

Parallel Control of Electrostatic Robotic Muscle

by

Mantas Naris

B.S.E., Case Western Reserve University, 2017

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy, Robotics
Robotics Program

2024

Committee Members:

J. Sean Humbert, Chair

Christoph Keplinger

Robert MacCurdy

Nikolaus Correll

Marco Nicotra

Eric Bogatin

Naris, Mantas (Ph.D., Robotics)

Parallel Control of Electrostatic Robotic Muscle

Thesis directed by Prof. J. Sean Humbert

Living things are capable of incredible dexterity, agility, robustness, and elegance in their interactions with their environment. For far longer than the field of robotics has even existed, humanity has aspired to capture even a fragment of that dynamic elegance – not merely out of curiosity, but also to inform the design of the artificial systems we hope to construct ourselves.

The field of soft robotics hypothesizes that by engineering comparably compliant, dynamically complex systems, some of this elegance can be translated to robotics. Thanks to recent strides in materials and fabrication, electrostatic actuators (ESAs) have emerged as a promising core technology for the field. These electrically-driven robotic muscles are powerful, fast, and can be made numerous geometries using a wide range of materials. However, electrostatic actuators are not without their own challenges. They operate at thousands of volts, but low currents – a niche where few technologies exist.

This work aims to address a fundamental problem for any electrostatic actuator-based soft robot, and for the future of our field: how to translate between the theoretical and computational tools used to study system dynamics and synthesize global control laws, and the high voltage signals needed to independently drive these robotic muscles – not one, or ten, but potentially hundreds to thousands of them. An answer to this question will open the door to the construction of highly performant soft robots with complex, coordinated motion, and facilitate the development of novel control techniques for these dynamically complex systems, taking us one step closer to the elegance of living things.

This work begins with a discussion on the current state of soft robotics, and identifies the limitations on the field imposed by a lack of actuators and accompanying drivers which are both performant and highly parallelizable. We identify the potential of electrostatic actuation (ESA)

to be a technology platform capable of resolving these limitations contingent on the development of parallelizable high voltage driver circuitry and feedback controls, and set our objective to do so.

In the second chapter, we build a dynamical model of an electrostatic actuator from first principles and geometry. We explore the physics and dynamics of these actuators along with the assumptions which can be used to simplify their modelling. The result is a linearized model of an ESA driven by a bridge-style switching amplifier which is tractable for control design.

In the third chapter, we introduce a parallelizable scheme for feedback control of electrostatic actuators. We introduce a parallelizable optoelectronic driver circuit suitable for driving high voltage capacitive loads such as ESAs. We present a procedure to refine our derived dynamical model through the system identification of physical actuators. Using this information, we design a control law for feedback regulation of actuator voltage, and demonstrate and validate its implementation on a benchtop system.

The fourth chapter, we demonstrate the utility of our approach by achieving parallel, independent control of soft actuators at unprecedented scales. We present the construction of a synthetic tissue – a soft, intelligent, robotic material with 100 independent channels of tightly integrated sensing, actuation, and control. Our system includes distributed magnetometers to leverage our insights from modelling ESA dynamics to implement both feedback control of actuator displacement and the sensing of externally applied forces. This chapter concludes with a discussion of the complex emergent behaviors that emerge as a result.

In the final chapter, we consider how to scale our approach by additional orders of magnitude. To do so, we present the design of a novel solid-state high-voltage switch based on series stacks of low-voltage transistors. Our design is more efficient, more performant, and better suited for mass manufacturing than existing alternatives. We introduce the theory of its operation, discuss its implementation, evaluate its performance against the state of the art, while also laying a framework for future studies.

*Naktį užkluptas žaibo nustėrau
suakmenėjau aštrios šviesos pervertas:
kas tai – meilė, neapykanta, pyktis, mintis
ar viename smūgyje – viskas?*

Justinas Marcinkevičius

Acknowledgements

This work would not have been possible without the aid of many people – more than I can possibly include – and I thank you all. First and foremost, I would like to thank my advisor, Prof. Sean Humbert, for his guidance and support throughout my graduate studies. Sean, you have been an incredible mentor, and I am grateful for the countless hours you have spent helping me navigate the world of dynamics and controls. Whiteboarding applied mathematics and controls theory together will forever remain a highlight of my studies.

Next, I would like to thank the members of my committee – each of whom has fundamentally shaped the direction of my research, perhaps more than they realize. Professors: I am deeply grateful for your willingness to dissect and discuss the details of my work, to send me down new paths of discovery whenever I hit a dead end, and to offer your wisdom and guidance when I needed it. Most of all, I thank you for your enthusiasm for doing so.

Thank you to my labmates and colleagues, Angie, Lluís, Leo, Zoe, Shakheeb, Beth, Gene, and Mike, for your friendship, for everything you've taught me, and for the many adventures along the way. Thank you to my students: Jonny, Andrew, Matt, Tyler, and the members of the CU Robotics Team. I have learned as much from you as I hope you have learned from me. Finally, I would like to express my deepest gratitude to the Faculty of the Robotics Program at the University of Colorado, Boulder for welcoming me as their first transfer student.

I would also like to acknowledge the scientific contributions of my colleagues which are presented in this work – specifically the members of our s-tissue team whose contributions are detailed in section 4.9.1. I would especially like to thank Dr. Shane Mitchell, who's development

of high voltage power optocouplers was a necessary foundation for most of the results presented in this work. I would also like to thank Dr. Mitchell for introducing me to the world of high voltage power electronics, and for the many, many hours spent arguing over how best to implement them. These conversations were among the most engaging and insightful of my studies.

This document was prepared with the assistance of a copilot (GPT-4, OpenAI, late 2023). All copilot-assisted writing has been verified, and I stand behind the entirety of this document.

Contents

Chapter

1	Introduction: The Current State of Soft Robotics	1
1.1	The Case for Soft Robotics	1
1.2	The State of the Art in Soft Transducers	2
1.3	The Problem of Controls	4
1.4	A Note On High Voltage Safety	7
2	The Dynamics of Electrostatic Actuators	8
2.1	Peano-HASEL Actuators	8
2.2	ESA Equivalent Circuit Model	9
2.3	Geometric Relations	11
2.4	Lagrangian Formulation	14
2.4.1	Insights Arising from Equilibrium	18
2.5	ESA Driver Topologies	19
2.6	Half Bridge Driver Model	22
2.6.1	Zero-Input Dynamics	24
2.6.2	Zero-Leak Dynamics	25
2.6.3	Linearized SISO Model	27
3	Scalable Feedback Control of Electrostatic Actuators	29
3.1	Optoelectronic Bridge Drivers	29

3.1.1	Optocoupler Fundamentals	29
3.1.2	Optocoupler Drive Circuitry	30
3.1.3	Optocoupler Alignment	31
3.1.4	Optocoupler Isolation	32
3.1.5	Optocoupler Thermal Management	33
3.2	High Voltage Sensing Circuitry	34
3.3	Half Bridge System Identification	35
3.3.1	Transfer Function	35
3.4	Frequency Domain System Identification	37
3.5	Loopshaping Control Design	38
3.5.1	Steady-State Error	39
3.5.2	Tracking Error	40
3.5.3	Maximum Overshoot	40
3.5.4	Maximum Gain	40
3.5.5	Noise Rejection	41
3.5.6	Closed Loop Bandwidth	41
3.5.7	Controller Design	42
3.6	Digital Control Implementation	44
3.6.1	Controller Discretization	44
3.6.2	Actuator Fault Detection	46
3.7	Summary, Validation, and Performance	46
3.7.1	Voltage Tracking	48
3.7.2	Closed Loop Frequency Response	50
3.7.3	Limitations of Voltage Control	52
4	A Case Study in Sensorimotor Systems Integration	55
4.1	System Architecture	56

4.1.1	Sensorimotor Unit Cell	56
4.1.2	Modular Backbone	57
4.1.3	Hardware Interface	59
4.1.4	Thermal Management	60
4.1.5	Structural Elements	61
4.1.6	Silicone Skin and Fabrication	63
4.1.7	System Assembly	63
4.2	Power Supply Stack	64
4.2.1	USB-C Power Delivery	65
4.2.2	High Voltage Boost Stage	66
4.2.3	High Voltage Capacitor Bank	66
4.2.4	High Voltage Lockout	67
4.2.5	High Voltage Indicator	68
4.2.6	The Missing Trip Circuit	69
4.3	Computational Architecture	70
4.4	Motion Capture Videography	73
4.5	Distributed Magnetometry	75
4.5.1	Magnetometer Fabrication	76
4.5.2	Magnetometer Circuitry	77
4.5.3	Magnetometer Calibration	77
4.5.4	Magnetometer Efficacy	78
4.6	Displacement Control of ESAs	79
4.6.1	Performance of Displacement Control	82
4.7	Force Sensing	84
4.7.1	Force Curve Fit	85
4.7.2	Distributed Force Sensing	86
4.8	Emergent Behaviors	87

4.8.1	Dynamic Conveyance	88
4.8.2	Distributed Audio and Haptics	89
4.8.3	Multimodal Interaction	90
4.9	Conclusion	91
4.9.1	Contributions	92
5	Solid State Drivers for Electrostatic Actuators	93
5.1	Switching High Voltages	93
5.1.1	Stacked Series MOSFETs	95
5.1.2	Electrons and Holes	97
5.1.3	Our Approach	99
5.2	A Floating Optoelectronic Gate Driver	100
5.2.1	High-Side Gate Driver Circuit	102
5.2.2	Low-Side Gate Driver Circuit	103
5.2.3	Gate Driver Performance	104
5.3	MOSFET Synchronization by Capacitive Gate Coupling	107
5.3.1	Choice of Coupling Capacitance	110
5.3.2	Voltage Balancing	111
5.3.3	Half-Bridge Voltage Balancing	115
5.3.4	Failure Modes And Current Limiting	118
5.4	Circuit Simulation	119
5.4.1	Transient Gate Behavior	121
5.5	Instrumentation	125
5.6	Manufacturing	126
5.7	Performance	128
5.7.1	Leakage Current	128
5.7.2	Slew Rate	128

5.7.3 Efficiency and Thermals 133

5.7.4 Pulse Width Modulation 135

5.7.5 ESA Step Response 136

5.8 Conclusion 138

Bibliography **140**

Figures

Figure

2.1	A 2x12 Peano-HASEL expanding actuator stack.	8
2.2	An electrical symbol representing an electrostatic actuator along with an equivalent circuit model.	9
2.3	Geometry of HASEL actuator deformation parametrized in terms of the zip angle α	11
2.4	A model of HASEL actuator deformation dynamics.	15
2.5	The internal topology of the typical HV switched-mode power supply (HV-SMPS).	19
2.6	Several inherently parallelizable bridge-style switching amplifier topologies suitable for driving ESAs.	21
2.7	Expanded equivalent circuit of a half-bridge-driven ESA.	23
2.8	Electrodynamics of the zero-input case.	24
2.9	Electrodynamics of the zero-leak case.	26
3.1	The basic principles of photoconductive power optocoupler operation.	30
3.2	Our implementation of a photoconductive power optocoupler and accompanying driver circuit.	31
3.3	Quasi-cross-sectional view of a high-performance high-voltage switching optocoupler. For simplicity, some features which are not in-plane are shown as such. Approximate scale.	33
3.4	A simple and effective HV sensing circuit for use with a microcontroller.	34

3.5	The complete half bridge driver implementation.	35
3.6	Maximum sew rate of our optoelectronic half-bridge as a function of input duty cycle.	36
3.7	Transfer function fit for 100 Peano-HASEL actuators. Inset figures show the input (black) and one standard deviation around the mean of the output (blue) at selected frequencies.	38
3.8	A block diagram representing a generalized single-input, single-output (SISO) system with feedback control.	39
3.9	Voltage controller design via loopshaping with superimposed performance requirements.	43
3.10	A block diagram of the "inner loop" voltage control scheme, highlighting interactions with the decomposed model of ESA dynamics.	47
3.11	Quasistatic voltage regulation.	48
3.12	Quasistatic displacement control via proxy of voltage regulation.	49
3.13	Normalized quasistatic displacement control via proxy of voltage regulation.	50
3.14	Frequency response of system under closed loop voltage control.	51
3.15	Characteristic hysteresis loop for a Peano-HASEL actuator stack.	52
4.1	The s-tissue: a soft shape display with simultaneous multifunctionality.	55
4.2	A sensorimotor unit cell from the s-tissue, exploded view.	57
4.3	A 1x10 s-tissue module and the subsystems found within.	58
4.4	Overview of the main elements of the s-tissue system.	59
4.5	Front panel subassembly and interface.	60
4.6	Rear panel subassembly with integrated exhaust fan.	61
4.7	Front panel subassembly, exploded view.	62
4.8	Topology of the s-tissue power supply stack.	64
4.9	A simple, effective USB-C power delivery sink circuit.	65
4.10	Dual stage high voltage boost converter.	66

4.11	Triply-redundant high voltage lockout system.	67
4.12	High voltage indicator circuit.	68
4.13	The computational architecture of the s-tissue, which is well suited for the demands of other highly parallelized soft robots.	71
4.14	A frame of motion capture videography data and accompanying 3D surface fit. . . .	74
4.15	Functional diagram of the magnetometry system and characterization of its perfor- mance.	75
4.16	Accurate, high-frequency distributed measurement of deformation.	78
4.17	Block diagram of the nested displacement and voltage control loops highlighting their interactions with the decomposed ESA dynamical model.	81
4.18	Comparison of voltage control and height control schemes under quasistatic conditions.	82
4.19	Aggregated photos demonstrating a scrolling text display.	83
4.20	Shape display by voltage control only.	84
4.21	Block diagram of the force estimation scheme within the s-tissue.	85
4.22	Force sensing data and polynomial fit.	86
4.23	The s-tissue, banana force scale.	87
4.24	A demonstration of computer vision-aided interaction with objects with the s-tissue.	88
4.25	A demonstration of the s-tissue interacting with liquids.	89
4.26	A demonstration of the force-sensing capabilities of the s-tissue.	90
4.27	Drawing on the s-tissue using a magnet-tipped wand. The system instantly renders any detected magnetic fields as surface deformations.	91
5.1	Transformer-based synchronization of series MOSFETs. A current pulse through the shared primary of each transformer creates simultaneous gate drive pulses on each individual secondary.	96
5.2	A CMOS half-bridge, constructed from a pair of PMOS and NMOS devices.	97
5.3	Perhaps the most common PMOS load switch driver topology.	98

5.4	A novel high voltage load switch based on an optoelectronic gate driver and a capacitively-coupled MOSFET stack. (US Provisional Patent No. 63/608,794)	99
5.5	A half-bridge ESA driver implemented using two MOSFET stack switches.	100
5.6	The minimal implementation of an optically-powered, optically-controlled gate driver.	101
5.7	High side of the optoelectronic gate driver circuit. Both the control signal and the power supply are transmitted as light, providing robust electrical isolation to the low-side circuitry.	102
5.8	Low side of the optoelectronic gate driver circuit.	103
5.9	Timing changes introduced by our gate driver to the switching of a BSS127 MOSFET.	105
5.10	A profile of the response time of our gate driver circuit.	106
5.11	Degradation in gate driver performance at higher frequencies.	107
5.12	The mechanics behind the capacitive coupling of two series MOSFETs, highlighting the parasitic gate capacitance C_g at the core of the method.	108
5.13	Capacitive coupling of N MOSFETs. Current scales as we descend down the ladder.	110
5.14	The induced change in gate-source voltage as a function of change in drain-source voltage shown for various values of coupling capacitance C_c	111
5.15	Effect of balancing resistance R_b on the drain voltage of each MOSFET within a coupled series stack, shown switching a resistive load.	113
5.16	Effect of balancing resistance R_b on the drain-source potential across each MOSFET within a coupled series stack, shown switching a resistive load.	114
5.17	Expected voltage and output current profiles of a stacked series MOSFET half-bridge amplifier used to drive an ESA.	116
5.18	The drain-source potential ΔV_{ds} across each MOSFET within a half-bridge composed of two coupled series stacks, driving a simulated ESA load at assorted values of R_b and C_c	117
5.19	Equivalent models of MOSFETs and Zener diodes used for simulations in LTSpice.	119
5.20	A comparison of the simulated and observed switching behavior of a BSS127 MOSFET.	120

5.21	Effect of variations in coupling capacitance C_c on gate turn-on timing throughout a capacitively-coupled series MOSFET stack.	123
5.22	Effect of variations in balancing resistance R_B on gate turn-off timing throughout a capacitively-coupled series MOSFET stack.	124
5.23	High speed, 4-channel high voltage instrumentation system. For additional measurement channels, multiple instruments can be synchronized with a trigger pulse.	125
5.24	Schematic and render of the prototype switch implementation.	126
5.25	Experimental setups for measuring the slew rate of direct-drive from an HV-SMPS (left) and of parallelizable switching circuits (right).	129
5.26	Characteristic trace of a slew rate evaluation test pulse delivered to a 2x12 Peano-HASEL actuator stack by a half-bridge amplifier circuit constructed from stacked series MOSFETs.	131
5.27	Maximum slew rate achieved using different switching technologies as a function of load capacitance.	132
5.28	Thermal imagery highlighting power delivery disparities using optocoupler (left) and solid-state (right) load switches to drive a resistive load.	134
5.29	Effect of PWM duty cycle variation on the output voltage and current to an ESA driven by a solid-state half-bridge amplifier.	135
5.30	Electromechanical step response of an ESA driven by a PWM-controlled solid-state half-bridge amplifier constructed with series MOSFET stacks.	137

Chapter 1

Introduction: The Current State of Soft Robotics

1.1 The Case for Soft Robotics

Traditional robotic architectures – built with rigid linkages and electric servomotors – excel at precise, repetitive tasks in well defined environments. However, in spite of numerous advancements, they struggle to adapt to changes in uncertain and dynamic environments, especially in comparison to the dynamic agility, dexterity, adaptability, and elegance of biological systems [80, 85, 86, 129]. Furthermore their heavy, inflexible structures driven by powerful motors necessitate a trade-off between performance and safety. This can be solved eg. by separating humans and robots in an industrial setting, but such a solution places fundamental limitations on the contexts where high-performance robots can be used. This is most sorely felt in the context of medical assistance and rehabilitation, prosthetics and orthotics, search-and-rescue, and human-robot collaboration, but extends to any context where robots must interact with delicate objects, such as the automation of agriculture.

Although robotics has long aimed to mimic the functionality of natural organisms, the comparative details of their structure, dynamics, and control could not be more different. Organisms have a highly hierarchical structure with trillions of individual elements. Cells form tissues, tissues form organs, and organs form organ systems which form the organism itself. Every level of this hierarchy demonstrates emergent properties arising from its components, and independent transducers (sensors or actuators) can exist at even the lowest levels of the hierarchy (eg. at the cell or tissue level). As a result, organisms contain several orders of magnitude more transducers

than modern day robots. These transducers are distributed throughout their structure, providing hyper-redundant actuation and extremely rich sources of sensory information.

Most importantly, like animals, soft robots leverage their inherent compliance in interactions with their surroundings to achieve unparalleled adaptability to uncertain environments [53,54,121]. In the limit case, an octopus arm demonstrates great feats of manipulation and mobility using no skeleton at all; it relies on a complex 3D arrangement of muscle capable of bending, shortening, elongating, and twisting called a muscular hydrostat [77–79], which is dexterously coordinated by patterns of neural activation [42–44].

The tradeoff is an explosion of dynamic complexity. However, if properly harnessed, this dynamic complexity is likely a key source of biology’s continued ability to drastically outperform current robots in agility, dexterity, adaptability, and elegance [129]; the field of soft robotics hypothesizes that by engineering comparably compliant, dynamically complex systems, some of this elegance can be translated to robotics [112,120,122].

To engineer such systems, progress must be made in two domains. The first is the development of soft transducers which form the basic building blocks of a soft robot. The second is in the development of effective controls techniques for the complex, high-dimensional dynamics which soft robots exhibit.

1.2 The State of the Art in Soft Transducers

Soft transducers – sensors and actuators – form the fundamental building blocks of soft robots, and their design remains a highly active area of research. Soft actuators can be driven by various means [52] including fluidic pressures [102,131,151], electrostatic forces [1,106,107], magnetic fields [66], piezo-response [92], temperature changes [46], or even the contraction of living muscle cells [9,82,83,104]!

Perhaps the most ubiquitous of these, fluidic actuators deform in response to hydraulic or pneumatic pressures. They can be designed in a variety of morphologies, often aided by automated design approaches, can integrate sensing [47], achieve high actuation force and large strokes [21], and

are particularly safe for human interaction [85]. They have been used for a variety of applications, ranging from grippers [47, 102], to hybrid robots [8, 19], to active surfaces [132], to continuum robotic tentacles [13]. However, they are not particularly efficient [96], and force engineers to make a fundamental tradeoff between performance and the inclusion of rigid and bulky compressors, reservoirs, and valves [148, 151].

One of the most promising technologies in soft robotics, electrostatic actuators are flexible capacitors which undergo deformation in response to Maxwell stresses induced by an applied high voltage field [1, 28, 38, 40, 41, 75, 123, 124, 130, 160]. They are performant, efficient, and versatile [107, 124], exhibiting large deformations, fast actuation, and high energy densities [38, 40], and are directly driven by electrical stimuli, which is convenient for systems integration. Some types can produce force densities 5–10 times higher than those of biological muscles [106]. They have been used in numerous soft robotics applications [41, 68, 124, 138], and even in micro air vehicles [48]. However, they rely on the maintained integrity of a dielectric layer to function properly. This dielectric is repeatedly subjected to large electrical potentials, and is prone to failure by electrical aging and dielectric breakdown, resulting in limited operational lifespans [75].

Materials with the strongest piezo-response are typically rigid ceramic crystals, and cannot truly be considered soft actuators. Thermally-driven actuators such as shape-memory alloys and thermoactive polymers are generally too slow to be of much practical use, especially when relaxation is driven by passive cooling.

Robots made from living cells represent an exciting new forefront in robotics. However, the work in this area represents only very preliminary steps. Living tissues are difficult to work with and integrate relatively poorly with existing robotics technologies. Simple engineered tissues cannot survive outside of nutrient baths, and tissue engineering the vasculature needed to circumvent this is no trivial feat [27, 63, 95, 133].

Electrofluidic actuators such as hydraulically amplified self-healing electrostatic (HASEL) actuators aim to bring the benefits of fluidic actuation into to the domain of electrostatic actuators by combining electrostatic delivery of power and the use of dielectric hydraulic fluids to both

amplify and direct force [1, 73, 75, 124]. A variety of materials can be used for their elastomeric shells and liquid dielectrics, and the actuators themselves can be made in numerous geometries using accessible manufacturing techniques, keeping their design flexible to the needs of a given application. Their flexible structure is inherently resilient to mechanical damage, and their liquid dielectric confers the ability to self-heal after electrical breakdown events [64] thus significantly increasing HASELs operational lifespan and reliability. They have demonstrated electromechanical conversion efficiencies up to 21 % [124], with a theoretical limit of about 80 % [111]. HASELs have a wide range of motion, contractile strains up to 24 % and expanding strains up to 124 % have been shown; they are also extremely fast, and have been shown to reach roll-off frequencies of 126 Hz [124]. HASEL actuators are generally an excellent choice for myriad soft robotics applications, and thus are used as the basis of the work presented here.

Like soft actuators, soft sensors have been developed based on a variety of techniques and sensing modalities, detecting changes in resistance or capacitance [2, 47, 76, 91], magnetic fields [138], or optical intensity [4, 163]. Perhaps most interesting of all is the ability of ESAs to self-sense their own strain state, much like natural muscle, reducing the dependence on external sensors for state feedback [75, 94, 124, 158]. Many of these methods have been successfully demonstrated to provide sufficient state feedback to achieve closed-loop control of a ESAs [69, 70, 94, 138].

1.3 The Problem of Controls

From the perspective of feedback control design, the characteristics of soft robots create significant challenges for current theory and methods. The derivation of precise models to describe and consequently control the dynamics of complex soft robots is difficult, as they lack a natural coordinate system and possess infinite degrees-of-freedom (DOFs) due to their elastic nature. In particular, conventional analysis and control synthesis approaches translate poorly to systems governed by nonlinear partial differential equations [84]. These difficulties are compounded as soft actuators typically exhibit nonlinear dynamics which may span multiple dynamic regimes [126], and often exhibit time-variance in their dynamics due phenomena such as the accumulation of

space charges in ESAs [33, 51, 70, 158]. Approaches to solve the problem of controls can loosely be grouped into model-based and model-free (data-driven) approaches. However, many strategies combine elements of both.

Whereas controls engineers typically start with linear methods, progress to nonlinear or model predictive control (MPC) approaches, and only then turn to machine learning in desperation, the opposite has occurred in soft robotics, where some of the first controllers developed were based on machine learning or other model-free, data-driven approaches [147] such as adaptive self-tuning control [87]. These approaches are powerful and convenient, and can simultaneously learn both the model and a suitable controller. Although machine learning approaches are considered by many to be problematic for controls as they introduce a "black box" into the system and make it difficult to prove stability or robustness, they are nonetheless capable of approximating mappings on complex, high-dimensional manifolds, and thus are potentially very well suited to the development of controls for soft robotics. The recent explosion of interest in this area suggests that machine learning approaches will play a significant role in the future of the field.

In terms of model-based approaches, open-loop control using approximate models that assume a steady-state force equilibrium have shown some success [147]. However, these models neglect significant aspects of the dynamics, and fail in situations requiring fast, accurate motion, or in the presence of applied external forces. Nevertheless, progress continues to be made towards the definition of soft robotic models which are tractable for controls design and still represent a useful approximation of the system [3, 23–26, 149]. Fractional order control is a model-based approach based on fractional calculus, which generalizes the differentiation operator \mathcal{D} to fractional powers \mathcal{D}^a . Recent developments in fractional order controllers allow the implementation of robust and adaptive control schemes with superior flexibility to their integer counterparts, and make these traditional methods better suited for soft robotics applications [101].

In high dimensional systems it is often observed through simulations and experiments that the dominant dynamics evolve in some subspace of the larger state space. Reduced order descriptions may capture these intrinsic dynamics in such a way that current theory and methods can be applied

to synthesize effective feedback controllers which rely on relatively few states [119].

Nervous systems have evolved to form useful reductions of high dimensional data in the presence of uncertainty, resulting in simple representations of sensory rich information streams. An underlying principle that is ubiquitous is sensorimotor convergence, in which sets of pooling neurons differentially weight arrays of sensory input to extract low-dimensional representations of complex data [61, 62, 65]. Likewise, complex motions using redundant actuators are composed through the combination of low-dimensional elements called motion primitives [30, 54–56]. It has been hypothesized that organisms encode their state and generate motor commands in terms of low-dimensional representations, and that mimicking how nature matches perception to dynamics and actuation could serve as an effective framework through which to approach infinite DoF systems.

The Koopman operator provides a framework through which to formalize this idea. The Koopman operator is an infinite-dimensional linear operator that evolves observables – functions of the state of a dynamical system [81, 98, 127]. It allows one to transform a nonlinear (finite or infinite) dynamical system into an infinite dimensional linear system, where eigenfunction expansion techniques can then be applied to analyze and reconstruct solutions. Its extension to input-output systems [115, 116] also allows linear control synthesis techniques to be applied, with the resulting linear control law inducing an equivalent nonlinear control law that is implementable in the original system. In terms of model reduction, it provides a pathway to finite dimensional linear representations of high dimensional nonlinear systems by restricting the Koopman operator to an invariant subspace spanned by specifically chosen observable functions. This approach has been demonstrated to be effective in numerous hyper-redundant soft robotic systems [12–14, 45, 71, 89].

A number of automated design techniques, such as genetic algorithms or compositional pattern-producing networks have been demonstrated as an effective tool to generate structures with respect to an objective function [18, 57–59, 145], and may provide insights into the design of both controls and system morphology.

