but as expected this resulted in overshoot and vulnerability to noise. On the other hand, reduced proportional gain results in waveforms not meeting amplitude specifications (to a lesser degree pitch or DC offset error as well), and not meeting specified second harmonic content for yaw. The electrically noisy environment of the PEU combined with the lightweight, unshielded, and very high-impedance voltage dividers were very susceptible to noise. Additional filtering, analog or digital, risks phase time delay and software filtering must be minimal to avoid computational overhead. This work focused on digital feedback control of the measured output voltage, and in the last section proposes an open-loop flight controller directly modulating a central pulse frequency corresponding to reference flapping waveforms, using rate-gyro measurements only.

Recently, progress has been made in on-board power systems on FIRs such as the UW Robofly [30] and the Harvard Robobee [27]. These robots were *powered* wirelessly. However, their flights were in open loop, that is, with *uncontrolled* flapping, resulting in very short flights, all of which were less than a second. Here, we report an advance in this area by demonstrating a system that is small enough ($\approx$ 90 mg not including controller) to fly on an FIR that can also modulate its output waveform in a way that can, repeatably, vary the thrust output of the flapping wings. The waveform it produces is close to a pure sinusoid. This avoids exciting high-frequency self-bending modes of the piezoelectric actuators that can induce high strains reduce efficiency. Of particular importance, this can reduce cracking [16].

Because roll torque for the UW Robofly [18, 11, 10, 9] is based on body geometry and wing placement, a linear actuator for thrust force given commanded flapping amplitude is prerequisite for attitude and position control of the robot in flight. This waveform generator can also be used to drive the DC offset and additional harmonic content of V_{sig} needed for both pitch and yaw torques, as discussed in [46, 20, 9].

The advance discussed here is a new waveform generation control strategy in which a simple model is used to to generate a lookup table for actuator pulse frequency throughout the flapping cycle. This simplified model makes use of the clear structure of the lookup tables. This basic structure is learned by repetitive control in sec.5.4, confirmed by simulation using continuous time proportional control with unity negative feedback (i.e. having zero noise and perfect instantaneous voltage measurements) in sec.5.2.2, and also computed by modeling of actuator and load in sec.5.3. Regardless of how this feedforward PFM LUT is obtained, simple linear operations on the row-entries k of LUT achieve the roll and pitch control control, while indexing the appropriate column of LUT based on measured yaw-rate results in the LUT shape corresponding to varying differential stroke speed for yaw control. Lateral force control can be achieved similarly to yaw control, since net fore-aft forces are similarly obtained by differential stroke speed of the wings, as described in chapter 1.

5.2 Digital feedback controlled waveform generation

The waveform generator control laws implemented on microcontrollers have been designed to discretize a flapping cycle into $N_{wfg} = 26$–$N_{wfg} = 100$ iterations. These values are chosen as a middle-ground between fine resolution of output voltage regulation vs computational burden of faster control. With $N_{wfg} = 26$ the waveform generator control law operates at ≈ 5 kHz with a period of ≈ 200 μs, for flapping frequencies of 180–200 Hz, as given by Eq. 5.1.

Because the digital outputs used in either PWM or PFM modes are operated in the hundreds of kHz, many individual pulses can occur within a single T_{wfg}, and only the duty-cycle is incremented once during the waveform generator control period of T_{wfg}. Thus, it is the slewing rate which is discretized into 26 possible values and has relatively coarse resolution, while the PEU output signals go with *the continuous time* integral of the slewing rates of the charge/discharge drive stages.

$$T_{wfg} = \frac{1}{f_{flap} N_{wfg}}, f_{wvg} = \frac{1}{T_{wfg}} \tag{5.1}$$

In this work for all control laws it is preferred to implement the N_{wfg} PFM/PWM value sequences by writing the sequences to a RAM buffer which is incremented every T_{wfg} seconds by a hardware timer, resulting in the MCU digital output pulsing at the frequency and period specified by the k^{th} element of the PFM sequence. This achieves timing accuracy for digital control, and minimizes the CPU intervention.

5.2.1 The feedback: voltage divider design

The sensor for the output voltage feedback is a simple voltage divider implemented with the lightest components possible and preferably without the complication of a surface mount buffer circuit. For low voltage sensing, a voltage divider is a trivial task, but for the FIR we wish to avoid the weight of buffers and the piezo voltages employed make this a somewhat tricky tradeoff between power consumption, equivalent source impedance, ADC recommendation for source impedance, ADC sample time, resolution and sensor gain, full-scale of

ADC input constrained to a maximum voltage, signal-to-noise ratio, weight, component count, and component voltage ratings. Appendix B discusses the these design choices.

5.2.2 Classical control approaches

Proportional control for the waveform generators was generally found to be vulnerable to noise and any phase-time delay incurred by filtering, corresponding to low gain and phase margins. Emperical design methods were used in LTSpice simulations and MATLAB Simscape Electrical simulation. Emperical tuning was also performed on physical tapped-inductor boost converter + halfbridge driver PEU's, with both emulated RC bimorph on breadboard, as well as driving actual Robofly wings with 2.7nF 'single-wide' actuators. While hardware implementation was costly due to the frequent breakage of actuators due to voltage transients during experimentation, the circuit simulations and hardware-in-loop RC load were useful to explore the design and issues. Quality feedforward alleviates this poor control by reducing the need for large proportional gain, and limited continuous-time or digital filtering is possible with careful design to fit within SSWaP budget of FIRs.

Notable onboard waveform generators from literature

The HAMR family of cockroach inspired piezo actuated robots [35, 33, 21] and the Harvard Robobee [28, 46] are very useful to study, as they have similar problems to solve in wireless power and control.

5.2.3 HAMR

The HAMR robots, like the Robofly and Robobee typically use a wire tether for piezo signals when it is convenient to do so. Unlike the Robobee or Robofly, the HAMR robots are ground robots and have therefore seen much more work with wireless operation under battery power. Thus, onboard PEU's have been used to generate piezo drive signals much more so than FIRs. Since the HAMR platforms feature many more actuators (two piezos per leg), the onboard systems have quite a difficult task of generating many piezo signals. The onboard systems of the HAMR have typically been the same two-stage boost converter and half-bridge driver topology as discussed in chapt.4 due to well-established high efficiency and high step ratio of the switch tapped-inductor boost converter. The HAMR, having many more actuators, has to have the same number of drivers. Simplicity of the half-bridge driver is desirable and presents a relatively easy, if power inefficient, solution. Given the large number of independent piezo signals to drive (each needing 2 control signals) and a limited number of digital outputs, the HAMR robots have favored smoothed square waves achieved by multiplexing a fewer number of distinct digital outputs at much lower frequency than the PFM employed here. This approach can be seen in the figures of [21], but would not result in the repeatable DC offset, differential stroke speed, and variable amplitude needed for the Robofly flight controller.

5.2.4 Robobee

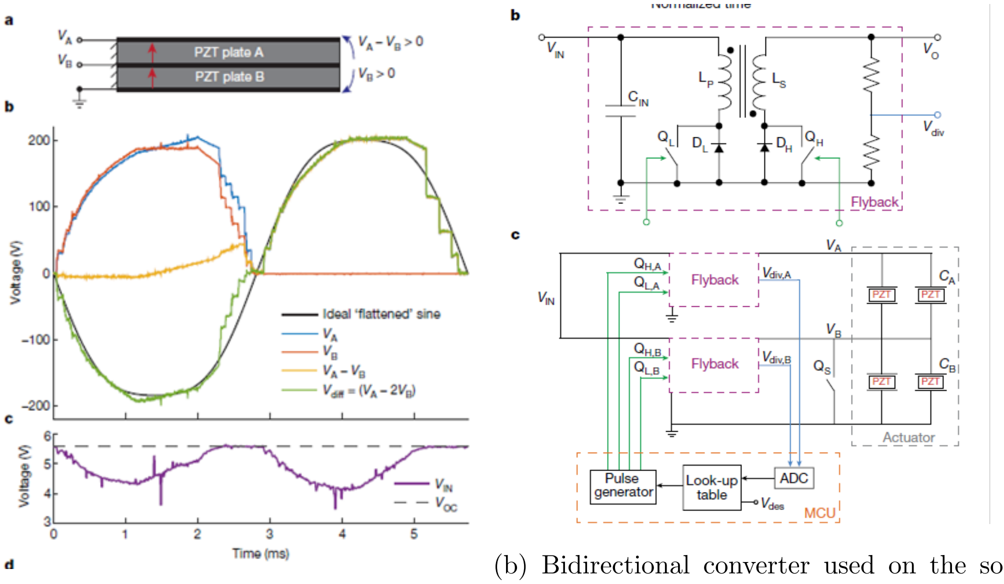

(a) Robobee piezo drive waveforms.

(b) Bidirectional converter used on the solar-powered Robobee.

Figure 5.1: Figures from piezo drive signals and power electronics circuit schematic figures for review. Figures from [28]

Briefly, Fig. 5.1 depicts the drive signals used for the two bimorph piezo actuators on the Robobee X-wing [28]. Significant harmonic distortion is clear due to the hsyteretic bang-bang control law and excessively long switch 'on' times for the Q_H. The first pulse of the discharging phase after 3ms, for example, appears to reduce the actuator voltage by nearly 90 volts instantaneously. This waveform control is considered unacceptable in this work due to actuator damage.

Since the tip deflection of the bimorph depends on *the differential* strain between both layers to the neutral fiber in the passive layer, maintaining a voltage across the least strained layer in the bimorph is actually wasteful. This is clear from comparing a simple ratios: consider two separate bimorph actuators and imagine we desire zero differential strain. This could be achieved with any number of voltages, we could for example apply $V_b = 2$V, $V_{sig} = 1V$, and on the second actuator we could apply 200V to one layer and 100V to the other. The second actuator, being at much higher voltage, would experience much higher losses and so it's better to get the same strain ratio with minimal input. The logical continuation of this is to drive only one of the two layers of the bimorph at a time, leaving the other layer to be unstrained, thereby getting the most strain out of the range of electric field. This insight was actually reported in[52], and [44] before [27] and was not *new* as claimed. While this is an efficient drive scheme, the particular implementation in[27] required a high voltage switch whose only 'job' was to hold one layer of the pzt to ground during half the cycle.

5.3 Model based feedforward waveform control

Since the piezoelectric actuator is operated well below its self resonant frequency, it is considered to be in quasi steady-state and as such can be considered to be a predominantly capacitive load [1]. Given the high electric fields of FIR application, the dielectric losses represented by R_0 as discussed in chap.1 should not be neglected, although the hysteretic losses can be in our FIR usage due to low amplitude current. Because the impedance of the attached load is small (but not negligible), we do not consider it here just for simplicity. Without considering the resonant mode, a model of the bimorph actuator depicted in Fig. 5.2 is connected to a bias capacitor at the top electrode and any bias charging stage is effectively disconnected by a reverse biased diode (therefore representing multiple PEU topologies in this work). To determine a feedforward actuation signal over time in order to drive the bimorph signal V_{sig} to reference V_{ref} during charging phase of the signal ($+ve\frac{dV}{dt}$), an energy balance can be derived. Considering the energy balance terms for the resonant modes improves accuracy. This makes use of:

a) the relationship of energy to the voltage across a capacitor given by Eq. 4.6.

b) the energy delivered to the actuator through the charging converter's inductor, given by Eq. 4.5 multiplied by the conversion efficiency

c) The energy lost by the parallel dielectric loss R_0 of each actuator layer, given as the integral of $\frac{V^2}{R}$ power dissipation over each of the resistors.

- Simplifying assumptions: small time-steps and small δV simplifies the energy dissipation terms; and assuming that the $C_b >> C_0$, the capacitance of the series (C_0-

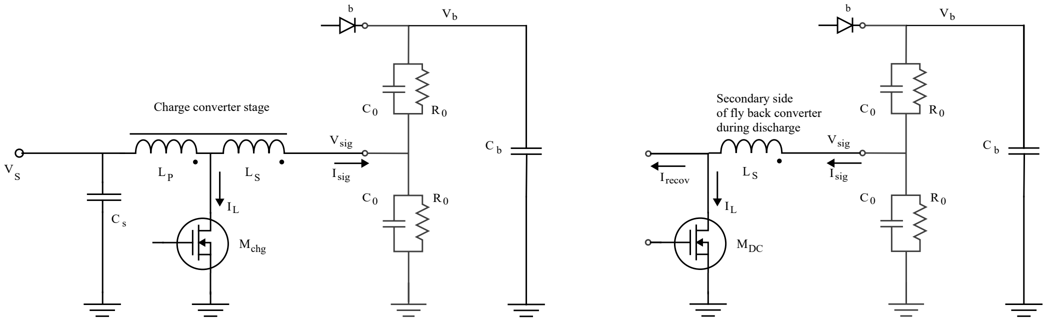

Figure 5.2: (Left) Simplified actuator model with principal capacitance and dielectric loss, connected to the charging stage (here shown as a switch-tapped inductor boost converter). The bias diode is reverse biased, and the bias rail capacitance is in series with the bimorph's top PZT principal capacitance.
(Right) The simplified actuator model during discharging phase. Here, A secondary inductor is shown an actuator discharging conversion stage with current I_L when switch M_{DC} conducts and output current I_{recov} when the inductor's stored energy collapses at the opening of M_{DC}. This discharging converter could take several forms, e.g. a flyback converter returning energy to the low voltage source V_s, or this inductor can also be a boost converter (as in the muxed PEU or two stage topology with energy recovery from chap.4), which would instead discharge the energy taken from the actuator back into the bias rail with typically 50% efficiency. The method of charge/discharge stages just changes the LHS of the energy balance equations below.

$C_b) \approx C_0$; the discharge conversion does not return energy to the bias rail with losses, although this too can be modeled in this case we simplify.

Using the models shown in Fig. 5.2, during actuator charging stage the energy balance is given by Eq. 5.2, and during discharging phase given by Eq. 5.3

$$\frac{0.5T_{on}^2V_s^2\eta}{L_p} = 0.5C_{bot}\left(-V_{sig}^2 + r_{sig}^2\right) + 0.5C_{top}\left(-\left(V_b - V_{sig}\right)^2 + \left(V_b - r_{sig}\right)^2\right) + \frac{T_k\left(\frac{V_{sig}}{2} + \frac{r_{sig}}{2}\right)^2}{R_0} - \frac{T_k\left(V_b - \frac{V_{sig}}{2} - \frac{r_{sig}}{2}\right)^2}{R_0} \tag{5.2}$$

$$\frac{0.5T_{on}^2 V_{sig}^2}{L_s} = 0.5C_{bot}\left(V_{sig}^2 - r_{sig}^2\right) + 0.5C_{top}\left(-\left(V_b - V_{sig}\right)^2 + \left(V_b - r_{sig}\right)^2\right) + \frac{T_k\left(\frac{V_{sig}}{2} + \frac{r_{sig}}{2}\right)^2}{R_0} - \frac{T_k\left(V_b - \frac{V_{sig}}{2} - \frac{r_{sig}}{2}\right)^2}{R_0} \tag{5.3}$$

Where

T_{on} is the on-time of the switch, less than the discrete control period T_k corresponding to the switch duty-cycle, this is the desired control variable.

T_k is the time period of the discrete control interval, over which the dielectric loss R_0 is dissipating power.

V_{sig} is the initial voltage of the actuator signal layer.

r_{sig} is the reference voltage of the actuator signal layer.

These energy balances are solved for the switch on-times, faciliated by symbolic computation utilities, to yield pulsing frequencies for a given waveform generator control period. Equations for the solutions to eqs. 5.3, 5.2 are in appendix A.2.1. The resulting pulse trains drive a circuit simulation of the corresponding model as intended, and this is what is prescribed in literature for obtaining these pulse sequences as in [38]. However there is still a big issue: model features influencing the phase angle between voltage and current to the actuator means that the choice of which stage between charge vs discharge to utilize at any given time, isn't just as simple as looking at the sign of the desired change in reference voltage at discrete time k.

Here, the phase shift visible in the experimental power calculations shown in Fig. 2.3 is apparent in the phase of the control signal relative to the reference waveform. In this case, charge and discharge control signals u_{chg}, u_{disCG} are zeroed when they disagree with the sign of the voltage change over time δV_{ref}. This is to make clear that approaches such as [37] which choose the timing of monotonically charging vs monotonically discharging phases based solely on the sign of the change in reference signal, are inadequate on their own to determine which of the conversion stages to activate at any given time, complicating the use of such a feedforward signal in addition to a proportional control law which demonstrates phase lag.

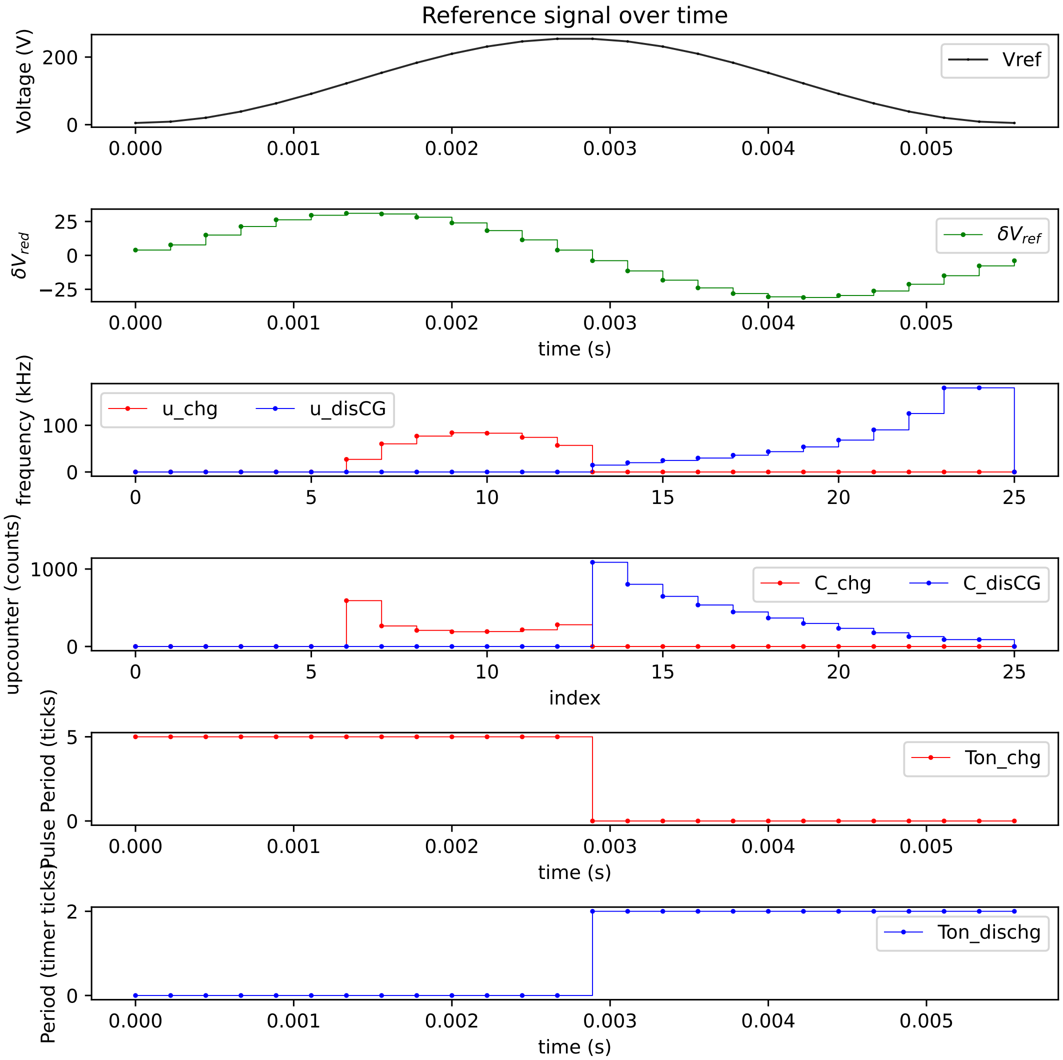

Figure 5.3: A model based PFM sequence (holding the pulse period constant) for charge and discharge phases is obtained from model inversion.