Numerous other works have approach these questions from a primarily biological perspective, examining a multitude of disparate neural architectures spanning from arthropods to cephalopods.

Studies into the structure and function of nervous systems have provided valuable insights for controls design, as well as the development of neuromorphic controllers which implement control laws with simulated neurons [42–44, 103, 128, 134–136, 139–143, 154–157, 161, 162].

Nonetheless, the problem of controls in soft robotics remains an open question. Progress towards solutions will require access to rich data streams and substantial computational power to facilitate the implementation of cutting-edge controls-theoretic and data-driven algorithms.

1.4 A Note On High Voltage Safety

This work enables the construction of soft robots with unprecedented levels of parallelization and performance, and consequently with unprecedented levels of high voltage power delivery. The very nature of electrostatic actuators and the applications for which they are used brings these high voltages in close interaction with their environment. Predicting all of the potential failure modes of hundreds of HV capacitors in contact with an unpredictable environment is at best extremely difficult, and at worst, impossible.

It is generally understood that the worst-case catastrophic discharge of a single HASEL actuator to the human body is painful, but safe. This work demonstrates scaling that potential discharge 100-fold, and opens the door to scaling it much, much further. A broader discussion discussion of safety is sorely needed in our field.

Attempting to replicate these results could kill you.

Reading and understanding [113] is **mandatory** prior to attempting to replicate any of the work presented in this document. Follow appropriate safety precautions and stay safe, dear reader.

Chapter 2

The Dynamics of Electrostatic Actuators

This chapter examines the dynamics of electrostatic actuators (ESAs) and their implications for control. ESAs are non-conservative transducers which convert between electrical and mechanical energy, and can also store energy internally in the form of both electrical and elastic strain energy. Thus their two inputs or outputs are voltage and force, and their two internal states are charge and geometry. These ideas are expanded on in the following sections.

2.1 Peano-HASEL Actuators




Figure 2.1: A 2x12 Peano-HASEL expanding actuator stack.

The work presented in this chapter will readily generalize to any type of electrostatic actuator with adjustments for that actuator’s specific geometry and mechanisms of action. However, of the various classes of electrostatic actuators, HASELs are a particularly good choice for the practical implementation of dynamic, performant soft robots.

Consequently, the reader should note that virtually all of the results presented in this work are derived and demonstrated using specifically 2x12 expanding stacks of Peano-HASEL actuators [74, 75, 100, 126], which are shown in figure 2.1. These actuators are fabricated according to the procedure presented in [69]: actuator shells are made from 20 μm thick polyester film (L0WS, Multi-Plastics), filled with a liquid silicone dielectric (PSF-5cSt, Clearco), and electrodes are screen printed with conductive ink (CI-2051, Engineered Materials System).

2.2 ESA Equivalent Circuit Model

The electrical dynamics of an ESA are often modeled with the equivalent circuit shown in figure 2.2 [69]. Also shown is a symbol created to represent ESAs in circuit diagrams; it is a version of the capacitor symbol modified to allude to the physical structure and deformation mechanisms of these actuators.




Figure 2.2: An electrical symbol representing an electrostatic actuator along with an equivalent circuit model.

The load capacitance $C_{\text{load}}(\alpha)$ is the predominant element of this model, and models the

charge stored on the actuator's electrodes. Its value will change due to changes in geometry, increasing as the actuator's electrodes zip together. It is therefore represented as a variable capacitance which is a function of the actuator's instantaneous geometry, represented by the generalized coordinate α . The expression for the load capacitance will depend on the geometry of the actuator and on the dielectric strength of the materials used to construct it, and can become relatively complex. However, in its simplest case, the capacitance can be approximated as a parallel-plate capacitor comprising of the zipped portion of the actuator. The capacitance of the separated electrodes in the unzipped portion is not considered, as it will be much smaller than that of the zipped portion, where electrodes are separated by a distance of only twice the material thickness of the elastomeric shell. This approximation is given by:

$$C_{load}(\alpha) \simeq \frac{\varepsilon_r \varepsilon_0}{2t_m} A_e(\alpha) \quad (2.1)$$

Here, ε_r is the relative permittivity of the dielectric material, ε_0 is the permittivity of free space, t_m is the thickness of the elastomeric shell (2 layers are used), and $A_e(\alpha)$ is the effective area of the zipped portion of the electrodes at a given geometric state α . The model for A_e will depend on the geometry of an individual actuator.

A parallel resistance R_{par} models the leakage between actuator electrodes through the body of the actuator. This parameter mainly affects the zero-input dynamics of the actuator, and we would expect R_{par} to be on the order of $1 \text{ G}\Omega$ or more.

A series resistance R_{ser} models the impedance of the actuator electrodes, which are typically made with conductive inks and exhibit impedances which can become significant. This parameter is governed by the electrode ink selection, geometry, and quality of connection to external circuitry. It affects actuator performance and efficiency, and we would expect R_{ser} to be on the order of $1 \text{ M}\Omega$ or less.

2.3 Geometric Relations

ESAs have 2 states, which we represent as generalized coordinates: one electrical and one geometric. The generalized electrical coordinate is charge q , however the generalized geometric coordinate is not immediately obvious. In fact, geometry of the actuator's deformation can be parametrized in three different ways: the pouch angle α , or the displacement in the x or z directions. However, thanks to a holonomic¹ conservation of volume, all geometry can be described in terms of any one of them. Our first objective is to build the geometric relations that allow this, shown in figure 2.3.

Our approach follows the geometric analysis for contractile Peano-HASEL actuators [74,126], with a few modifications to more conveniently represent expanding actuator stacks. The actuator is represented as two conjoined circular segments parametrized by the angle α , with geometry constrained by a holonomic conservation of volume. The actuator is zipped when $\alpha = \alpha_f = \frac{\pi}{2}$, and unzipped when $\alpha = \alpha_0$.




Figure 2.3: Geometry of HASEL actuator deformation parametrized in terms of the zip angle α .

As discussed in section 2.2, the equivalent capacitance of the actuator is modeled as a parallel-plate capacitor with plates defined by the zipped portion of the electrode with area $A_e(\alpha)$. For a

¹ technically, scleronomic

constant electrode width W , this can be represented in terms of the effective length of the electrode $l_e(\alpha)$, which corresponds to the zipped portion of the electrode. As before, we assume as t_m is very thin compared to the displacement of the electrodes, and thus only consider this zipped portion towards the effective area of the capacitor. Equation 2.1 becomes the constitutive relation:

$$C(\alpha) = \frac{\varepsilon_r \varepsilon_0 W}{2t_m} l_e(\alpha) \quad (2.2)$$

Since the actuator pouches are sealed and contain a fixed volume of incompressible dielectric fluid, the volume of the actuator is constant, and can be represented as a function of two other constants: the width of the actuator W , and the cross-sectional area A of the pouch. The latter follows directly from the area of a circular segment:

$$A = \frac{L_p^2}{2\alpha_0^2} (\alpha_0 - \sin(\alpha_0) \cos(\alpha_0)) \quad (2.3)$$

The axial length of the pouch l_a is the length of the cord of the circular segment:

$$l_a(\alpha) = \frac{l_p(\alpha)}{\alpha} (\sin(\alpha)) \quad (2.4)$$

The pouch length l_p is the length of the arc of the circular segment:

$$l_p(\alpha) = 2r(\alpha)\alpha = \sqrt{\frac{2A\alpha^2}{\alpha - \sin(\alpha) \cos(\alpha)}} \quad (2.5)$$

The radius $r(\alpha)$ which defines the circular segment is:

$$r(\alpha) = \sqrt{\frac{A}{2\alpha - 2\sin(\alpha) \cos(\alpha)}} \quad (2.6)$$

The effective electrode length l_e follows from the portion of the actuator which is not zipped:

$$l_e(\alpha) = L_p - l_p(\alpha) \quad (2.7)$$

The height of the actuator $h_s(\alpha)$ is nominally represented by twice the saggita of the circular segment:

$$h_s(\alpha) = \frac{l_p(\alpha)}{\alpha} (1 - \cos(\alpha)) \quad (2.8)$$

Alternatively, the actuator's height can be represented by the average height $h_c(\alpha)$, which is defined by the centroid of the circular segment:

$$h_c(\alpha) = \frac{2}{l_a(\alpha)} \int_{-\alpha}^{\alpha} r(\alpha) (\cos(\theta) - \cos(\alpha)) r(\alpha) \cos(\theta) d\theta \quad (2.9)$$

$$h_c(\alpha) = \frac{2r(\alpha)^2}{l_a(\alpha)} (\alpha - \sin(\alpha) \cos(\alpha)) \quad (2.10)$$

This representation is mainly useful for calculating a more accurate representation of viscous damping due to fluid flow using a Poiseuille flow model as done in [126]. However, it could also be used to more accurately represent the displacement of an actuator operating under large compressive loads.

Some of the above equations directly relate to a constant parameter at one extreme of the actuator's motion, which we will refer to as L_p , L_a , H_s , and L_e . Most of these parameters are defined relative to the resting state of the actuator α_0 , whereas L_e is determined by optimal filling of the actuator as discussed in [74].

$$L_p = l_p(\alpha_0)$$

$$L_a = l_a(\alpha_0)$$

$$H_s = h_s(\alpha_0)$$

(2.11)

$$L_e = l_e(\alpha_f) = L_p - \sqrt{\pi A}$$

It is advantageous to build compound actuators with multiple repeating pouches. Let N_x be the number of pouches repeated in the x direction, and N_z be the number of pouches repeated in the z direction. For contractile actuators (acting in the x direction), N_x will scale the displacement, and N_z will scale the force output. For expanding actuators (acting in the z direction), the inverse is true. N_x will scale the force output, and N_z will scale the displacement. As discussed above, the Peano-HASELS used predominantly in this work are expanding actuators with $N_x = 2$ and $N_z = 12$. However, our derivation of dynamics is generic to this.

With these relations in place, the displacements of the actuator in the x and z directions can be written entirely in terms of α and constant parameters:

$$x(\alpha) = (L_a - l_a(\alpha) - l_e(\alpha)) \cdot N_x \quad (2.12)$$

$$z(\alpha) = (h_s(\alpha) - H_s) \cdot N_z \quad (2.13)$$

Consequently, so long as these functions are bijective, it is also possible to find complementary expressions $\alpha(x)$ or $\alpha(z)$, and therefore in principle use any of these variables as the generalized coordinate for the actuator. For example, $C(\alpha)$ can be rewritten as $C(z)$, and so forth. However, we will refrain from finding an explicit expression for $\alpha(z)$ as it gets a bit complex, is not generalizable, and ultimately proves unnecessary for the conclusions we wish to draw. The expression for $\alpha(x)$ is given indirectly in [74]. It is important to note that the range of motion is still constrained to $\alpha \in [\alpha_0, 2\pi]$, and consequently $z \in [0, z(2\pi)]$.

2.4 Lagrangian Formulation

We use the Lagrange formulation to derive the equations of motion for an expanding Peano-HASEL actuator. This formulation is typically approached from the context of mechanical systems and generalized displacements. However, it can be elegantly extended to electrical and electromechanical systems [22,67,114], as energy is the common currency between the two domains. It should be noted that [126] also uses a Lagrangian formulation to derive the dynamics of a Peano-HASEL actuator. However, that work is done for contractile, rather than expanding actuators, and limits its analysis to a single generalized coordinate α . We consider the derivation of the dynamics of an expanding actuator and examine both geometric and electrical generalized coordinates.

The generalized coordinates λ_i we select for this analysis are: displacement $\lambda_1 = z$ and charge $\lambda_2 = q$. As shown in figure 2.4, the actuator itself is considered to be massless, and the only massive element is an external mass which rests on top of the actuator. Viscoelastic effects of the actuator materials and viscous effects of the fluid flow are represented by a generic spring (constant K) and damper (constant B) respectively. A more physical definition of the viscous damping term based on Poiseuille flow between parallel plates can be found in [126]. A more accurate definition of viscoelastic material effects would add a substantial degree of complexity to the model. We find a generic spring and damper sufficient for this analysis.




Figure 2.4: A model of HASEL actuator deformation dynamics.

Electrical dissipation is given by $R = R_{ser}$, the total resistance of the actuator's electrodes. As drawn in figure 2.4, R is split between the two resistors. Our model also includes an external voltage source $v(t)$ and an external force $f(t)$ which perform non-conservative work on the system. They are represented as virtual work terms $\delta\mathcal{W}_{nc}$ acting in the direction of virtual displacements $\delta\lambda_i$, which correspond to the generalized coordinates λ_i , and define the respective virtual forces Q_i :

$$\delta\mathcal{W}_{nc} = v\delta q - f\delta z \quad (2.14)$$

The Lagrangian \mathcal{L} is the difference between the kinetic co-energy \mathcal{T}^* and the potential energy \mathcal{V} of the system. The potential energy consists of mechanical potential energy \mathcal{V}_m , viscoelastic strain energy stored within the actuator's materials \mathcal{V}_v , and electrical energy \mathcal{V}_e stored on the plates of the capacitor:

$$\mathcal{L} = \mathcal{T}^* - \mathcal{V}_m - \mathcal{V}_v - \mathcal{V}_e \quad (2.15)$$

The kinetic co-energy² \mathcal{T}^* of the system is the kinetic co-energy of the mass M , given by:

$$\mathcal{T}^* = \frac{1}{2}M\dot{z}^2 \quad (2.16)$$

The mechanical potential energy \mathcal{V}_m of the system is the potential energy of the mass M displaced by z :

$$\mathcal{V}_m = Mgz \quad (2.17)$$

² You may be more familiar describing this term as kinetic energy. These ideas are identical in the context of Newtonian mechanics. The difference becomes important at relativistic speeds. Please see [114] for a more thorough discussion.

The viscoelastic strain energy V_v is represented by a spring with constant K :

$$\mathcal{V}_v = \frac{1}{2}Kz^2 \quad (2.18)$$

Finally, the electrical potential energy V_e is the energy stored on the plates of the capacitor:

$$\mathcal{V}_e = \frac{q^2}{2C(z)} \quad (2.19)$$

The Lagrangian is therefore:

$$\mathcal{L} = \frac{1}{2}M\dot{z}^2 - Mgz - \frac{1}{2}Kz^2 - \frac{q^2}{2C(z)} \quad (2.20)$$

We also wish to consider the dissipative effects \mathcal{D} within the actuator, which consist of viscous damping \mathcal{D}_v and resistive dissipation \mathcal{D}_e :

$$\mathcal{D} = \mathcal{D}_v + \mathcal{D}_e \quad (2.21)$$

Viscous dissipation \mathcal{D}_v is represented by a damper with constant B :

$$\mathcal{D}_v = \frac{1}{2}B\dot{z}^2 \quad (2.22)$$

Resistive dissipation \mathcal{D}_e is represented by a resistor with constant $R = R_{ser}$:

$$\mathcal{D}_e = \frac{1}{2}R\dot{q}^2 \quad (2.23)$$

The equations of motion can now be found using Lagrange's equation:

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{\lambda}_i} \right) - \frac{\partial \mathcal{L}}{\partial \lambda_i} = \mathcal{Q}_i - \frac{\partial \mathcal{D}}{\partial \dot{\lambda}_i} \quad (2.24)$$

Differentiation of the respective terms gives:

$$\frac{\partial \mathcal{L}}{\partial z} = -Mg - Kz + \frac{C'(z)q^2}{2C(z)^2} \quad (2.25)$$

$$\frac{\partial \mathcal{L}}{\partial \dot{z}} = M\dot{z} \quad (2.26)$$

$$\frac{\partial \mathcal{L}}{\partial q} = -C(z)^{-1}q \quad (2.27)$$

$$\frac{\partial \mathcal{L}}{\partial \dot{q}} = 0 \quad (2.28)$$

$$\frac{\partial \mathcal{D}}{\partial \dot{z}} = B\dot{z} \quad (2.29)$$

$$\frac{\partial \mathcal{D}}{\partial \dot{q}} = R\dot{q} \quad (2.30)$$

which result in the following equations of motion:

$$M\ddot{z} + B\dot{z} + Kz = \frac{q^2 C'(z)}{2C(z)^2} - Mg - f \quad (2.31)$$

$$q = C(z)(v - R\dot{q}) \quad (2.32)$$

Without loss of generality, we can consider the voltage source to act directly on the capacitive element of the electrostatic actuator instead of the actuator itself, in essence bypassing R_{ser} and setting $R = 0$. The resistive dissipation term such an assumption neglects will reappear in a more nuanced manner during our discussion of the driver circuit and its dynamics in section 2.6. To draw a connection the terminology of that section, equation 2.32 is in terms of v_{out} , and equation 2.33 is in terms of v_{load} . The inclusion of resistive dissipation is not necessary for the validity of the remaining analyses in this section, so we can simplify equation 2.32 to:

$$q = C(z)v \quad (2.33)$$

Equations 2.31 and 2.33 can be consolidated into a single equation of motion:

$$M\ddot{z} + B\dot{z} + Kz = \frac{1}{2}C'(z)v^2 - Mg - f \quad (2.34)$$

or, if we nonetheless wish to include the resistive dissipation in this model, the full version of equation 2.34 can be rewritten using 2.32 as:

$$M\ddot{z} + B\dot{z} + Kz = \frac{1}{2}C'(z)(v^2 - 2R\dot{q}v + R^2\dot{q}^2) - Mg - f \quad (2.35)$$

Note that $C'(z)$ is the derivative of the capacitance with respect to z , and only depends on z :

$$C'(z) = \frac{\partial}{\partial z} [C(z)] = \frac{d}{dz} [C(z)] \quad (2.36)$$

2.4.1 Insights Arising from Equilibrium

At equilibrium, $\ddot{z} = \dot{z} = 0 = \dot{q}$, which reduces either formulation (2.34, 2.35) of the equation of motion to:

$$Kz = \frac{1}{2}C'(z)v^2 - Mg - f \quad (2.37)$$

which leads to some interesting conclusions:

1. If we neglect the mass and the elastic strain of the material, and assume that force output is constant over the actuator's stroke from rest to a displacement \bar{z} , the force output is proportional to the square of the voltage. Such a formulation essentially relegates the work done by the actuator on the mass to the external force. While not directly used in this work, this result is consistent with the common understanding in the field:

$$f = \frac{1}{2}C'(\bar{z})v^2 \quad (2.38)$$

Had we been considering the contracting actuator case, and using a $C(x)$ representation, this equation would exactly match the results from [74, 126, 158].

2. With no external applied force and a constant mass, then equilibrium displacement of the actuator can be set purely by an applied voltage \bar{v} assuming $C'(z)$ is monotonic. This means that by controlling the voltage, we have a proxy for controlling the displacement of the actuator, and suggests voltage regulation as a useful control strategy for ESAs at equilibrium or under quasi-static conditions.

$$z = \frac{C'(z)}{2K}v^2 - \frac{Mg}{K} \quad (2.39)$$

This equation is affine, and will evaluate to negative values of z . Because z cannot be negative by definition, this suggests there is a minimum voltage threshold that must be reached before the actuator begins to move. This is consistent with the literature, and clearly visible in figures 3.12, 3.13, and 3.15. The minimum pull-in voltage v_0 can be found at the point $z = 0$ where the voltage is sufficient to overcome the mass:

$$v_0 = \sqrt{\frac{2Mg}{C'(z)}} \quad (2.40)$$

3. Under a regulated voltage \bar{v} , if $C'(z)$ is known, then the magnitude of the externally applied force can be inferred by measuring the displacement of the actuator. This is a useful property for force sensing, and the design of a force estimator $E_f(\bar{v}, z)$ is explored further in 4.7.

$$f = \frac{1}{2}C'(z)\bar{v}^2 - Kz - Mg \quad (2.41)$$

Equation 2.41 is valid for $z \in [0, z(2\pi)]$ and $\forall \bar{v} \geq 0$.

2.5 ESA Driver Topologies

While ESAs have great potential for use in soft robots, the design of their drive circuitry presents significant challenges. To drive them, it is necessary to control the strength of the electric field applied across their electrodes. However, they operate at high voltages ($10^2 - 10^4$ V) and low currents ($< 10^{-1}$ A) – a niche where few technologies exist. The simplest method to drive an ESA is by using a low voltage control signal to set the output voltage of a high voltage amplifier. Commercially available HV switched-mode power supplies (HV-SMPS) with an amplifier mode, such as a Trek 50/12 (Advanced Energy) are commonly used in laboratory settings [1, 75, 158], whereas untethered solutions often call for custom implementations [99]. While this approach is well suited for the evaluation of individual actuators, it introduces limits on performance and does not scale well. When choosing a HV-SMPS, engineers must make tradeoffs between voltage range, current output, and system size. The aforementioned Trek 50/12 can provide 12 mA at 50 kV, but is the size of a refrigerator. Smaller amplifiers sacrifice voltage range, current capacity, or both.

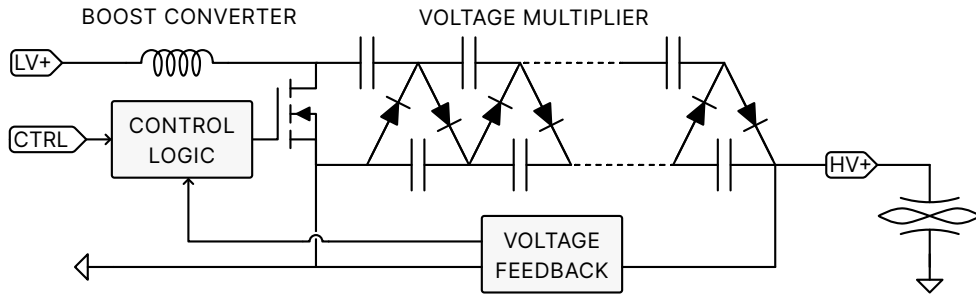


Figure 2.5: The internal topology of the typical HV switched-mode power supply (HV-SMPS).

This is particularly problematic because as predominantly capacitive loads, ESA power draw is highly variable. When holding a constant voltage, ESAs draw minimal power, mainly to counteract leakage through the actuator. This is highly dependent on actuator geometry, materials, and design, and specifically on R_{par} , but is typically in the range of 10^{-2} W to 10^{-1} W. However, when quickly changing voltage, the peak power draw of ESAs can be several orders of magnitude higher, in the range of 10^0 W to 10^3 W. For example, the gripper presented in [158] draws 0.65 W in a holding state, but jumps 50-fold to a peak of 37 W to achieve its fastest actuation speeds. Unpublished laboratory results show this difference can be substantially larger still, and although short, these peaks are necessary for maximal actuator performance. When driven directly from an amplifier, voltage regulation performance is limited by the slew rate and peak current output of the amplifier, forcing a tradeoff between actuator performance and unused power supply capacity – a precious resource. Commercially available amplifiers with appropriate voltage and current output are several orders of magnitude bulkier and more expensive than the actuators themselves, making this 1:1 pairing impractical for applications requiring the coordinated actuation of multiple ESAs. For example, an array of 100 HASEL actuators would cost under 100 \$ in raw materials and weigh about 1.1 kg, while the required high power driving electronics to independently address each actuator using the direct-drive approach (EMCO A Series, XP Power) would cost over \$25,000 and weigh over 7 kg. This estimate uses a power supply which can only reach 6 kV. Some actuator types would require even larger, heavier, and more expensive supplies to achieve maximal performance at 8 to 12 kV.

The implementation of a parallel switching amplifiers, an approach first demonstrated for ESAs in [99], provides an elegant solution to scaling. In this method, a HV SMPS is set to provide a fixed output voltage and charge a HV capacitor bank. Together, they provide a HV supply rail capable of delivering the large current spikes drawn by the capacitive ESAs without requiring a SMPS with the capacity to supply these peaks directly.

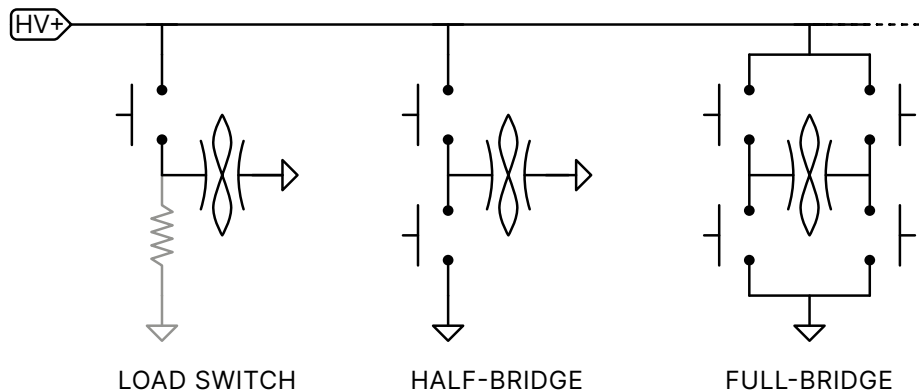


Figure 2.6: Several inherently parallelizable bridge-style switching amplifier topologies suitable for driving ESAs.

Electrically controlled switches in one of several commonly used topologies like the ones shown in figure 2.6 are used to rapidly toggle connections between the ESA and the supply and ground rails, causing current to flow into or out of the actuator. Given sufficiently fast switching, the electrical dynamics of the ESA act as a low-pass filter, and the actuator responds smoothly to this input. The simplest version, a load switch, only requires a single HV-capable switch, but limits the performance of the actuator as discharge rate is limited by the RC time constant of the actuator and an optional low-side leakage resistor. Maximizing performance is a question of balancing leakage-opposed charging with leakage-driven discharge. As discussed in [99], a much better approach is to add a second low-side switch and thus gain independent control of both charge and discharge rates. This approach is demonstrably parallel [99, 138], and more complex switch topologies, such as the full-bridge (H-bridge), can be used to reverse polarity and counteract the accumulation of space charges in the actuator, a phenomenon known to reduce actuator performance over repeated charge/discharge cycles [33, 51, 158]. Alternative topologies could be used to further improve efficiency, or recover some power from switching the actuator [17, 60]. We will be considering a half-bridge topology for this work, which generalizes to the more complex topologies.

2.6 Half Bridge Driver Model

Consider the equation of motion (2.33) with respect to $v = v_{load}$ for the generalized electrical coordinate q of the HASEL, along with its derivative:

$$q = C(z)v \quad (2.42)$$

$$\dot{q} = \frac{d}{dt} [C(z)]v + C(z)\dot{v} \quad (2.43)$$

We understand the dynamics of the generalized coordinate governing the mechanical dynamics (\dot{z}) to be several orders of magnitude slower than the dynamics of the generalized coordinate governing the electrical dynamics (\dot{q}). This quasistatic assumption allows us to approximate the capacitance as constant at the relevant timescales, and thus decouple the electrical dynamics. With $C(z) \simeq C_{load}$, and therefore $\dot{C} \simeq 0$, we can write the electrical dynamics directly in terms of the actuator voltage v and the net actuator current:

$$\dot{q} = C_{load}\dot{v} = i_{net} \quad (2.44)$$

To facilitate analysis, we will consider the net actuator current i_{net} as the sum of three components: the charging current i_{chg} , the discharge or drain current i_{drn} , and the leakage current i_{leak} :

$$\dot{q} = C_{load}\dot{v} = i_{chg} - i_{drn} - i_{leak} \quad (2.45)$$

Where these currents show up may not immediately be obvious, and requires a closer look at the expanded electrical model of a half-bridge-driven ESA shown in figure 2.7.




Figure 2.7: Expanded equivalent circuit of a half-bridge-driven ESA.