These equations come from a simplified model but still provide a useful starting point for a feedforward term, however they are very hard to compute with available microcontroller resources especially since positive and negative solutions drop out and one must be chosen from multiple, and the solutions are vulnerable to modelling inaccuracies from varying load or actuator properties and/or time-varying control authority of the DC-DC conversion stages e.g. from changing supply voltage, drifting transistor gains, parasitic capacitance, inductance, and resistance of traces, etc. The most confounding factors for a purely feedforward model based control sequence in practice are 1) the large variation in capacitance when different actuators are driven in FIR voltage ranges $V = 150 - 300$ V, such as studied in [38] 2) difficulty to measure or predict variable wing loading during PEU experiments and from vehicle to vehicle.

For these reasons, a repetitive controller was envisioned and implemented to learn the control sequence instead. Subsequently, a motional feedback scheme to modulate feedforward pulse trains based on either the full model or a simplified model is demonstrated.

5.4 Repetitive control design and implementation

Here we describe how we created the desired waveform, measured the output of the PEU, and measured thrust from the flapping wing of an FIR. In all experiments and depicted waveforms, the V_{Sig} is set to have DC-offset equal to half the V_{Bias} command, which corresponds to a zero pitch-torque command [46, 9].

5.4.1 Robofly implementation of a repetitive controller

We fabricated a benchtop version of the PEU (Fig. 5.4 and implemented the controller in C on a ST Microelectronics STM32F4 microcontroller, chosen because it provides many timers, a small (8 mg) package, and floating-point math. The voltage divider outputs from the PEU sensing circuitry were measured by the ADC of the STM32F4 microcontroller at the controller sampling frequency of 8.1 kHz. The code was debugged and tuned on a passive resistor-capacitor (RC) circuit with impedance similar to the load of the FIR's piezo actuator-wing system. The lumped-parameter coefficients modeling the FIR's wing system were found to be highly variable due to wear and varying payload, making model-based control design difficult. The RC test load, therefore, was valuable for empirical tuning even though it does not exhibit the full dynamics of the target load.

Voltages V_{Bias} and V_{Sig} were directly measured by oscilloscope (Rigol DS1054Z) at 10 kHz. Simultaneously, a running log of preceding 10 flapping cycles was stored in the MCU until the end of the experiment and then transmitted via UART to a custom-written NI LabVIEW software interface. Additionally, a custom circuit used isolation amplifiers (AMC1302, Texas Instruments) to measure the highside currents I_S, I_{Bias}, and I_{Sig} depicted in Fig. 3.10, and these measurements were logged by an NI USB digital acquisition system (DAQ) concurrently with oscilloscope measurements. Oscilloscope, MCU, and DAQ measurements were synchronized in time by measuring an additional digital timing signal produced by the MCU. Fig. 5.5 shows an example learned waveform driving the RC load.

To measure thrust, a single wing in a Robofly similar to [11] was driven by the PEU while mounted to a capacitive force sensor similar to that presented in [17]. This sensor measures small displacements, read-out as changes in capacitance, of a flexible beam. This configuration provides a high-bandwidth, low-error measurement of the beam position. The beam was configured so that thrust force from the wing induced deformation along its flexible mode. The output of the capacitance meter (Microsense 8810) varies linearly with displacement and therefore applied force, and was also measured by the NI DAQ and LabVIEW software interface (detailed in §A.3.1). To eliminate sensor drift, the electrical current and force sensors were zeroed by taking a time-averaged measurement after settling for several seconds after wing flapping ceased.

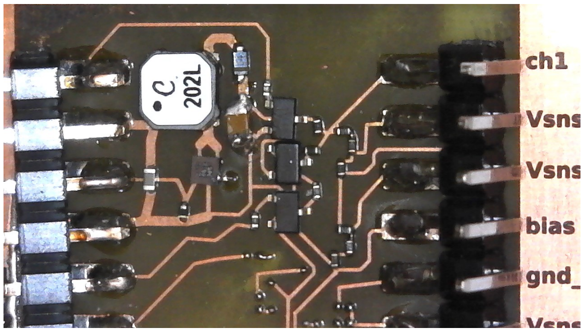

Figure 5.4: Benchtop version of the PEU. Compared to the flight-weight PEU in Fig. 1.5, it has smaller dimensions, is on a heavier non-flex FR-4 substrate that is easier to prototype but is functionally equivalent, includes additional connections for characterization, uses a GaNFET as the the boost converter switch, uses a larger bias capacitor C_2, and the coupled inductor's shield was not removed. The MCU is on an external PCB for ease of development. The authors previously demonstrated these components on a single flex circuit in [30].

All thrust experiments were done at 240 V bias, near the 250 V maximum voltage rating of the PEU's components. This maximizes potential thrust and control authority but reduces boost converter efficiency.

Insofar as end-to-end efficiency, because the half-bridge driver loses reverse-power (i.e. dumps to ground any negative I_{Sig}), we have zeroed the reverse power as in [10, 18] and do not consider negative power to be 'returned'. Taking measurements of current and voltage of channel 1 and bias, as indicated in Fig. 4.1, with $\hat{I}_{Sig} = \max(0,\ I_{Sig})$, and $\overline{x}$ as the time average, taken over 10-cycles, end-to-end efficiency is given by:

$$\eta = \frac{\overline{(V_{Sig}\hat{I}_{Sig} + V_{Bias}I_{Bias})}}{\overline{(V_sI_s)}} \tag{5.4}$$

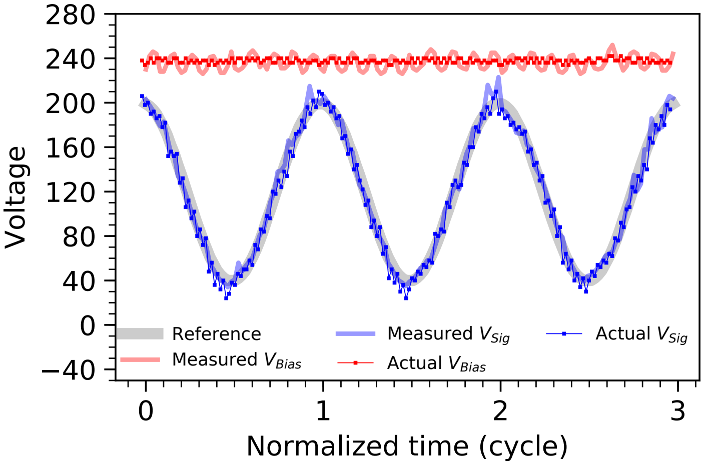

Figure 5.5: Typical voltage over time normalized to the period of one wing-stroke, from the PEU driving the RC ‘dummy’ load, showing that it follows the reference signal well. Settings: frequency 160 Hz, 160 V_{pp}, 120 V_{DC} offset from gnd, and bias signal $V_{Bias} = 240$ V. Oscilloscope measurements (“Actual”) assumed to be ground truth are overlain, showing agreement with online ADC measurements (“Measured”) made by the MCU. The abscissa is time normalized to the period of one wing stroke.

5.5 *Results*

The estimated weight of our PEU excluding interchangeable microcontroller but including flex PCB, larger bias storage capacitor, and dual drivers is 90 mg, fitting within the payload budget of a < 500 mg flying insect robot (FIR).

Fig. 5.7 shows that the PEU and repetitive controller is capable of producing signals across a range of desired amplitudes despite electrical noise in feedback measurements significantly exacerbated by the test apparatus. Relatively good control of V_{Bias}, V_{Sig} can be observed as a relatively small deviation from reference and a nearly constant DC bias voltage, despite the large range of instantaneous boost converter power throughout the flapping cycle.

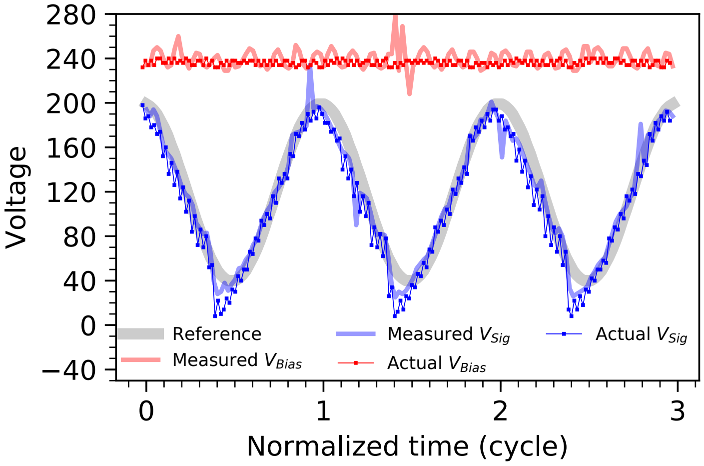

Figure 5.6: Typical waveform from the PEU driving a Robofly wing to follow a sinusoidal reference signal, showing slightly diminished performance relative to the RC load. Waveform settings are the same as in Figure 5.5.

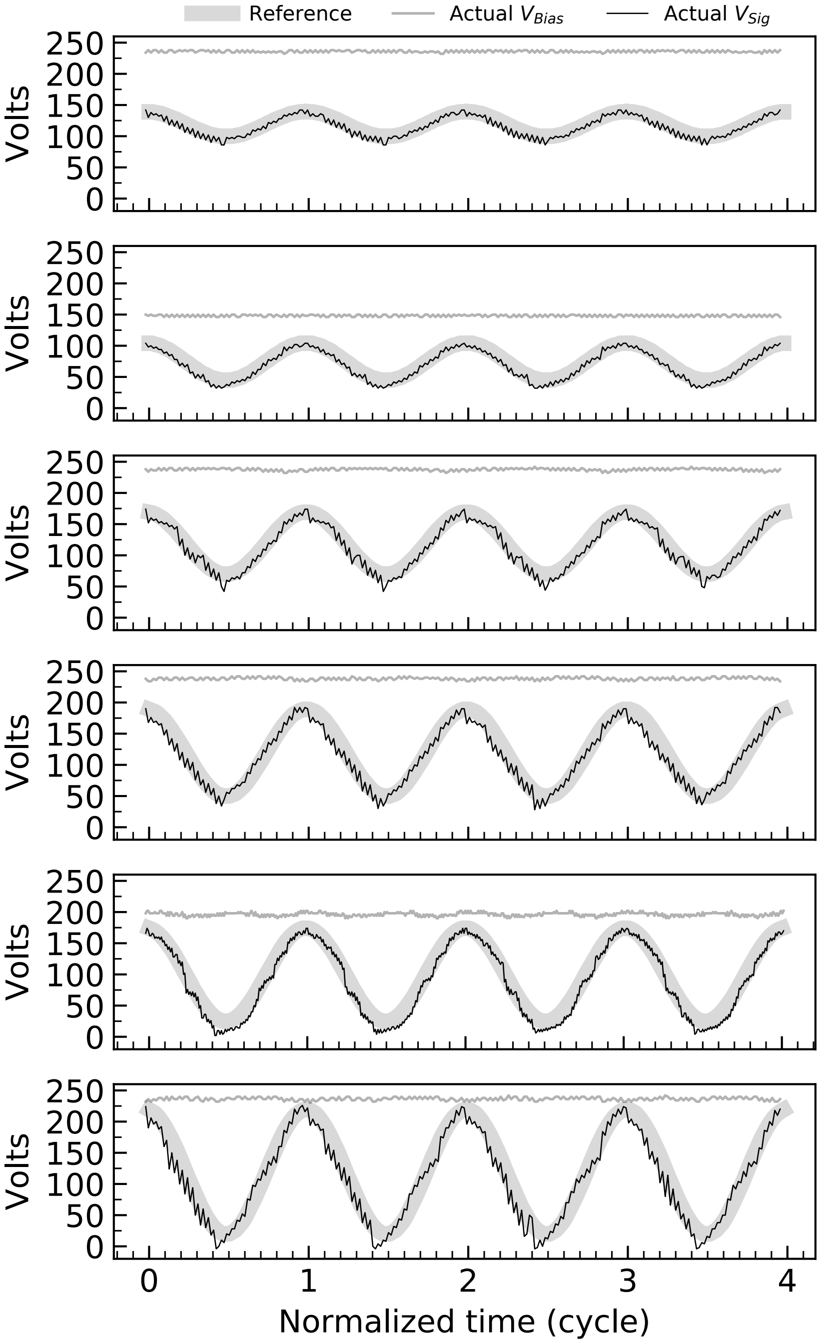

Figure 5.7: The PEU can consistently produce waveforms across a range of desired amplitudes and V_{Bias}, while maintaining a nearly constant V_{Bias}. From top-to-bottom, the V_{Bias}/V_{Sig} set-points are respectively: 240/40; 150/60; 240/100; 240/140; 200/150; 240/200.

The PFM output of a single driver channel including high-side and low-side transistors is shown in Fig. 5.9 over one complete cycle of wing-flapping. An example of the PFM *LUT* input to the high-side and low-side driver transistor gates is shown in Fig. 2.6.

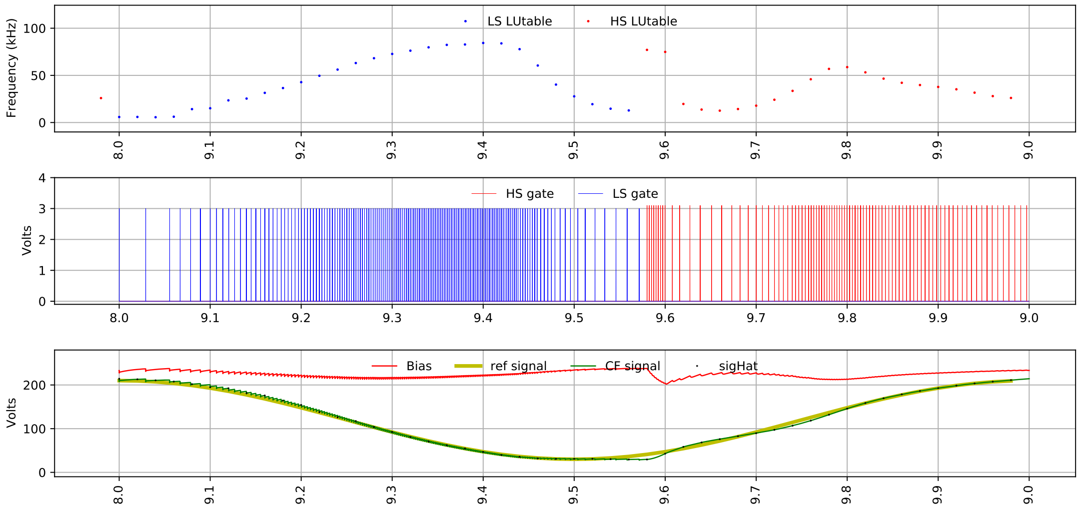

Figure 5.8: (Duplicate of Fig. 2.6) The Lookup table (converted to frequency) for the driver highside and lowside outputs. [Mid] The pulse trains corresponding to the commanded frequency set at each sample of the lookup table (50 controller samples per cycle). Note that there are many pulses possible within a single control sample period, since the digital timers are running at *much* higher frequency than is the digital controller. Each pulse period is an integer multiple of the system clock period, typically 0.3–5 μsec. The pulse width is adjustable to change driver authority. [Bot] The output waveform from the simulated circuit with a high-impedance bias voltage source representing the boost converter (which is faster to simulate). The unaided repetitive controller can be seen to compensate for periodic disturbances in the bias rail which supply the drive stage.

The nominal voltage over time waveform for a balanced Robofly in hover would be a unipolar sinusoid exactly centered between ground and V_{Bias} (zero pitch) and having no second harmonic content (zero yaw). This nominal waveform is depicted in Figs. 2.6, 5.8 as the reference signal, while the control signals plotted in Fig. 5.9 are the controller signals learned from scratch in order to drive the output of the half-bridge toward reference voltage.

The control signals generated by the repetitive controller are very informative regarding the structure of the control torque generation. Fig. 5.9 shows two principal dimensions of the high-dimension lookup table.