Here, R_{on} is the on-state impedance of each switch. R_{off} is the off-state impedance of each switch. R_{ext} is an external leakage resistance, which represents the collective action of leakage currents from neither the bridge circuit, nor the actuator itself. Most significantly, this would include the leakage through the sensing circuit discussed in section 3.2. R_{ser} and R_{par} are the series and parallel resistances of the actuator, respectively. C_{load} is the load capacitance of the actuator, which as discussed above, is assumed to be constant for the purposes of this analysis.

v_{out} is the measurable output voltage of the bridge. v_{load} is the voltage across the actuator electrodes, which cannot be measured directly, but is the true voltage across the actuator (unless specified otherwise, v is taken to be v_{load}). However, if we assume that the voltage drop across the actuator leads and electrode material is small relative to everything else, that is to say $R_{ser} \ll R_{par}$, $R_{ser} \ll R_{ext}$ and $R_{ser} \ll R_{off}$, then we can approximate both voltages as $v_{load} \simeq v_{out}$ when implementing hardware. This proves to be a reasonable assumption in practice.

Switches are controlled with digital signals from a microcontroller, typically using pulse width modulation (PWM). For the half-bridge, these input signals are ρ_{chg} and ρ_{drn} , which represent the

duty fractions of the high-side (charge) and low-side (drain) switches, respectively. The duty fraction $\rho_i \in [0, 1]$ is the fraction of time spent in the on-state, and is a direct analog for duty cycle, which is typically specified as the percentage of time spent in the on-state. For example, a 15% duty cycle on the high-side switch corresponds to $\rho_{chg} = 0.15$. This representation is chosen purely for mathematical convenience. We design the system under the assumption that the charge and drain switches should never be on at the same time, as this would short the rail and simply waste power.

To build a model for this system, let us consider its two fundamental modes of operation. The first only considers the leakage current in the absence of inputs, and therefore charge and drain currents, and is discussed in section 2.6.1. The second considers the charge and drain currents in the absence of leakage, and is discussed in section 2.6.2. A good model of the system is then given by the sum of these two cases, as discussed in section 2.6.3.

2.6.1 Zero-Input Dynamics

First, let us isolate the contributions of the leakage current i_{leak} . This is the zero input case, where both switches are turned off. In this case, the actuator is effectively disconnected from the bridge circuit, and the only dynamics correspond to the leakage current i_{leak} .

$$\dot{v} C_{load} = -i_{leak} \quad (2.46)$$

The equivalent sub-circuit is shown in figure 2.8.




Figure 2.8: Electrodynamic of the zero-input case.

We can consolidate the various leakage paths into a single leakage resistance R_{leak} , according to:

$$R_{leak} = R_{par} || (R_{ser} + R_{ext}) = \frac{R_{par} (R_{ser} + R_{ext})}{R_{par} + R_{ser} + R_{ext}} \quad (2.47)$$

Note that this analysis neglects the leakage through R_{off} of each switch. This proves to be a reasonable assumption, so long as $R_{off} \gg R_{leak}$, which is typically the case. The dynamics will be more complex if particularly leaky switches are used. This is not considered in this work.

Following these simplifications, the dynamics of the zero-input case can be represented by the linear dynamics of a RC discharge circuit:

$$\dot{v} C_{load} = -i_{leak} = \frac{-v}{R_{leak}} \quad (2.48)$$

2.6.2 Zero-Leak Dynamics

Next, let us isolate the contributions of the charge and discharge currents i_{chg} and i_{drn} . This is analogous to a zero-state response, although would be a slight stretch of the terminology to call it that. We assume R_{par} , and R_{ext} are infinite, and thus $i_{leak} = 0$. As in section 2.6.1, we also assume R_{off} is infinite.

$$\dot{v} C_{load} = i_{chg} - i_{drn} \quad (2.49)$$

If we assume that only one switch is ever on, we can consolidate the resistances into a single variable according to: $R_{load} = R_{on} + R_{ser}$, as shown in figure 2.9.




Figure 2.9: Electrostatics of the zero-leak case.

The nonlinear dynamics of this system are given by:

$$C_{load}\dot{v} = i_{chg} - i_{drn} = \rho_{chg} \frac{v_+ - v}{R_{load}} - \rho_{drn} \frac{v}{R_{load}} \quad (2.50)$$

Nonlinear dynamics are comparably difficult to work with, especially in the context of controls design. If we are willing to accept some loss in fidelity close to the rails, we can linearize this system around a nominal operating point, \bar{v} .

Consider the ~~twitter logo~~ state vector \mathbf{z} , which is represented as a perturbation $\tilde{\mathbf{z}}$ from the nominal operating point $\bar{\mathbf{z}}$, which should be an equilibrium point of the system:

$$\mathbf{z} = \bar{\mathbf{z}} + \tilde{\mathbf{z}} = \begin{bmatrix} v \\ \rho_{chg} \\ \rho_{drn} \end{bmatrix} = \begin{bmatrix} \bar{v} + \tilde{v} \\ \bar{\rho}_{chg} + \tilde{\rho}_{chg} \\ \bar{\rho}_{drn} + \tilde{\rho}_{drn} \end{bmatrix} \quad (2.51)$$

Because we have assumed there is no leakage (and thus no zero-input dynamics), any voltage is an equilibrium point when both switches are turned off. We choose the center voltage $\bar{v} = \frac{1}{2}v_{rail}$ for simplicity and greatest accuracy across the entire operating range. Thus, the equilibrium point $\bar{\mathbf{z}}$ is given by:

$$\bar{\mathbf{z}} = \begin{bmatrix} \bar{v} \\ \bar{\rho}_{chg} \\ \bar{\rho}_{drn} \end{bmatrix} = \begin{bmatrix} \frac{1}{2}v_{rail} \\ 0 \\ 0 \end{bmatrix} \quad (2.52)$$

and the perturbation $\tilde{\mathbf{x}}$ is given by:

$$\tilde{\mathbf{x}} = \begin{bmatrix} \tilde{v} \\ \tilde{\rho}_{chg} \\ \tilde{\rho}_{drn} \end{bmatrix} = \begin{bmatrix} v - \frac{1}{2}v_{rail} \\ \rho_{chg} \\ \rho_{drn} \end{bmatrix} \quad (2.53)$$

We linearize using a Taylor series expansion of $\dot{v} = f(v, \rho_{chg}, \rho_{drn}) = f(\mathbf{x})$ around $\bar{\mathbf{x}}$:

$$f(\tilde{\mathbf{x}}) \simeq f(\bar{\mathbf{x}}) + \sum_i \left. \frac{\partial f}{\partial x_i} \right|_{\bar{\mathbf{x}}} \cdot \tilde{x}_i \quad (2.54)$$

$$\begin{aligned} \left. \frac{\partial f}{\partial v} \right|_{\bar{\mathbf{x}}} &= \frac{\rho_{drn} - \rho_{chg}}{R_{load}C_{load}} & \left. \frac{\partial f}{\partial v} \right|_{\bar{\mathbf{x}}} &= 0 \\ \left. \frac{\partial f}{\partial \rho_{chg}} \right|_{\bar{\mathbf{x}}} &= \frac{v_{rail} - v}{R_{load}C_{load}} & \left. \frac{\partial f}{\partial \rho_{chg}} \right|_{\bar{\mathbf{x}}} &= \frac{v_{rail}}{2R_{load}C_{load}} \\ \left. \frac{\partial f}{\partial \rho_{drn}} \right|_{\bar{\mathbf{x}}} &= \frac{-v}{R_{load}C_{load}} & \left. \frac{\partial f}{\partial \rho_{drn}} \right|_{\bar{\mathbf{x}}} &= \frac{-v_{rail}}{2R_{load}C_{load}} \end{aligned}$$

Thus, the linearization around $\bar{\mathbf{x}}$ is given by:

$$f(\tilde{\mathbf{x}}) \simeq (\rho_{chg} - \rho_{drn}) \frac{v_{rail}}{2R_{load}C_{load}} \quad (2.55)$$

Since these dynamics are independent of the actuator voltage v , and because:

$$\dot{v} = \dot{\tilde{v}} + \dot{\tilde{v}} = 0 + \dot{\tilde{v}} = \dot{\tilde{v}}, \quad (2.56)$$

we note that $f(\tilde{\mathbf{x}}) = \dot{\tilde{v}} = \dot{v}$, and thus we can write the linearized model for the zero-leak actuator voltage as:

$$\dot{v} = (\rho_{chg} - \rho_{drn}) \frac{v_{rail}}{2R_{load}C_{load}} \quad (2.57)$$

or equivalently:

$$\dot{v} C_{load} = i_{chg} - i_{drn} = (\rho_{chg} - \rho_{drn}) \frac{v_{rail}}{2R_{load}} \quad (2.58)$$

2.6.3 Linearized SISO Model

We combine the zero-input and zero-leak dynamics and thus the contribution of all three currents to get the complete, linearized model of the half bridge electrodynamics:

$$\dot{v} C_{load} = i_{chg} - i_{drn} - i_{leak} = (\rho_{chg} - \rho_{drn}) \frac{v_{rail}}{2R_{load}} - \frac{v}{R_{leak}} \quad (2.59)$$

Having designed the system assuming that the charge and drain switches should never be simultaneously closed, we can consolidate ρ_{chg} and ρ_{drn} into a single variable ρ according to:

$$\rho_{chg} = \begin{cases} \rho & \text{if } \rho > 0 \\ 0 & \text{otherwise} \end{cases} \quad \rho_{drn} = \begin{cases} -\rho & \text{if } \rho < 0 \\ 0 & \text{otherwise} \end{cases}$$

Note that this form assumes symmetric rates³ of charge and discharge. This proves to be an adequate assumption for this work, but may not hold in general. This consolidation results in a single-input, single-output (SISO) system with linear dynamics:

$$\dot{v} C_{load} = \rho \frac{v_{rail}}{2R_{load}} - v \frac{1}{R_{leak}} \quad (2.60)$$

Finally, we solve for \dot{v} and group all parameters into a , the inverse of the leakage time constant, and b , the maximum expected slew rate of the bridge:

$$a = \tau_{leak}^{-1} = \frac{1}{R_{leak} C_{load}}$$

$$b = \dot{v}_{max} = \frac{i_{max}}{C_{load}} = \frac{v_{rail}}{2R_{load} C_{load}}$$

This results in the final form of our model:

$$\dot{v} = b\rho - av \quad (2.61)$$

A linear SISO model is an excellent starting point for closed loop controls. Furthermore, the parametrization gives insights and lets us make predictions about actuator performance directly from parameters which we can tune during the design process.

³ To model asymmetric rates, a corrective factor can be added to either the charge or discharge duty cycle.

Chapter 3

Scalable Feedback Control of Electrostatic Actuators

This chapter discusses the fundamental basic building block of scalable feedback control for ESAs: the half bridge driver and voltage sensor. We will discuss the design and implementation of these components, the system identification and control design process, and finally, the performance of the resulting closed loop system. This chapter demonstrates the efficacy and limitations of a voltage control approach, and provides a foundation for the design of the s-tissue hardware platform discussed in chapter 4.

3.1 Optoelectronic Bridge Drivers

The practical implementation of the switching amplifier circuits presented in chapter 2 requires HV switches capable of handling the high voltages required to drive ESAs. These switches must be controlled by low voltage digital signals from a microcontroller, while protecting the microcontroller from high voltages. Furthermore, the switches must be compact, and be manufacturable in large quantities. An optoelectronic approach, first presented in [100], provides an effective solution.

3.1.1 Optocoupler Fundamentals

Optocouplers are devices which transfer signals as light, and in doing so, provide a means to link two circuits without a direct electrical connection. They are typically used to provide electrically isolated interfaces for high-voltage equipment, or on communication lines between systems

with different ground potentials.




Figure 3.1: The basic principles of photoconductive power optocoupler operation.

Generally, an optocoupler consists of a photodiode or phototransistor optically coupled to an LED. In our implementation shown in figure 3.1, a high voltage photodiode (OZ100SG, Voltage Multipliers, Inc.) is used in photoconductive mode. As this device is used directly as a load switch, we will call it a power optocoupler. When the LED is off, the reverse-biased photodiode blocks current flow. When illuminated, the photoelectric effect causes current to flow in response to absorbed photons. Thus, the photodiode conducts current in proportion to the radiant intensity of light emitted by the LED, which in turn is proportional to the LED input current. Although this relatively simple circuit allows direct control of high voltage current from low voltage signals, its implementation presents several practical challenges: driver circuitry, alignment, isolation, and thermals.

3.1.2 Optocoupler Drive Circuitry

Driving the optocoupler is a matter of driving a high power infrared LED (L1IZ-094000, Lumileds) Any number of driver circuits can be used, the simplest being a low-side MOSFET and a current limiting resistor shown in figure 3.2. We pair a $7.5\text{ k}\Omega$ resistor with our MOSFET (NDC7002N, onsemi) for the results in chapters 3 and 4, but use variable resistors paired with a different MOSFET (PMV20EN, Nexperia) for the results in chapter 5. Experiments in this latter set control for LED current. The resistor sets the current flow through the LED, and the

MOSFET is used to switch the current on and off. The average brightness of the LED is the in effect a function of the average current. Transient voltage suppression (TVS) diodes (SP4021-01FTG, Littelfuse) are included to protect the low-side circuitry from any transients or charge build-up on the high-side.

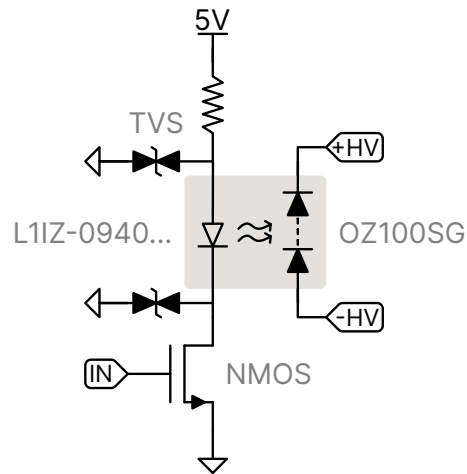


Figure 3.2: Our implementation of a photoconductive power optocoupler and accompanying driver circuit.

It should be noted that as the LED heats up, its impedance decreases, and it draws more current. This will effectively increase the gain of the optocoupler, and can lead to thermal runaway and failure as discussed further in section 3.1.5. For higher thermal stability, a constant current source such as a current mirror can be used to drive the LED.

3.1.3 Optocoupler Alignment

The first challenge is that of aligning the optical centers of the LED and photodiode during assembly. The magnitude of the photoelectric effect is proportional to the radiant intensity of the light, which falls off with the square of distance. Furthermore, the 150° beam angle of the LEDs must be aligned with the photodiode. In practice however, the radiant intensity starts to fall off at a beam angle of about $\pm 30^\circ$. Together, these two factors make this system particularly sensitive

to changes in alignment. Small variations of these parameters can introduce large variations in the performance of individual optocouplers.

These facts are mitigated by milling a small $0.055" \times 0.105"$ slot into one face of each diode's transparent epoxy packaging, which fits over the surface mounted LED to correctly position and align the photodiode during assembly. Slots are milled to a precise depth of $0.015" \pm 0.005"$ using a manual mill (Grizzly Industrial), but are visually aligned to the optical center of the photodiode, as the axial positioning of the diode body within its epoxy substrate was observed to vary substantially between individual components received from the manufacturer, and proved challenging to measure both quickly and precisely by other means. There is no such geometric variability in the LEDs.

3.1.4 Optocoupler Isolation

The electrical potential between the LED and the high side of the photodiode will regularly reach 8kV, with mere millimeters of separation between them.. Any discharge between these two points will destroy the low voltage circuitry of optocoupler, and potentially damage any connected circuits such as microcontrollers. Individual optocouplers must also be shielded from optical interference, both from neighboring optocouplers and from the environment.

One approach is to mount an opaque, FDM 3D-printed (eg. 1.75mm PETG Filament, Overture; MK3S, Prusa) plastic shell to the PCB over each optocoupler to provide optical isolation, while simultaneously serving as a reservoir to fill with a polymerizing dielectric gel (Sylgard 527, Dow Chemical) to provide electrical isolation. A flexible FDM 3D-printed (eg. 1.75mm TPU95A Filament, Overture; MK3S, Prusa) gasket helps form a seal between the shell and the PCB. This is the approach shown in figure 3.3.

The other approach is to simultaneously achieve optical and electrical isolation by encapsulating the high-voltage portion of the PCB in an opaque insulating epoxy (Quickset White, Pratley). This requires the construction of suitable moulds to hold the PCB and epoxy in place during curing, and all but impossible to inspect and repair. However, it is more compact.




Figure 3.3: Quasi-cross-sectional view of a high-performance high-voltage switching optocoupler. For simplicity, some features which are not in-plane are shown as such. Approximate scale.

Regardless of approach, where possible, high- and low-voltage components are placed on opposite sides of the PCB. Transient voltage suppression (TVS) diodes are added to the most exposed and vulnerable parts of the circuitry to protect low-voltage components by clamping any transients or accumulated charges to safe voltages.

To test insulation, all ground and low voltage connections were grounded, and the encapsulated driver was tightly wrapped in a grounded wire mesh¹. The high voltage portions of the circuit were then brought to rail voltage (8 kV), and any insulation failures were immediately reflected in the load current. External failures could furthermore be located by corona discharge.

3.1.5 Optocoupler Thermal Management

Finally is the challenge of thermal management. Both elements of the optocoupler will generate substantial heat. The LED in the form of a large current draw, and the photodiode because of resistive losses during operation. As discussed in section 5.7.3, these devices are not particularly efficient, and generate substantial heat during operation.

If left unchecked, heat will cause catastrophic failure to the device. LEDs typically exhibit a positive temperature coefficient², meaning that their current draw increases with temperature.

¹ Although at inception, this test was somewhat improvised due to the COVID lockdown – the wire mesh was cut from a screen door on my house – it proved to be very effective and particularly useful for informing design iterations. The grounded mesh closely and uniformly hugs the surface of the component under test, and if there is any weak point in the insulation, the location of breakdown can clearly be seen.

² Such LEDs are typically listed as having a negative thermal coefficient with respect to forward voltage.

Increased current causes increased luminant intensity, which in turn causes increased conduction through the photodiode, and consequently, additional heat generation. If the device gets hot enough, this will create a positive feedback loop that can lead to thermal runaway, catastrophic failure, and permanent destruction of the optocoupler. However, even smaller amounts of unmanaged heat can be dangerous in the long term through the mechanism of thermal degradation and eventual dielectric breakdown of the insulating epoxy.

One approach is to mount a heatsink on the PCB opposite the optocoupler like shown in figure 3.3. This is easier when using the dielectric gel approach, which only takes up one side of the PCB. The other option is to limit the current, and therefore the brightness of the LED. This will both reduce the heat generated by the LED and by the photodiode, but at the cost of limiting the maximum achievable slew rate of the actuator. Even better would be to use a constant-current LED driver, but this increases the complexity of the drive circuitry, has not been attempted, and is rendered unnecessary by the results in chapter 5. Due to its compact nature, we use the epoxy approach for the s-tissue, and the dielectric gel approach shown in figure 3.3 for designs aiming to maximize performance.

3.2 High Voltage Sensing Circuitry




Figure 3.4: A simple and effective HV sensing circuit for use with a microcontroller.

To measure kilovolt scale signals with a microcontroller, we use the design shown in figure 3.4. A $1\text{ G}\Omega$, high impedance voltage divider steps down the voltage by a ratio of 10000:3. A voltage follower operational amplifier (TLV271) provides a high input impedance. The output of

this buffer circuit is then measured with the microcontroller’s internal ADC (Teensy 4.0, PJRC). The microcontroller filters the data, and can then use this information to provide closed loop control of the actuator. Improvements to this design are discussed in section 5.5.

3.3 Half Bridge System Identification

A half bridge driver for ESAs can be constructed using the fundamental components discussed so far. The resulting circuit is shown in figure 3.5, and represents a fundamental building block of scalable feedback control for ESAs.

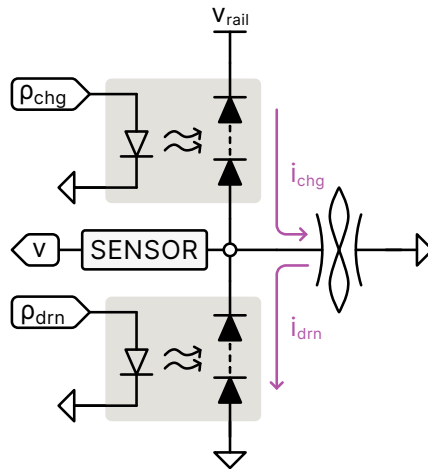


Figure 3.5: The complete half bridge driver implementation.

3.3.1 Transfer Function

A transfer function representation of our system is found using the Laplace transform, which is defined as:

$$\tilde{f}(s) = \mathcal{L}\{f(t)\} = \int_0^{\infty} f(t)e^{-st} dt \quad (3.1)$$

We convert equation 2.61 to a phase space representation by taking the Laplace transform of both sides:

$$s\tilde{v}(s) = b\tilde{\rho}(s) - a\tilde{v}(s) \quad (3.2)$$

simplifying:

$$\tilde{v}(s + a) = b\tilde{\rho} \quad (3.3)$$

and solving for the transfer function:

$$G(s) = \frac{\tilde{v}(s)}{\tilde{\rho}(s)} = \frac{b}{s + a} \quad (3.4)$$

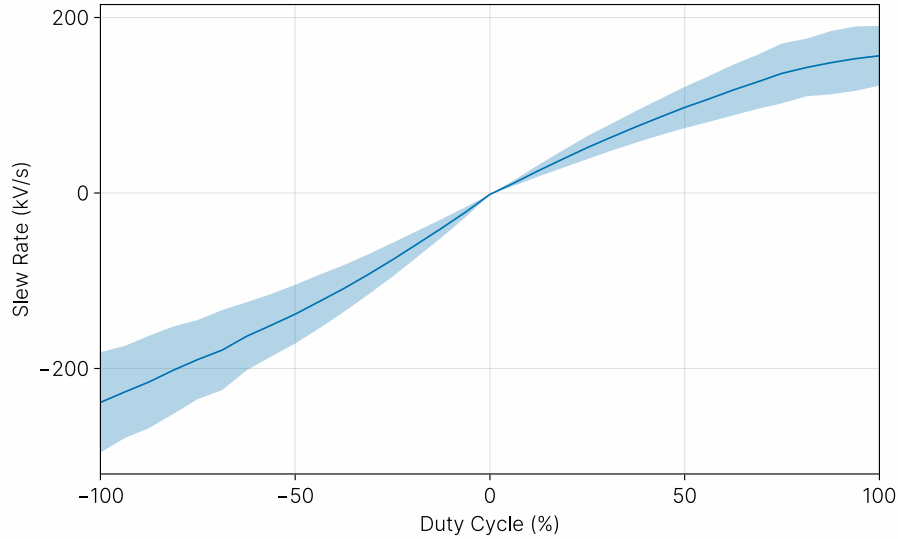


Figure 3.6: Maximum sew rate of our optoelectronic half-bridge as a function of input duty cycle.

As shown in figure 3.6, we can expect the maximum slew rate for the full hardware platform introduced in chapter 4 to be approximately 300 kV/s (0.3 V/ μ s), so we can approximate $b \simeq 300$. Likewise for this hardware, the leakage current is dominated by the 1 G Ω sense resistor, and the actuator capacitance is approximately 500 pF, giving: $a \simeq 2$. Thus, our theoretical transfer function is approximately:

$$G(s) = \frac{300}{s + 2} \quad (3.5)$$

Although this model uses loosely approximated parameter values, we expect that the model resulting from our system identification process will closely resemble this form.

3.4 Frequency Domain System Identification

To obtain a more accurate model of G (specifically $G(j\omega)$) than estimated by equation 3.5, we perform system identification in the frequency domain. The procedure is as follows:

1. Select a sequence of frequencies ω_n , or $f_n = 2\pi\omega_n$ to excite the system which are integer ratios of the sampling frequency f_s , and are roughly logarithmically spaced through the domain of interest. This is to ensure that the signal energy for a given test is contained in a single peak in the discrete frequency domain, and to ensure that the analysis covers the logarithmically spaced domain of interest. If it is not realistic to use integer ratios of the original f_s , the input signal can be zero-padded to artificially increase its effective resolution in phase space. A chirp signal can also be used, as done in [70]. However, this approach provides a comparably poor signal-to-noise ratio, captures digital artifacts, and ultimately leads to less accurate results.

2. Generate a sequence of inputs ρ_n for each ω_i in the set of ω_n .

$$\rho_i(t) = a \sin(\omega_i t) \quad (3.6)$$

The magnitude a of these sample inputs should be adjusted to have the best possible signal-to-noise ratio without saturating the actuator. In practice, this means smaller amplitudes at lower frequencies.

3. Measure actuator voltage response to each input $v_n(t)$ For our results, we use the circuit described in section 3.2 along with the built-in ADC of a Teensy 4.0 microcontroller set to a sample frequency f_s of 1 kHz.

4. Take the Discrete Fourier Transform (DFT) of each v_n and ρ_n using a DFT algorithm such as FFTW [34].

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi jkn/N} \quad (3.7)$$

5. Compute the value of G for each ω_n to build the transfer function data set \hat{G} .

$$\hat{G} = \left\{ \frac{\tilde{v}_i(j\omega_i)}{\tilde{\rho}_i(j\omega_i)} \quad \forall i \right\} \quad (3.8)$$

6. Fit a transfer function $G(s)$ to the data set \hat{G} , for example using the BGFS optimizer from `ControlSystemsIdentification.jl` [16, 20].

The results of this procedure repeated across 100 Peano-HASEL actuators on the s-tissue hardware platform presented in chapter 4 are shown in figure 3.7. Each point in \hat{G} is the result of a single 15 s test at a specific frequency f_n . This analysis gives:

$$G(s) = \frac{434}{s + 11} \quad (3.9)$$

which is plotted as G_{fit} along with G_{est} , the estimate from our model in equation 3.5.

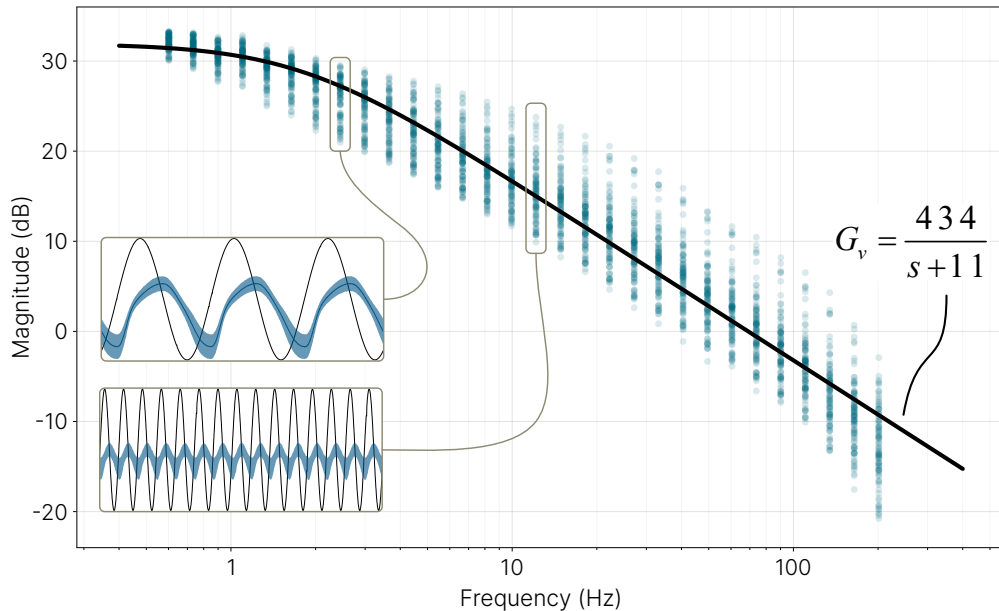


Figure 3.7: Transfer function fit for 100 Peano-HASEL actuators. Inset figures show the input (black) and one standard deviation around the mean of the output (blue) at selected frequencies.