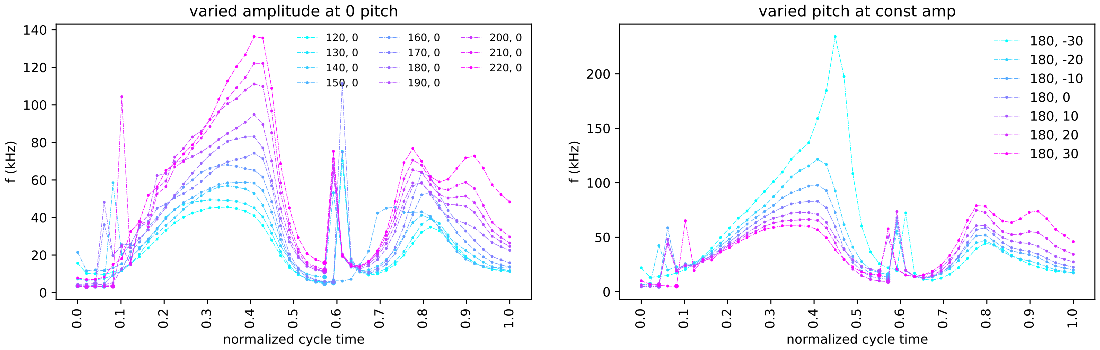

(a) [Varied amplitude with constant pitch]

(b) [Varied pitch with constant amplitude]

Figure 5.9: Control signals for two principle dimensions (pitch and roll) of the lookup table. This is the simulation output, which has good fidelity to the physical robot but is considerably cleaner since the data from physical experiments was collected with varied PEU and controller system settings, where the simulation was held constant. The frequency of the PFM is the control variable, which is muxed to the high-side or low-side of the driver depending on the voltage error. **Note** the low-side is active during the first half of each control cycle, and the high-side is active during the latter half of the flapping cycle. The transition from high-side to low-side occurs at $\approx 0.1\&0.6$ cycle time.

What is clear, in studying Fig. 5.9, is that pitch and roll would be achieved by different actions from the driver. They are uncoupled. As seen in Fig. 5.9b, in order to cause a roll (brought about by modulating flapping amplitude of one wing relative to the other) the amplitude of the signal to one wing must be increased, so the driver must increase the activation of *both* the high-side *and* the low-side.

In order to achieve pitch instead of roll, the operation is different. For example, to apply a pitch command of -30 volts as in the cyan trace in Fig. 5.9b, the entire waveform must be reduced in voltage. The waveform generator learned to do this by significantly *increasing* the low-side activation, and *decreasing* the high-side activation.

Additionally, it appears in both Fig. 5.9a and Fig. 5.9b that the relationship is fairly consistent over the range of commands in amplitude and pitch, suggesting that linear interpolation is feasible once the waveform generator has learned the two extreme cases.

The problem however, is that there is some nonlinearity apparent especially in the pitch modulation of Fig. 5.9b when the controller is resorting to greatly increasing the PFM to the lowside at $t = 0.45$ cycles in the cyan trace. This is a nonlinear increase from the less extreme pitch command cases plotted by the other traces. This occurred because the controller began to saturate because the low-side driver lacked the authority to 'pull' the output down adaequately. In other words, the waveform generator was trying to reduce the output voltage at that time in the cycle (close to zero volts) as much as it possibly could, but was unable to do enough. Studying the plots of the reference and measured voltage (to see if measurement error plays a role) was informative in diagnosing closed loop system characteristics and quirks. In this particular case, it would appear that increasing the low-side driver authority just slightly by either 1) increasing the pulse period of the hardware timer would grant the controller more authority, or 2) adjusting the output or gate resistances of the physical circuit.

It's worth noting that the high-side and low-side stages of the driver have different authority but can be balanced. The highside authority is tuned by selecting good pulse period of the digital timer as well as by physical resistor values in the driver circuit itself. The high-side portions of all traces are fairly well balanced in the center of the y-range in Figs. 5.9, indicating that the waveform generator does not appear to saturate the high-side either at high or low pulse frequency.

5.5.1 Waveform quality

The flight mechanics of FIRs and insects alike require smooth wing motion to maximize lift and minimize wear on the mechanism. piezoelectric actuators are extremely high-bandwidth which is advantageous for control [65]. But it also means that distortion of the V_{Sig} is transmitted directly to the actuator, and the wings. It is desirable to avoid high frequency content of the waveform because the half-bridge driver topology is inefficient with reactive power, and the piezoelectric actuators can be easily damaged by discontinuous V_{Sig}, V_{Bias} or frequency content at the actuator's resonant frequency [16]. Kinematic analysis of the wing motion at some radius from the wing root in [27], indicates that harmonic distortion is substantially filtered by the transmission and wing mechanism. Nevertheless, it is unknown how thrust and efficiency are impacted, if at all, by harmonic distortion.

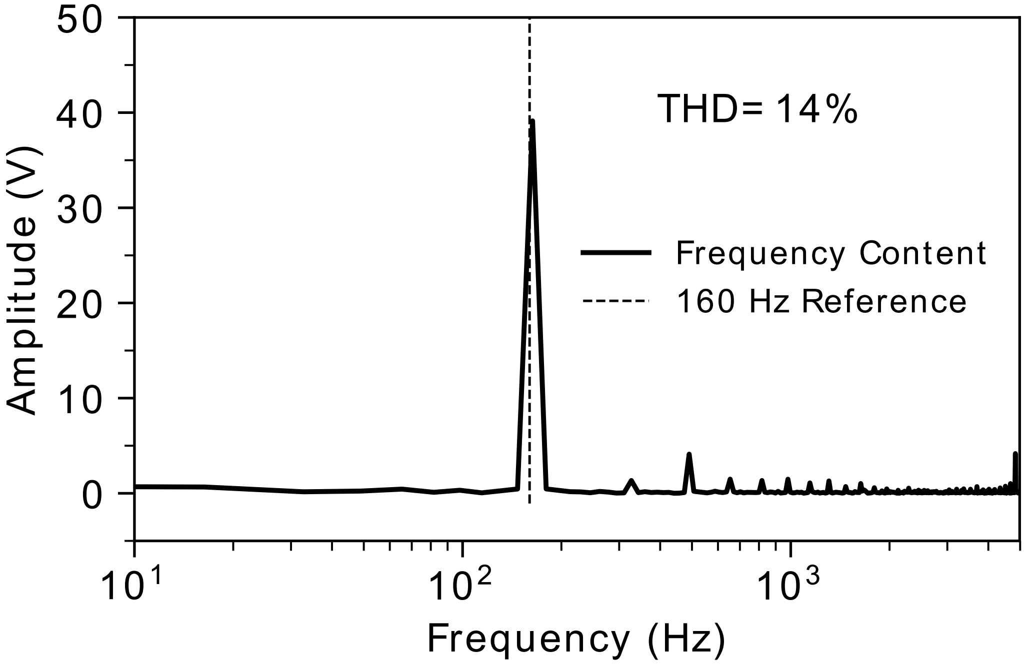

Figure 5.10: Frequency content of measured V_{Sig} from Fig. 5.6 over 10 wing-flapping cycles.

Fig. 5.10 shows a Fourier decomposition of the 160 V_{pp} commanded waveform shown in Fig. 5.6 for 10 cycles. It shows moderate frequency content far away from the fundamental frequency of the reference signal, with a total harmonic distortion (THD) of 14%, which includes quantization effects. The waveform applied to the RC test load depicted in Fig. 5.5, had the benefit from more cycles of controller training, and is visibly superior with a lower

THD of 8%. The contribution of measurement noise is difficult to know with certainty, given the measurement apparatus, but it is conjectured to represent a non-negligible component of the waveform's harmonic content. Excessive low-side actuation at $t = 1.4$ cycles in Fig. 5.6 results in V_{Sig} error, V_{Bias} droop due to loading, and self-induced measurement noise.

Assuming the oscilloscope's direct measurements of V_{Bias} and V_{Sig} are ground truth, the microcontroller's measurements are noisier (Figs. 5.6 and 5.5). On FIRs, software or hardware filtering must be carefully considered due to trade-offs in computational expense and time-delays or potential loading effects on already high-impedance on-board voltage dividers for V_{Bias}, V_{Sig}.

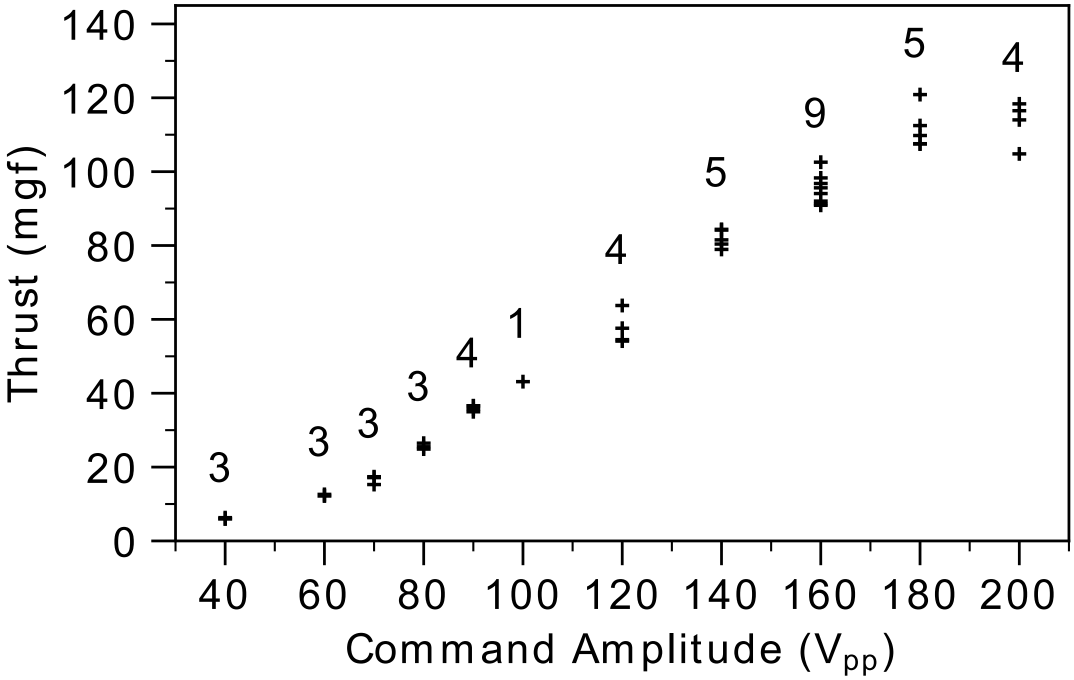

Figure 5.11: Measured thrust in milligram-force is approximately linear over commanded signal amplitude. Sample size at each commanded amplitude is labeled above the data.

5.5.2 Controlled Thrust

Fig. 5.11 shows that wing thrust varies essentially linearly with voltage amplitude (with exceptions at very large and small amplitude). This matches previous results showed on a FIR driven with a standard benchtop piezo amplifier [49], indicating that our PEU's output is not significantly different from a controls perspective.

As commanded amplitude decreases, the slope of the $T(V_{cmd})$ levels off as expected due to the correspondingly small stroke velocity of the wing and resultant decrease of the wing's passive angle-of-attack and thrust generation. We conjecture that decreasing thrust per command $T(V_{cmd})$ at higher commanded amplitude > 200 V may be due to limitations of other parts of the transmission-wing system, such as wing flexing or transmission nonlinearities. Although the specific slope of the $T(V_{cmd})$ and maximum thrust achieved by this half-fly is expected to vary dramatically for robots of different design and state of repair, the strong linearity of the measured thrust over commanded flapping amplitude is vital to linear controllers such as used for roll angle and altitude in [18]. This result is therefore an important step for achieving autonomous flight for FIRs.

Conclusions and Future Work

This work reports the first PEU at a weight compatible with an FIR that is capable of modulating wing thrust. This capability is essential for flight control, because it is needed to vary loft and/or roll torque on a two-winged robot. Results were validated with thrust measurements on an FIR.

Because the system can produce arbitrary yet precisely-controlled and repeated waveform, it provides a direct path to creating the signals needed to actuate pitch and yaw torques, which are needed for attitude and position control [46, 20, 18]. Pitch and yaw torques are driven by modulating V_{Sig}'s DC offset from ${}^{V_{Bias}}/_{2}$, and adding a 2^{nd} harmonic content [46].

Achieving lift-to-weight ratio > 1 for our FIR carrying PEU and anticipated sensor suite [18] will require $V_{Bias} > 200$ V, $V_{Sig} > 180$ V_{pp}. The bottom waveform of Fig. 5.7 exceeded this, and with stroke-averaged power to the actuators of $P_{net} = 85$ mVA had an end-to-end efficiency of 20%, and boost converter efficiency of 35% when supplied with an expected supply voltage of $V_s = 7\ V$ as from onboard photovoltaic power sources [30]. our measured end-to-end efficiency of 20% under full load is less than the 28% reported in [27], which is expected for our half-bridge topology.

In our experience, FIRs driven by standard piezo amplifiers require trimming to eliminate bias torques [11], and demonstrate variation in thrust from robot to robot and due to mechanical wear. We anticipate that PEU calibration would entail an additional learning phase while tethered to the ground. The PEU would learn and store a number of waveforms in lookup tables (e.g. Fig. 5.7). In controlled flight, the PEU could then interpolate between stored tables to produce arbitrary commanded waveforms, achieving a high control bandwidth.

Our simple linear controller was able to learn how to make sinusoids with low distortion. Future improvements will reduce the weight of the output capacitor, improve the waveform, broaden the range of load impedances that can be accommodated, speed up calibration, improve signal amplitude range closer to the bias for greater control authority, use the controller look-up table adjustments to monitor load conditions, and allow the system to adapt to gradual changes from mechanical wear.

Creating arbitrary repeatable waveforms given varying loads under the constraints of an FIR-compatible PEU benefits from a quality feedforward term. In this section I present a method to obtain lookup tables online for appropriate periodic pulse trains for both the driver and boost converter. I devised a repetitive control similar to iterative learning strategies using an LT-Spice based simulation.This simulation implements the circuit and lumped parameter model for actuator and load shown in Fig. 2.2. The actuator load (equated to series RCL elements) is the Van Dyke model described in IEEE piezoelectricity standards [1] and also detailed in [38]. The simulation results using this model were validated against the physical circuit and robot mechanism during this work. In this approach, the output of the previous cycle is stored in memory, compared to the reference signal, and linear corrections to the PFM look-up tables are made online. This repetitive learning occurs of a period of approximately 200 cycles in the results plotted below.

This system is radically different than the systems which appear to be used in other

FIRs discussed in §5.2.2, in that the *control signal lookup table* is updated in a memory loop, instead of a traditional controller computing a one-off scalar control signal at every control sampling instance k by comparing feedback to a reference obtained by indexing a reference lookup table. The simplest version of the control law is stated in Eq. 5.5. At each time k within the cycle, which ranges from 1 to n, where n is the number of samples per cycle (50–100), the look up table is updated online in response to the most recent voltage measurement according to:

$$L(k-1) = \begin{cases} L(k-1) + K_\ell L(k-1)\left(R(k) - \hat{v}\right) & \text{if } s(k) = -1 \\ L(k-1) + K_h L(k-1)\left(R(k) - \hat{v}\right) & \text{if } s(k) >= 0 \end{cases} \tag{5.5}$$

where

k is the cycle index wrapped to $[1, n]$
L is the driver pulse-period lookup table, vector of length n
$s(k)$ is the k^{th} active output, either high-side or low-side.
K_h is the high-side controller gain
K_ℓ is the low-side controller gain
$\hat{v}$ is the most recent scalar voltage measurement from the ADC.
R is the reference voltage, vector of length n
n is the length of the lookup table

K_h, K_ℓ have opposite signs due to the former pulling the output up in voltage, while the low-side instead sinks the output down toward ground. The K_h, K_ℓ may have different magnitudes due to differing control authority of high-side compared to low-side driver stages. These gains are also a function of the voltage differentials $V_{bias} - V_{Sig}$ and $V_{Sig} - gnd$.

This simple formulation can benefit from increasing heuristic or filtering sophistication if onboard computation allows. For example, since the control law in Eq. 5.5 can be executed in a loop after the completion of an entire sequence L, noncausal filters can possibly be used on the vector of measurements $\overline{\hat{v}}$ - so long as the system is in quasi steady-state and the correction terms are sufficiently small so the assumption holds that the current cycle dynamics are similar to the last.

The parameter s determines which driver output is active throughout the cycle, either the high-side or the low-side. Initial values for s can be pre-computed just by considering the rate of change of the reference signal: $s = \frac{\dot{R(k)}}{|\dot{R(k)}|}|_{k=1}^{n}$. In other words, the driver should be using the high-side when we're supposed to be increasing the output voltage, and use the low-side when the reference voltage is decreasing. This is the approach suggestd in [38], however this does not account for the phase angle of the time-varying complex load of the FIR piezo-wing system. In particular, if a naive feedforward signal which does not account for this phase is combined with a proportional control element as depicted in Fig. 5.12, the feedforward input is out of phase with the proportional controller which results in harmonic distortion. Due to this deficiency in the approaches taken in literature, it became clear that s should also be a learned parameter.

Slow adaptation of s has been implemented successfully in simulation in order to learn a phase angle between input PFM and the reference. This phase of charge vs. discharge action corresponds to the phase angle between the voltage and current of the complex actuator-wing system. §2.3.3. s can be learned by a heuristic which monitors the LUT for the condition that the PFM is saturated, and that the sign of the error indicates that the alternate driver output should be activated at that point in the cycle (i.e. it should switch to the low-side from using the high-side or vice-versa). In addition to greatly improving the control of the proportional plus repetive control scheme, this simple heuristic provides a low frequency online phase angle measurement of the complex load.

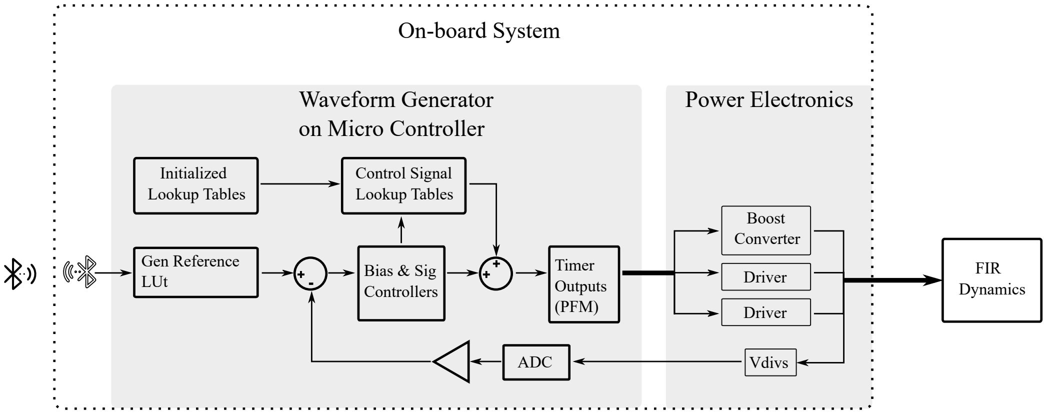

Figure 5.12: Waveform generator subsystem. The memory loop is formed between the "Control Signal Lookup Tables" and "Bias & Sig Controllers" blocks. The "Initialized Lookup Tables" shows the initialization of the control signal lookup tables, which only needs to happen once at the commencement of a new command incoming wirelessly from the flight controller. In order to initialize these LUts, the microcontroller can interpolate between existing lookuptables in a library of learned waveforms, or it can generate a guess naively or based upon a model.