3.5 Loopshaping Control Design

Our approach is to first translate practical performance requirements for the closed loop system into constraints on the open loop transfer function $L = GK$. This approach - called loopshaping - allows us to directly design a K for a given G , and thus "shape" the closed loop dynamics $\frac{L}{1+L}$ to meet our requirements.




Figure 3.8: A block diagram representing a generalized single-input, single-output (SISO) system with feedback control.

Central to this approach is the fact that several important closed loop transfer functions (CLTFs) can be approximated in terms of the open loop transfer function (OLTF) with some assumptions. These include the complementary sensitivity function T :

$$\frac{y}{r} = -\frac{y}{v} = T = \frac{L}{1+L} \quad (3.10)$$

and the sensitivity function S :

$$\frac{e}{r} = S = \frac{1}{1+L} \quad (3.11)$$

At low frequencies, $|L| \gg 1$, thus:

$$S = \frac{1}{1+L} \simeq \frac{1}{L} \quad (3.12)$$

$$T = \frac{L}{1+L} \simeq 1 \quad (3.13)$$

At high frequencies, $|L| \ll 1$, thus:

$$S = \frac{1}{1+L} \simeq 1 \quad (3.14)$$

$$T = \frac{L}{1+L} \simeq L \quad (3.15)$$

3.5.1 Steady-State Error

Given a step input $R(s) = \frac{1}{s}$, by the final value theorem, the steady state error will be:

$$e(\infty) = \lim_{s \rightarrow 0} sE(s) = \lim_{s \rightarrow 0} sS(s)R(s) = \lim_{s \rightarrow 0} S(s) = S(0) \quad (3.16)$$

Thus, to achieve a desired steady-state error $e(\infty) < e_{ss}$, we require that $|S(0)| < e_{ss}$. Since $S(0) \simeq \frac{1}{L(0)}$, this can be rewritten as $|L(0)| > \frac{1}{e_{ss}}$, resulting in an open-loop constraint on gain:

$$|L(0)|_{dB} > 20 \log \left(\frac{1}{e_{ss}} \right) \quad (3.17)$$

It should be noted, however, that the slope of magnitude of L at low frequencies will have a more significant effect on the steady-state error. To track a step input, the slope of $|L|_{dB}$ should be greater than -20 dB/decade in the low frequency limit.

3.5.2 Tracking Error

For tracking performance, we desire a low tracking error gain $S = \frac{e}{r}$ at low frequencies. Thus we require that $|S(s)| < e_{tr}$. Since $S \simeq \frac{1}{L}$ at low frequencies, this can be rewritten as $|L(j\omega)| > \frac{1}{e_{tr}}$, resulting in an open-loop constraint on gain:

$$|L(j\omega)|_{dB} > 20 \log \left(\frac{1}{e_{tr}} \right) \quad (3.18)$$

3.5.3 Maximum Overshoot

One important metric for transient performance is the maximum overshoot M_o . Assuming second order dynamics, this can be translated into a representation of the damping ratio ζ as:

$$\zeta = \sqrt{\frac{\ln(M_o)^2}{\pi^2 + \ln(M_o)^2}} \quad (3.19)$$

Assuming the system is underdamped ($\zeta < 0.7$), which is necessary for any discussion of overshoot, we can approximate the phase margin as: $\phi_m \simeq 100^\circ \zeta$, resulting in an open-loop constraint on phase margin:

$$\phi_m > 100^\circ \sqrt{\frac{\ln(M_o)^2}{\pi^2 + \ln(M_o)^2}} \quad (3.20)$$

3.5.4 Maximum Gain

The overall quality of the transient behavior is also affected by the maximum gain M_r of the closed loop system. We can limit the gain of the $\frac{y}{r}$ transfer function.

$$|T(s)| < M_r \quad (3.21)$$

First, this gain is expressed in terms of the damping ratio ζ , assuming a second order system:

$$M_r = \frac{1}{2\zeta\sqrt{1-\zeta^2}} \quad (3.22)$$

Then, we solve this equation for ζ in terms of M_r :

$$\zeta = \sqrt{\frac{1}{2} \pm \frac{\sqrt{M_r^2 - 1}}{2M_r}} \quad (3.23)$$

Which translates to an open-loop constraint on phase margin:

$$100^\circ \sqrt{\frac{1}{2} - \frac{\sqrt{M_r^2 - 1}}{2M_r}} < \phi_m < 100^\circ \sqrt{\frac{1}{2} + \frac{\sqrt{M_r^2 - 1}}{2M_r}} \quad (3.24)$$

3.5.5 Noise Rejection

Noise typically manifests at high frequencies. To reject the effects of noise on our output, we wish to limit the gain of the transfer function $T = -\frac{y}{v}$ at high frequencies.

$$|T(s)| < M_n \quad (3.25)$$

Using the approximation $T \simeq L$, this results in an open-loop constraint on gain:

$$|L(j\omega)|_{dB} < 20 \log(M_n) \quad (3.26)$$

3.5.6 Closed Loop Bandwidth

The bandwidth is the highest frequency for which the control system can effectively track and regulate a dynamic process. The closed loop bandwidth will closely match the open loop bandwidth, which is defined as the crossover frequency ω_c , the frequency at which the open loop gain is -3 dB (half the power).

$$\omega_c = \{\omega : |L(j\omega)|_{dB} = -3\} \quad (3.27)$$

Thus, the desired closed loop bandwidth ω_{bw} can be directly translated into an open-loop constraint on crossover frequency:

$$\omega_c > \omega_{bw} \quad (3.28)$$

The following test can be used to determine if a system meets this requirement for a desired bandwidth ω_{bw} :

$$|L(j\omega_{bw})|_{dB} > -3 \text{ dB} \quad (3.29)$$

3.5.7 Controller Design

To achieve feedback control over the voltage of an ESA governed by the electrical dynamics modelled by G_v , we would like to design a control law K_v , such that the closed loop T_v has the following performance specifications:

- Less than 1 % steady-state error.
- A tracking error under 10 % for all frequencies up to 10 Hz and under 20 % for all frequencies up to 20 Hz.
- The rejection of at least 90 % of noise at all frequencies beyond 250 Hz.
- A bandwidth of at least 100 Hz ($5\times$ the bandwidth of G_z).
- Less than 3 dB of overall maximum gain for well-conditioned transient behavior.

The loopshaping approach presented in this section provides a means for selecting K_v and designing the behavior of the closed-loop transfer function $T_v = \frac{L_v}{1+L_v}$ based on the shape of the open-loop transfer function $L_v = K_v G_v$. This transfer function is shown in figure 3.9 along with the limits imposed on it by our performance specifications. Also shown in this figure are G_v , the plant transfer function of the electrical dynamics of the actuator found by frequency domain system identification in section 3.4, and K_v , the voltage control law which we design. Please refer to figure 3.10 for a summary of these transfer functions and how they interact.

We design the control law K_v to satisfy our performance requirements. K_v is given in terms of kV to PWM duty fraction $\rho \in [0, 1]$:

$$K_v(s) = \frac{470}{s + 150} \quad (3.30)$$

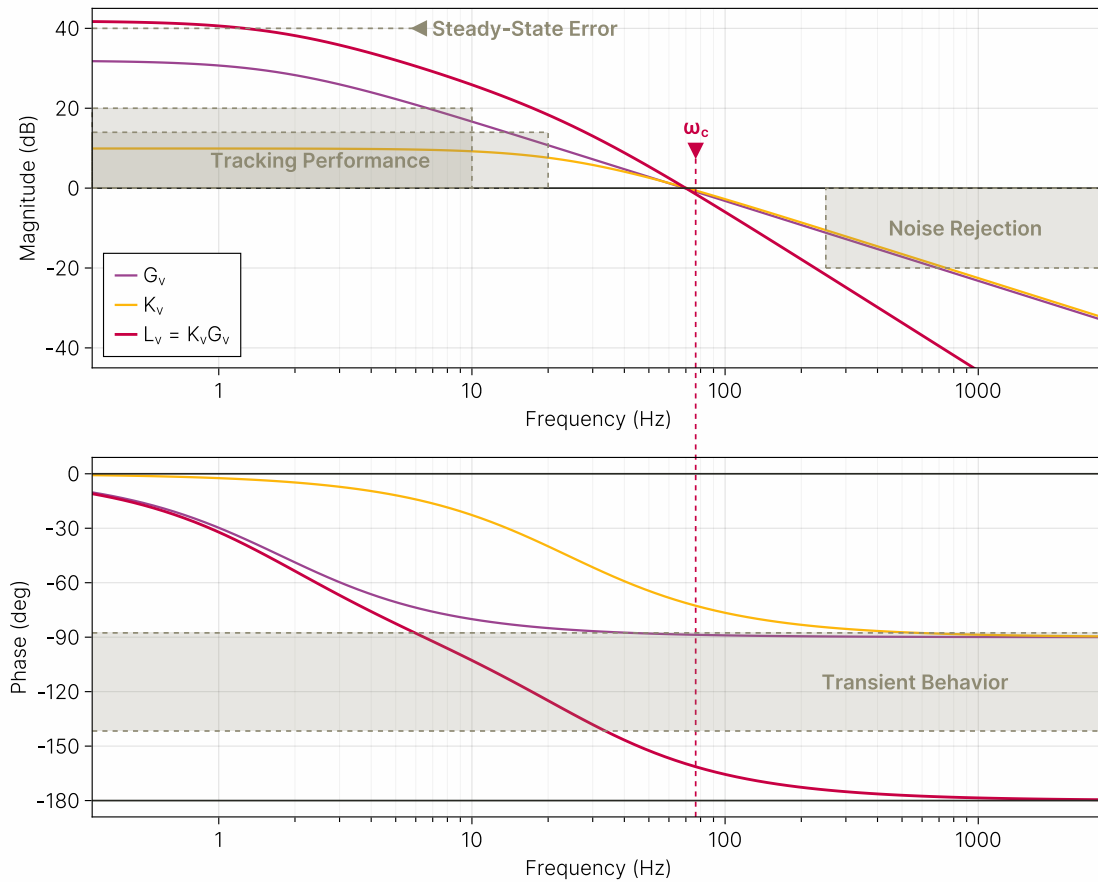


Figure 3.9: Voltage controller design via loopshaping with superimposed performance requirements.

There remains substantial room for improvement in the design of K_v . Although this controller satisfies the performance requirements in terms of steady-state error, tracking error, and noise rejection, it fails to satisfy the requirements on phase margin. In fact, the phase margin is only approximately 19° , which suggests that this controller will struggle with transient behavior, overshoot, and possibly even stability. Observations of controller performance (which are not presented directly) are consistent with this determination.

Based on the curvature of the open-loop transfer function L_v in figure 3.9, we would expect the closed-loop transfer function T_v to have a bandwidth of approximately 80 Hz. Interestingly however, based on the results shown in figure 3.14, the bandwidth of the closed-loop system T_v is approximately 200 Hz, exceeding our expectations by more than a factor of 2. Such a discrepancy

suggests the possibility of inaccuracies or significant, unaccounted-for factors within our methodology, and warrants further analysis.

3.6 Digital Control Implementation

Practically, it is often easiest to implement a controller using a microcontroller. Although physics operates in continuous time,³ and the controller is designed in continuous time, the microcontroller operates in discrete time. Thus, we must discretize the controller before implementing it in firmware.

3.6.1 Controller Discretization

The Z-Transform lets us quickly translate any controllers we design from the continuous time domain to a discrete time domain suitable for implementation on a microcontroller.

$$\tilde{x}(z) = \mathcal{Z} \{x[k]\} = \sum_{n=0}^{\infty} x[k]z^{-n} \quad (3.31)$$

The delay operator z^{-1} represents a delay of one time step, such that $z^{-1}u_k = u_{k-1}$, or more generally:

$$z^{-n}u_k = u_{k-n} \quad (3.32)$$

A continuous time transfer function $H(s)$ can be discretized using zero-order hold:

$$H(z) = \left(\frac{z-1}{z}\right) \mathcal{Z} \left\{ \frac{H(s)}{s} \right\} \quad (3.33)$$

This can be solved manually by using transform tables, or facilitated by using software. It should be noted that other discretization methods exist (eg. direct substitution, first-order hold, etc.), each with it's own trade-offs in terms of complexity, stability, and accuracy.

We rewrite equation 3.30 scaled in terms of the raw representations used on the MCU: an 8-bit encoding of voltage and a 16-bit encoding of PWM duty fraction.

$$K_v(s) = \frac{3e5}{s + 150} \quad (3.34)$$

³ Until Stephen Wolfram proves otherwise: <https://www.wolframphysics.org/>

Numerous tools exist to facilitate the discretization of systems. For example, to discretize the continuous time controller $K_v(s)$ using zero-order hold at 600 Hz in Julia:

```

julia
1  using ControlSystems
2  K = tf(3e5, [1, 150])
3  Δt = 1/600
4  Kd = c2d(K, Δt; method=:zoh)

```

The resulting discrete transfer function should then be rearranged in the form:

$$H(z) = \frac{U(z)}{E(z)} = \frac{b_0 + b_1 z^{-1} + \dots + b_m z^{-m}}{1 + a_1 z^{-1} + \dots + a_n z^{-n}} \quad (3.35)$$

which lets us solve for u_k in a form which is easily implemented in firmware:

$$u_k = b_0 e_k + \sum_{i=1}^m b_i e_{k-i} - \sum_{i=1}^n a_i u_{k-i} \quad (3.36)$$

Using values from the example above, our firmware implementation would look as follows:

```

c pseudocode
1  const int N = 1;
2  const int M = 2;
3  float a[N] = {-0.7788};
4  float b[M] = {0.0, 442.4};
5
6  void controller() {
7      uint8_t v_ref = update_reference();
8      uint8_t v_mon = read_voltage();
9      for (int i = N; i > 0; i--) u[i] = u[i-1];
10     for (int i = M; i > 0; i--) e[i] = e[i-1];
11     e[0] = v_ref - v_mon;
12     u[0] = b[0]*e[0];
13     for (int i = 1; i <= M; i++) u[0] += b[i]*e[i];
14     for (int i = 1; i <= N; i++) u[0] -= a[i]*u[i];
15     int16_t pwm = round(u[0]);
16     (pwm > 0) ? write_chg(pwm) : write_drn(-pwm);
17 }

```

3.6.2 Actuator Fault Detection

If an actuator fails due to dielectric breakdown, it enters into a low impedance state. Thus, attempting to charge a failed actuator essentially shorts the high voltage rail through the high-side switch. This is not consequential for brief periods of time, but prolonged attempts to charge a failed actuator generate substantial amounts of heat and will eventually cause damage to switching and controls hardware. In a worst case scenario, such conditions can lead to insulation failure and cascading damage to multiple components.

With large numbers of actuators, it should be expected that some actuator failures will eventually occur. Consequently, failure monitoring and auto-shutdown are implemented as part of the driver firmware.

```

c pseudocode
1  if (v_mon < V_TH && e[2] > ERR_TH) {
2      badness++;
3  }
4  if (hasel_ok && badness > badness_frames(3.0)) {
5      hasel_ok = false;
6      send_status(STS_HASEL_ERR);
7  }

```

3.7 Summary, Validation, and Performance

We have modeled the dynamics of the actuator as two linear transfer functions, G_v and G_z , representing the decoupled electrical and displacement dynamics discussed in chapter 2 respectively. The transfer function G_v includes the dynamics of the half-bridge drive circuit shown in 3.3 and discussed throughout chapter 2, and maps from the input duty fraction ρ to the voltage across the actuator v . Meanwhile, the transfer function G_z models the mapping from actuator voltage v to actuator displacement z .

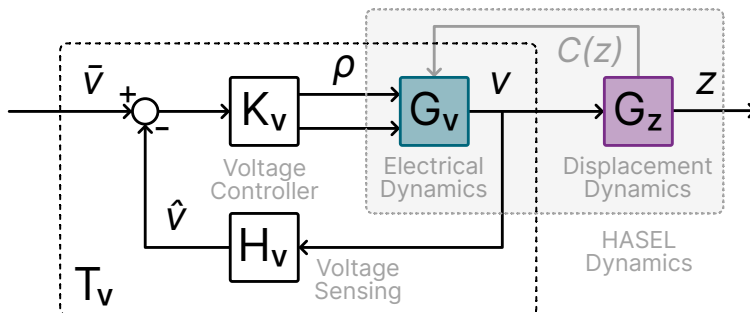


Figure 3.10: A block diagram of the "inner loop" voltage control scheme, highlighting interactions with the decomposed model of ESA dynamics.

As discussed in chapter 2, we are able to decouple these due to their different timescales. Figure 3.10 shows this model, along with the neglected coupling effects of displacement-dependent capacitance $C(z)$.

Also shown are the component transfer functions of T_v , the closed voltage feedback loop which maps from a reference or setpoint voltage \bar{v} , to v , the voltage across the actuator. As well as well as its various component transfer functions, which we have constructed over the course of this chapter. A transfer function model for G_z is estimated from first principles models in section 2.6.3, and parameter values are refined using a frequency domain system identification procedure on benchtop hardware in section 3.4. The control law K_v is designed via loopshaping in section 3.5.7. The voltage sensor H_v , which is introduced in section 3.2, is assumed to have a unit transfer function.

3.7.1 Voltage Tracking

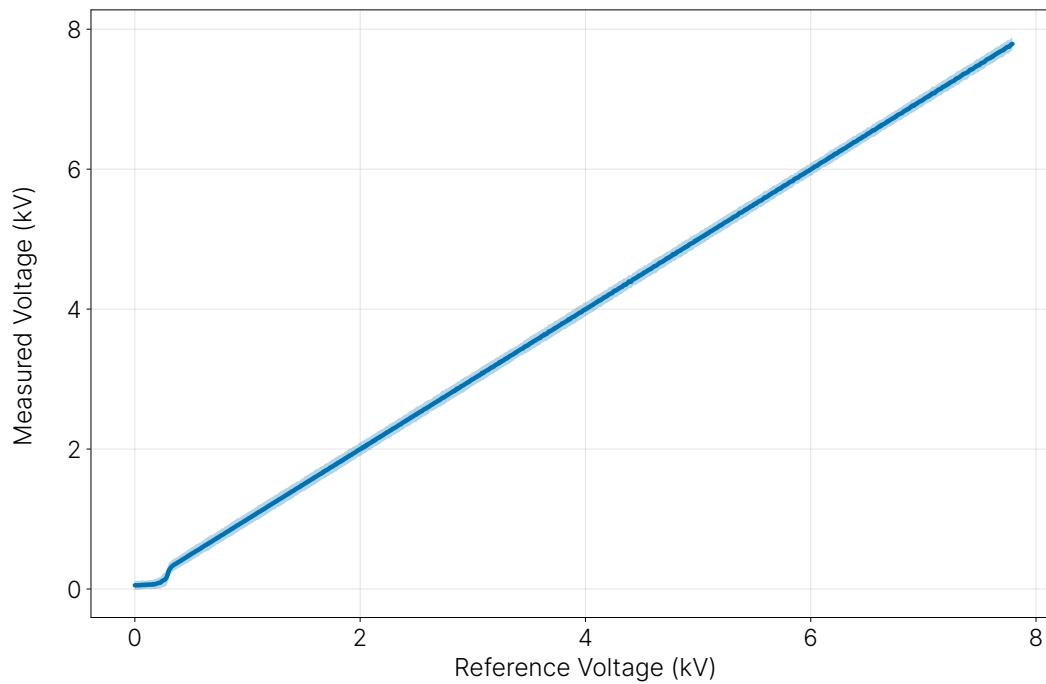


Figure 3.11: Quasistatic voltage regulation.

Figure 3.11 validates the tracking performance of voltage regulation for 100 independent 100 Peano-HASEL actuators, and driven by 100 independent optoelectronic half bridge drivers. In other words, this shows the quasistatic behavior of the voltage control loop T_v , the mapping from desired setpoint voltage \bar{v} to the actual voltage across the actuator v . The reference input to each channel ramps from zero to the 8kV rail voltage over 10 seconds. We consider this slow enough to be considered a quasistatic test. As suggested by equation 2.39, we would expect voltage regulation to be an adequate proxy for rudimentary displacement control under such quasistatic assumptions. The results shown in figure 3.12 confirm this – to a degree.

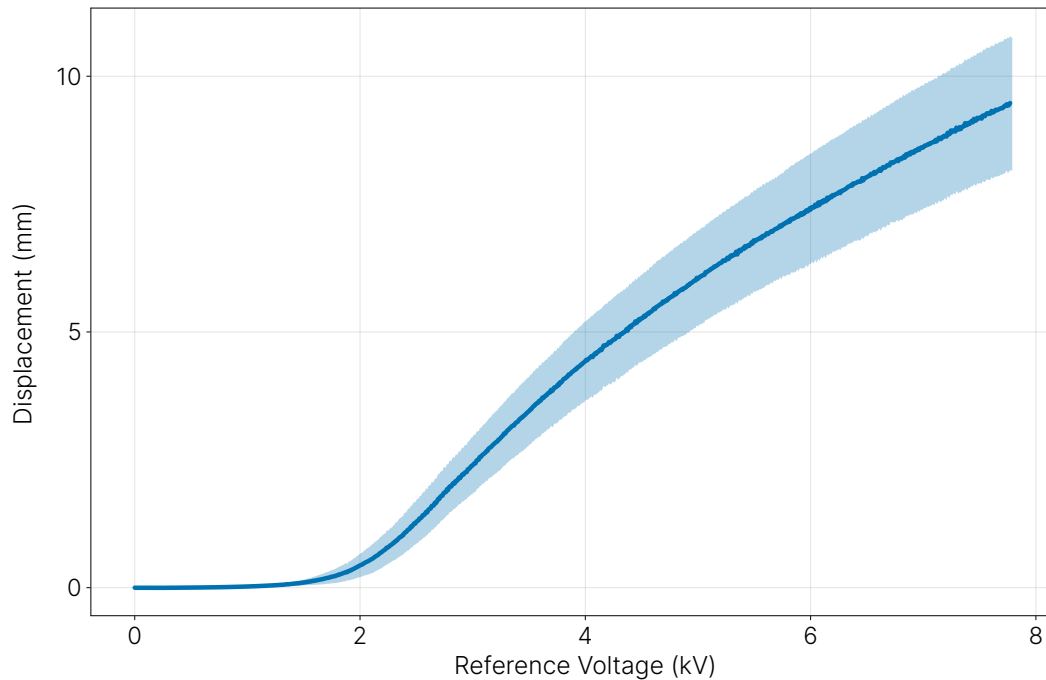


Figure 3.12: Quasistatic displacement control via proxy of voltage regulation.

Figure 3.12 shows the quasistatic behavior of G_z , the mapping from voltage across the actuator v to actuator displacement z . The 100 actuators used to collect this data exhibited a margin of variability in their construction, as they were constructed by hand by 5 different people over a span of 2 years. If we account for this variation by normalizing the displacement to each actuators respective maximum displacement, we can see in figure 3.13 that the results are quite consistent. This suggests that for at least quasistatic positional control, we only need to measure this voltage/displacement curve to achieve an effective result.

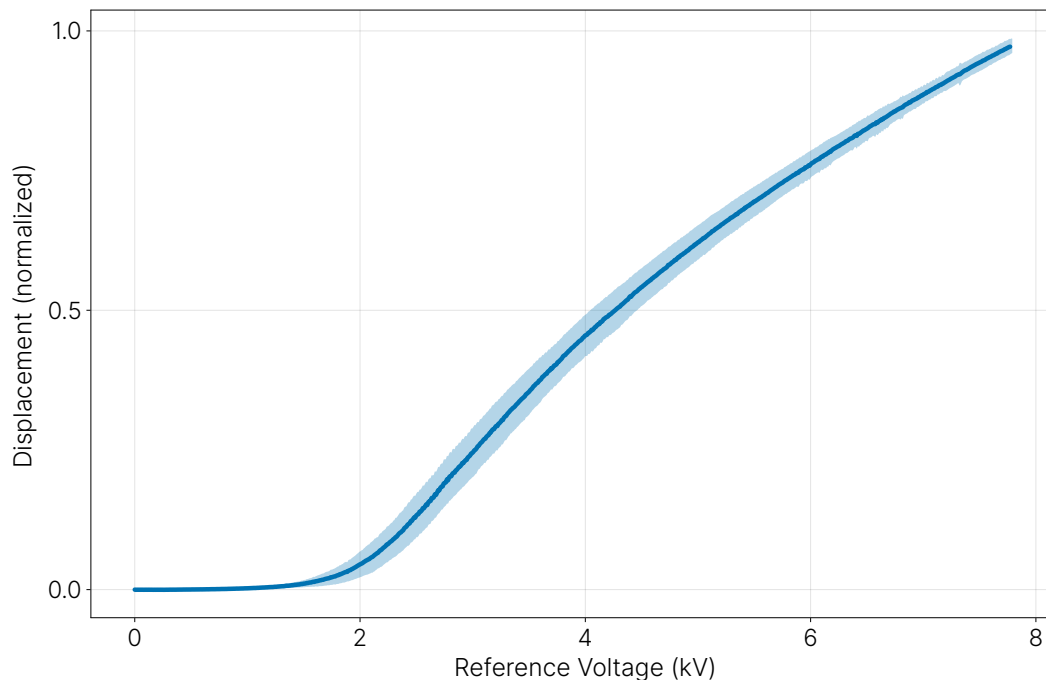


Figure 3.13: Normalized quasistatic displacement control via proxy of voltage regulation.

3.7.2 Closed Loop Frequency Response

To verify and validate the performance and design of our voltage control scheme, we consider the closed loop frequency response. We evaluate the performance of the control scheme on 100 individual actuators running simultaneously within the s-tissue hardware stack presented in chapter 4. The mean ± 1 standard deviation of the transfer functions G_z and T_z is shown in figure 3.14.

The transfer function G_z models the mechanical dynamics of the actuator, specifically the transfer function from voltage v across the actuator to its displacement z . The process of identifying G_z mirrors the process of identifying G_v presented in section 3.4, except that voltage is regulated using the controller K_v presented in this chapter, and actuator displacements are measured using the motion capture (MoCap) videography system (OptiTrack) discussed in section 4.4. The sensor limit shown in figure 3.14 represents the sample rate of the MoCap system, and only applies to the identification of G_z . Note that G_z models the transfer function from the measured voltage v , not the reference voltage \bar{v} asked of the controller. Voltages are measured with the circuit described in

section 3.2. As measured, the open-loop displacement dynamics modeled by G_z exhibit a bandwidth of approximately 20 Hz.

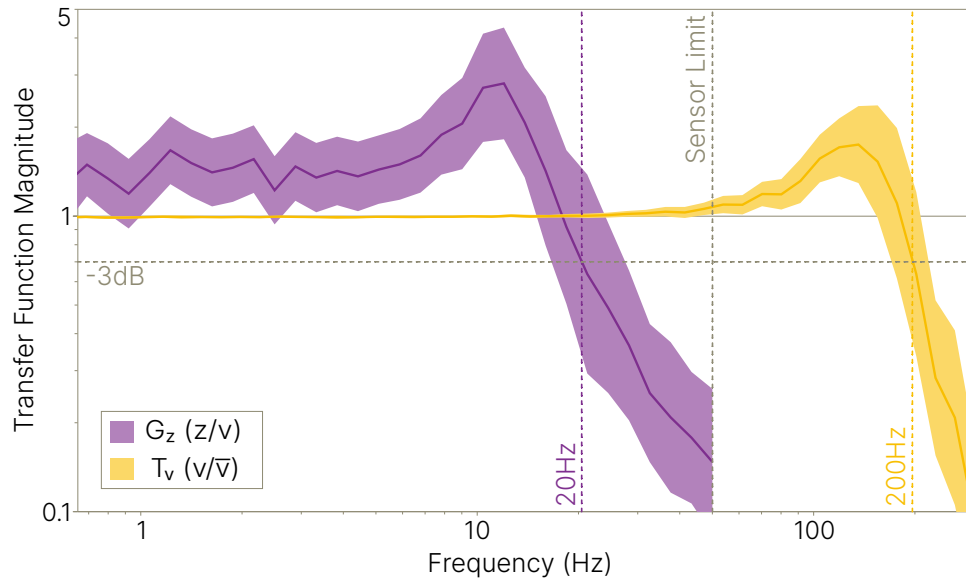


Figure 3.14: Frequency response of system under closed loop voltage control.

The transfer function T_v models the combined closed-loop dynamics of the voltage controller K_v , the electrical dynamics G_v , and the voltage sensor H_v . The process of identifying T_v mirrors the process of identifying G_v presented in section 3.4, except that the analysis is performed by comparing the reference voltage \bar{v} sent to the controller K_v , to the true voltage v measured across the actuator. An overview of the various transfer functions under consideration is given in section 3.7 and more broadly in section 4.6.

The closed-loop dynamics of T_v demonstrate a bandwidth of 200 Hz, and constant gain over the entire bandwidth of the actuator's displacement dynamics G_z . The bandwidth of T_v is larger than the bandwidth of G_z by a full order of magnitude. In other words, our controller is sufficient to control voltage at or above the entire range of speeds we expect the actuator to move, and our design is validated, along with our decision to decouple the electrical and mechanical dynamics.

3.7.3 Limitations of Voltage Control

Taking full advantage of the high-speed dynamics electrostatic actuators are capable of attaining means operating them well beyond a quasistatic regime. Although displacement control via the proxy of voltage regulation is effective in the quasistatic case, its efficacy is limited in general. These limitations are perhaps best illustrated by figure 3.15, which shows 60, 1 Hz charge-discharge cycles of a Peano-HASEL actuator at a non-quasistatic rates.

A hysteresis loop is clearly visible, which is in part caused by the dissipative elements of the dynamics playing a non-trivial role. It is also partially due to the inherent asymmetry of the displacement-dependent capacitance. Because the mechanical dynamics responsible for the value of $C(z)$ will lag behind the faster electrical dynamics governed by it, the effective capacitance of the circuit will be lower during the charge phase than the discharge phase.

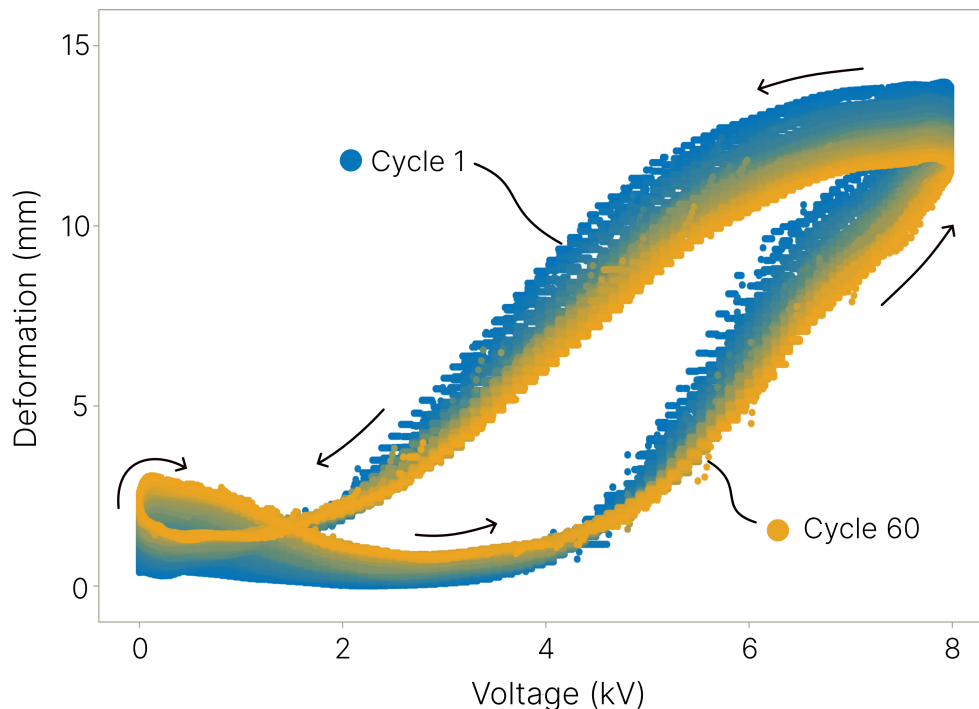


Figure 3.15: Characteristic hysteresis loop for a Peano-HASEL actuator stack.

The existence of this hysteresis loop is further compounded by the property of dielectric

absorption, sometimes referred to as charge retention. When a capacitor is charged, the molecules in its dielectric become polarized, aligning their electric dipoles with the applied field. When discharged, most of these dipoles return to their original random orientations, releasing the stored energy. However, some dipoles may become trapped or immobilized due to imperfections in the dielectric material, preventing a complete and immediate return to the initial state. As a result, even after the capacitor is supposedly discharged, a residual charge remains in the dielectric, which slowly leaks out over time. This phenomenon is commonly observed in certain capacitors, especially those with liquid dielectrics.

In effect, this introduces a time-variant component to the dynamics of the system, with the capacitance $C(z)$ becoming dependent on voltage history as well: $C(z, v(t))$. This presents a significant challenge for control design, as the system is no longer time-invariant, and the majority of control design techniques rely on this assumption. One option is to use adaptive control techniques. Another is to treat the time variance as an uncertainty in the dynamics, and to leverage techniques of robust control to design a suitable controller for the family of systems parametrized by this uncertainty. These approaches are promising, but they are left for future research.

However, dielectric absorption fundamentally impedes the performance of the actuator, and its outright prevention would be preferable to any compensation from the control system. Dielectric absorption can be mitigated by using a full-bridge driver shown in figure 2.6 and reversing the voltage polarity at each zero crossing. This approach is used by [159] and others, but doing so exactly doubles the component requirements of the driver circuitry, and is not always practical, especially at larger scales.

The solution to this phenomenon is ultimately a question of materials science, and remains an active area of research [37, 144]. It's severity is highly dependent on the type, source, and even batch of liquid dielectric and elastomeric film materials used to construct the actuator, as well as environmental controls and procedures employed during the manufacturing of the actuator. Further research on the materials of these actuators is needed to better understand, model, and eventually minimize the prevalence of this effect.

We note that charge retention effects are asymptotic, and so we tolerate them by characterizing our system after a warm-up period of 60-100 cycles, and more importantly, by adding direct measurement and regulation of displacement as an outer control loop. This work is discussed in chapter 4.

Chapter 4

A Case Study in Sensorimotor Systems Integration




Figure 4.1: The s-tissue: a soft shape display with simultaneous multifunctionality.

The entire point of the technologies presented so far is to enable the construction of the next generation of highly parallelized soft robotic platforms. To truly evaluate their efficacy for such a goal, we must integrate them into a cohesive robotic system which is capable of sensing and responding to its environment, and accomplishing complex tasks within it. In this chapter, I describe the design and implementation of the s-tissue¹: a 10×10 array of electrostatic actuators, optoelectronic bridge drivers, magnetic displacement sensors, with integrated software and controls. The s-tissue serves as a case study in sensorimotor integration through which to evalu-

¹ Let's say this stands for "synthetic tissue". The original name tried a little too hard to be an acronym.

ate the technologies described in the preceding chapters in an application-oriented context. More fundamentally, this robotic platform showcases the rich multifunctionality that arises from tightly integrating highly parallelized soft robotic actuation and sensing technologies with feedback control at previously undemonstrated scales.

This work results in a numerous emergent capabilities directly relevant to a variety of scientific and industrial contexts, including: high-fidelity shape display, sensing of surface deformations and external forces, multimodal user interaction, spatially distributed audio, dynamic manipulation of objects and liquids, and more; some of which are shown in figure 4.1. For a comparison of the s-tissue to other shape displays, please see [69]. The s-tissue demonstrates the promise of using soft robotic materials to create high-DoF, high-performance, and sensor-rich robotic systems, and provides a glimpse into the exciting future of our field.

4.1 System Architecture

4.1.1 Sensorimotor Unit Cell

A HASEL actuator described in chapter 2, and driver circuit described in chapter 3 are combined with a soft magnetic displacement sensor [138] to form a sensorimotor unit cell shown in figure 4.2. With appropriate supporting hardware, each cell is capable of sensing voltage and 3-axis magnetic field at 1 kHz each, measuring surface displacements to a resolution of 0.1 mm and external forces to a resolution of 50 mN, and achieving voltage regulation up to 200 Hz and displacement control up to 20 Hz.




Figure 4.2: A sensorimotor unit cell from the s-tissue, exploded view.

4.1.2 Modular Backbone

Ten unit cells are combined into a module along with a pair of microcontrollers (MCUs), a power supply stack, thermal management system, structural components, and a basic user interface as shown in figure 4.3. The two MCUs are responsible for sensing and actuation, and are called the sensor control unit (SCU) and driver control unit (DCU), respectively. Modules are functionally independent, providing the majority of supporting systems needed to run their unit cells, and requiring only a low voltage DC power supply (approx 40 W), earth ground connection, and two USB 2.0 ports on a host computer for interfacing and control.




Figure 4.3: A 1x10 s-tissue module and the subsystems found within.

The full s-tissue system is constructed from 10 modules mechanically linked with threaded rods and covered with a silicone skin, and optionally with retroreflective markers for motion capture videography (MoCap). A central PC (AMD Ryzen 7 5800X CPU, 32 GB RAM, Linux 5.14) provides a user interface, and is used to coordinate the modules and perform high level control using the software described in [69]. A central AC/DC power supply (Ion SFX 650-G, Fractal Design) provides power to all modules. Each module is individually grounded to a ground plate which has been directly connected to the steel superstructure of our laboratory building² to provide the largest possible reservoir for dissipation of accumulated charges. The full s-tissue is shown in figure 4.4.

² On the recommendation of our facility’s electrician. This is equivalent to grounding to the earth prong of an electrical outlet, but with a much more direct, transparent path.




Figure 4.4: Overview of the main elements of the s-tissue system.

This hardware architecture is highly modular by design, and with the exception of some high voltage electronics, all components can be manufactured with a combination of off-the-shelf parts, readily available raw materials, and accessible rapid-prototyping manufacturing processes such as 3D printing or laser cutting. This allowed me to separately evaluate and iterate on the design of individual subsystems or components – most of which underwent one or two revisions. Perhaps even more importantly, it enabled incremental repair and maintenance of a system which contains hundreds if not thousands of points of potential high voltage discharge in a tightly packed volume. If (when) discharge events occur, the resulting damage is likely to cascade across multiple components in highly unpredictable ways. The modular design makes it relatively easy to quickly identify and repair faults as they occur, and to replace worn or damaged components. The result is an iterable hardware platform which continues to operate reliably after approximately 3 years and several hundred hours of use.

4.1.3 Hardware Interface

The majority of user interaction with the hardware is handled via a software interface. However, some low-level hardware interactions are broken out to a hardware interface in the front panel of each module illustrated in figure 4.5:




Figure 4.5: Front panel subassembly and interface.

- HV Status – an amber LED which illuminates whenever the HV rail is energized. It’s function and implementation are detailed in section 4.2.5.
- SCU/DCU Status³– white LEDs which are tied to a GPIO pin on each microcontroller, are controlled by firmware, and indicate firmware state.
- Power Status – a white LED which illuminates if the module has power.
- SCU/DCU Reset – tactile switches to put individual microcontrollers into a low power state. They are rarely used in practice, but invaluable when needed.
- HV Lockout – a toggle switch which disables the high voltage rail when in the downward position. It is part of the high voltage lockout system described in section 4.2.4.

4.1.4 Thermal Management

No individual component within this system generates a tremendous amount of heat on its own. However, the combined operation of 100 parallel channels adds up to several hundred watts of heat dissipated within an enclosed space. The main sources of heat are, in descending order: high voltage photodiode conduction losses, relatively high-current optocoupler LEDs, microcontroller CPUs, and the various power supplies, converters, and regulators. As described in sections 5.7.3

³ The value of a single clearly visible LED tied to a GPIO pin on every microcontroller cannot be overstated for debugging and troubleshooting when bringing up a new hardware stack.

and 3.1.5, if left unchecked, heat can cause catastrophic failure of optocoupler LEDs or insulation. In the latter case, this will lead to high voltage discharge events and cascading electrical damage to multiple components.




Figure 4.6: Rear panel subassembly with integrated exhaust fan.

To prevent heat buildup, a low-noise, low-power, 12 V, 60 mm fan is added to the back of each module to ensure some degree of airflow throughout the system. Fans are oriented to exhaust air from the module to minimize interference with the skin, and operate continuously whenever the module is powered.

4.1.5 Structural Elements

The structural design characteristics of the s-tissue can be represented by the front panel subassembly shown exploded in figure 4.7 below:




Figure 4.7: Front panel subassembly, exploded view.

Panels are made from acetal sheets (various dimensions, ePlastics.com), which are laser cut (Speedy 360 flexx, Trotec) and then sanded at 200 grit to reduce surface charge accumulation and minimize reflections which may interfere with motion capture videography. Other structural components are 3D printed (MK3S, Prusa) from composite reinforced PETG (XT-CF20, ColorFabb), and often make use of brass heat-press inserts to robustly fasten components together.

Wherever possible, brass fasteners are used internally to minimize ferrous sources of magnetic interference. Black anodized steel fasteners are used externally to minimize reflections for motion capture videography. Most of the empty space is filled with wires connecting the various electrical components, and the geometry of internal spacers creates paths and tie points for wire management.

Standoffs mount the panel to the PCB at the correct distance for access to PCB mounted buttons and switch. Polycarbonate “light pipes” channel and diffuse light from the variously sized LEDs to the plane of the front panel.

4.1.6 Silicone Skin and Fabrication

The silicone skin is fabricated by mixing 200 g of Ecoflex 00-30 (Smooth On) in equal parts of A and B with 3% by mass black coloring agent added (Silc-Pig, Smooth On). The mixture is cast using a 62 cm x 62 cm square frame mold on a tabletop and left to settle to an average thickness of 550 μm . To achieve a matte surface finish, we mix equal parts A and B of Psycho Paint (Smooth On), add 400% by mass of solvent (NOVOCS Matte, Smooth On), 3% by mass Silc-Pig, and spread the mixture evenly over the cured EcoFlex surface. After drying, the silicone skin is carefully stretched over the 10x10 display. Panels secured by screws along each edge of the array hold the skin in place.

4.1.7 System Assembly

Two acetal panels (top and bottom) serve as the primary structural backbones of each module. All sensorimotor unit cells (drivers, sensors, and actuators) are bolted to the top panel. All supporting electronics (power stack and microcontrollers) are bolted to the bottom panel. The top subassembly is designed to be easily removable from individual modules for access and replacement of channel-specific circuits, without requiring complete disassembly of the full array. To achieve this, the top and bottom panels are connected by 4 \times 3D printed support pillars, which also serve as wire management channels. Each support pillar has two parts, one permanently affixed to each of the two panels. Each support pillar is joined by a single bolt which is accessible from the top panel.

Threaded cross bars passed through internal brackets provide a simple means to hold multiple modules together. The general assembly procedure consists of assembling and wiring the top and bottom subassemblies separately, then connecting the electrical connections between them, and finally by mating the subassemblies together with bolts.

4.2 Power Supply Stack




Figure 4.8: Topology of the s-tissue power supply stack.

The hardware stack requires power supplied at multiple voltage levels. A high voltage supply (8 kV) is needed to drive the actuators themselves. A 5 V supply is needed for USB communication, and a 3.3 V supply is needed for the remaining low voltage components.

Nominally, power to all 10 modules is provided by a centralized AC/DC power supply (Ion SFX-L 650W, Fractal Design) as shown in figure 4.8. Although it tethers the hardware to a wall outlet, this approach has several advantages and is well suited for a large, stationary, high-power system such as the s-tissue. It avoids the headaches of battery management, and directly provides 12 V (primary), 5 V, and 3.3 V rails which conveniently satisfy the low voltage power requirements of the system, removing the need for dedicated step-down regulators. To provide the high voltage rail, the 12 V primary supply is stepped up to 8kV as described in section 4.2.2.

A single centralized power supply also facilitates the implementation of a global emergency stop (e-stop) system. An e-stop kill switch wired to control the AC/DC power supply and thus cut all power to the system is always kept within reach of the operator. This approach is simple, reliable, and transparent in its operation.

The downside of this approach is that it kills all power to the system, including the power source for the indicator LEDs. This can create scenarios where the hardware still holds dangerous amounts of charge, but does not advertise this. Numerous other indicators exist to immediately distinguish total power loss conditions, which in practice tend to occur because the operator has

intentionally pressed the emergency stop button. Ultimately, leakage currents will discharge the entire system even from a fully charged state within a few minutes, rendering it safe for disassembly and troubleshooting. Future kill switch implementations should only cut hazard power supplies.

Each module is designed with a power budget of 48 W. This can be broken down to 30 W of output from the high voltage rail (average 3 W per actuator), 8 W for low voltage power rails, and 10 W of headroom to account for any voltage conversion losses and inefficiencies.

4.2.1 USB-C Power Delivery

The power supply stack presented can be modified to take advantage of the USB-C power delivery (PD) standard. This allows the system to be powered at the module level by universally available 45 W USB-C phone chargers. This is a convenient and portable alternative to the centralized AC/DC power supply, and is well suited for smaller systems sharing a similar hardware architecture to the s-tissue. When using USB-C PD, additional step down converters may be needed to source the logic-level low voltage rails.




Figure 4.9: A simple, effective USB-C power delivery sink circuit.

The simplest way to implement a USB-C PD sink is to use a PD "trigger" IC, such as the IP2721 (Injoinic Technology) shown in figure 4.9. These simple ICs implement a minimal set of PD vocabulary to negotiate a fixed voltage from a power supply. Once power delivery has been negotiated, provided, and stabilized, the IC enables a NMOS load switch and connects the bus to the rest of the system.

The inclusion of a load switch is important, because the bus will only provide a current limited 5 V supply on startup until negotiation is complete. Without a load switch, the current drawn by downstream components - namely boost converters - can load the bus voltage and destabilize the negotiation process, leaving the system stuck in an endless startup loop.

4.2.2 High Voltage Boost Stage

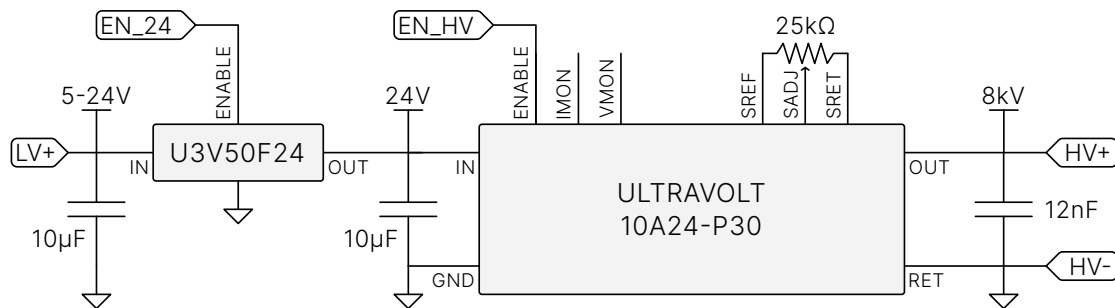


Figure 4.10: Dual stage high voltage boost converter.

To provide a high voltage rail, we implement a dual stage high voltage boost converter. The first stage (U3V50F24, Pololu) steps up the primary (nominally 12 V) low-voltage rail to a 24 V intermediate level. Although for our system, the primary supply is nominally 12 V, this part can accept a broad range of input voltages (2.9-24 V).

The second stage (Ultravolt 10A24-P30, Advanced Energy) steps up the intermediate voltage into the kilovolt range to provide the high voltage rail. The rail output voltage is set via a 25 k Ω potentiometer, we use 8 kV for the s-tissue. This supply is rated for 30 W of output power at 10 kV, thus we expect it to be able to sustain up to 3 mA of high voltage supply current.

4.2.3 High Voltage Capacitor Bank

To stabilize the rail voltage, as well as to provide higher peak current for the capacitive actuators, a 13.2 nF capacitor bank is added to the HV rail. This capacitor bank is comprised of 4×3.3 nF capacitors (615R Series, Vishay) in parallel. However, capacitors of this type must be

derated to 30% of their nominal capacitance when operating near their maximum voltage, so the effective capacitance of the bank is closer to 4 nF under normal operating conditions.

As discussed further in chapter 5, this capacitance is likely insufficient for peak actuator performance, and should instead be much larger than the total ESA load capacitance. However, such a strategy brings diminishing returns in terms of performance, and quickly runs into safety concerns as discussed in [113]. The physical size of the high voltage capacitor bank may also become impractical for larger systems. Ultimately, the question of how much capacitance should be added to the rail warrants further study.

4.2.4 High Voltage Lockout

To ensure the safety of the operator and the hardware, a triply-redundant high voltage lockout system is included. First, a lockout switch is added to the PCB mounted on the Ultravolt module⁴. Second, a lockout switch is added to the front panel of each module, and serves as the primary user-facing lockout switch. Third, a GPIO pin on the DCU is connected directly to the enable circuitry, providing a software-based lockout. Schottky diodes are added as shown in figure 4.11 to prevent backfeeding between the three sub-circuits. If any of these three sub-circuits pulls the enable line low, the HV power supply will shut down.




Figure 4.11: Triply-redundant high voltage lockout system.

⁴ This switch is not easily accessible, and was primarily used to guarantee that power supplies was disabled while working with partially disassembled hardware.

4.2.5 High Voltage Indicator

The high voltage indicator is an amber LED (HSMA-C150, Broadcom) which illuminates whenever the high voltage rail is energized. While seemingly trivial, this indicator circuit has repeatedly proven to be an invaluable, if not critically necessary feature, and is therefore discussed in detail. Importantly, this LED is not controlled by software, but rather by the circuit shown in figure 4.12 which connects directly to the high voltage rail. A fully hardware-based implementation ensures that software failures will not compromise the efficacy of the indicator.

We implement the circuit shown in 4.12 using a N-channel MOSFET and a $1.2\text{ k}\Omega$ resistor to set the LED current. The resistors $R1$ and $R2$ should be selected so that the light turns on for any value of v_{rail} satisfying:

$$v_{rail} \frac{R2}{R1 + R2} \geq v_{th} \quad (4.1)$$




Figure 4.12: High voltage indicator circuit.

An important shortcoming of this indicator design is that it nonetheless relies on a low voltage supply to operate, as the high voltage rail cannot provide enough current to power the LED directly. Consequently, it would show a false negative if the system suddenly lost all low voltage power with charged rails and actuators.

Another shortcoming is that this design will always show as negative if the HV rail is shorted, even if the HV boost stage is running at full power. This is a correct representation of the rail

voltage⁵, but it is not necessarily a useful indicator of the state of the HV power supply. A better design would consider both, possibly by tying into either the enable pin or low voltage input of the HV power supply.

4.2.6 The Missing Trip Circuit

Trip detection and auto-shutoff are safety systems commonly present in high voltage systems which are notably absent from our hardware. A trip detection system is meant to detect short conditions across the output of a power supply. Typically, these systems will shut down the power supply if a maximum current threshold is exceeded, indicating a short condition somewhere downstream. These circuits are therefore implemented using a high-side current sensor, as a low-side current sensor would not be capable of detecting short conditions, as these by definition bypass the return path.

It is challenging, but certainly possible to implement such a high-side current sensor for our 8kV system. Likewise, auto-shutoff could be implemented using the enable pin of the power supply. Such a system should also short the rail to dissipate any energy stored on the rail capacitors in the event of a trip condition or loss of low voltage power. This could be implemented with a load resistor in series with normally-closed HV switch, such as certain reed relays.

However, as we are primarily interested in repeatedly charging and discharging multiple capacitive loads in parallel, in some cases we can expect our power supply to operate at, near, or even above 100% nominal load. Consequently, it becomes tricky to even distinguish a trip condition from normal operation. Perhaps integrating the high-side current, in essence detecting a sustained short, would be effective approach. Nonetheless, defining a trip condition for such a system remains an open question and addressing this important shortcoming is left for future work.

⁵ Nonetheless, it is an instance of "the danger light turns off during a fault state" which should be avoided.

4.3 Computational Architecture

Highly parallelized soft robots, by their very nature, require a substantial amount of data transfer between their individual hardware elements and likewise to and from a multitude of centralized processing, coordination, and decision making tasks which operate in a global context. Their data transfer rates, context switching frequencies, and latency requirements are far more demanding than those of most conventional robots, and at least today, their development necessitates working across two computational domains with diametrically opposing requirements.

In the first, the objective is generally to coordinate the action of distributed hardware elements to achieve emergent behaviors. Additional objectives may include the study of functional principles of sensorimotor systems or the development and application of theoretical frameworks for the control of dynamically complex robots. Success in achieving these objectives is facilitated by access to the the full spectrum of cutting-edge tools available to science, whether they be mathematical, computational, or based on machine learning. Tasks in this domain must interact with rich data streams from highly parallelized hardware, and interpret complex spatiotemporal data for decision making in a global context. These tasks demand abundant computational resources (eg. CPU, RAM, GPU), often benefit from parallelization, and are best implemented on a computer which can leverage the full power of operating systems and high-level programming languages and the convenient abstractions they provide.

In the second, the objective is to create interfaces which support the unique functionalities of novel hardware in a manner which can be scaled to large numbers of parallel hardware channels. Certain tasks in this domain, such as SPI communication, demand control of MHz-frequency signals with timing precision requirements sometimes reaching the order of nanoseconds. While not all tasks are this demanding, their timing requirements are nonetheless comparably stringent, and implementing these tasks as firmware on an embedded microcontroller is typically the logical choice. Embedded systems such as these operate with limited computational resources, including processing power, memory, and storage. Consequently, they are often programmed in low-level, statically-

compiled languages like C or Rust, and virtually never run a full operating system. However, to facilitate the implementation of required functionality, the use of other tools like a real-time operating system (RTOS) or a hardware abstraction layer (HAL) such as the Arduino framework is common.

The computational architecture of a highly parallelized soft robot must support operation across both of these disparate domains, and provide a means of linking them together. We implement the s-tissue system with the topology shown in figure 4.13, which is well suited for this purpose. It consists of a centralized computer responsible for high-level processing and coordination connected to a distributed array of microcontrollers responsible for parallelized, low-level hardware interaction. These subsystems are linked via high-speed, low-latency communication over USB.