In Fig. 5.12, the output signal of "control signal lookup table" is basically a learned feedforward term. This feedforward control signal is responsible for the bulk of the controller effort and current feedback has little result on the current output because *the Bias & Sig controllers are very weak/low gain.* This system is therefore shockingly simple and given time to learn performed substantially better than my attempts with a more traditional controller. The potential downside is the bandwidth of repetitive control which would fail at higher frequency R. Since the repetitive control would not be 'starting from scratch' when online, learning a slightly different waveform when the reference changes at typical frequencies during FIR flight experiments appears to converge within a couple of cycles in simulations of slewing R. This analysis of stability and bandwidth for this scheme is out of scope for this work, in which repetitive control gain is recommended to be reduced as possible.

5.5.3 Waveform Generator design in simulation

This section outlines the simulation design methods and simulated performance for learning a single waveform.

Python + LT Spice method:

Algorithm 1 Python + LT Spice waveform generator design simulation

```
procedure Simulate system to learn control signal lookup table. (circuit simulation file, command amplitude, flapping frequency, DC offset, or arbitrary reference waveform )
    In Python: Modify LTSpice sim file to accept initial conditions and PWL files for the control inputs from python script.
    Define reference waveform (i.e. sinusoid for hovering, offset sinusoid for rolling, or arbitrary signal if specified by input args). Lookup tables are arrays of length 50–100 controller samples per cycle.
    Initialize lookup tables for control inputs - these are the same length as above, and have units of clock ticks. The PFM output frequency is simply equal to $\frac{1}{period}$.
    for $N$ learning iterations do
        Modify the LTSpice simulation file for current initial conditions.
        Write the lookup tables to PWL files ( tuples with digital voltage and timestamp in a *.txt) which are read by LTSpice. Requires emulating the working of the MCU timers.
        Call LTSpice from operating system shell
        Read LTSpice simulation output in python
        Generate corrections to lookup tables based on the simulation output
    end for
    Plot results
    Optional: save end-result lookup table, add this to a library of waveforms.
end procedure
```

Fig. 5.13 shows an example of a waveform learning exercise in which the waveform generator controller was constrained to 8-bit fixed point math starting from a lookup table initialized to a constant low frequency. The pulse width parameters had to be carefully tuned to get good results.

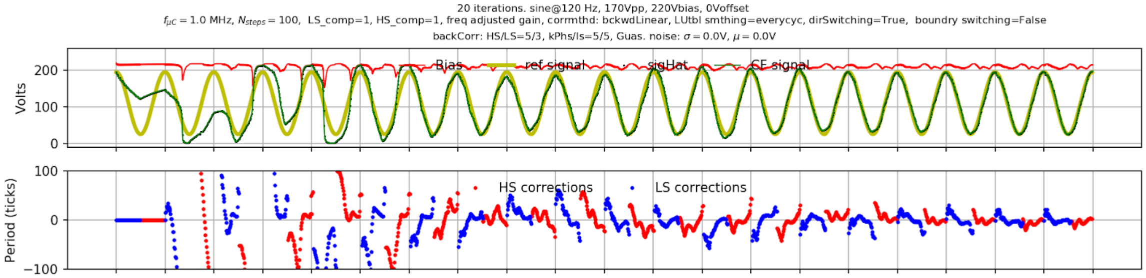

Figure 5.13: [Top] Waveform of the first twenty learning cycles from an 8-bit fixed point version of the repetitive waveform generator controller. [Bottom] The corrections made to the lookup table in response to the error of the previous cycle compared to the reference sinusoid.

5.6 Motional Feedback

To extend waveform generation for the muxed PEU into flight control, this work provisionally explored a waveform generation strategy utilizing motional feedback from an onboard IMU to complement and possibly even replace voltage feedback for waveform generation. Similar to the repetitive controller designed above, the feedforward input described here can complement a traditional voltage feedback based control. The basic idea is to use a simplified model based feed-forward look-up table (or equivalently obtain one via repetitive control) and simple transformations of this LUT for flight control. At the time of writing there is insufficient experimental validation of this technique to claim its effectiveness, however provisional results are shown below with expected functionality following the same logic used in [27] to claim that the demonstrated waveform generation extends to flight control by modulating amplitude, DC offset, and harmonic content.

Full model-based feedforward terms are computationally intensive to compute online but can possibly be computed at low priority in spare MCU cycles, and precomputed 4-dimensional lookup tables would be too large for all control dimensions of roll, pitch, yaw. e.g. a flapping cycle segmented into 26 steps would require an array of 26^3 elements. If amplitude were a dimension the LUT would be even larger. While repetitive control was shown to be able to achieve pitch and roll authority after a handful of flapping cycles, learning time may result in reduced control bandwidth. If repetitive control is done in offline training phase, full-cycle length lookup tables for pitch, roll/altitude, and yaw dimensions would similarly result in enormous arrays, eating into the full RAM allotment of the NRF52840, STM32, and AVR core microcontrollers utilized in this work. Additionally, the voltage dividers used for voltage measurement feedback are wasteful with power, high-impedance, and by necessity for weight significantly exceed the voltage ratings of SMD resistors as discussed in appendix B. Therefore motional feedback is suggested to improve waveform generation under voltage feedback, and possibly even replace the need for voltage feedback by achieving waveform generation suitable for flight control without an inner control loop for voltage regulation.

Motional feedback is also utilized for high-fidelity audio amplification, where audio fidelity is improved by feedback control of the actual motions of the speaker cone which result from the amplifier's output signal, as opposed to the voltage delivered to the coils [40] - thereby circumventing audio output distortion caused by the action of unknown transfer functions of realistic speaker systems. FIR waveform generation has the same need for high fidelity.

It was found in [11] that the Robofly exhibits little coupling between roll, pitch, and yaw. This is advantageous for a linear controller, and a feature which should be unchanged for this waveform generator as well.

LUT generation

It was postulated in [27] that PEU generated waveforms demonstrating modulation of amplitude, DC offset, and harmonic content (for differential slew rate) are sufficient to achieve flight control. Therefore this work proceeds to generate and demonstrate such modulation. The primary goals of this are to address the muxed PEU's need for a different waveform basis than the standard on Eq. 1.1 to achieve yaw torque, the desire to generate initial LUT without or before repetitive control, and the goal to implement direct motional feedback acting on the PFM LUTs instead of reference voltage waveforms. As with many processes, small nonlinearities in the PEU transfer function relating PFM input to voltage output could cause undesired coupling between amplitude, roll, pitch, yaw in this case, but are assumed to be 'fixable' by feedback control and can be tuned.

The PFM generation employed is intended to result in the "flattened sine' found by [27] to be a more energy efficient way to flap because since the extreme voltages contribute less to the lift while costing more energetically, so the waveform can be 'flattened' a bit to reduce the voltage extremes. The 'flattened' sine is achieved by this control signal, as can be seen in the nominal waveforms and the oscilloscope measurements.

For the switching mode charge and discharge converter stages of the muxed PEU, the total δV from commencement to completion of a charging or discharging phase is proportional to square root of the total energy transferred, which is in turn proportional to the area of the PFM LUT over the charge/discharge phase. For our purposes this just means that we should preserve LUT total area when changing the PFM control signal for yaw torques. For yaw control, total thrust, DC offset, and amplitude of the resultant output voltage should be unchanged throughout yaw commands. This is easy to see graphically by examining the PFM signals of Fig. 5.17, but is described mathematically below.

To ensure that a yaw modulation (increasing powers of the $\sin^2$ base) is an area preserving transformation, the yawing LUT columns are just re-scaled after the initial generation according to normalization factor Q computed as the ratio of the definite integrals of the yawing PFM LUT over the nominal LUT PFM coeffs. This normalization is expressed below in Eq. 5.6 and is computed in a for-loop offline during LUT generation. Applying this re-normalization was observed in simulation to restore the peak-to-peak amplitude and DC offset of the output waveforms through range of yawing actuation.

Nonlinearity of the charging and discharging conversion stages over PFM inputs is expected at the higher output power of the PEU, this is likely due to variation of efficiency and the loaded supply voltage (uncompensated yet) over the output voltage range. Thus, some undesired coupling between the roll, pitch, and yaw actions remains, but in simulation plots below this appears to be small enough to be corrected by feedback control. This is similar to the unintended coupling between longitudinal force actuation and roll/pitch observed in [12].

$$Q = \frac{\sum_0^{k=N_h} \sin\left(\frac{k}{N_h}\pi\right)^2 = \frac{1}{2}N_h}{\sum_0^{k=N_h} LUT[k,m]} \tag{5.6}$$

- $N_h = \frac{N_{wfg}}{2}$ is the number of discrete iterations of the wfg controller in one half of the complete flapping cycle.
- m is the column index of the LUT $\in [-(N_{yawCols}-1)/2,\ N_{yawCols}-1)/2]$ during lookup table generation, corresponding to the discretized range of negative to positive yaw control signal.
- k is the index of the controller used to select the k^{th} row of the LUT.

- $LUT[k,m]$ is a single column of the yawing lookup table, computed in Eq. 5.7

- $N_{yawCols}$ is the integer number of distinct columns to negative and positive of zero yaw torque command. For more resolution, this value is increased.

In words, the numerator represents the area of the first π radians of the function $\sin(k)^2$. The factor Q simply gets applied to the entire LUT column after it is computed, to re-scale it so that it has the same area as the original (zero yaw) nominal base. This basis function, $\sin(\cdot)^2$, is based on that the pulsing frequency $f \propto$ duty-cycle DC for constant pulse period. The numerator of Eq. 5.6 simplifies analytically to be simply half of number of waveform generator iterations per half flapping cycle, or $\frac{1}{2}\frac{N_{wfg}}{2}$.

This normalization factor can be scaled proportionally if the pulse width changes.

Voltage measurements at the extrema of the output waveforms can be used to perform low frequency adaptation of the PFM amplitude constants. Thus, if the output is too low in voltage after a discharging half-cycle then the corresponding PFM feedforward amplitutde factor f_{DC} is reduced proportionally to the error. This is similar to the repetitive control algorithm, except that a single term is adjusted based on a single measurement as opposed to the entire lookup table adjusted online by much higher frequency feedback $\hat{v}$.

PFM coefficients for precomputed cycle LUT

Yawing torque is applied by differential wing stroke speeds resulting in a net moment. This is easily achieved by the addition of a second harmonic to the output voltage at twice the flapping frequency and with amplitude set by the yaw control signal, as expressed in Eq. 1.1. For the muxed PEU in chapter.4, it is convenient to enforce equal charge and discharge times, with deadtime to ensure break-before-make mux switching. Thus for this implementation a different equation described in Eq. 5.7 was used to result in differential stroke speeds, using powers of the sinusoid instead of second harmonic content. To avoid computational expense of computing of complicated expressions (especially containing trigonometric functions or sqrt) online, Eq. 5.7 was used to precompute a look-up table for the full range of expected yaw rate error signal. The operation is described below in Eq. 5.7.

$$LUT[k,m] = (\sin(\pi(k - N_h)/N_h))\ \cos(((k - \phi_k)/N_h)\pi)^m \tag{5.7}$$

The pulsetrain formulation in Eq. 5.7 has a structure based on the flattened sine reference innovation of [27], the energy stored on the capacitor, and the varying voltage across the discharging inductor, the time derivative of the reference voltage, terms to smooth the start and stop of the pulse-train as well as incorporate the 'flattened sine'.

Online Flight control operations on PFM LUT

When the robot experiences yaw rate error during flight, the appropriate column of a lookup table such as found by repetitive control, generated by Eq. 5.7, or generated by model inversion is simply indexed by interpolation. Roll rate, and pitch rate errors are scaled by control gains and multiplied into the k^{th} element of the indexed LUT column.

Once the LUT is in memory, it costs very little to merely index the columns m for yaw and apply scaling and offset terms for roll and pitch. This online operation is what is used during flight, at every half-cycle or every full cycle. This is described in Eq. 5.8

$$LUT[k,m] = (f + \omega_x k_{rx} + \omega_y k_{px}) LUT_1[k,m] \tag{5.8}$$

5.6.1 Control authority demonstration for roll, pitch, and yawing waveforms

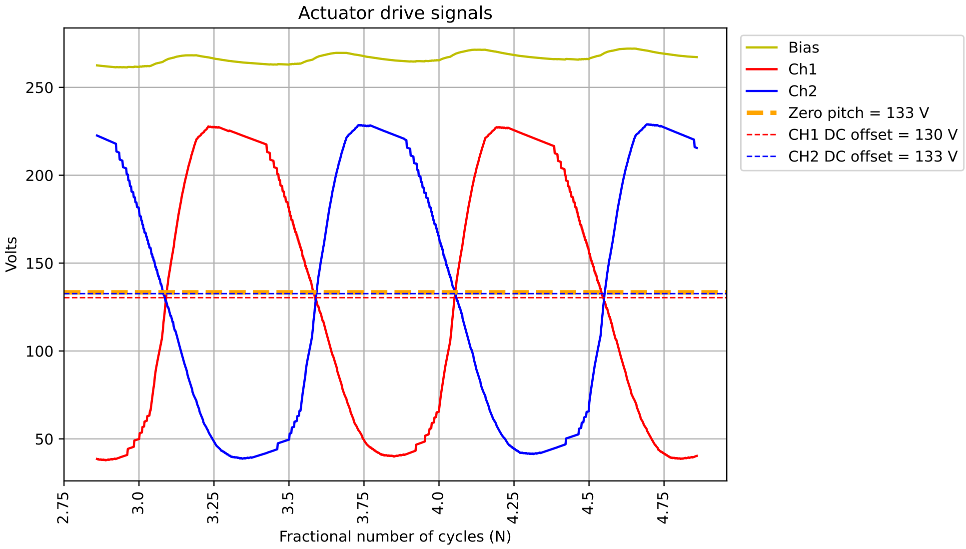

Figure 5.14: An example of the new muxed PEU driving two independent channels to achieve yaw in Spice simulation.

Roll

Roll authority of the driver is achieved by varying the amplitude of one wing relative to the other. How the driver achieves this is shown in Fig. 5.15. The actuator signals necessary to achieve this are shown in the bottom plots, and represent distinct transformations of the control signal, thereby minimizing coupled pitch, total thrust, and yaw.

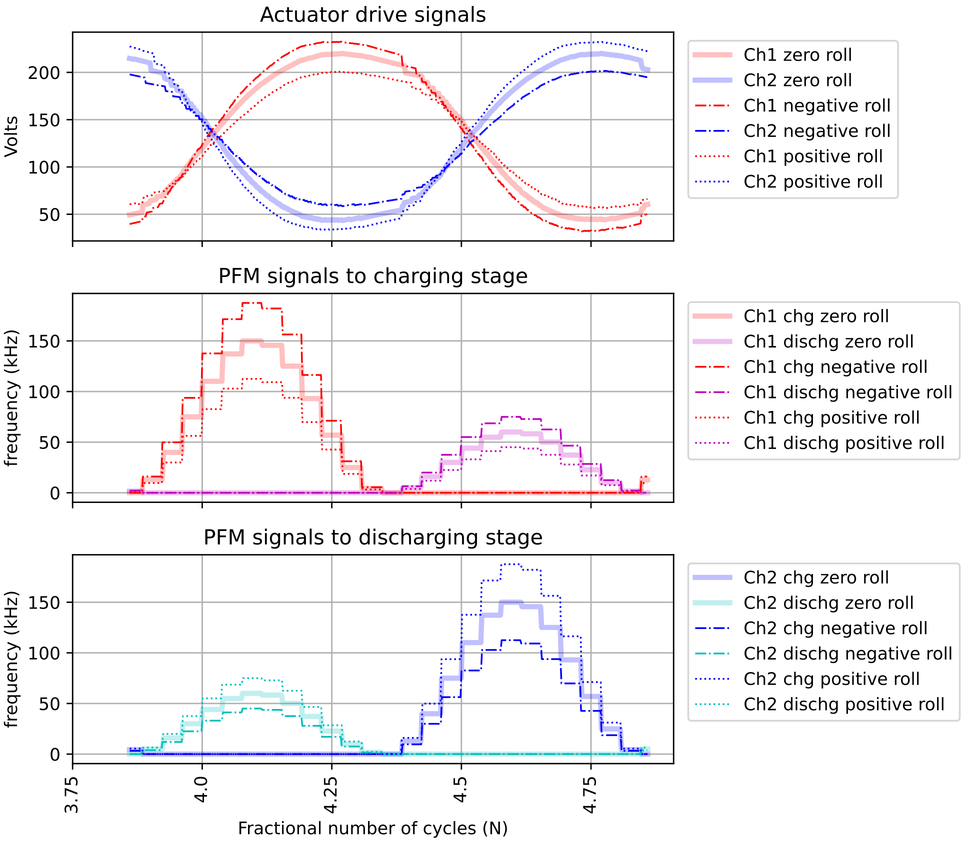

Figure 5.15: An example showing control signal and resultant waveform of the new muxed PEU driving two independent channels to achieve roll in Spice simulation. PFM signals for the charging and discharging stages are generated using eqs. 5.7, 5.8.

Pitch

Pitch authority of the driver is demonstrated by the ability to execute variable amplitude diference between left and right wing, indepentently. This is shown for separate pitch commands of negative, neutral, and positive pitch torque in Fig. 5.16. The actuator signals necessary to achieve this are shown in the bottom plots, and represent distinct transformations of the control signal, thereby minimizing coupled roll, total thrust, and yaw.