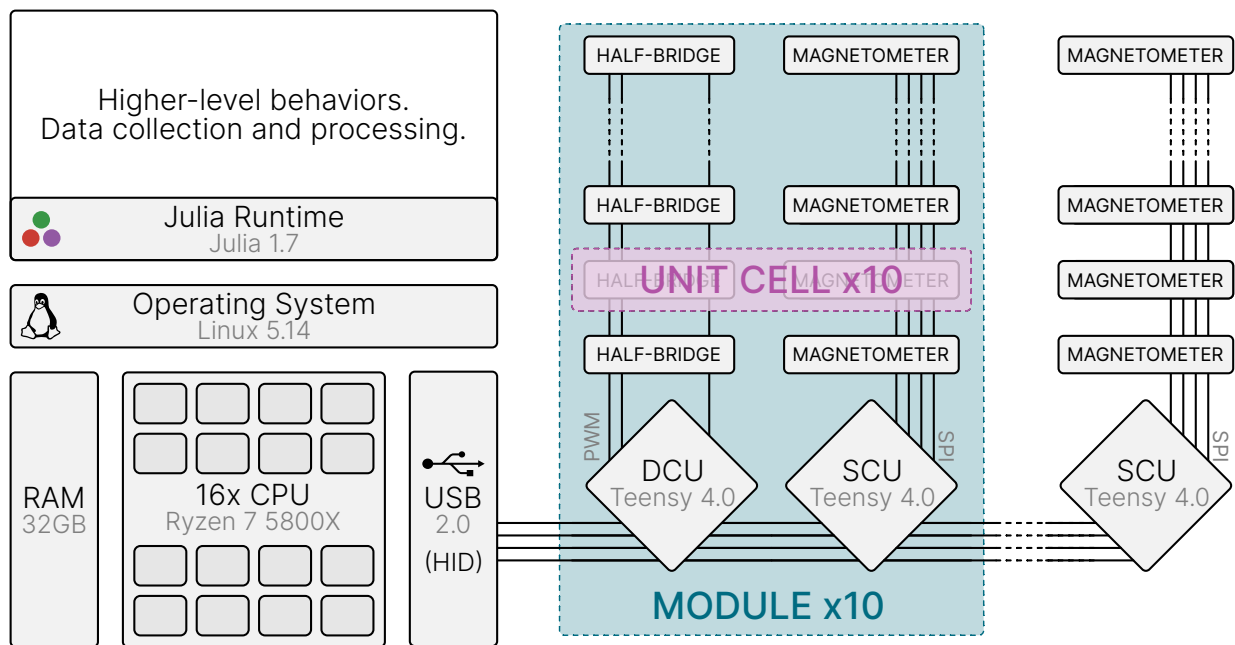


Figure 4.13: The computational architecture of the s-tissue, which is well suited for the demands of other highly parallelized soft robots.

The central compute system is based on a PC (Ryzen 7 5800X, Advanced Micro Devices, 32GB RAM) running Linux (Kernel 5.14). We write the software stack for the central system in

Julia [7], a high-level language for scientific computing. We chose to use Julia because it provides an elegant syntax and best-in-class tooling for the mathematical representation and analysis of dynamical systems [118], natively supports a system of multi-threading and task synchronization primitives useful for the asynchronous demands of highly parallelized robots, provides an elegant means to wrap and call into C libraries and thus enable low-level interactions with system hardware, and compiles highly performant code. Unlike other robotics software frameworks, our julia-based software stack allowed us to write, compile, and include additional code at system runtime, which proved incredibly convenient, especially for the development of control laws.

Each module of the s-tissue described in section 4.1.2 includes two microcontrollers, named the Driver Control Unit (DCU) and Sensor Control Unit (SCU). The entire s-tissue includes both 10 DCUs and 10 SCUs, all of which are based on the Teensy 4.0 MCU (PJRC).

DCUs are responsible for the voltage control loop of the 10 ESAs within their module. For each ESA, the DCU monitors voltages from the sensors described in section 3.2, runs an instance of the control law described in section 3.6, and generates the PWM signals which drive the optocoupler half-bridge described in section 3.3, and thus drive the the actuator. Each DCU also monitors it's respective ESAs for short circuit conditions and failures.

SCUs are responsible for the magnetometry system described in section 4.5. Each of the sensorimotor unit cells described in section 4.1.1 contains one magnetometer, one ESA, and one optoelectronic drive circuit. Thus, each SCU is likewise responsible for 10 magnetometers, with which it communicates over a Serial Peripheral Interface (SPI) Bus with a clock speed set to 1 MHz.

Both types of MCU asynchronously stream relevant data to or from the central computer in real-time over the Universal Serial Bus (USB). The USB protocol is highly flexible, and it is possible to configure a device as any number of multiple endpoints which behave in different ways.

Perhaps the most ubiquitous form of communication with microcontrollers, USB-virtualized Serial communication uses bulk transfer endpoints. These endpoints have large bandwidth at the cost of unbounded latency, meaning this system makes no guarantees on timing whatsoever; the packet is sent whenever there is unallocated bandwidth on the bus. The transmission latency will

vary from packet to packet, and is dependent on the other traffic on the bus, and packets may be grouped together at the discretion of the host controller. For a system like ours, which implements hierarchical control loops in parallel across multiple devices, low communication latency is critical, and serial-based communication proves problematic.

A better option, used in our implementation, is the Human Interface Device (HID) protocol, which is based on interrupt transfers. Unlike the bulk transfers used by serial communication, the USB protocol guarantees that interrupt transfers are serviced within a bounded time frame at the cost of overall bandwidth. We configure SCUs and DCUs for a 1 ms polling interval, or in other terms, a 1 kHz polling frequency.

4.4 Motion Capture Videography

The ground truth reference for the displacement of the hardware system is provided by a motion capture videography system (MoCap). A grid of 100 circular retroreflective markers (Scotchlite 7610, 3M) are placed in the center of each cell. Their position is captured by an array of 8 infrared cameras (OptiTrack Prime 13 W), and their 3D position is reconstructed by by software (Motive 3, Optitrack) to sub-millimeter precision. The videography system is calibrated with a wand (CS-W500, OptiTrack) and calibration square (CS-200, OptiTrack) using the standard procedure described in the OptiTrack/Motive documentation. Four markers in each corner of the 10×10 display are used to offset the ground plane, setting the average z-coordinate of the marker positions to zero. Data frames are then streamed as UDP multicast packets to a local network.

At the time of this work, no official client packages supported depacketization of the raw point cloud data available in the NatNet data stream. Therefore, I wrote a client package to receive and interpret NatNet data streams from Motive directly in Julia, which we do at 240 Hz.

Once received and unpacked, outliers are removed, a polynomial surface is fit to the data, and then resampled at the desired spatial resolution - typically matching the position of the sensorimotor unit cells. The resulting surface samples are then used as the ground truth reference for the displacement of the hardware system. Figure 4.14 shows a frame of point cloud data along with

the fitted surface.




Figure 4.14: A frame of motion capture videography data and accompanying 3D surface fit.

Our software framework enables us to receive and process MoCap data streams quickly enough to render these plots in real time (240 Hz), which is plenty given the refresh rate of our monitor (60 Hz). Because the soft surface is not perfectly flat, displacement is measured relative to the displacement of each actuator at rest, which can be captured with a quick calibration script.

4.5 Distributed Magnetometry

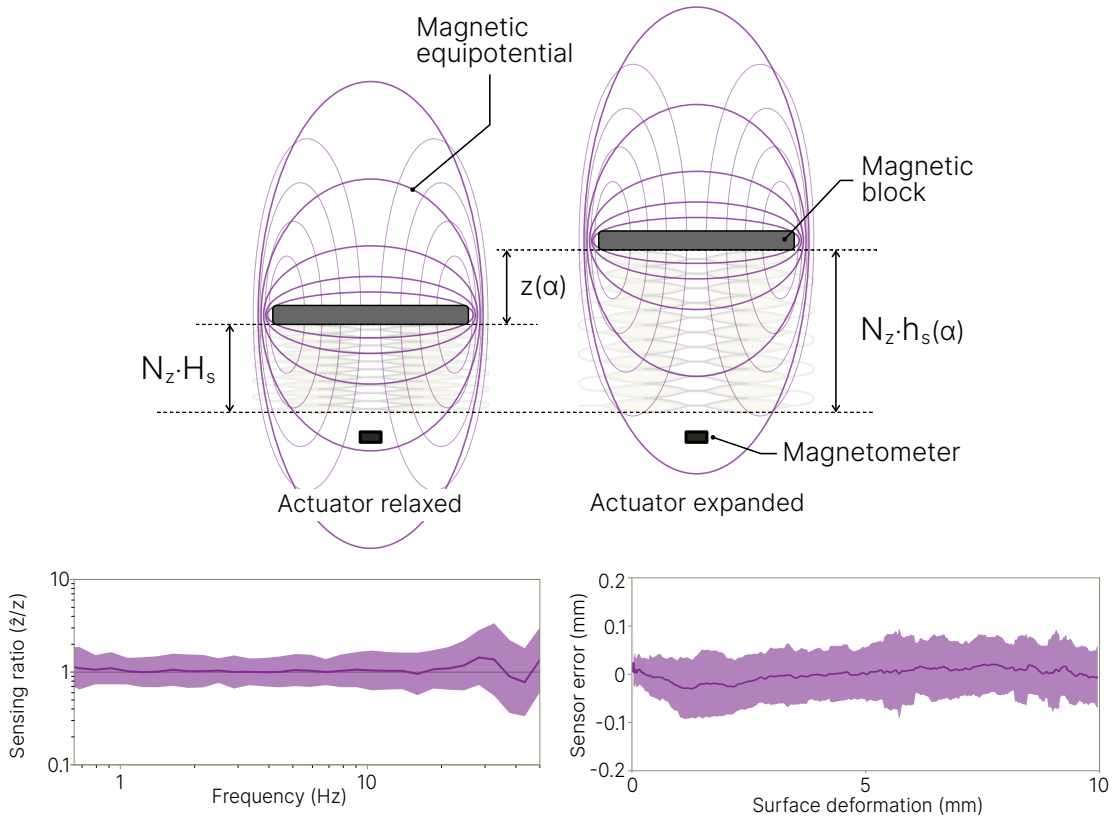


Figure 4.15: Functional diagram of the magnetometry system and characterization of its performance.

Each cell uses the combination of a soft magnetic block and a magnetometer [138] to transduce actuation to measured deformation as shown in figure 4.15. Changes to the displacement z of actuator caused by deformation of the surface – whether from actuator movement or externally applied forces – results in a displacement of the magnetic block. In turn, this causes a change in the 3D magnetic flux density $\{b_x, b_y, b_z\}$ measured by the magnetometer. A distributed magnetic sensor array is the result of implementing such sensors within each cell of the surface.

A third-order polynomial mapping

$$\hat{z} = p_0 + p_1 b_z + p_2 b_z^2 + p_3 b_z^3 \quad (4.2)$$

maps from a raw magnetometer sensor reading b_z (in mG) to the estimated deformation \hat{z} (in mm) of the respective cell using a set of 4 coefficients $\{p_0, \dots, p_3\}$ unique to that cell. The magnetic data is sampled at 600 Hz and processed through a digital first-order low-pass filter with a 50 Hz cutoff frequency. The polynomial coefficients for each cell are determined using a calibration procedure based on MoCap data as discussed in section 4.5.3.

As shown in the performance characterization plots at the bottom of figure 4.15, using this method, each cell in the surface can measure deformation with a mean error below 0.1 mm for quasi-static deformation, which results in a sensing resolution of 0.8 % given a 12 mm range of actuator deformation. Frequency response characterization also shows the sensors can accurately track HASEL deformations up to 30 Hz. A gradual decrease in accuracy can be observed above this point, as the amplitude of deformation begins to approach the sensing resolution.

This is more so a limitation of our test methodology – which depends on, and is limited by, the displacement dynamics of the actuator – than it is a fundamental limitation of our sensing approach. The magnetometers have a sample rate of 1 kHz, and there is no reason to suspect a frequency dependence in their operation. The magnetometer measurements in each cell are not substantially influenced by the magnetic flux of adjacent cells; however, the mounting of the skin restricts the movement, and thus increases the measurement error, of those cells comprising the edges of the 10×10 surface.

4.5.1 Magnetometer Fabrication

The magnetic sensing system consists of two components: a soft, magnetized silicone block, and a commercially available magnetometer. This design, along with its fabrication process, are based directly on [138]. Soft magnetic blocks consist of flexible silicone (EcoFlex 00-30, Smooth On) suffused with bonded neo-powder (NQP-B+ 20441, Neo Magnequench). To create them, first, a batch of EcoFlex 00-30 is hand-mixed using the required 1:1 ratio of parts A:B. Next, 1.67 mass % of neo-powder is added and hand-mixed into the silicone. Once most of the powder is roughly distributed in the silicone, the mixture is placed into a planetary mixer (ARV-310, Thinky), and

subjected to the following procedure.

The mixture is first degassed in a vacuum (~ 0.2 kPa) for 60 s, then mixed at a speed of 2000 rpm for 30 s, before the speed is decreased to 200 rpm and the mixture is mixed at 10.5 kPa for an additional 60 s. The prepared mixture is poured into a 42 mm x 42 mm x 4 mm acrylic mold and cured under the strong magnetic field generated by two ND42 magnets (CMS Magnetics) separated by 4 cm for about 2.5 hr before being removed from the mold. As the silicone cures, this process aligns the magnetic particles within the mixture, thus magnetizing the silicone block. Once cured, the orientation of the magnetic particles (and therefore the magnetism of the block) is fixed.

4.5.2 Magnetometer Circuitry

A low cost, off-the-shelf 3-axis magnetometer (LIS3MDL, ST Electronics) completes the magnetometry system by recording changes in the magnetic flux $\{b_x, b_y, b_z\}$ induced by the motion of the magnetized silicone block.

Magnetometers are purchased pre-mounted on breakout boards (LIS3MDL Carrier, Pololu) which are then mounted directly underneath the HASEL actuator and magnetic block. All 10 of the magnetometers within a module interface to that module's SCU via a shared SPI bus, with a clock speed set to 1 MHz using SPI Mode 0. Due to the physical length of the SPI bus, we found that substantial noise could manifest and interfere with the quality of sensor measurements from the 10th actuator in some modules. Future improvements to this system could include topological changes and the inclusion of terminating resistors on the SPI bus.

4.5.3 Magnetometer Calibration

A mapping from magnetic flux measurements to deformation is generated empirically for each HASEL actuator in the array. To do so, each HASEL in the array is driven with a voltage ramp from 0 kV to 8 kV over 8 s. The z-coordinates of both the magnetometer reading and of the MoCap data are sampled during this test. Afterwards, these samples are fed into a least-squares polynomial fitting function to generate a unique third-order polynomial fit for each cell.

This calibration procedure takes approximately 10 min and is performed one cell at a time. As a result, inaccuracies can be introduced when adjacent cells actuate due to the mechanical coupling of adjacent cells through the surface skin. To account for this, the estimated deformation for a given cell $\hat{z}_{i,j}$ is modified to be:

$$\hat{z}_{i,j} = (p_0 + p_1 b_z + p_2 b_z^2 + p_3 b_z^3) + a (\hat{z}_{i-1,j} + \hat{z}_{i,j-1} + \hat{z}_{i+1,j} + \hat{z}_{i,j+1}) \quad (4.3)$$

with scaling factor $a=0.05$. For edge cases (e.g., where $\hat{z}_{i,j}$ does not point to a valid cell), the adjacent \hat{z} values are set to zero.

4.5.4 Magnetometer Efficacy

The combined action of the magnetic sensors within all 100 cells results in a magnetic sensor array capable of accurate deformation mapping at rapid time scales, as shown in figure 4.16. The only limit to the resolution at which the s-tissue is capable of resolving surface deformations is the density of cells.




Figure 4.16: Accurate, high-frequency distributed measurement of deformation.

When compared to computer vision, the primary advantage of our magnetic sensing approach is that it functions equally well in dimly-lit or object-occluded environments. The downsides of our magnetic sensing approach include the disruptive influences of external magnetic materials (which can lead to inaccurate deformation mappings), and the fact that the display must be recalibrated to the magnetic signature of its environment whenever it is moved to a different location. Such a limitation could be crippling for highly mobile robots – however, additional research is needed to gauge the full significance and magnitude of these disruptions.

4.6 Displacement Control of ESAs

As discussed in section 3.7.3, the voltage-displacement curve of a HASEL actuator creates a time-variant hysteresis loop. This is a problem for voltage control, as it means that the voltage required to achieve a given displacement will vary between actuators, and will also vary over time for a given actuator due to charge retention effects.

We hypothesize that this hysteresis loop can be compensated for using a closed-loop feedback system which measures the displacement of the actuator and adjusts the voltage accordingly. Specifically, our objective is to compensate for charge retention effects and variability in actuator construction, as well as to provide a means of compensating for external force disturbances.

Recall equation 2.34, the equation of motion for a HASEL actuator:

$$M\ddot{z} + B\dot{z} + Kz = \frac{1}{2}C'(z)v^2 - Mg - f \quad (2.34)$$

This equation is both nonlinear and control-affine, and thus we don't expect linear controls to work particularly well. Nonetheless, we attempt to implement a linear control loop for displacement control, and see how it performs. Using the same methodology as described in chapter 3 modified to include the previously described voltage controller and measuring displacement using the MoCap system described in section 4.4, we use frequency domain system identification on the full set of 100 actuators to estimate a linear transfer function from voltage to displacement:

$$G_z(s) = 0.014 \frac{(24\pi)^2}{s^2 + 12\pi s + (24\pi)^2} \quad (4.4)$$

A bode plot of the data used to generate this transfer function is shown in figure 3.14. The transfer function G_z demonstrates second order dynamics, as we'd expect from the form of equation 2.34. Values of the damping ratio ζ and natural frequency ω_n can be extracted from the canonical form:

$$H(s) = k \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \quad (4.5)$$

Using the same loopshaping approach described in chapter 3.5, we use the model of G_z to design a linear displacement controller:

$$K_z(s) = 30 \frac{s + 70}{s} \quad (4.6)$$

This controller results in an open-loop gain $|K_z G_z| > 5$ under 20 Hz for tracking performance, and a gain $|K_z G_z| < 0.1$ over 300 Hz for noise rejection. We also include the prefilter:

$$F_z(s) = \left(\frac{200}{80}\right)^2 \frac{(s + 80)^2}{(s + 200)^2} \quad (4.7)$$

Following the procedure described in section 3.6, both the control law and the prefilter are discretized via ZOH and implemented at 200 Hz. In contrast to the voltage controller, which is implemented on the embedded microcontroller that is the DCU, the displacement controller is implemented within the high-level julia code on running on the central computer as discussed in section 4.3.

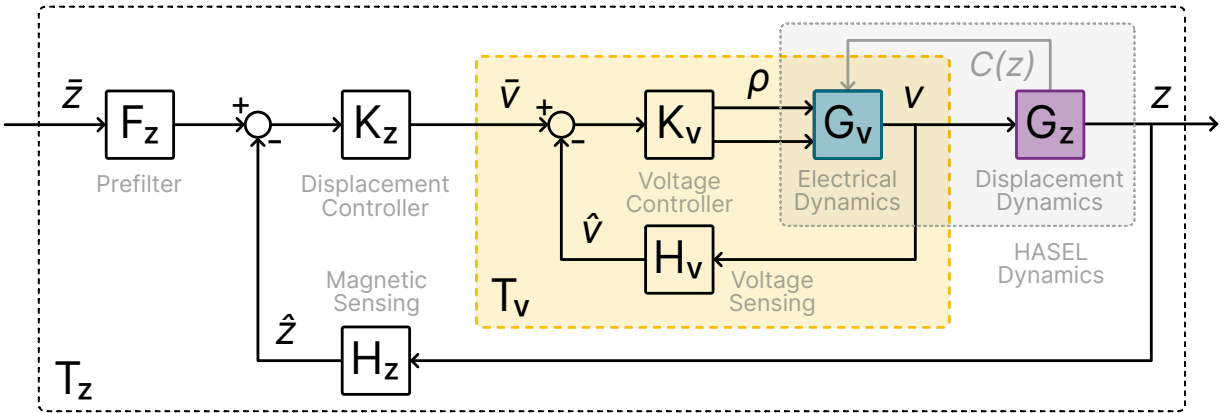


Figure 4.17: Block diagram of the nested displacement and voltage control loops highlighting their interactions with the decomposed ESA dynamical model.

Taken together with our previous results, we create the hierarchical control scheme shown in figure 4.17. The voltage control loop T_v introduced in chapter 3 becomes an inner loop surrounded by an outer displacement control loop T_z . The loop T_v has the same structure and components as described in section 3.7. From the perspective of the outer loop, it can be considered to be part of the plant.

4.6.1 Performance of Displacement Control




Figure 4.18: Comparison of voltage control and height control schemes under quasistatic conditions.

As expected, under quasistatic conditions, the position controller regulates the surface geometry to a greater degree of steady-state accuracy when compared to actuator voltage regulation alone, as shown in figure 4.18. The increased accuracy proves useful for applications requiring high-fidelity shape display, like the rendering of scrolling text shown in figure 4.19.




Figure 4.19: Aggregated photos demonstrating a scrolling text display.

However, the controller does not offer much benefit in the dynamic regime, and we can conclude that a linear control loop for displacement control provides little benefit beyond compensation for actuator variability under quasistatic conditions. This compensation is still useful, but it is not a particularly interesting result. Future implementations can likely skip the system identification and use a simple proportional controller to achieve the comparable results.

To improve dynamic performance, a nonlinear controller based on the full, nonlinear dynamics presented in chapter 2 would likely be necessary. For such a task, adaptive, sliding mode, or model-predictive control approaches could prove effective. This is left for future work.

However, as evidenced by figure 4.20, rapid control of surface geometry can just as easily be achieved with voltage control alone if a high degree of accuracy is not needed. This proves to be a common use case in soft robotics, confirming voltage control as a fundamental technology for electrostatic actuators.



Figure 4.20: Shape display by voltage control only.

4.7 Force Sensing

Recall equation 2.41, which relates the force applied to an ESA to the deformation and voltage of the actuator.

$$f = E_f(\bar{v}, z) = \frac{1}{2}C'(z)\bar{v}^2 - Kz - Mg \quad (2.41)$$

It suggests that the magnitude of an external force applied to an ESA under voltage regulation (\bar{v}) can be determined from a voltage measurement \hat{v} in conjunction with a displacement measurement \hat{z} . Thus, in addition to deformation sensing, the magnetometer also enables external force measurement when combined with HV sensor measurements from the HASEL driving circuit presented in 3.2, as shown in figure 4.21.

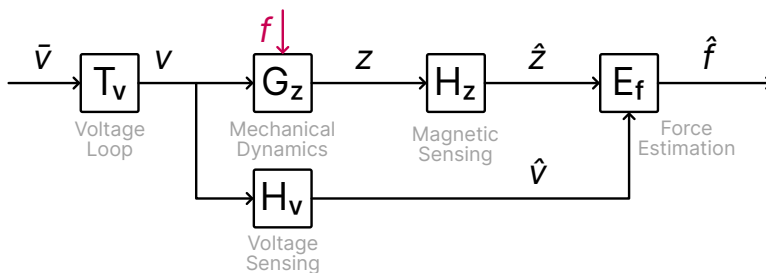


Figure 4.21: Block diagram of the force estimation scheme within the s-tissue.

4.7.1 Force Curve Fit

Rather than deriving an analytical expression for the force estimator $E_f(\hat{z}, \hat{v})$, we approximate it directly from experimental data using a high order polynomial fit. The estimator takes the form:

$$\hat{f} = E_f(\hat{v}, \hat{z}) = \sum_{i,j} p_{ij} \hat{v}^i \hat{z}^j \quad (4.8)$$

The force sensing map shown in 4.22 is created using a dynamical mechanical analyzer (DMA) (800E2, TestResources), which applied incremental deformation to a Peano-HASEL stack at a rate of 0.5 mm/s – slow enough to maintain quasi-static conditions. Each actuator produces approximately 2.5 N for small (sub-mm) deformations. The resolution of the force mapping is within 50 mN – which roughly corresponds to adding a 5 g mass to the surface.




Figure 4.22: Force sensing data and polynomial fit.

Prior to adding the force, the actuator receives a constant voltage of 4 kV and the baseline deformation is measured. The change in deformation and force are recorded from the DMA until a maximum of 25 N. This test is repeated for a sequence of HASEL input voltages ranging from 4 to 8 kV using three different actuators.

The mean data from all test actuators at each voltage is used to determine the polynomial surface fit using the MATLAB Curve Fitting Toolbox (MathWorks, 2021b) and resulting in a 15-term, 5th-order polynomial. The resulting coefficients are given in [69]. The force map is assumed to be the same for all cells.

4.7.2 Distributed Force Sensing

The display can thus self-sense distributed forces, for example to detect the size and shapes of various objects on the surface. Due to the continuous elastic skin, forces applied across two adjacent cells are averaged across the cells. One can thus estimate the total force of an object by summing the distributed forces measured by each cell.

The utility of this approach is perhaps best demonstrated in parallel with operation of the surface as a text display: by implementing an interactive scale which displays an object's mass in

realtime (figure 4.23). One region of the surface acts as the scale using embedded force sensing, and the other region acts as a text display which displays the mass (in grams) of detected objects based on the summation of forces across the scale region.




Figure 4.23: The s-tissue, banana force scale.

This simple implementation of a novel functionality also highlights the our system’s ability to operate in parallel; independent cells are capable of simultaneous, independent behavior. Our approach uses only z-axis magnetometry measurements. Force sensing capabilities could be further expanded by incorporating data from the additional magnetometer axes to measure surface shear forces [29, 153].

4.8 Emergent Behaviors

As repeatedly observed in the study of biology, sensorimotor integration at large scales tends to result in emergent behaviors [30, 31, 44, 85]. This hardware platform provides a case study for a similar phenomenon in soft robotics.

As shown by figure 4.23, 4.1, and others throughout this section, the combined actuating, sensing, and control capabilities of the s-tissue enable a rich range of behaviors which are not possible on systems with a lesser degree of sensorimotor integration. It is important to note that

although some of these behaviors were discussed during the early design phase as potential demos, the s-tissue hardware was not explicitly designed to perform any of them.

4.8.1 Dynamic Conveyance

Perhaps one of the best example of emergent behaviors resulting from the complex sensorimotor interactions of the s-tissue is shown in figure 4.24. Here, a camera mounted above the display is used to detect objects by color. The surface of the shape display then morphs to dynamically convey these objects across its surface, sorting them by color. A more detailed discussion of the hardware and software implementation of this demo is given in [69].




Figure 4.24: A demonstration of computer vision-aided interaction with objects with the s-tissue.

Due to its soft, continuous structure, such interaction is not limited to the discrete manipulation of objects, but can also be used to interact with liquids⁶, as shown in figure 4.25. This type of interaction is rather uncommon amongst robotic platforms.

⁶ Yes, we poured water on a 8 kV, 650 W, \$30,000 hardware stack which took 4 years to build.




Figure 4.25: A demonstration of the s-tissue interacting with liquids.

4.8.2 Distributed Audio and Haptics

Due to the high speed dynamics of the HASEL actuators, the s-tissue can generate audible sound up to 500 Hz, (just shy of the note C5), which is entirely limited by the polling rate of the firmware and communication stack. Changes to the firmware could theoretically increase this to about 10 kHz, and thus cover the majority of the human hearing range. Each actuator can independently play a unique waveform, allowing the system to play spatially distributed sounds. Unfortunately, the author has not found an elegant way to represent this result via pdf. Using the same mechanism, the s-tissue can generate haptic feedback, which, for example can be combined with the touch-based interaction shown in figure 4.26 to create highly interactive interfaces.

4.8.3 Multimodal Interaction




Figure 4.26: A demonstration of the force-sensing capabilities of the s-tissue.

Because the s-tissue can react to multiple modalities of external stimuli, including applied forces, displacements, and magnetic fields, it is possible to implement the force-based touch detection scheme shown in figure 4.26; forces detected by the right half of the surface are mirrored as actuator displacements on the left. Such capabilities have applications in haptics, object interaction, and teleoperated human interaction [32].

Likewise, the s-tissue can react to changes in external magnetic fields, as demonstrated by the demo shown in figure 4.27. In this demo, the fields generated by moving a magnet-tipped wand across the surface are detected, measured, and represented as actuator displacements, allowing a user to interactively draw on the shape display. In this operating mode, the s-tissue acts as a distributed magnetic flux sensor which can simultaneously form a physical representation of the flux density to create an intuitive visualization of unseen forces.