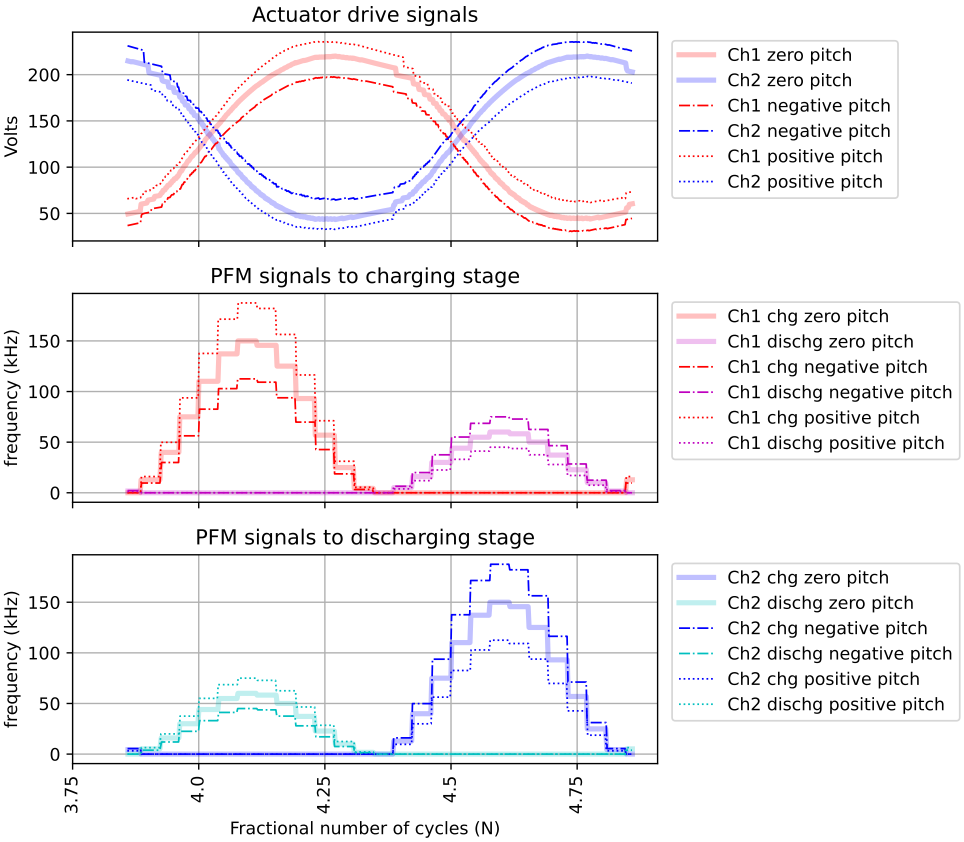

Figure 5.16: An example showing control signal and resultant waveform of the new muxed PEU driving two independent channels to achieve pitch in Spice simulation. PFM signals for the charging and discharging stages are generated using eqs. 5.7, 5.8.

Yaw

Yaw authority of the driver is not required for the minimal attitude stabilization of roll and pitch, but is demonstrated by the ability to execute variable differential $\frac{dV}{dt}$ of left and right wing simultaneously. Simply put, the waveforms should each rise and fall in voltage at different rates, thus creating a yaw moment due to differential wing drag during flapping. This is shown for separate yaw commands of negative, neutral, and positive yaw torque in Fig. 5.17. The actuator signals necessary to achieve this are shown in the bottom plots, and represent distinct transformations of the control signal, thereby minimizing coupled roll, total thrust, and pitch.

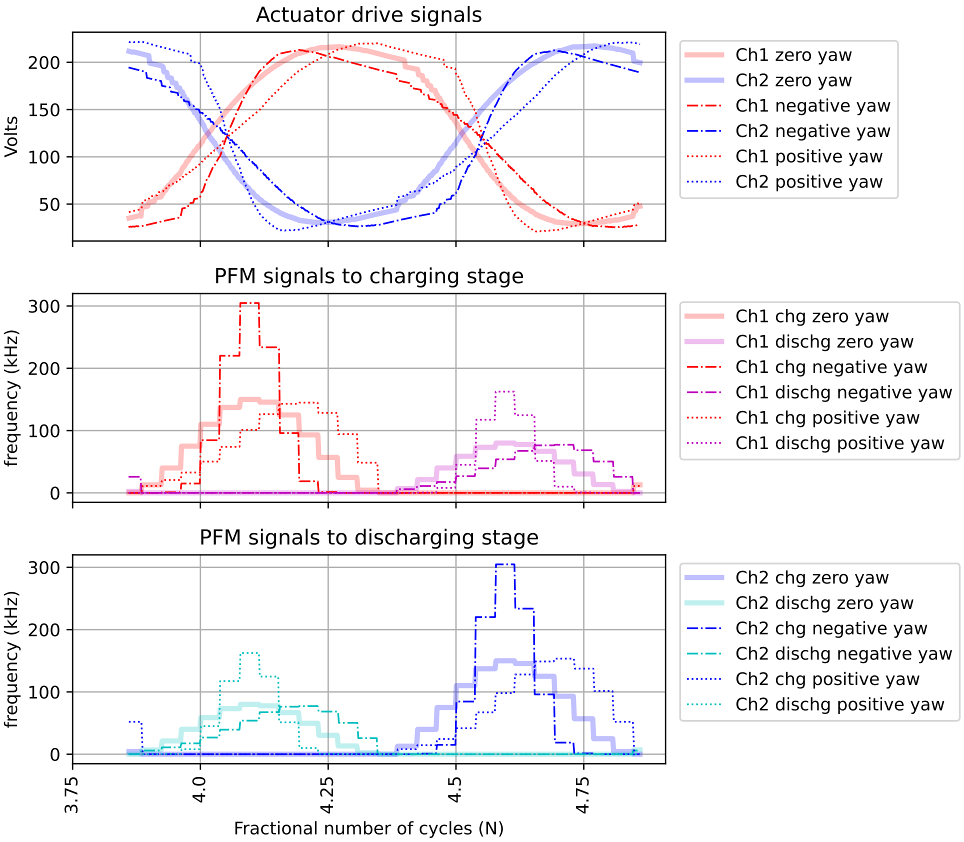

Figure 5.17: Yawing waveforms generated in the high-fidelity Spice simulation. PFM signals for the charging and discharging stages are generated using eqs. 5.7, 5.8.

Lateral force

An example of a non-equal charge/discharge implementation of a signals resulting in differential stroke-speed to generate net lateral force is shown below. Notably, using the muxed PEU this would either require different actuator orientation on the back-to-back Robofly. A degree of undesired coupling to DC offset is observed due to this particular PFM generation scheme's failure to account for differences in authority between charge and discharge converters. The *LUT* in this case is just generated by a model-free equation simpler than Eq. 5.7 jsut for the purposes of this demonstration, but can easily be improved by better modeling or adaptive control to reduce the coupling to amplitude and DC offset in the event of changed actuator configuration in the robot.

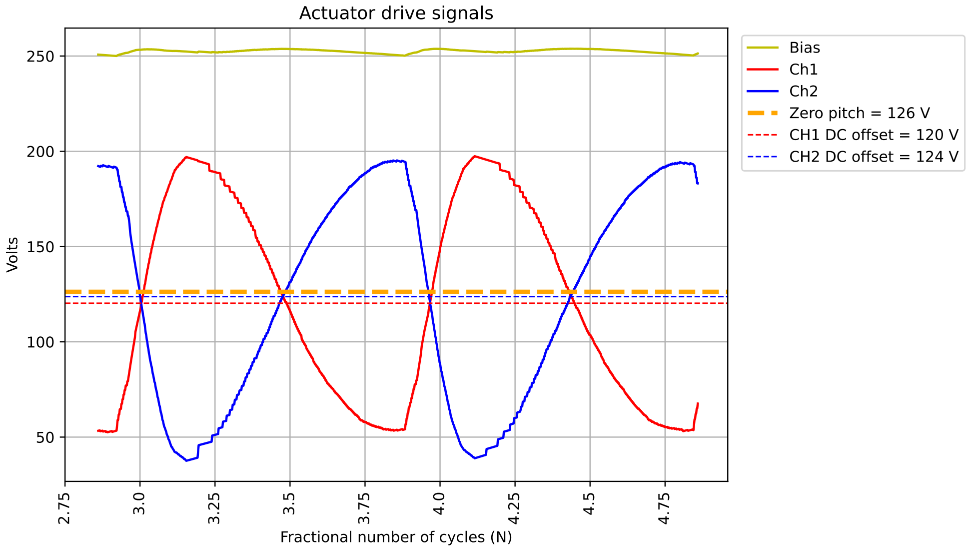

Figure 5.18: An example of the muxed PEU generating two independent waveforms for net longitudinal force in Spice simulation.

Combined actuation of roll, pitch, and yaw

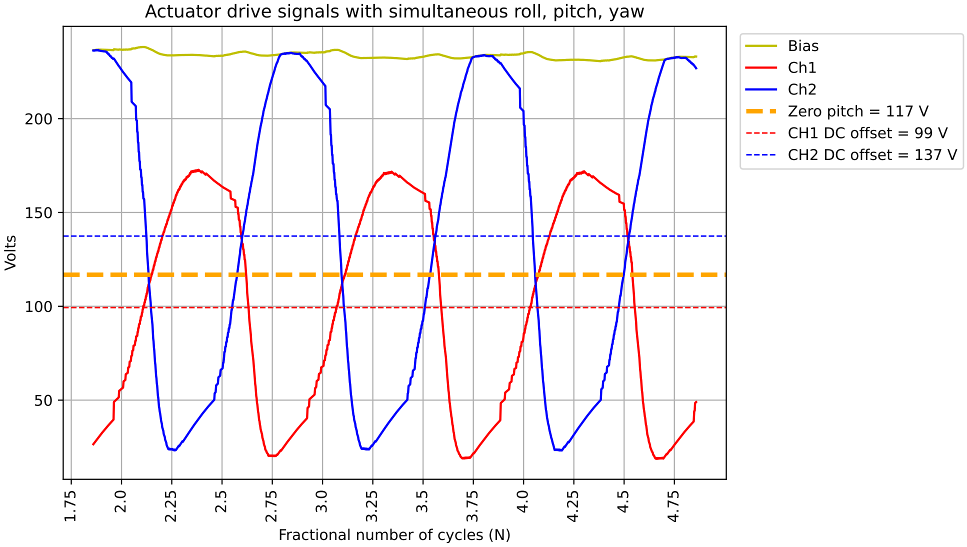

Figure 5.19: An example of the muxed PEU design generating two independent waveforms for combined roll, negative pitch, and yaw flight control signals.

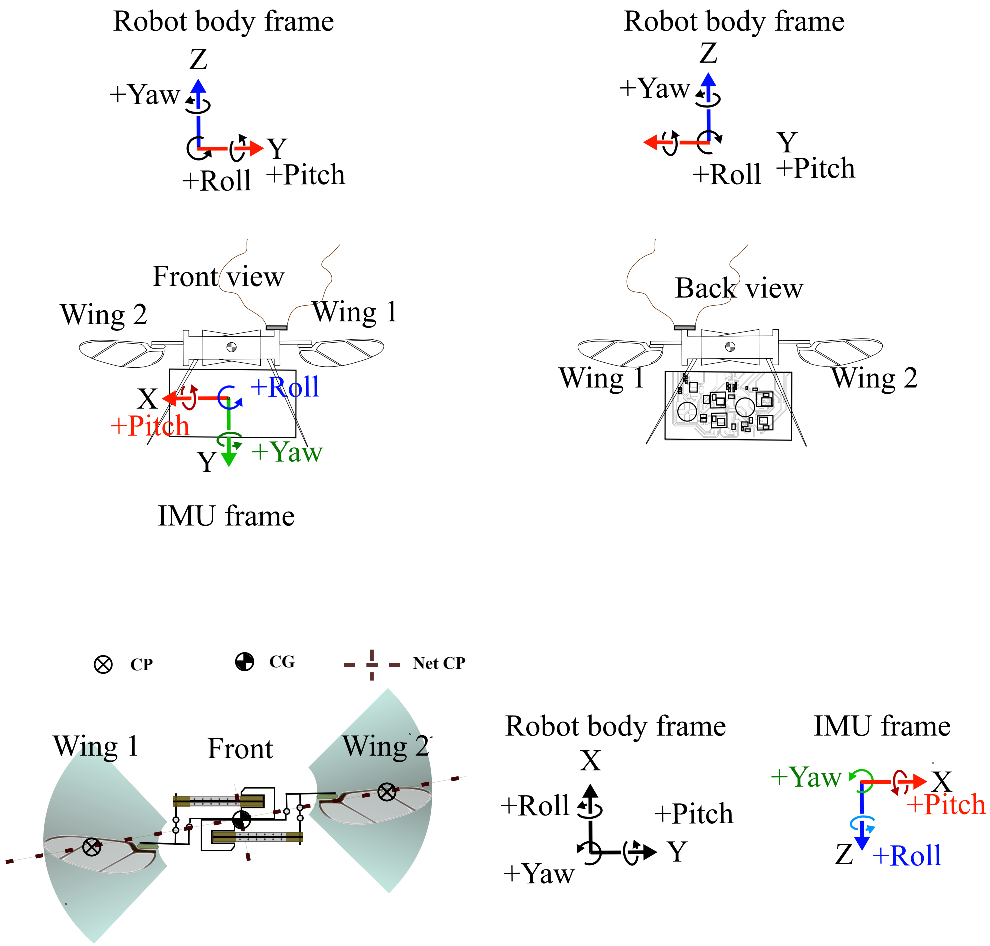

Figure 5.20: A full diagram of coordinate frames, control gestures, physical layout, and actuators for the 2-wing RoboFly.

5.7 *Implementation of Motional Feedback LUT Control*

5.7.1 IMU placement and roll, pitch, and yaw

The IMU was oriented on the robot as close to the CG as possible, as indicated on Fig. 5.20. The views of the robot aid in determining the correct coordinate transformation from IMU to robot body frame. The top-down view is used by inspection to determine the correct control gains in response to angular rate, refering to the drawn actuator orientation with PZT top layer (bias) positioned furthest from the airframe centerline.

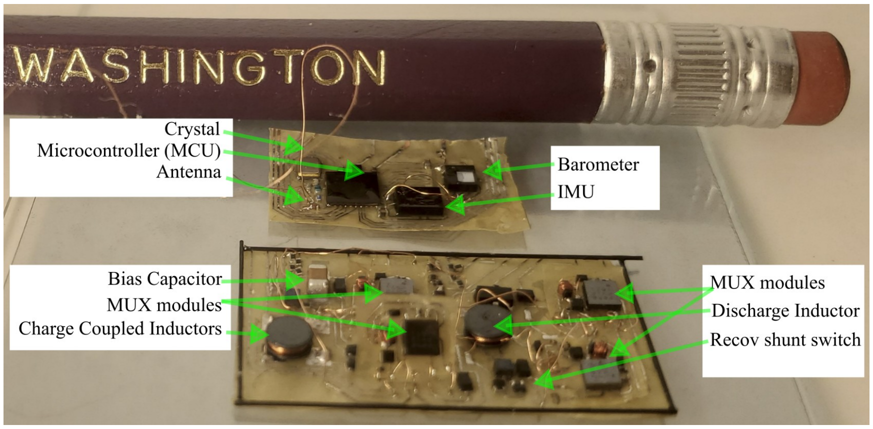

Figure 5.21: An instantiation of the muxed PEU from chap.4 using 4 LFBD modules in the foreground, and in the background is a microcontroller circuit with IMU, barometer, voltage regulator, and NRF52840 with RF antenna, next to a pencil for scale.

Hardware demonstration

Fig. 5.21 shows the hardware instantiation used in the tests of the muxed PEU from chap.4 and a microcontroller implementing the motional feedback control attitude stabilizer described above in sec.5.6.

Figure 5.23 shows an initial static displacement test of the robot PEU charging stage, proving that the PEU can charge the actuators to full voltage as shown by the wing displacement. Robot damage prevented lift greater than weight.

Despite the cracked actuators preventing lift greater than weight, the muxed PEU was able to flap the wings and generate the predicted output signals to the piezo actuators. This is significant because it shows the design has actually matured through the fraught initial stages of a new idea.

During the experiment in Fig. 5.22, the robot experienced four cycles of rolling to approximately ±15°, and flapped 90 times at 180 Hz. Telemetry was transmitted via wireless communication protocol “Enhanced Shockburst” (ESB) to a receiver unit connected by serial to the host PC. A total of 7-Bytes of data was transmitted once per flapping cycle, consisting of the current flapping cycle counter, and the gyro $\omega_x, \omega_y, \omega_z$. Notably, the

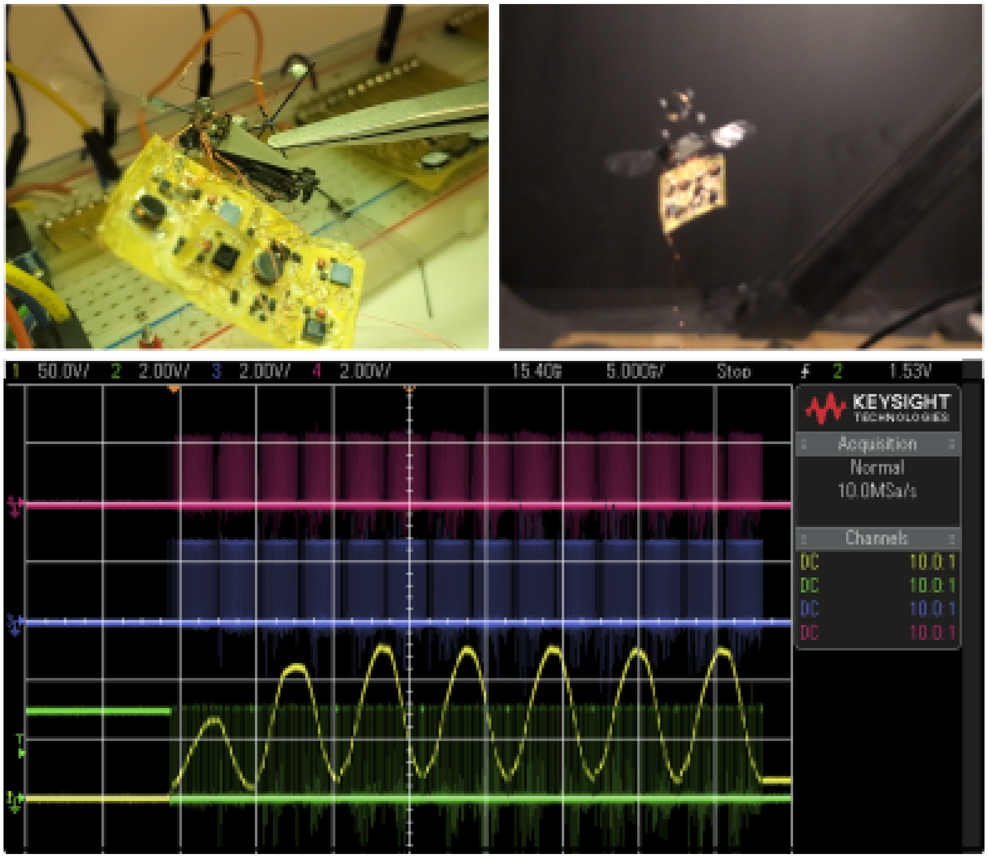

Figure 5.22: (Top left) Left and right wing motion is shown by the blurred wing outlines while the muxed PEU drives both wings independently. Actuator and circuit damage required limiting of the PEU output below the voltage rating of the components. (Top Right) Unsuccessful flight demonstration in which the onboard gyro is used to modulate the output voltages and flap the wings independently, but due to damage was insufficient for lift greater than the weight of the laden robot. (Bottom) Oscilloscope capture of the output voltage to ch1 actuator is plotted in yellow as it starts from rest. Bias and ch2 are not measured. The magenta and blue traces are the pulse trains to the discharge and charge stages (respectively) of the PEU. The output voltage of each stage is muxed to both wing actuators separately during the first and then second half wing strokes. The green trace is a digital output which is set and then cleared at the start and stop of the 5 kHz controller ISR, plotted to verify the timing and successful modulation of the pulse frequencies.