By interpolating between sensor values, the position of the wand tip can also be determined at a resolution greater than the physical resolution of the cells, suggesting a wide range of potential applications for this technology as a human interface device, for example as an input tool for digital

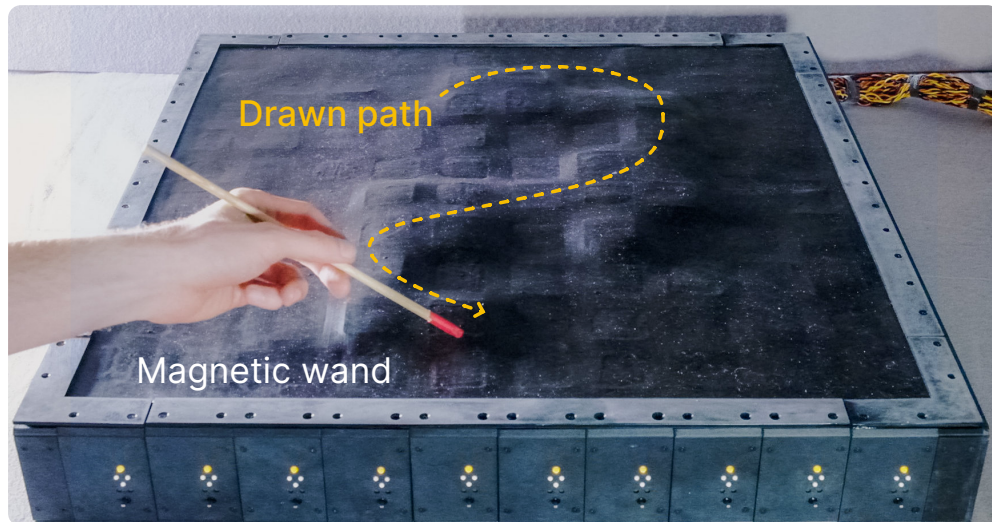


Figure 4.27: Drawing on the s-tissue using a magnet-tipped wand. The system instantly renders any detected magnetic fields as surface deformations.

illustration, or as a dynamically-reconfigurable array of buttons.

4.9 Conclusion

The s-tissue system presented in this chapter is a high-speed soft robotic shape display which combines electrostatic actuation, magnetic sensing, and feedback control in a highly scalable, parallelized architecture. This system demonstrates an implementation of feedback control at an unprecedented scale in soft robotics and validates our approach for parallelizing the control of electrostatic actuators.

An active surface is but a topology. Fundamentally, the s-tissue is an ESA-driven soft robot with 100 channels of parallel feedback control closed for both actuator state variables. As such, the hardware stack on which it is based and the lessons learned in its design could also be readily applied to the design of other high DoF soft robots like continuum manipulators [150] or bioinspired systems [125].

The integration of soft robotic sub-systems at this scale results in a multitude of emergent capabilities including high-fidelity shape display, traveling waves which can transport both solids and liquids across the surface, and complex multimodal user interactions; together, they hint to

the rich multifunctionality that highly parallelized soft robots can achieve.

4.9.1 Contributions

The s-tissue system presented in this chapter is the result of the work and intellectual contributions of a large team of individuals. Our final results are presented in [69], and an outline of each author's contributions follows. M.N. led the design of the shape display; all authors contributed towards its design and conceptualization. S.K.M., E.A., and N.K. designed the HASEL actuator. M.N. and S.K.M. designed the voltage control scheme, HASEL driver circuit, and power supply system. V.S. and K.L. designed the sensor system and circuitry. M.N. wrote the firmware, software, and communications framework. B.K.J., M.N., and A.V. designed the closed loop control laws. M.N. designed the mechanical components of the shape display and integrated all subsystems. M.N., B.K.J., V.S., K.L., and A.V. built the shape display. M.N., B.K.J., V.S., K.L., and A.V. developed the experimental methodology and analyzed data; M.N., B.K.J., V.S., A.V., and K.L. wrote code for data collection and analysis; B.K.J., M.N., V.S., K.L., and A.V. implemented demos and experiments. M.N., B.K.J., and V.S. designed the figures; B.K.J. performed the photography and videography; M.N. and B.K.J. created the illustrations. M.E.R., C.K., N.C., and J.S.H. supervised and administrated the project. The paper was written by B.K.J., M.N., K.L., V.S., and A.V.; all authors contributing to editing.

Chapter 5

Solid State Drivers for Electrostatic Actuators

Having demonstrated the requirements and benefits of controlling ESAs in parallel we turn to a more practical question: how to adapt our approach for mass manufacturing. The missing link for widespread use of ESAs with closed-loop control are compact, efficient, high-voltage, low-current switches. To date, no such design has been demonstrated – let alone designs which are amenable to standard mass manufacturing techniques – placing hard practical limits on the construction of ESA-based multi-actuator robots. This chapter aims to rectify that.

5.1 Switching High Voltages

Switching amplifiers can be almost perfectly efficient if the right conditions are met, a particularly important trait when constrained by a tightly limited HV power budget imposed by the HV-SMPS – the bulkiest, most expensive part of the system. An ideal switch does not consume power when open, as it blocks voltage without passing current; nor does it consume power when closed, as it passes current without a voltage drop. The only time such an idealized switch could consume power is during transitions between the two states. Thus, the critical parameters for maximizing the efficiency of a practical switch are low on state impedance, high off state impedance, and fast switching speed.

The most commonly used switch is a HV optocoupler, which works by using a low voltage LED to vary the reverse conductance of a reverse-biased HV photodiode via the photoelectric effect [99, 138, 158]. These switches can be highly compact, and have been used throughout this work.

However, those switches suitable for switching ESAs are made from rare, expensive components which are only produced by a single manufacturer (Voltage Multipliers Inc.), and require time consuming, labor-intensive assembly. Furthermore, they require optical isolation in addition to electrical isolation, and their relatively high on-state impedance reduces efficiency below 50%, as discussed further in section 5.7.3. Combined with significant low voltage power draw for the LED, this inefficiency generates enough excessive heat to cause runaway thermal failures, forcing a tradeoff between performance and the inclusion of bulky heatsinks. This phenomenon is discussed further in sections 3.1.5 and 5.7.3.

One alternative is to use electromechanical switches, such as reed relays. These operate by using an electromagnet to pull a spring-loaded switch open or closed. They can be highly efficient, but are relatively bulky. Furthermore, the magnetic fields they generate mandate minimum spacing between switches to prevent interference, further increasing their effective footprint. They can be damaged by the surge currents commonly seen when driving capacitive loads, such as ESAs. The author is not aware of any works in the literature which have used these switches in the context of soft robotics.

Had this work been written 50 years earlier, the author would explore the use of vacuum tubes such as krytrons, thyratrons, or planar triodes [5]. But these are no longer manufactured at large scales, and are bulky, expensive, and have limited operational lifespans.

Another option is to use solid-state transistors, such as metal oxide semiconductor field effect transistors (MOSFETs), insulated gate bipolar transistors (IGBTs), or burlington junction transistors (BJTs). One could argue that the transistor is the foundation of modern day electronics, and they can be found in just about any electronic device made in the past few decades. However, transistors are typically designed and intended for applications requiring much lower voltages and much higher currents. Among transistor types, MOSFETs achieve the highest blocking voltages, but even the latest silicon-carbide (SiC) MOSFETs are limited to about 4kV in reasonably compact form factors. These high voltage MOSFETs are expensive, specialty devices. With further advancements in materials science, peak MOSFET switching voltage continues to increase, and

ESA operating voltage continues to decrease [37, 144]. In the future, MOSFETs should become the obvious choice for driving ESAs, as their switching speed, form factor, and efficiency are hard to match. In the present, individual MOSFETs are not a viable option.

5.1.1 Stacked Series MOSFETs

Although the blocking voltage of individual MOSFETs is limited for a given package size, MOSFETs have been stacked in series to achieve higher blocking voltages in a variety of applications ranging from instrumentation [5], to buck converters [11, 108], to high speed pulse generators [6, 93], to plasma containment systems for nuclear fusion research [10].

While varied in their implementation, these approaches all share the same fundamental requirements to prevent damage to the MOSFETs in the stack: all devices must be coupled to switch with precisely synchronized timings, and voltage must be evenly distributed across each device in the stack at all times, including during switching transients. This can be done either with gate-side techniques – which mainly aim to synchronize the switching of devices, load-side techniques – which directly balance voltages, or a combination of both.

One such hybrid technique is active balancing. This involves measuring closing the loop around load-side measurements of voltage or rate of voltage change, and compensating via subtle adjustments to the timing of each gate [35, 90, 137]. This method is the by far the most complex, and is perhaps best suited for high power applications with extremely limited tolerance for imbalance. Driving ESAs does not fall into this category. For most other applications, if devices are precisely synchronized, reliable load-side balancing can be achieved entirely using passive components, such as an RC ladder [36, 39, 50]. Assuming the addition of such a ladder, we can focus mainly on the problem of synchronization using gate-side techniques.




Figure 5.1: Transformer-based synchronization of series MOSFETs. A current pulse through the shared primary of each transformer creates simultaneous gate drive pulses on each individual secondary.

Perhaps the most straightforward gate-side approach is transformer-based gate coupling [10, 88]. In this method, each MOSFET gate is connected to the secondary of a transformer, and all primaries are connected together in series. A trigger pulse sent through this chain of primaries appears simultaneously on the gate of each device, and they switch in near-perfect synchronization. However, this approach requires bulky, high isolation voltage transformers, and is therefore better suited to applications such as high energy physics where size is not a predominant constraint.

A particularly attractive method is capacitive gate coupling [5, 6, 39, 49, 50, 97, 108]. In this method, external capacitors are added to create a voltage divider with the parasitic capacitance of each MOSFET gate, thus coupling them together. This method is simple, effective, and requires minimal components. It has been demonstrated both with discrete devices [5, 49, 50], and as a custom IC implemented on a silicon-on-insulator (SOI) CMOS process [11, 97, 108]. Recently, this method has been demonstrated¹ to work for femtosecond lasers in the voltage ranges needed to drive ESAs [93]. This is the method we use, and is discussed further in section 5.3. It should be noted that numerous other methods exist, but are less relevant to this context [15, 72, 90, 109, 110, 117, 146, 152].

¹ This paper's authors neglect to cite **any** of the existing literature even though they clearly and exactly copy Hess' & Baker's original design. Shame on them.

5.1.2 Electrons and Holes

Typically, bridge circuits are built with complimentary MOSFET (CMOS) pairs: one P-Type (PMOS) and one N-Type (NMOS). This, as shown in figure 5.2 is done for convenience given the requirements for driving these two types of device.

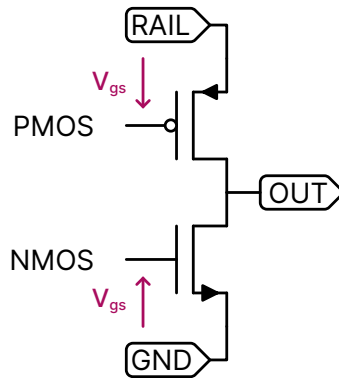


Figure 5.2: A CMOS half-bridge, constructed from a pair of PMOS and NMOS devices.

From the perspective of gate design, it is easier to use PMOS devices for the high-side switch(es) in the bridge. To activate, a PMOS needs a negative gate-source potential. Thus, driving the PMOS is a matter of pulling the gate potential low relative to the source, which is typically the positive rail. This is typically done using a AC-coupling capacitor, or using a pull-down circuit like the one shown in figure 5.3.

In contrast, a NMOS device needs a positive gate-source potential to activate. Since the source of the NMOS is typically the negative rail (or ground in single-ended applications like ours), driving a low-side NMOS is generally trivial. This is especially true when their gate threshold voltage is low enough to be driven directly from a logic-level input. However, push-pull gate driver circuits are often used to either improve switching speeds, or to provide higher gate drive voltages.




Figure 5.3: Perhaps the most common PMOS load switch driver topology.

However, these two types of MOSFET are not equivalent opposites. In the realm of semiconductor manufacturing, it is well-established that NMOS devices tend to offer higher blocking voltages and lower on-state resistances when compared to their PMOS counterparts. This phenomenon can be primarily attributed to the inherent disparities in the mobility of charge carriers in the silicon substrates of these devices: electrons, which are the majority charge carriers in NMOS devices, tend to have higher mobility in silicon than holes, which are the majority charge carriers in PMOS devices.

In spite of the advantages to NMOS devices, CMOS architectures are nonetheless ubiquitous, as substantial challenges arise when trying to build bridge-style circuits using purely NMOS devices. The source of the high-side NMOS becomes connected to the output of the bridge, which varies in time, thus requiring a gate driver which is always capable of providing a positive potential relative to this rapidly varying source voltage.

The typical approach is to use a charge pump. However, the design of such a circuit becomes prohibitively difficult when the source voltage is expected to violently swing from 0 to almost the HV rail, several thousand volts possibly several thousand times per second. A transformer could be used, but this approach is bulky, expensive, and doesn't provide complete isolation. A battery could be used, but is also bulky and would eventually need to be replaced.

A circuit capable of generating a small positive potential relative to this varying source voltage would allow the construction of bridge circuits using entirely N-channel MOSFETs which could take advantage of their superior performance characteristics - namely increased blocking voltages. This is especially relevant to soft robotics due to the high voltages required to drive ESAs. Our design to accomplish this is presented in the following sections.

5.1.3 Our Approach

To enable the control of arrays of high voltage capacitive loads, and thus the construction of complex, ESA-driven, multi-actuator soft robots we introduce a novel high voltage switch shown in figure 5.4.

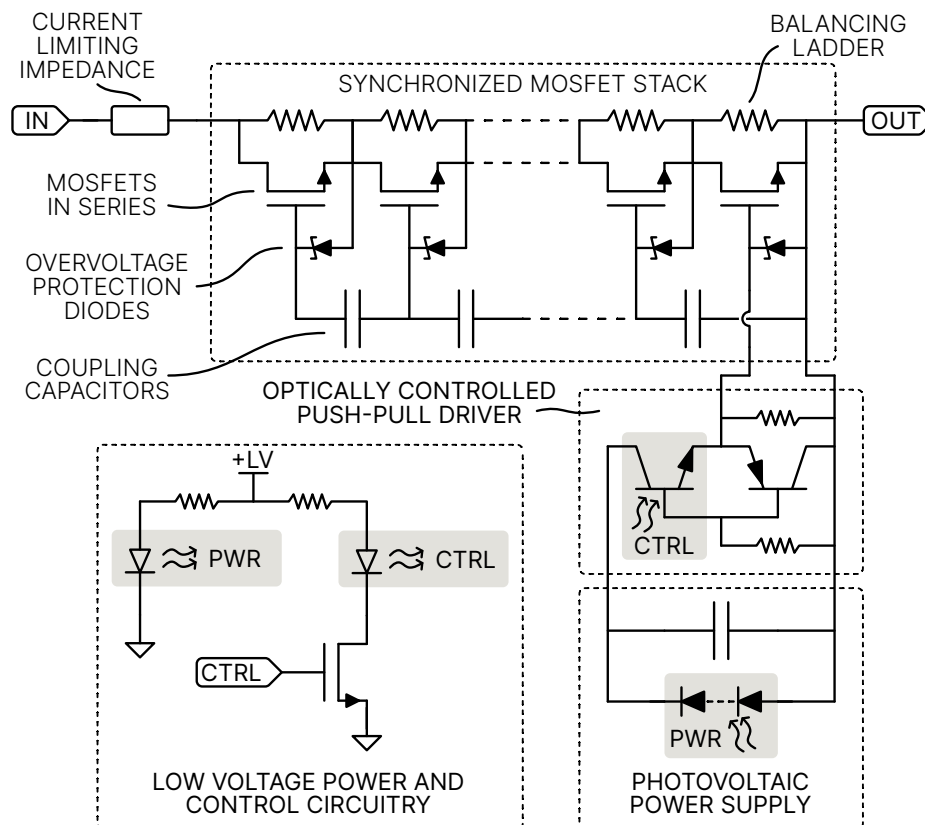


Figure 5.4: A novel high voltage load switch based on an optoelectronic gate driver and a capacitively-coupled MOSFET stack. (US Provisional Patent No. 63/608,794)

It is constructed from a series stack of N-channel MOSFETs which are synchronized by capacitive gate coupling. A resistive balancing ladder in series with the MOSFET stack passively balances load-side voltages during switching transients.

The entire stack is controlled using a novel floating optoelectronic gate driver, which is both powered and controlled using light for complete electrical isolation. This driver floats along with the output of the switch, allowing us to use identical switches (made entirely with N-channel MOSFETs) for both the high- and low-side switches of bridge-style amplifiers, such as the half-bridge ESA driver shown in figure 5.5. With the addition of a voltage sensor such as the one introduced in section 3.2, this circuit would be well-suited for closed-loop voltage control of ESAs.




Figure 5.5: A half-bridge ESA driver implemented using two MOSFET stack switches.

5.2 A Floating Optoelectronic Gate Driver

The optoelectronic gate driver shown in figure 5.6 allows a floating load switch to be controlled by a low voltage digital control signal while maintaining isolation from the high voltage switching circuitry, which is necessary for the practical implementation of bridge-style amplifiers.

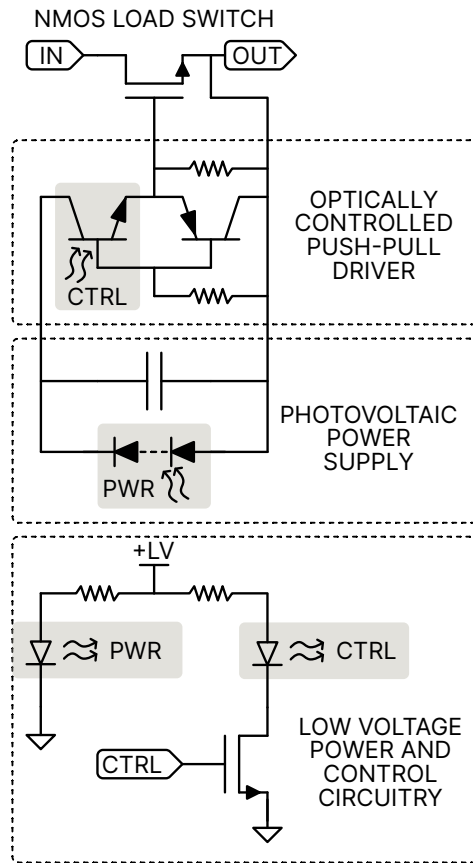


Figure 5.6: The minimal implementation of an optically-powered, optically-controlled gate driver.

As the energy requirements to switch a single MOSFET (especially a small, signal-level MOSFET as opposed to a power MOSFET) are comparably small, a photovoltaic is an excellent choice for wirelessly providing power. Referencing the photovoltaic to the bridge's output achieves a fully isolated power supply which floats along with the output of the bridge, exactly as desired. This approach is compact, effective, and inexpensive. The following sections discuss the various subcircuits of this gate driver in further detail.

5.2.1 High-Side Gate Driver Circuit




Figure 5.7: High side of the optoelectronic gate driver circuit. Both the control signal and the power supply are transmitted as light, providing robust electrical isolation to the low-side circuitry.

A photovoltaic optocoupler (TLP3910, Toshiba) and a capacitor ($0.1 \mu\text{F}$) provide a small local power supply which floats along with the low-side potential of the switch, and thus along with the source pin of M_x , the bottom MOSFET in the stack. However, this photovoltaic optocoupler has relatively slow turn-on and turn-off times. Therefore, a push-pull driver constructed from a phototransistor optocoupler (TCLT1103, Vishay), a PNP transistor, and a pair of pull-down resistors ($200 \text{ k}\Omega$, $1 \text{ M}\Omega$) is used to drive the gate of the first MOSFET. This design takes advantage of the high speed photodiode broken out by the base pin of the TCLT1103 to achieve much faster switching than could be achieved by switching the photovoltaic optocoupler directly.

The photovoltaic optocoupler needs to provide enough energy to fully charge and discharge the gate of the bottom MOSFET in the stack at a switching frequency f_s . The energy stored in the gate to turn the MOSFET on is calculated from the gate charge Q_g and the gate voltage V_g according to the relation:

$$E_g = \frac{1}{2} Q_g V_g \quad (5.1)$$

Thus, the minimum power output of the photovoltaic is given by:

$$P = \frac{f_s Q_g V_g}{2\varepsilon} \quad (5.2)$$

where ε is a loss coefficient, which is 1 for a perfectly efficient system, or less than 1 to account for

by a NMOS-based LED driver linked to the input pin of the gate driver. This part is designed to receive 40 mA of current in the on-state, also less than the maximum 60 mA allowed by the datasheet. Increased current would likewise increase the high-side current when activated, leaving some headroom for design optimization. It should also be noted that the TCLT110X series includes comparable transistor-output optocouplers with alternate current transfer ratios.

A wide range of NMOS devices could be used to complete the circuit shown in figure 5.8, we use (PMV20EN, Nexperia). This design provides a high input impedance, allowing this circuit to be driven from virtually any device capable of outputting 3.3 V or 5 V logic. To facilitate debugging, an indicator LED is driven in parallel with the LED of the TCLT1103, and activates whenever the MOSFET stack is intended to activate.

All nodes of this circuit exposed to the high-side circuitry include transient voltage suppression (TVS) diodes (SP4021-01FTG, Littelfuse) to clamp any transients or charge buildups which may occur during operation.

5.2.3 Gate Driver Performance

We next look to quantify the performance of the presented gate driver. Of primary interest is its ability to drive a MOSFET to saturation at sufficiently fast frequencies for PWM control.

Figure 5.9 shows the effects of the gate driver on the switching dynamics of a N-channel MOSFET (BSS127, Diodes Inc.). It compares the behavior of a MOSFET driven through our gate driver with the behavior of the same MOSFET driven directly by an idealized signal. For testing, the MOSFET is connected as a low-side switch for a 200 k Ω load resistor on a 300 V rail. The rail is provided by a HV-SMPS (610E, Trek) set to supply mode and augmented with a 15 nF capacitor (FKP 1, WIMA) for enhanced stability. Voltage data is measured directly by an oscilloscope (SDS1202X, Siglent) using a 10 \times probe. Current data is calculated from the voltage drop across the load resistor.

For both tests, the input waveform is a 10 kHz, 50% duty cycle square wave generated by a function generator (33500B Series, Keysight). 10 kHz is selected to correspond to a realistic

frequency range for PWM control. For the directly-driven test set, the function generator output is set to range from 0 V to 10 V, and directly connected to the MOSFET gate. For the gate driver test, the function generator output is set to range from 0 V to 3.3 V, comparable to the digital output of a microcontroller, and connected to the gate driver input. Because the nominal voltage of the input signal varies between tests, figure 5.9 shows a normalized input signal.

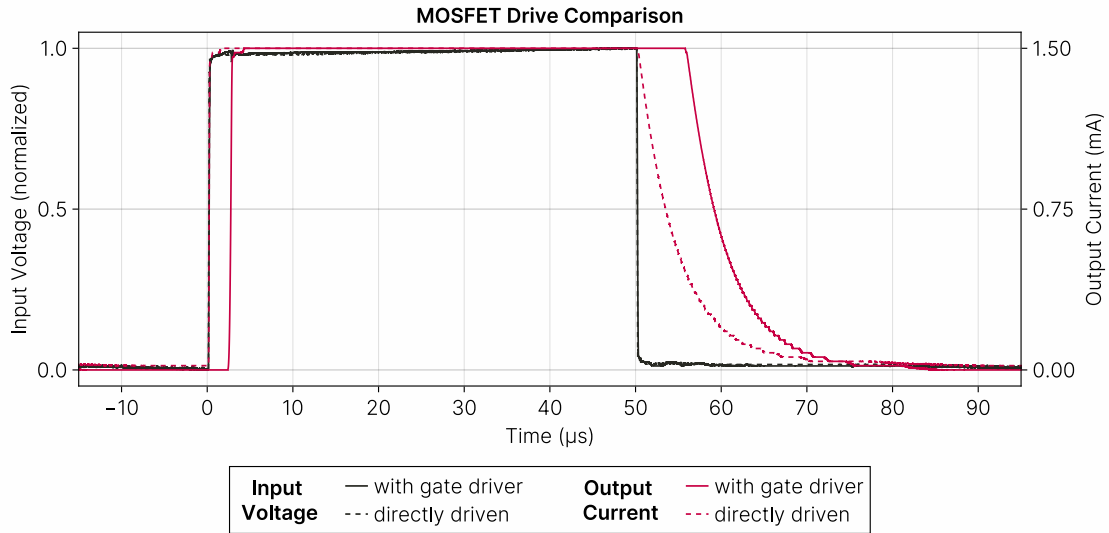


Figure 5.9: Timing changes introduced by our gate driver to the switching of a BSS127 MOSFET.

An RC-like time constant appears in the turn-off of the MOSFET. This is a function of the MOSFET's C_{ds} and the impedance of the load path, which for this test, is dominated by the 200 kΩ load resistor. This time constant is not significantly affected by the gate driver, and will not necessarily manifest for a predominantly capacitive load such as an ESA.

The gate driver introduces an approximately 3 μs delay in turn-on, and an approximately 6 μs delay in turn-off. Such delay will not have a substantial effect for PWM control in the 10 kHz range. Likewise, the slight (~6%) extension of the pulse width introduced by the driver is fairly straightforward to mitigate in software. Thus, based on this result, the gate driver is suitable for PWM control.

Next, we take a more detailed look at the timing of gate driver operation, specifically, by

comparing the input signal to the voltage output of our gate driver circuit. Other than measurement location, the test setup is identical to the previously described test. As shown in figure 5.10, the gate voltage reaches approximately 11.5 V in response to a 10 kHz square wave with 50% duty cycle. This is sufficient for a BSS127 MOSFET, which enters saturation above approximately 5 V, while remaining comfortably below this device's maximum gate-source voltage of 20 V.




Figure 5.10: A profile of the response time of our gate driver circuit.

The delay time t_{dly} is the time between the rising edge of the input pulse and the point at which the gate voltage begins to rise, defined as the point it reaches 5% of its on-state value. The rise time t_{rise} is defined as the time it takes the gate voltage to rise from 5% to 95% of its on-state value. The turn-on time t_{on} is the sum of t_{dly} and t_{rise} . Each of these metrics has an analogue for turn-off. The storage time t_{strg} is the time between the falling edge of the input pulse and the point at which the gate voltage begins to fall, defined as the point it reaches 95% of its on-state value. The fall time t_{fall} is defined as the time it takes the gate voltage to fall from 95% to 5% of its on-state value. The turn-off time t_{off} is likewise the sum of t_{strg} and t_{fall} .

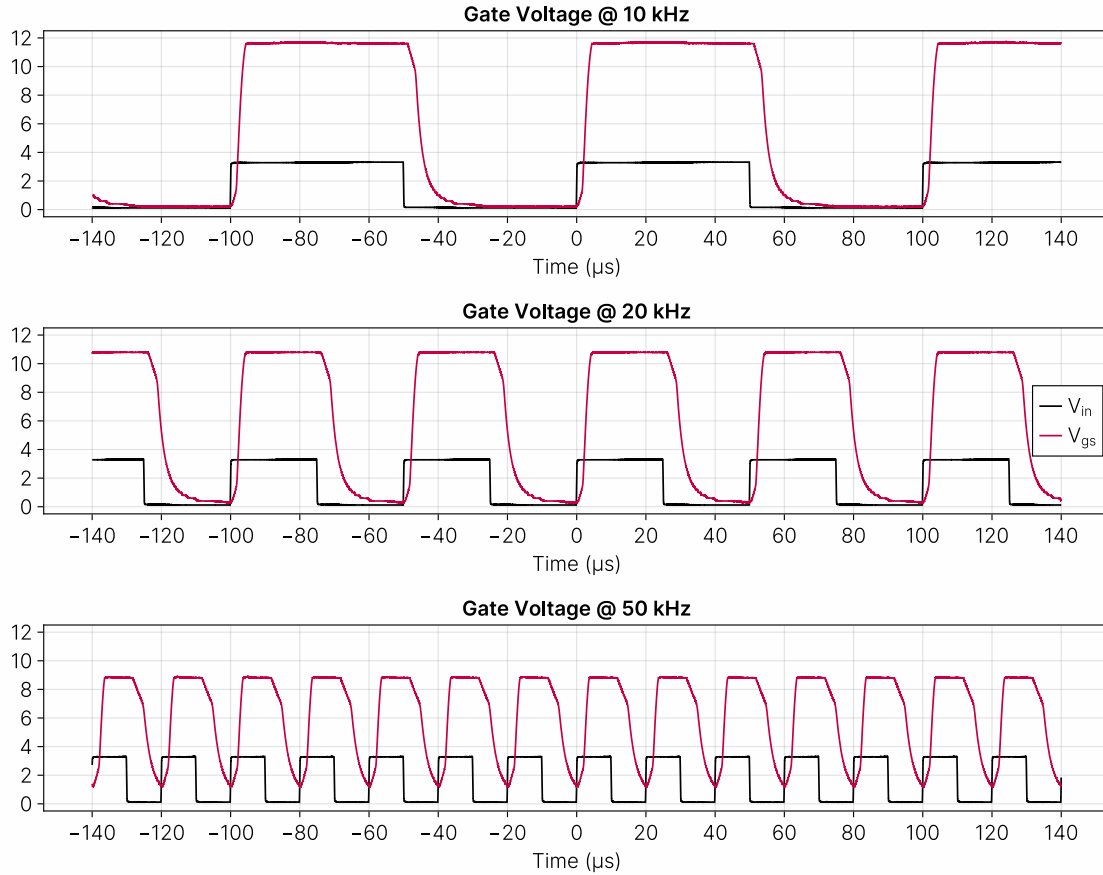


Figure 5.11: Degradation in gate driver performance at higher frequencies.