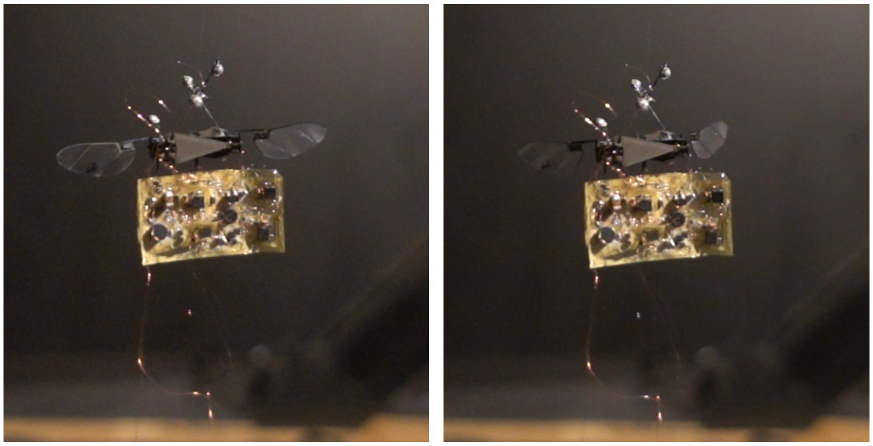

Figure 5.23: Despite cracked actuators unable to flap with full amplitude at flying frequency, earlier video and the still frames from a step input test of the charging stage show the wings moving as expected to their full stroke amplitude, from zero stroke angle with $V_b = V_{sig} = 5$V at rest (left), to maximum static displacement when bias was charged to the full voltage rating of the bias capacitor, $V_b = 250$ V.

flapping index incremented reliably during this data collection, proving the microcontroller and RF link was operating sufficient to send ESB datagrams uninterupted at 180 Hz.

5.8 Conclusions

It was already demonstrated in [20] that an onboard rate gyro can stabilize the attitude of an FIR, by effectively adding damping to the closed-loop dynamics of the robot by using the same control torque inputs resulting from variable DC offset and relative amplitude of the two wings. It has been shown in [46] that such modulation of the unipolar sinusoidal flapping waveforms work to control the robot. In this work the contribution is to show the missing link between those two results, which is the ability of a PEU to generate waveforms which result in the DC offset and differential amplitude originating those roll and pitch torques.

Feedforward terms over the discretized flapping cycle are stored in lookup tables which are then scaled linearly for roll and pitch control dimensions, but only a single lookup table for yaw is needed, thus reducing the LUT memory requirements. Linear transformations of a baseline library of LUTs provide a strategy for angular velocity control, but can readily be extended to full flight control because the yaw/lateral force indexing and linear operations modulating the LUT map directly to the modulation methods used in roll, pitch, yaw/lateral force controllers in the standard waveform generator equation of Eq. 1.1.

Suggestions for further development are included in chap.6.

Chapter 6 Conclusions

This work has presented solutions for key outstanding problems in wireless power, control, and actuation autonomy. This work has resulted in several milestone results for the field and notable publications such as [30] which was was the first ever wireless lift-off of an FIR actuated soley by by onboard systems, and subsequent work [31] has shown an additional wireless power solution, while in [29] I was the first to actually demonstrate in hardware that the onboard PEU design can indeed generate waveforms of sufficient quality to linearly modulate the thrust required for roll torque and altitude control. In the same work I also demonstrated that an unaided repetitive control law can learn an unknown periodic control signal within the span of a few wing strokes, thus providing a valuable feedforward term to any waveform generator control used since and going forward.

Still, the primary aim of this work is to advance the technology needed for wireless operation of the FIRs which faces a major bottleneck in power autonomy. Control, sensing, and aerodynamics research problems have attracted a great deal of quality research effort but these endeavors now greatly benefit from this work to cut the cord. This work presents an exploration of state of the art solutions, and presents new power electronics and waveform generation control to enable such advancement. Specifically, the waveform generation and power electronics work together to result in a doubling of the effective power capacity of our FIRs by improving end-to-end efficiency over the previous state of the art, from 28% to $\approx$ 54% with estimated 75% the weight and expected further reduction as the prototype matures.

This new design is ≈ 75 % the weight of previous SoTA, capable of higher voltage, and to our knowledge is the first to demonstrate independent wing actuation by in an onboard PEU. The PEU design methods developed in this work shed light on other possible PEU topologies which are expected to build on these results. These contributions help the field by unlocking applications and experiments which have been precluded by being literally tethered to the desk.

Continuations

This work suggests continued testing of the motional feedback attitude stabilization using proportional control of measured roll rate, to directly modulate the PFM LUTs instead of an inner waveform generator voltage control loop. This simpler approach is in agreement with the author's sense that nature and bioinspired engineering favor simplicity and lightweight feedback, model-free, and adaptive strategies. Initial model based formulations of the PFM LUTs can possibly be static as simulations and early tests indicated, although further experimentation should verify this. However, to reduce saturation or potential coupling of roll, pitch, or yaw caused by naively generated PFM sequences such as in Fig. 5.19, it is recommended to add a low gain repetitive control element.

The next step beyond roll-rate control is to control robot attitude and then position in $[x, y, z, \Psi]$. A simple complementary filter or Kalman filter is recommended for state estimation given the high frequency vibration observed in FIR flight measurements which would impact inclination estimates relying on the accelerometers. Altitude estimation and control z can be improved using the 6.5mg barometer IC visible on the instantiated microcontroller board in Fig. 5.21.

Other topologies made possible by the design methods and highside switches based on pulse transformers should be explored. In particular, the heavy bias capacitor is undesired. There are options which could possibly avoid the bias capacitor, and implement the 'bridged bipolar' drive scheme presented in [52], also similar to the alternating drive topologies suggested in literature. Besides just weight savings of nearly 17 mg, separating the top electrodes of independent actuators would be *very* desirable for waveform generation. A realistically sized bias capacitor is only between $3 - 10\times$ the actuator capacitance and MLCC capacitor technology is likely too matured to yield substantial improvements, and so the actuators are coupled in undesirable ways which may could complicate flight control down the road. Possible configurations and control strategies to compensate or avoid this potential problem should be kept in mind.

Practically speaking, revisions to the muxed PEU should better route traces, reduce unpopulated pcb area, and integrate the highside switches instead of attaching separate modules. While useful for testing, manufacture, and occasional replacement after accidents, the added weight and variability of overall CG given modular muxes is significant, and the jumper wires are difficult and risk incorrect connections. While this work resorted to jumper wires where necessary on a single-layer flex PCB, '1.5' layer boards with vias are possible using narrow strips of flex pcb made at the same time as the circuit and then placed on the underside of a single layer board, which could reduce the failure rate of jumper solder joints and damage/tangling/mistaken connections made during assembly. Additionally, it may be possible to achieve mechanical lowpass filtering of the onboard IMU if the flex-PCB is cut so as to achieve a prescribed stiffness and form a mass-spring LPF with the mass of the IMU.

Even with the same muxed PEU topology, reduction of operating voltage to < 200 V would allow the substitution of lighter EPC2012 GaNFET for the 350 V EPC2050 used here, which would further improve the basic trade made for less inductor weight in exchange for highside switches. Switching to the lighter GaNFET and reducing the weight of mux modules and extra solder/flux, and jumpers is expected to reduce the weight of the muxed PEU to $\approx 115 - 120$ mg.

Multiple forks of waveform generator control strategies are open. The most promising direction for waveform generator control in the authors opinion is proportional + repetitive control starting from either a) a model-based feedforward to provide a precomputed lookup table $LUT[k, m]$ or b) the PFM functions developed in chap. 5. Either way, the dimensions of the precomputed LUT are the discrete time cycle index k and index m obtained by interpolating yaw torque to the appropriate column m. Reduced dimensionality of the LUTs is desirable for memory requirements and should be valid given the linearity and uncoupled roll, pitch, and amplitude terms of the waveform generator. Online, LUT should be adapted with low gain using the repetitive control algorithm developed in this work.

Additionally, the dynamics of the repetitive control signal provides information about the load such as the phase angle between the real and reactive components of the actuator-wing system, providing interesting avenues of system identification.

This work indicates that actuator design and choice of operating voltages and frequencies affects the dielectric losses of the actuators which can significantly increase power requirements. The ratio of dielectric loss to ‘useful’ aerodynamic damping may provide a optimization path for actuator frequency, applied voltage, and wing properties.

BIBLIOGRAPHY

[1] Ieee standard on piezoelectricity. *ANSI/IEEE Std 176-1987*, pages $0_1 - -$, 1988.

[2] Palak Bhushan and Claire J Tomlin. Milligram-scale micro aerial vehicle design for low-voltage operation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–9. IEEE, 2018.

[3] Rodney A. Brooks and Anita M. Flynn. Fast, cheap, and out of control: A robot invasion of the solar system. *Journal of the British Interplanetary Society*, 42:478–485, 1989.

[4] Simon Chaput, Martin Renaud, Rémi Meingan, and Jean-François Pratte. A 3.7 v to 200 v highly integrated dc-dc converter with 70.4% efficiency for portable electrostatic mems applications. In *2014 IEEE 12th International New Circuits and Systems Conference (NEWCAS)*, pages 381–384. IEEE, 2014.

[5] Yufeng Chen, Hongqiang Wang, E Farrell Helbling, Noah T Jafferis, Raphael Zufferey, Aaron Ong, Kevin Ma, Nicholas Gravish, Pakpong Chirarattananon, Mirko Kovac, et al. A biologically inspired, flapping-wing, hybrid aerial-aquatic microrobot. *Science Robotics*, 2(11), 2017.

[6] Yufeng Chen, Huichan Zhao, Jie Mao, Pakpong Chirarattananon, E Farrell Helbling, Nakseung Patrick Hyun, David R Clarke, and Robert J Wood. Controlled flight of a microrobot powered by soft artificial muscles. *Nature*, 575(7782):324–329, 2019.

[7] Andreas Christ, Mark G Douglas, John M Roman, Emily B Cooper, Alanson P Sample, Benjamin H Waters, Joshua R Smith, and Niels Kuster. Evaluation of wireless resonant power transfer systems with human electromagnetic exposure limits. *IEEE Transactions on Electromagnetic compatibility*, 55(2):265–274, 2012.

[8] Yogesh M Chukewad and Sawyer Fuller. Yaw control of a hovering flapping-wing aerial vehicle with a passive wing hinge. *IEEE Robotics and Automation Letters*, 6(2):1864–1871, 2021.

[9] Yogesh M Chukewad, Johannes James, Avinash Singh, and Sawyer Fuller. Robofly: An insect-sized robot with simplified fabrication that is capable of flight, ground, and water surface locomotion. *IEEE Transactions on Robotics*, 37(6):2025–2040, 2021.

[10] Yogesh M Chukewad, Avinash T Singh, Johannes M James, and Sawyer B Fuller. A new robot fly design that is easier to fabricate and capable of flight and ground locomotion. In

2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4875–4882. IEEE, 2018.

[11] Daksh Dhingra, Yogesh M Chukewad, and Sawyer B Fuller. A device for rapid, automated trimming of insect-sized flying robots. *IEEE Robotics and Automation Letters*, 5(2):1373–1380, 2020.

[12] Daksh Dhingra, Kadierdan Kaheman, and Sawyer B Fuller. Modeling and lqr control of insect sized flapping wing robot. *arXiv preprint arXiv:2406.20061*, 2024.

[13] R Dudley and Charles P Ellington. Mechanics of forward flight in bumblebees: I. kinematics and morphology. *Journal of Experimental Biology*, 148(1):19–52, 1990.

[14] RONALD S FEARING. Toward micromechanical flyers. In *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2001 NAE Symposium on Frontiers of Engineering*, page 21. National Academies Press, 2002.

[15] R.S. Fearing, K.H. Chiang, M.H. Dickinson, DL Pick, M. Sitti, and J. Yan. Wing transmission for a micromechanical flying insect. In *Robotics and Automation (ICRA), 2000 IEEE Int. Conf.*, volume 2, pages 1509–1516. IEEE, 2000.

[16] Ben M. Finio, Nestor O. Perez-Arancibia, and Robert J. Wood. System identification and linear time-invariant modeling of an insect-sized flapping-wing micro air vehicle. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ Int. Conf.*, pages 1107–1114, September 2011.

[17] Benjamin M Finio, Kevin C Galloway, and Robert J Wood. An ultra-high precision, high bandwidth torque sensor for microrobotics applications. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 31–38. IEEE, 2011.

[18] Sawyer B Fuller. Four wings: An insect-sized aerial robot with steering ability and payload capacity for autonomy. *IEEE Robotics and Automation Letters*, 4(2):570–577, 2019.

[19] Sawyer B Fuller, Michael Karpelson, Andrea Censi, Kevin Y Ma, and Robert J Wood. Controlling free flight of a robotic fly using an onboard vision sensor inspired by insect ocelli. *Journal of The Royal Society Interface*, 11(97):20140281, 2014.

[20] SB Fuller, EF Helbling, P Chirarattananon, and RJ Wood. Using a mems gyroscope to stabilize the attitude of a fly-sized hovering robot imav 2014: Int. In *Micro Air Vehicle Conf. and Competition*, 2014.

[21] Benjamin Goldberg, Raphael Zufferey, Neel Doshi, Elizabeth Farrell Helbling, Griffin Whittredge, Mirko Kovac, and Robert J Wood. Power and control autonomy for high-speed locomotion with an insect-scale legged robot. *IEEE Robotics and Automation Letters*, 3(2):987–993, 2018.

[22] MA Graule, P Chirarattananon, SB Fuller, NT Jafferis, KY Ma, M Spenko, R Kornbluh, and RJ Wood. Perching and takeoff of a robotic insect on overhangs using switchable electrostatic adhesion. *Science*, 352(6288):978–982, 2016.

[23] William G Hurley and Werner H Wölfle. *Transformers and inductors for power electronics: theory, design and applications.* John Wiley & Sons, 2013.

[24] Seiko Instruments Inc. Low esr, low leak current chip-type edlc cpx3225a752d, 2016.

[25] Vikram Iyer, Elyas Bayati, Rajalakshmi Nandakumar, Arka Majumdar, and Shyamnath Gollakota. Charging a smartphone across a room using lasers. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(4):143, 2018.

[26] Noah T Jafferis, Moritz A Graule, and Robert J Wood. Non-linear resonance modeling and system design improvements for underactuated flapping-wing vehicles. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3234–3241. IEEE, 2016.

[27] Noah T Jafferis, E Farrell Helbling, Michael Karpelson, and Robert J Wood. Untethered flight of an insect-sized flapping-wing microscale aerial vehicle. *Nature*, 570(7762):491–495, 2019.

[28] Noah T Jafferis, E Farrell Helbling, Michael Karpelson, and Robert J Wood. Untethered flight of an insect-sized flapping-wing microscale aerial vehicle. *Nature*, 570(7762):491–495, 2019.

[29] Johannes James and Sawyer B Fuller. A high-voltage power electronics unit for flying insect robots that can modulate wing thrust. In *2021 IEEE International Conference on Robotics and Automation (ICRA).* IEEE, 2021 [accepted].

[30] Johannes James, Vikram Iyer, Yogesh Chukewad, Shyamnath Gollakota, and Sawyer B Fuller. Liftoff of a 190 mg laser-powered aerial vehicle: The lightest wireless robot to fly. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018.

[31] Johannes James, Xingyi Shi, Joshua R. Smith, and Sawyer B. Fuller. Magnetically coupled resonators for wireless power transmission to insect sized flapping wing robots. In *Hilton Head Workshop 2024: A Solid-State, Sensors, Actuators, and Microsystems Workshop*, Hilton Head, SC, June 2024. ACM Press.

[32] Kaushik Jayaram, Noah T Jafferis, Neel Doshi, Ben Goldberg, and Robert J Wood. Concomitant sensing and actuation for piezoelectric microrobots. *Smart Materials and Structures*, 27(6):065028, 2018.

[33] Kaushik Jayaram, Jennifer Shum, Samantha Castellanos, E Farrell Helbling, and Robert J Wood. Scaling down an insect-size microrobot, hamr-vi into hamr-jr. *arXiv preprint arXiv:2003.03337*, 2020.

[34] Mattan Kamon, Michael J Tsuk, and Jacob White. Fasthenry: A multipole-accelerated 3-d inductance extraction program. In *Proceedings of the 30th international design automation conference*, pages 678–683, 1993.

[35] Michael Karpelson, Benjamin H Waters, Benjamin Goldberg, Brody Mahoney, Onur Ozcan, Andrew Baisch, Pierre-Marie Meyitang, Joshua R Smith, and Robert J Wood. A wirelessly powered, biologically inspired ambulatory microrobot. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2384–2391. IEEE, 2014.

[36] Michael Karpelson, Benjamin H Waters, Benjamin Goldberg, Brody Mahoney, Onur Ozcan, Andrew Baisch, Pierre-Marie Meyitang, Joshua R Smith, and Robert J Wood. A wirelessly powered, biologically inspired ambulatory microrobot. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2384–2391. IEEE, 2014.

[37] Michael Karpelson, Gu-Yeon Wei, and Robert J Wood. Milligram-scale high-voltage power electronics for piezoelectric microrobots. In *2009 IEEE international conference on robotics and automation*, pages 2217–2224. IEEE, 2009.

[38] Michael Karpelson, Gu-Yeon Wei, and Robert J Wood. Driving high voltage piezoelectric actuators in microrobotic applications. *Sensors and actuators A: Physical*, 176:78–89, 2012.

[39] Michael Karpelson, John P Whitney, Gu-Yeon Wei, and Robert J Wood. Energetics of flapping-wing robotic insects: towards autonomous hovering flight. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1630–1637. IEEE, 2010.

[40] JA Klaassen, SH de Koning, et al. Motional feedback with loudspeakers. *Philips Tech. Rev*, 29(5):148–157, 1968.

[41] Andre Kurs, Aristeidis Karalis, Robert Moffatt, John D Joannopoulos, Peter Fisher, and Marin Soljacic. Wireless power transfer via strongly coupled magnetic resonances. *science*, 317(5834):83–86, 2007.

[42] Younghoon Lee, Zhijian Ren, Yi-Hsuan Hsiao, Suhan Kim, Won Jun Song, Chengkuo Lee, and Yufeng Chen. Liftoff of a soft-actuated micro-aerial-robot powered by triboelectric nanogenerators. *Nano Energy*, 126:109602, 2024.

[43] Mario Lok, Elizabeth Farrell Helbling, Xuan Zhang, Robert Wood, David Brooks, and Gu-Yeon Wei. A low mass power electronics unit to drive piezoelectric actuators for flying microrobots. *IEEE Transactions on Power Electronics*, 33(4):3180–3191, 2017.

[44] Mario Lok, Xuan Zhang, Elizabeth Farrell Helbling, Robert Wood, David Brooks, and Gu-Yeon Wei. A power electronics unit to drive piezoelectric actuators for flying microrobots. In *2015 IEEE Custom Integrated Circuits Conference (CICC)*, pages 1–4. IEEE, 2015.

[45] H Lord. Pulse transformers. *IEEE Transactions on Magnetics*, 7(1):17–28, 1971.

[46] Kevin Y Ma, Pakpong Chirarattananon, Sawyer B Fuller, and Robert J Wood. Controlled flight of a biologically inspired, insect-scale robot. *Science*, 340(6132):603–607, 2013.

[47] Kevin Y Ma, Samuel M Felton, and Robert J Wood. Design, fabrication, and modeling of the split actuator microrobotic bee. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 1133–1140. IEEE, 2012.

[48] Mico Perales, Mei-huan Yang, Cheng-liang Wu, Chin-wei Hsu, Wei-sheng Chao, Kun-hsien Chen, and Terry Zahuranec. Characterization of high performance silicon-based vmj pv cells for laser power transmission applications. In *Proc. of SPIE Vol*, volume 9733, pages 97330U–1, 2016.

[49] Néstor O Pérez-Arancibia, Kevin Y Ma, Kevin C Galloway, Jack D Greenberg, and Robert J Wood. First controlled vertical flight of a biologically inspired microrobot. *Bioinspiration and Biomimetics*, 6(3):036009, September 2011.

[50] Abraham Pressman. *Switching power supply design*. McGraw-Hill, Inc., 1997.

[51] S.B. Fuller research group. Autonomous Insect Robotics Laboratory, 2018.

[52] Shannon A Rios and Andrew J Fleming. A new electrical configuration for improving the range of piezoelectric bimorph benders. *Sensors and Actuators A: Physical*, 224:106–110, 2015.

[53] Alanson P Sample, David T Meyer, and Joshua R Smith. Analysis, experimental results, and range adaptation of magnetically coupled resonators for wireless power transfer. *IEEE Transactions on industrial electronics*, 58(2):544–554, 2010.

[54] S Sherrit, HD Wiederick, and BK Mukherjee. Accurate equivalent circuits for unloaded piezoelectric resonators. In *1997 IEEE Ultrasonics Symposium Proceedings. An International Symposium (Cat. No. 97CH36118)*, volume 2, pages 931–935. IEEE, 1997.

[55] Xingyi Shi and Joshua R Smith. Reconfigurable and adaptive coupled relay resonator platform for a moving receiver. In *2019 International Workshop on Antenna Technology (iWAT)*, pages 182–185. IEEE, 2019.

[56] Vamsi Talla, Bryce Kellogg, Benjamin Ransford, Saman Naderiparizi, Shyamnath Gollakota, and Joshua R Smith. Powering the next billion devices with wi-fi. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, pages 1–13, 2015.

[57] Vamsi Talla, Bryce Kellogg, Benjamin Ransford, Saman Naderiparizi, Shyamnath Gollakota, and Joshua R Smith. Powering the next billion devices with wi-fi. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, pages 1–13, 2015.

[58] W. S. N. Trimmer. Microbots and micromechanical systems. *Sensors and Actuators*, 19:267–287, 1989.

[59] Qing-Ming Wang, Xiao-Hong Du, Baomin Xu, and L Eric Cross. Electromechanical coupling and output efficiency of piezoelectric bending actuators. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 46(3):638–646, 1999.

[60] Benjamin H Waters, Brody J Mahoney, Gunbok Lee, and Joshua R Smith. Optimal coil size ratios for wireless power transfer applications. In *2014 IEEE international symposium on circuits and systems (ISCAS)*, pages 2045–2048. IEEE, 2014.

[61] Benjamin H Waters, Alanson P Sample, and Joshua R Smith. Adaptive impedance matching for magnetically coupled resonators. *Prog Electromagn Res*, pages 694–701, 2012.

[62] John P Whitney, Pratheev S Sreetharan, Kevin Y Ma, and Robert J Wood. Pop-up book mems. *Journal of Micromechanics and Microengineering*, 21(11):115021, 2011.

[63] R. J. Wood. The first takeoff of a biologically inspired at-scale robotic insect. *IEEE Trans. Robotics*, 24(2):341–347, 2008.

[64] Robert J. Wood, Ben Finio, Michael Karpelson, Kevin Ma, Néstor O. Pérez-Arancibia, Pratheev S. Sreetharan, Hiroto Tanaka, and John P. Whitney. Progress on “pico” air vehicles. *Int. J. Robotics Research*, 31(11):1292–1302, 2012.

[65] Robert J Wood, E Steltz, and RS Fearing. Optimal energy density piezoelectric bending actuators. *Sensors and Actuators A: Physical*, 119(2):476–488, 2005.

[66] Yang Zou, Weiping Zhang, and Zheng Zhang. Liftoff of an electromagnetically driven insect-inspired flapping-wing robot. *IEEE Transactions on Robotics*, 32(5):1285–1289, 2016.

Appendix A Vehicle and circuit design and fabrication

This section is to provide a brief summary of just a few of the mechanical design and fabrication efforts required in this research endeavor.

A.1 Airframe

The old transmission was split to clear the winge hinge, weakening the transmission. Additionally, the actuators were covered by the airframe which impeded repair and electrode attachment. I worked on redesigning the airframe which became quite a rewarding effort, resulting in productive collaborations and the current airframe design which performs nicely.

snap joint fabrication

I made snap fit inserts and alignment features which make it possible to assemble an entire robofly without glue. The snap fits are cut out by the DPSS laser during airframe release cut, and consist of a barbed feature which guides into a target hole cut in the opposing material. One inserted through the laminate hole, the male barb returns to original width and holds the barb into position tight to the opposing material. This greatly eases fabrication and reduces manufacturing variability, unless you forget the glue.

A.1.1 CAD

I produced the PTC CREO airframe and flapping visualizations for old and new triple-wide airframe designs to aid design considerations and serve the lab in presentations. These files and PWL files for the kinematics are available by request.

A.2 Actuator

To improve actuator reliability, reduce electrode impedance for the switching PEU's, and reduce fabrication steps, electrode attachment was redesigned. Former methodology for attaching electrodes to the piezos used conductive epoxy placed over carefully insulated bridges from a solder point at the actuator base to both the top and bottom layers of the bimorph. To attach the signal electrode, an additional conductive epoxy joint was needed to access the middle layer of the bimorph through a window cut into the insulation of the actuator base. This conductive epoxy step required expensive material, additional insulation where bridging the gap between actuator base and pzt top and bottom, careful manual application to ensure clean but wide contact area, and an additional oven curing requiring careful handling and significant time. These additional steps often resulted in failure. The result, if correctly done, was an electrical contact measured to vary from 25-1000 Ω. This additional contact resistance and parasitic capacitance within the conductive epoxy matrix results in losses for switching-mode DC-DC converters. Additionally, the conductive epoxy frequently lost adhesion to the actuator due to flexing or handling, resulting in the effective loss of the robot. For these reasons a 25 μm copper electrode layer was added to the actuator design so as to achieve superior and easier electrical solder contact to all piezo bimorph layers via the conductive carbon fiber prepreg.

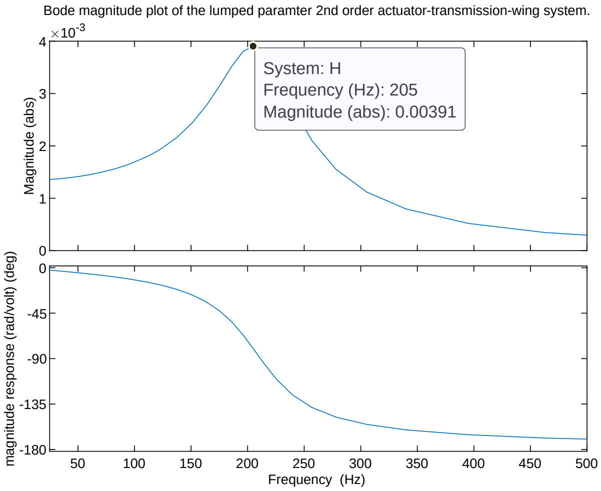

Figure A.1: Bode magnitude plot showing the wing-transmission-actuator system over input excitation frequency.

A.2.1 PFM duty-cycle for driving a simplified actuator model

As in chap. 5, Eq. 5.2 and Eq. 5.3 can be solved for the 'on' time of the charge and discharge switches over the period of time of the k^{th} control iteration. In order to facilitate model improvements, in this work this calculation was performed using symbolic computational utilities in Python, similar to the MATLAB symbolic toolbox. Full code is available by request or in the authors github repositories, but a shortened example code is below:

```
import numpy as np
import sympy as sp
from IPython.display import display

# capacitive and dielectric loss as parallel R0 to actuator C0
Etransferchgfcn = sp.Eq( (etas) * (1/2)*( ((Ton*(VS))/Lp)**2 * Lp), (1/2
                                       )*Cbot*(-Vsig**2 + Vref**2) + (1/2
                                       )*Ctop*(-(Vb-Vsig)**2 + (Vb-Vref)
                                       **2) + t_k*((Vref+Vsig)/2)**2/R0 -
                                        t_k*(((Vb-Vref)+(Vb-Vsig))/2)**2/
                                       R0 )

display(Etransferchgfcn)
print("latex for Etransferchgfcn: ", sp.latex( Etransferchgfcn  ))

#solve for Ton (which will be divied up across the span of the control
                                       period, we compute the PFM for
                                       this below.).
Tonchgf_fam = sp.solve(Etransferchgfcn, Ton)
TonDCf0 = Tonchgf_fam[0]
TonDCf1 = Tonchgf_fam[1]
print("there are " + str( len(Tonchgf_fam) ) + " solutions. choose the
                                       positive one.")
Tonchgfn0 = sp.lambdify( [Vb, VS, Vsig, Vref, Ls, Lp, Cbot, Ctop, etas,
                                       R0, t_k] , TonDCf0, "numpy")
Tonchgfn1 = sp.lambdify( [Vb, VS, Vsig, Vref, Ls, Lp, Cbot, Ctop, etas,
                                       R0, t_k] , TonDCf1, "numpy")
```

A.3 Misc Code

A.3.1 ICRA 2021 LABView

Fig. A.2 shows the front panel of the LABView VI used for [29]. The full VI, oscilloscope and serial interfaces, and highside current monitor calibration, and capacitive displacement sensor reading software is available by request.

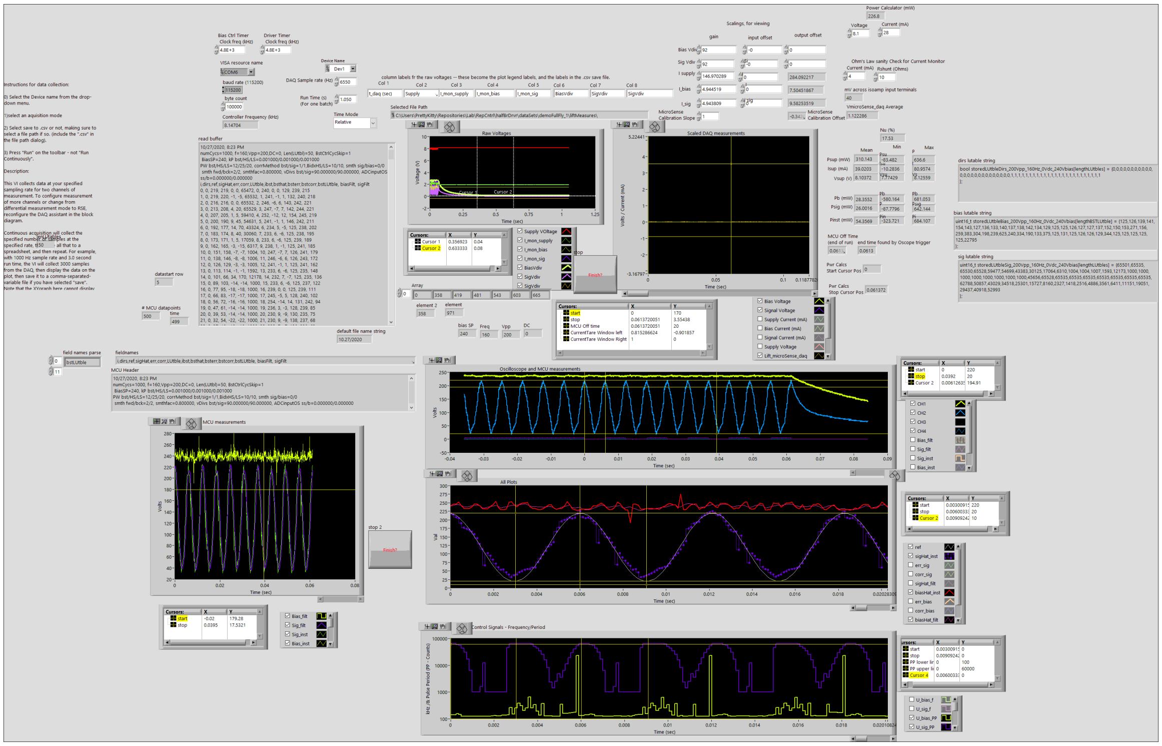

Instructions for data collection:

0) Select the Device name from the drop-down menu.

1)select an aquisition mode

2) Select save to .csv or not, making sure to select a file path if so. (include the ".csv" in the file path dialog).

3) Press "Run" on the toolbar - not "Run Continuously".

Description:

This VI collects data at your specified sampling rate for two channels of measurement. To configure measurement of more channels or change from differential measurement mode to RSE, reconfigure the DAQ assistant in the block diagram.

Continuous acquisition will collect the specified number of samples at the specified rate, t all that to a spreadsheet, and then repeat. For example, with 1000 Hz sample rate and 3.0 second run time, the VI will collect 3000 samples from the DAQ, then display the data on the plot, then save it to a comma-separated-variable file if you have selected "save". Note that the XYgraph here cannot display

Bias Ctrl Timer Clock freq (kHz): 4.8E+3
Driver Timer Clock freq (kHz): 4.8E+3
VISA resource name: COM6
baud rate (115200): 115200
byte count: 100000
Controller Frequency (kHz): 8.14704
Device Name: Dev1
DAQ Sample rate (Hz): 6550
Run Time (s) (For one batch): 1.050
Time Mode: Relative
lenLUtables: 50
datastart row: 5
MCU datapoints: 500
time: 499
default file name string: 10.27/2020

column labels fr the raw voltages -- these become the plot legend labels, and the labels in the .csv save file.

Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8
t_daq (sec)	Supply	I_mon_supply	I_mon_bias	I_mon_sig	BiasVdiv	SigVdiv	SigVdiv

Selected File Path: C:\Users\PrettyKitty\Repositories\Lab\RepCntrl\halfBrDrvr\dataSets\demoFullFly_1\liftMeasures\

read buffer

```
10/27/2020, 8:23 PM
numCycs=1000, f=160,Vpp=200,DC=0, Len(LUtbl)=50, BstCtrlCycSkip=1
 BiasSP=240, kP bst/HS/LS=0.001000/0.001000/0.001000
PW bst/HS/LS=12/25/20, corrMethod bst/sig=1/1,BidxHS/LS=10/10, smth sig/bias=0/0
 smth fwd/bck=2/2, smthfac=0.800000, vDivs bst/sig=90.000000/90.000000, ADCinputOS
ss/b=0.000000/0.000000
i,dirs,ref,sigHat,err,corr,LUtble,ibst,bsthat,bsterr,bstcorr,bstLUtble, biasFilt, sigFilt
0, 0, 219, 219, 0, 0, 65472, 0, 240, 0, 0, 129, 239, 215
1, 0, 219, 220, -1, -5, 65532, 1, 241, -1, 1, 132, 240, 218
2, 0, 216, 216, 0, 0, 65532, 2, 246, -6, 6, 143, 242, 221
3, 0, 213, 208, 4, 20, 65529, 3, 247, -7, 7, 142, 244, 221
4, 0, 207, 205, 1, 5, 59410, 4, 252, -12, 12, 154, 245, 219
5, 0, 200, 190, 9, 45, 54631, 5, 241, -1, 1, 146, 242, 211
6, 0, 192, 177, 14, 70, 43324, 6, 234, 5, -5, 125, 238, 202
7, 0, 183, 174, 8, 40, 30060, 7, 233, 6, -6, 125, 238, 195
8, 0, 173, 171, 1, 5, 17059, 8, 233, 6, -6, 125, 239, 189
9, 0, 162, 165, -3, -15, 6317, 9, 238, 1, -1, 125, 241, 185
10, 0, 151, 158, -7, -7, 1004, 10, 247, -7, 7, 126, 241, 179
11, 0, 138, 146, -8, -8, 1006, 11, 246, -6, 6, 126, 243, 172
12, 0, 126, 129, -3, -3, 1005, 12, 241, -1, 1, 125, 241, 162
13, 0, 113, 114, -1, -1, 1592, 13, 233, 6, -6, 125, 235, 148
14, 0, 101, 66, 34, 170, 12178, 14, 232, 7, -7, 125, 235, 136
15, 0, 89, 103, -14, -14, 1000, 15, 233, 6, -6, 125, 237, 122
16, 0, 77, 95, -18, -18, 1000, 16, 239, 0, 0, 125, 239, 111
17, 0, 66, 83, -17, -17, 1000, 17, 245, -5, 5, 128, 240, 102
18, 0, 56, 72, -16, -16, 1000, 18, 254, -14, 14, 131, 242, 94
19, 0, 47, 61, -14, -14, 1000, 19, 236, 3, -3, 128, 239, 85
20, 0, 39, 53, -14, -14, 1000, 20, 230, 9, -9, 130, 235, 75
21, 0, 32, 54, -22, -22, 1000, 21, 230, 9, -9, 138, 237, 68
```

field names parse: 0, 11, bstLUtble

fieldnames: i,dirs,ref,sigHat,err,corr,LUtble,ibst,bsthat,bsterr,bstcorr,bstLUtble, biasFilt, sigFilt

MCU Header

```
10/27/2020, 8:23 PM
numCycs=1000, f=160,Vpp=200,DC=0, Len(LUtbl)=50, BstCtrlCycSkip=1
 BiasSP=240, kP bst/HS/LS=0.001000/0.001000/0.001000
PW bst/HS/LS=12/25/20, corrMethod bst/sig=1/1,BidxHS/LS=10/10, smth sig/bias=0/0
 smth fwd/bck=2/2, smthfac=0.800000, vDivs bst/sig=90.000000/90.000000, ADCinputOS ss/b=0.000000/0.000000
```

Array: 0 | 0 | 358 | 419 | 481 | 543 | 603 | 665
element 2: 358
element: 971

bias SP	Freq	Vpp	DC
240	160	200	0

Cursors:	X	Y
Cursor 1	0.356923	0.04
Cursor 2	0.633333	0.08

Cursors:	X	Y
start	0	170
stop	0.0613720051	3.55438
MCU Off time	0.0613720051	20
CurrentTare Window left	0.815286624	-0.901857
CurrentTare Window Right	1	0

Scalings, for viewing

	gain	input offset	output offset	
Bias Vdiv	92	-0	0	
Sig Vdiv	92	-0	0	
I supply	146.970289	0		284.092217
I_bias	4.944519	0		7.50451867
I_sig	4.943809	0		9.58253519

MicroSense Calibration Slope: 1
MicroSense Calibration Offset: -0.342
Power Calculator (mW): 226.8
Voltage: 8.1
Current (mA): 28

Ohm's Law sanity Check for Current Monitor
Current (mA): 4
Rshunt (Ohms): 10
mV across isoamp input terminals: 40
VmicroSense_daq Average: 1.122286
Nu (%): 17.53

	Mean	Min	Max
Psup (mW)	310.143	-83.482	636.6
Isup (mA)	39.0203	-10.2836	80.9574
Vsup (V)	8.10372	7.77429	8.12559
Pb (mW)	28.3552	-580.164	681.053
Psig (mW)	26.0016	-87.7796	642.144
Pinst (mW)	54.3569	-323.721	684.107

MCU Off Time (end of run): 0.0613
end time found by Oscope trigger: 0.0613
Pwr Calcs Start Cursor Pos: 0
Pwr Calcs Stop Cursor Pos: 0.061372

dirs lutable string

```
bool storedLUtbleDirs_200Vpp_160Hz_0Vdc_240Vbias[lengthLUtbles] = {0,0,0,0,0,0,0,0,0,0,0,0,
0,0,0,0,0,0,0,0,0,0,0,0,0,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1,1
};
```

bias lutable string

```
uint16_t storedLUtbleBias_200Vpp_160Hz_0Vdc_240Vbias[lengthBSTLUtble] = {125,126,139,141,
154,143,127,136,133,140,137,138,142,134,129,125,125,126,127,127,137,152,150,153,271,156,
259,383,304,198,239,625,240,334,190,133,375,125,131,125,126,126,129,344,125,125,125,125,
125,22795
};
```

sig lutable string

```
uint16_t storedLUtbleSig_200Vpp_160Hz_0Vdc_240Vbias[lengthLUtbles] = {65501,65535,
65530,65528,59477,54699,43383,30125,17064,6310,1004,1004,1007,1593,12173,1000,1000,
1000,1000,1000,1000,1000,1000,45656,65528,65535,65535,65535,65535,65535,65535,65535,
62788,50857,43029,34518,25301,15727,8160,2327,1418,2516,4886,3561,6411,11151,19051,
29437,40918,52993
};
```

Cursors:	X	Y
start	0	220
stop	0.0392	20
Cursor 2	0.00612635	194.91

Cursors:	X	Y
start	-0.02	179.28
stop	0.0395	17.5321

Cursors:	X	Y
start	0.00300915	220
stop	0.00600333	20
Cursor 2	0.00909242	10

ursors:	X	Y
start	0.00300915	0
stop	0.00909242	0
PP lower lir	0	100
PP upper lii	0	60000
Cursor 4	0.00600333	0

Figure A.2: Front panel of LABView data collection software used for [29].

A.3.2 Other code

Matlab and python measurement tools for THD, AutoCorrelation, routines: MCR Embedded Capacitance Optimizer, Repetitive Controller Waveform Generator, STM, Atmel, and NRF codes, airframe and actuator design files, circuit PCB files, circuit simulations, and scripted symbolic computation for pulse trains and actuator system identification are all available by request to JohannesMJames@gmail.com.

A.4 High-side current sensor

Design of the PEU often required the ability to measure currents on non-grounded conductors. This 'high-side' current sensing is challenging because non-contact current sensors in various forms are generally insensitive, susceptible to noise, bulky, expensive, or impose loading effects on the measured system. Unfortunately, a simple series resistor element and Ohms law typically uses high-gain and finely tuned instrumentation amplifiers which usually need to be **ground referenced**. Since the series sensing resistor cannot be ground referenced for high-side measurements, it is convenient to reference the instrumentation amplifier to the highside, and pass it's output across a galvanic isolation[1], so that the data acquisition system or microcontroller recording this current measurement can be ground referenced and isolated from the high-side. This clever arrangement requires a floating power supply for the isolated high-side amplifier, which just 'rides' up and down in potential with the high-side. Conveniently, integrated circuits are available to achieve all this: 'isolation amplifier'. I therefore made use of these circuits, and designed a circuit with several such ICs, floating power supplies, charging circuitry for the floating power supples, and simple safe interface for user to connect inputs and outputs and switch between operation and charging.

For the floating power supplies, I found it convenient to simply use coincell rechargeable batteries. Using batteries as opposed to some other isolated power supply, such as transformer with as many secondary windings as wanted followed by rectification and conditioning circuitry, is not substantially harder and would be an improvement to the device.

A.4.1 circuit schematic

The circuit designed for the current sensing is simple at its core: it is just an series resistor and we're applying Ohm's law and amplifying the output using an isolated instrumentation amplifier. The isolation amplifier has a defined minimum and maximum input voltage

---

[1] this isolation gap can be achieved via transformer with complicated modulation/demodulation, or more modernly capacitively as is the IC employed here.

though, so the shunt resistor was sized so as to result in the voltage drop across the shunt resistor due to expected min/max current to range within 20–80% of the manufacture specc'd min and max. This resistor is made easier to swap out using SMD soldering tweezers, although future revisions may wish to make the range switchable.

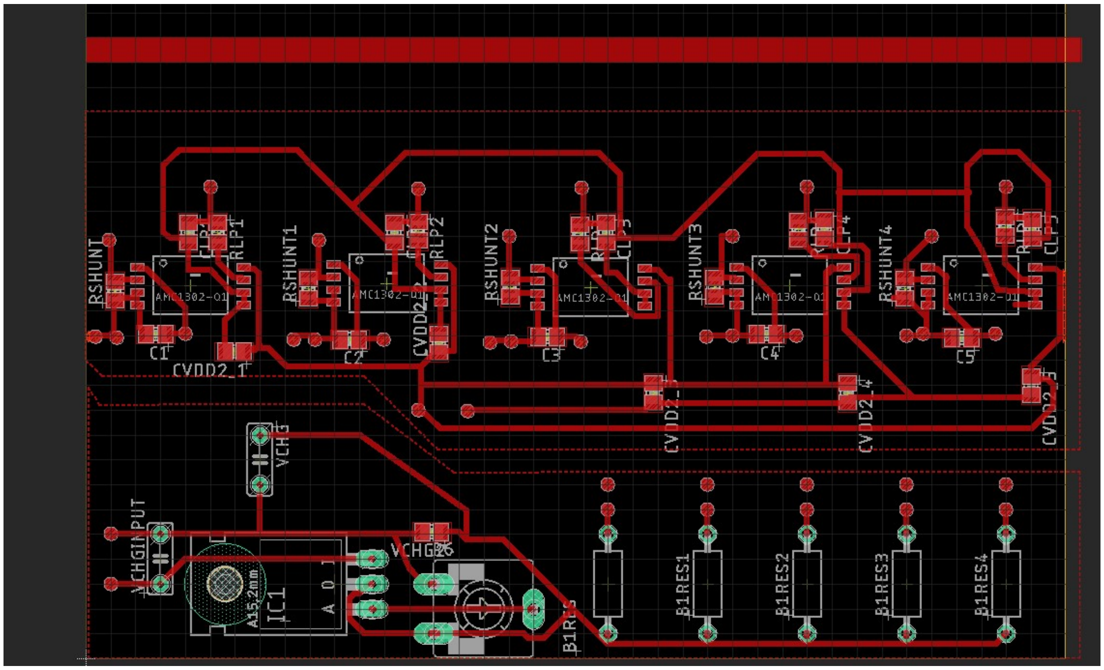

Figure A.3: Schematic of the custom high-side current monitor circuit.

A.4.2 calibration

Since this custom sensor was not exhaustively studied so that gain and drift can be documented, I decided it was safer to tare the instrument based on recorded data immediately before and after any usages. It was found that drift was negligible even at very different state of charge levels of the floating power supply.

Calibration was done as often as wanted but especially prior to experiments. Fortunately, repeated calibrations showed that the sensors' gains were quite constant over typical room temperature and humidity fluctuations as well. However, the output RC filter and shunt resistor values (being slightly different from channel to channel amongst the 5 channels), did require calibrating each channel. To simplify calibration, a known current source (Fluke process calibrator used for calibrating industrial instrumentation) and a LabVIEW calibration script, were used to easily generate calibration results.

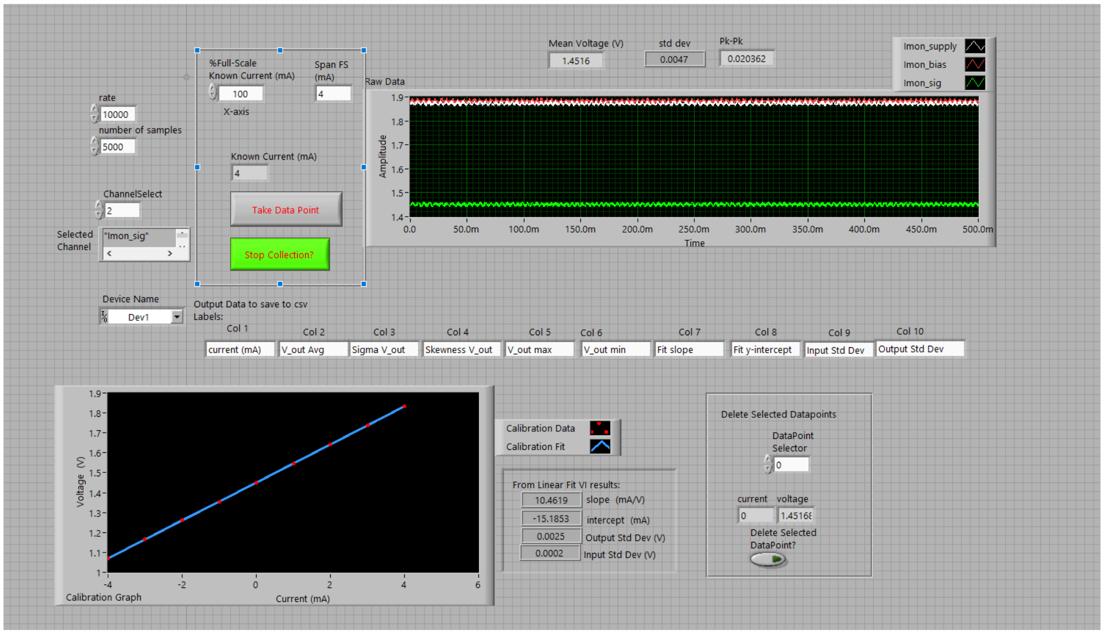

rate: 10000
number of samples: 5000
ChannelSelect: 2
Selected Channel: "Imon_sig"
Device Name: Dev1
%Full-Scale Known Current (mA): 100
Span FS (mA): 4
Known Current (mA): 4
Mean Voltage (V): 1.4516
std dev: 0.0047
Pk-Pk: 0.020362

Output Data to save to csv Labels:

Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8	Col 9	Col 10
current (mA)	V_out Avg	Sigma V_out	Skewness V_out	V_out max	V_out min	Fit slope	Fit y-intercept	Input Std Dev	Output Std Dev

From Linear Fit VI results:

10.4619	slope (mA/V)
-15.1853	intercept (mA)
0.0025	Output Std Dev (V)
0.0002	Input Std Dev (V)

Delete Selected Datapoints
DataPoint Selector: 0
current: 0
voltage: 1.4516

Figure A.4: High-side current monitor calibration front panel.

Appendix B Voltage feedback on ultra-light FIR PEUs

The voltage dividers selected for this work have $R_{eq} = 100-300\ k\Omega$ and dissipate $P_{diss} = 9-3$ mW at full input voltage of 300V input. This high R_{eq} violates manufacturer guidelines for $R_{eq} \leq 50$ kΩ but were observed to function anyway. Decreasing R_{eq} unacceptably increases power dissipation and worsens the SNR and resolution by not utilizing the full range of analog input. $R_{low} = 100 - 300$ kΩ was implemented using 0201 sized resistors, and $R_h = 10 - 30$ MΩ was implemented using 2-3 series 10 MΩ 0201 resistors. Larger resistors offer higher voltage ratings but are substantially heavier, and even though the voltage ratings on the PEUs are exceeded, experiments have shown the prototypes have survived this without failure. R_h is a series combination in order to distribute the voltage drop which exceeds component ratings. If payload increases allow for op-amp ICs, active filters are recommended to improve the impedance of the voltage feedback.

Eqns.B.1, B.2 describe competing considerations for FIR PEU voltage dividers which seek to avoid additional components such as operational amplifiers or buffer ICs. Analysis is simplified by assuming the current into the ADC sampler is negligible.

$$V_{ADC} = V_{sig}G, \;\; G = \frac{R_{low}}{R_h + R_{low}}, \;\; \hat{V_{sig}} = \frac{V_{ADC}}{G} \;\; R_{eq} = (\frac{1}{R_{low}} + \frac{1}{R_h})^{-1}, \;\; R_{tot} = R_{low} + R_h \tag{B.1}$$

$$\tau_{SH} = R_{eq}C_{SH} \;\; P_{diss} = \frac{V_{in}^2}{R_{tot}}, \;\; SNR = \frac{V_{FS}}{\nu}, \;\; \Delta V_{in} = \frac{GV_{in}}{2^n} \tag{B.2}$$

Where:

1. V_{out} is the voltage input to the ADC sampler, which must be less than device full-scale input, here 3.3 V.

2. V_{sig} is the voltage of the piezo drive signal being measured, assumed to be a maximum expected operation value $V_{sig} = 300$V for design purposes.

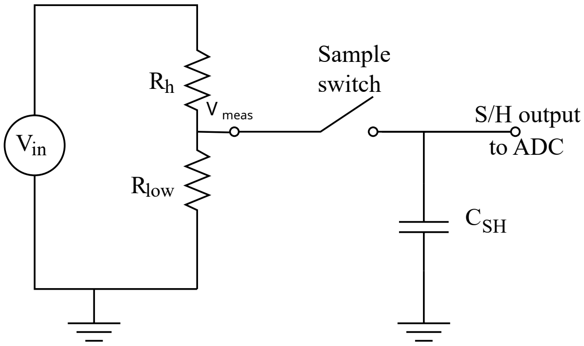

Figure B.1: PEU voltage sensor and the ADC.

3. G is the gain of the sensor.
4. $\hat{V_{in}}$ is the estimated voltage of the signal being measured.
5. R_{eq} is the Thevenin equivalent impedance of the voltage divider, as seen from the ADC input as indicated in Fig. B.1.
6. $\hat{V_{ADC}}$ is the estimated voltage at the ADC's sample-hold input.
7. τ_{SH} is the time constant of the 1^{st} order system formed by the voltage divider R_{eq} and the ADC's sample-and-hold capacitor C_{SH} which should be sufficiently charged until settled before acquisition.
8. P_{diss} is the power dissipated by the voltage divider.
9. SNR is the signal to noise ratio, here the full-scale voltage of the analog input V_{FS} (0-$\hat{V_{sig}}$ volts) compared to the amplitude of the noise ν
10. $n = 10$ is number of bits of the ADC.
11. ΔV_{in} is the resolution (in volts) of the voltage measurement after rescaling by the gain. This is the minimum amount of detectable change in the piezo drive signal.

The ADC resolution and accuracy is generally improved if the acquisition time T_{meas} is increased, most manufacturers recomment over 8 time constants. For digital control, $T_{meas} < T_{wfg}$ i.e. the sample acquisition must complete within the period of the waveform generator controller. However, the total impedance of the voltage divider must be quite large to avoid excessive power dissipation. The ADC sample and hold capacitor loads the voltage divider when the sample switch is closed, forming an RC system with time constant τ_{SH}, suggesting acquisition time to be as long as possible.

The sets of eqs. B.1 & B.2 indicate some competing goals. In particular, minimal τ_{SH} benefit by small R_{low} and proportionally larger R_h, however minimal P_{diss}, maximum SNR, and minimum ΔV_{in} benefit by larger R_{low} relative to R_h, just meaning that we should utilize the full range of allowable ADC voltage input which suggests larger R_{low}. A voltage follower inserted between the voltage divider and the ADC resolves all these problems and is tempting, but in this work was decided against due to increased flex pcb complexity and needing an additional ≈ 3 mg per channel for SMD opamp/buffer ICs, discrete components, and pcb area.