As presented, this gate driver proves suitable for PWM inputs at 10 kHz, with performance starting to degrade above 50 kHz as visible in figure 5.11. Also tested, but not shown in this figure is the ability of the gate driver to maintain a constant on-state at about 8 V. Switching delays could be reduced by changing the (nominally 200 k Ω and 1 M Ω) resistor values within the push-pull driver on the high-side of the overall gate driver circuit shown in figure 5.7. With additional parameter tuning, this design could likely operate at 100 kHz or more.

5.3 MOSFET Synchronization by Capacitive Gate Coupling

The switching of the stack is synchronized by capacitively coupling the MOSFET gates using a topology originally presented in [5, 50]. As shown in figure 5.12, this approach is a matter of

voltage division between the total equivalent parasitic gate-source capacitance C_g of the MOSFETs themselves and a coupling capacitance C_c added between MOSFET gates. The equivalent gate capacitance is found by:

$$C_g = C_z + C_{gs} + C_{gd}(A_v + 1) \quad (5.3)$$

Here, C_z is the forward capacitance of the Zener diode added to clamp any transient overvoltages on the MOSFET gate, $C_{gs} = C_{iss} - C_{rss}$ is the MOSFET gate-source capacitance, $C_{gd} = C_{rss}$ is the MOSFET gate-drain capacitance, and A_v is a voltage ratio used to approximate the Miller effect:

$$A_v = \frac{\Delta V_{ds}}{\Delta V_{gs}} = \frac{v_o}{v_d + v_g} \quad (5.4)$$

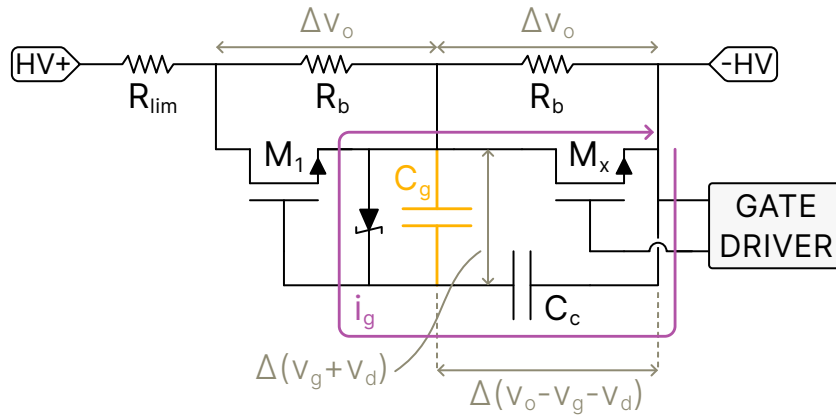


Figure 5.12: The mechanics behind the capacitive coupling of two series MOSFETs, highlighting the parasitic gate capacitance C_g at the core of the method.

Consider the two MOSFET stack shown in figure 5.12. When the switch is open, v_o is the potential blocked by each MOSFET. Assuming even voltage balancing, a switch consisting of a N -MOSFET stack will thus block a potential of $N \cdot v_o$. When the switch is closed, the potential blocked by each MOSFET is zero, assuming the current limiting resistor R_l is much larger than the $R_{ds,on}$ of each MOSFET. Thus, each MOSFET will see a change in drain-source potential of v_o over the course of a switching event.

When the switch is open, the Zener diode is forward-biased, and the gate voltage is $-v_d$, determined by the forward voltage of the diode v_d . The switch is driven fully closed at a desired

gate drive voltage of v_g . Thus, each MOSFET will see a change in gate-source potential of $v_g + v_d$ over the course of a switching event. When the switch is open, the two capacitors are in series, storing a total charge of:

$$Q_{off} = C_c (v_o - v_d) - C_g v_d \quad (5.5)$$

When the base MOSFET M_x is closed, the two capacitors are now connected in parallel, and coupling capacitor discharges through the gate of MOSFET M_1 , until both share a potential of v_g , with the charge stored:

$$Q_{on} = (C_c + C_g) v_g \quad (5.6)$$

By conservation of charge, $Q_{off} = Q_{on}$, and therefore:

$$C_c v_o - C_c v_d - C_g v_d = C_c v_g + C_g v_g \quad (5.7)$$

rearranging gives:

$$C_c (v_o - v_d - v_g) = C_g (v_g + v_d) \quad (5.8)$$

and finally:

$$C_c = C_g \frac{v_g + v_d}{v_o - v_d - v_g} \quad (5.9)$$

Thus, to synchronize a stack for a given C_g and A_v , the coupling capacitance C_c should be selected according to the ratio:

$$C_g = C_c (A_v - 1) \quad (5.10)$$

$$C_c = C_g \frac{1}{A_v - 1} \quad (5.11)$$

Now consider the N -device stack shown in figure 5.13. The conditions, requirements, and thus coupling capacitance $C_{c,1}$ needed to activate the top MOSFET M_1 is equivalent to the value of C_c given by equation 5.11. However, because the gate of each subsequent MOSFET undergoes the same voltage change, requires the same amount of charge, and consequently the same gate current i_g , it follows that each stage of the capacitive coupling ladder must provide proportionally more current. The coupling capacitance $C_{c,n}$ added to the n th MOSFET is therefore scaled to:

$$C_{c,n} = n \cdot C_c \quad (5.12)$$




Figure 5.13: Capacitive coupling of N MOSFETs. Current scales as we descend down the ladder.

Any voltage imbalance resulting from imperfections in synchronization is handled by the balancing ladder, which is discussed in 5.3.2.

5.3.1 Choice of Coupling Capacitance

This method was originally intended – and thus far has only been demonstrated – for resistive and inductive loads [50], where both the on- and off-state voltages are approximately constant between switching events. However, we are primarily interested in building bridge-style amplifiers for predominantly capacitive loads such as ESAs, where there is no such constant voltage. These voltages will vary between switching events dependent on the instantaneous state of charge stored in the load. Consequently, the coupling capacitance C_c must be selected for the switch to operate across a range of blocking voltages.

For a given blocking voltage, under-estimation of the coupling capacitance results in incomplete activation of some MOSFETs in the stack. If substantial, this can lead to voltage imbalances and damage to the stack. Over-estimation of the coupling capacitance results in excessive gate voltages, which could also damage the MOSFETs. However, the Zener diodes will bleed off these excesses, and the over-estimation will generally lead to a greater operating range at the cost of some added inefficiency in the form of energy loss through the Zener diodes. Thus, it is generally

better to over-, rather than under-estimate the added coupling capacitance. However, a conclusive answer to this issue will only be possible given more rigorous longevity and failure analysis.




Figure 5.14: The induced change in gate-source voltage as a function of change in drain-source voltage shown for various values of coupling capacitance C_c .

From the plot shown in figure 5.14, 1.9 pF is the optimal value of C_c for operating a $N \times 500$ V load switch. However, higher capacitances are helpful for operating at increasingly smaller blocking voltages, such as for half bridges operating closer to the rail. Given the benefit of over-estimation and the constraints imposed by the availability of components in convenient multiples, we use 2.7 pF as the base value of C_c for our prototype hardware.

5.3.2 Voltage Balancing

As visible in figures 5.4, 5.1, 5.12, and 5.13, series MOSFET stack implementations often implement a resistive balancing ladder in parallel with the MOSFETs. This serves two roles. It ensures even voltage division during any transients the switch may experience in the off-state. It

also corrects for any charge imbalances resulting from imperfections in switching synchronization.

The primary trade-off of a resistive balancing ladder is between the magnitude of the leakage current and the transient balancing speed. The use of a resistive ladder introduces a leakage current which may be undesirable when driving capacitive loads with limited power budget. As shown in figures 5.15 and 5.16, the larger the resistance, the slower the stack will self-correct any charge imbalances. However, the larger the resistance, the smaller the leakage current. The optimal value of the balancing resistance is therefore a matter of the specific application's requirements.

Figures 5.15 and 5.16 show two perspectives of the same test. Both tests are simulations conducted in LTSpice, using models and parameters detailed and validated in section 5.4. Likewise, both tests simulate a stack of $8 \times$ BSS127 MOSFETs acting as a high-side load switch for a $200 \text{ k}\Omega$ load resistor on a 4 kV rail simulated as an ideal voltage source. The gate driver discussed in 5.2 is not used for these tests. Rather, the MOSFET stack is activated by two $50 \mu\text{s}$, 10 V pulses, simulated as an ideal, floating voltage source.

Figure 5.15 shows $V_{d,i}$, the drain voltage of each MOSFET, M_i , in the stack. Also shown is V_{out} , the output voltage of the load switch, and equivalently, the source voltage $V_{s,x}$ of the base MOSFET in the stack, M_x . For this test, a coupling capacitance of $C_c = 1.9 \text{ pF}$ is used.

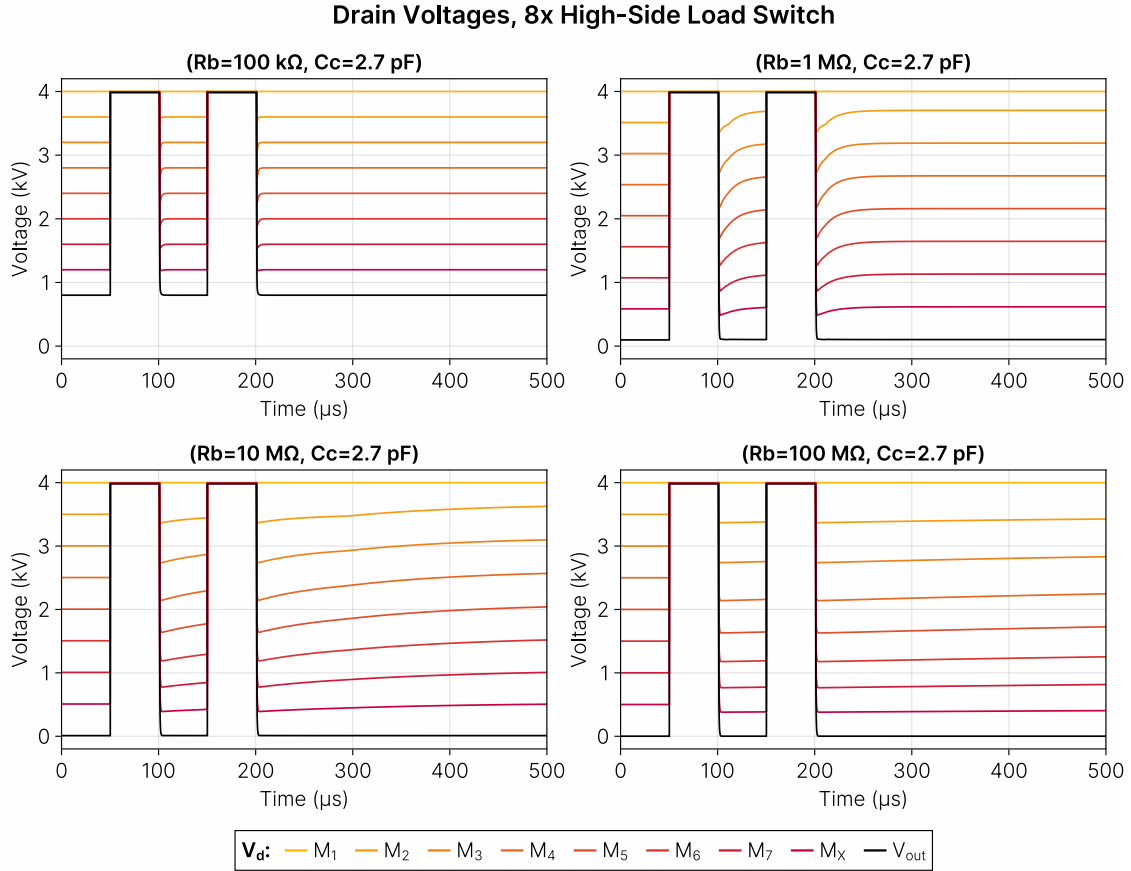


Figure 5.15: Effect of balancing resistance R_b on the drain voltage of each MOSFET within a coupled series stack, shown switching a resistive load.

Figure 5.16 shows $V_{ds,i}$ the drain-source potential across each MOSFET, M_i , in the stack. This perspective serves to highlight the voltage balance between the stacked MOSFETs. Ideally, these voltages would always identically divide the overall potential across the stack. In practice, the balancing speed depends on the magnitude of the balancing resistor R_b . For this test, a coupling capacitance of $C_c = 2.7$ pF is used. The different coupling capacitances are shown to demonstrate the efficacy of capacitive coupling across a relatively wide range of values.




Figure 5.16: Effect of balancing resistance R_b on the drain-source potential across each MOSFET within a coupled series stack, shown switching a resistive load.

The transient balancing speed is predominantly governed by the RC time constant created by the balancing resistors (R_b) and the effective output capacitance (C_{ds}) of each MOSFET. Thus, MOSFETs with minimal output capacitance are desirable, as they will achieve a given transient balancing speed with larger balancing resistors and therefore smaller leakage current.

Because the leakage current is also a function of the number of MOSFETs in the stack, it is beneficial to use many small MOSFETs instead of a few large ones. Many small MOSFETs will contribute smaller fragments to the overall resistance of the balancing ladder, and thus facilitate achieving both relatively fast transient balancing and high off-state impedance. These smaller MOSFETs also typically have smaller output capacitances, and can therefore support larger bal-

ancing resistors.

For our prototype hardware, we use $12 \times 100 \text{ M}\Omega$ resistors for an overall off-state impedance of $1.2 \text{ G}\Omega$. At a 6 kV blocking voltage, this results in a leakage current of $5 \mu\text{A}$, or a perfectly acceptable 0.1% of the maximum current output of a suitable power supply (eg. Ultravolt 6AA-24P30, Advanced Energy).

To further reduce the leakage current, it may be possible to use alternative balancing ladders. A capacitive ladder combines theoretically infinite DC impedance with very low transient impedance, but does not provide steady-state balancing. However, a parallel RC network could be used to leverage the advantages of both.

A diode-based ladder could also achieve an optimal combination of steady-state balancing and high impedance by leveraging the nonlinear voltage-dependent leakage effects of diodes. If the requirements are small, it may even be possible to achieve robust balancing using nothing other than the MOSFET body diode itself, and thus remove the need for a balancing ladder entirely. However, such an approach could just as easily irreparable damage the MOSFETs, and requires additional investigation.

5.3.3 Half-Bridge Voltage Balancing

We are explicitly interested in using this switch design for the construction of bridge-style amplifier circuits for capacitive loads, with output waveforms and behavior like that shown in figure 5.17. The performance of capacitively coupled MOSFET stacks under such operating conditions has not been explored by the existing literature. Nonetheless, it is critical that the synchronization and voltage balancing mechanisms remain robust for operation within bridge amplifier circuits.

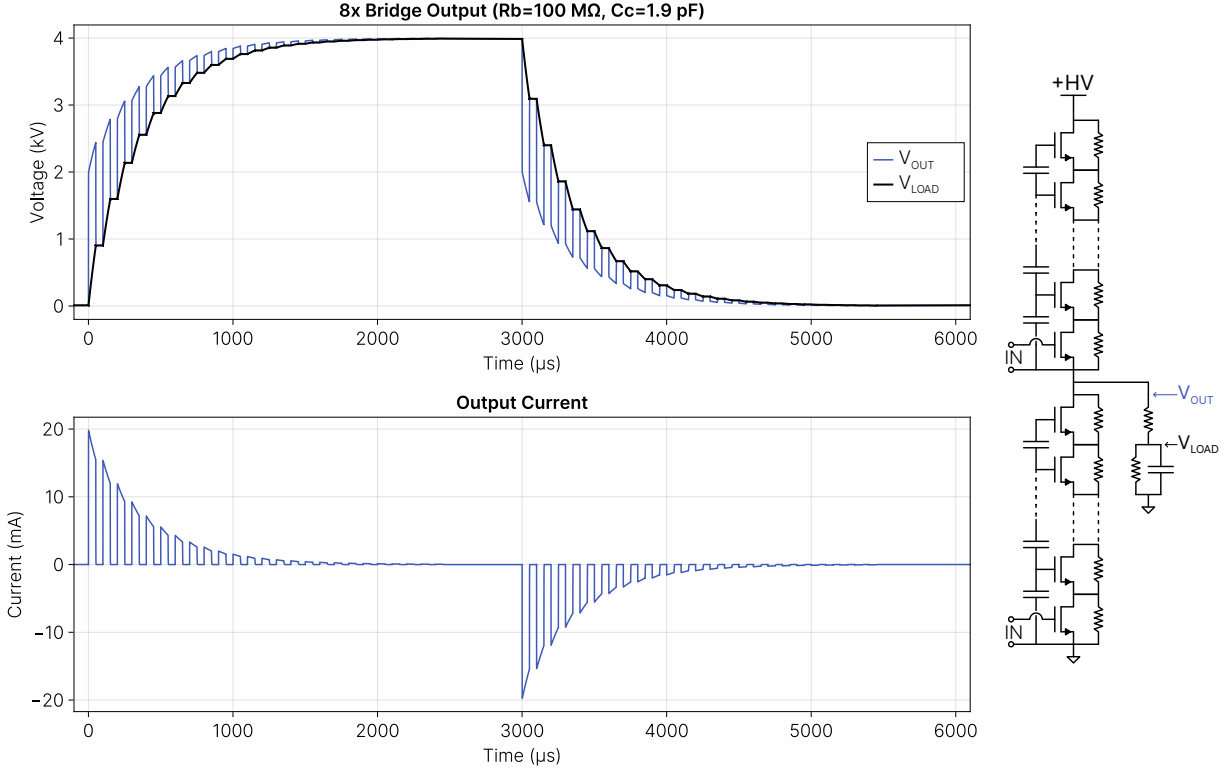


Figure 5.17: Expected voltage and output current profiles of a stacked series MOSFET half-bridge amplifier used to drive an ESA.

Here we can see that as discussed in section 5.3.1, the voltage blocked by each switch will not be constant between switching events. Furthermore, each switch will also see a violent change in blocking voltage whenever the opposing switch activates. We therefore consider the voltage balancing under these conditions, shown in figure 5.18.

Figures 5.17 and 5.18 show the simulated output of a half-bridge amplifier constructed from two stacks of $8 \times$ BSS127 MOSFETs along with the response of a connected ESA load with $C_{load} = 1 \text{ nF}$, $R_{ser} = 100 \text{ k}\Omega$, and $R_{par} = 1 \text{ G}\Omega$. MOSFET stacks are capacitively coupled with a base coupling capacitance C_c , and voltages are balanced with a resistive balancing ladder of resistance R_b .

Gate drivers are simulated as ideal voltage sources in parallel with a $1 \text{ M}\Omega$ resistor. The rail is a 4 kV ideal voltage source. Input signals are a train of 10 V , $50 \mu\text{s}$ pulses every $100 \mu\text{s}$: in other

words, a 10 kHz square wave at a 50 % duty cycle. To periodically charge and discharge the load, the input signal is applied to either the high- or low-side stack, alternating every 3 ms. Both figures are generated from results produced by the same test.

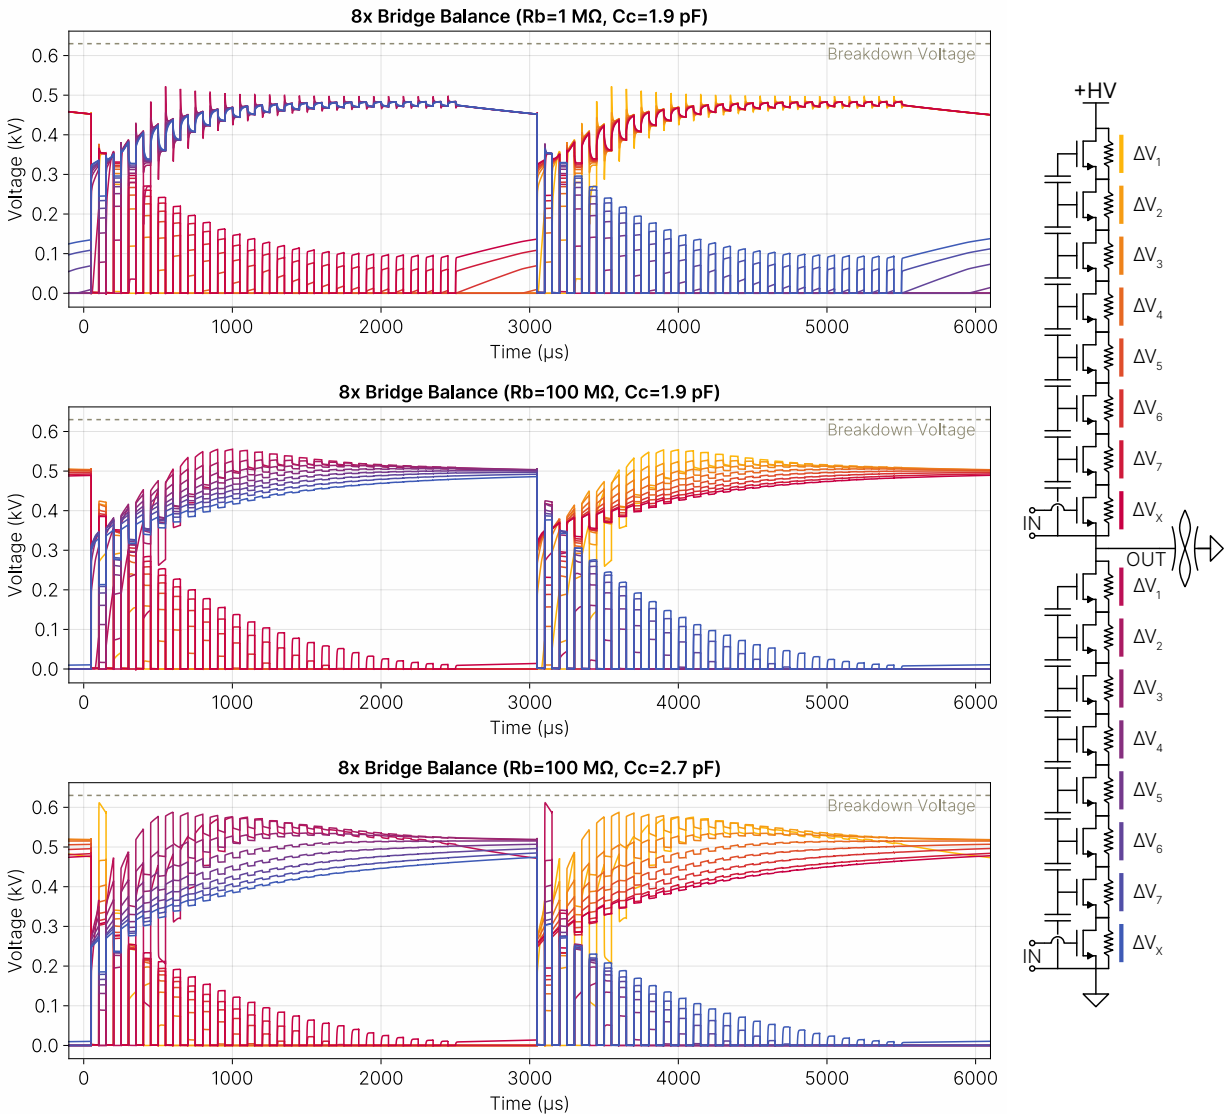


Figure 5.18: The drain-source potential ΔV_{ds} across each MOSFET within a half-bridge composed of two coupled series stacks, driving a simulated ESA load at assorted values of R_b and C_c .

5.3.4 Failure Modes And Current Limiting

Individual MOSFETs in the stack may briefly experience voltages greater than their breakdown voltage and enter avalanche conditions due to insufficiently fast voltage balancing. The trade-off, as discussed in section 5.3.2, is to allow greater leakage current through the switch by using smaller values of balancing resistance R_b . Ultimately however, due to the nature of our design, such individual overvoltage events will be very brief and are unlikely to damage the MOSFETs. In this case, the avalanche effects serve as a secondary, more aggressive, voltage balancing mechanism: MOSFETs which see a greater voltage will also exhibit increased leakage current until voltage balance is restored.

However, the entire stack may also experience overvoltage conditions simultaneously. Perhaps the most likely way for this to occur is by the ESA load entering "generator mode" – increasing in voltage in response to the capacitance reduction caused by mechanical deformation while charged, sometimes raising the actuator voltage above even the rail voltage [99]. Increased current flow during prolonged avalanche conditions will cause sharp local temperature increases within the material of the MOSFET. This will damage the device, and can quickly lead to failure. Repeated prolonged avalanche events may also degrade the performance of the MOSFET over time [105]. This incremental degradation manifests as an increase in $R_{ds,on}$, which could provide a metric to evaluate the health of the stack.

One option for reducing the likelihood, duration, and intensity of any overvoltage conditions is to add a voltage safety factor to the design of the switch. Understanding that the breakdown voltage of our MOSFETs is 600 V (nominally, in practice this value is closer to 630 V), we design our prototype switches to use a nominal blocking voltage of 500 V per MOSFET. A safety factor is further beneficial to the design of the stack, as it allows individual MOSFETs to fail without causing a cascading failure of the entire stack. Ultimately however, the primary mode of damage created during overvoltage conditions is caused by heating resulting from sustained overcurrent.

Therefore, the single most effective method for preventing damage to the switch is to pre-

vent prolonged overcurrent conditions through the device. A current limiting resistor (nominally $100\text{ k}\Omega$ - $200\text{ k}\Omega$) is added in series with the stack, and doubles as a current sense resistor for the instrumentation described in section 5.5. For testing prototype hardware, we aim to limit current to approximately 20 mA , substantially less than the 70 mA maximum for each individual MOSFET specified by the manufacturer (BSS127, Diodes Inc.). However, current limiting resistors ultimately dissipate energy – increasingly so with higher currents through the switch – and thus eliminate some of the efficiency benefit to using solid state switches in the first place. A better solution could be to use a resonant inductor or LC tank for current limiting, as used for zero current switching (ZCS) or zero voltage switching (ZVS) schemes in switching power supplies.

5.4 Circuit Simulation

Circuit simulation proves to be a valuable tool in the design, study, and optimization of this system. It allows exploration and iteration faster than building prototype hardware, including the ability to sweep through sets of design parameters. It also provides insight into those behaviors of the system which are prohibitively difficult to measure in the lab. Unless otherwise noted, all simulations are run using LTSpice (Analog Devices) with default settings.

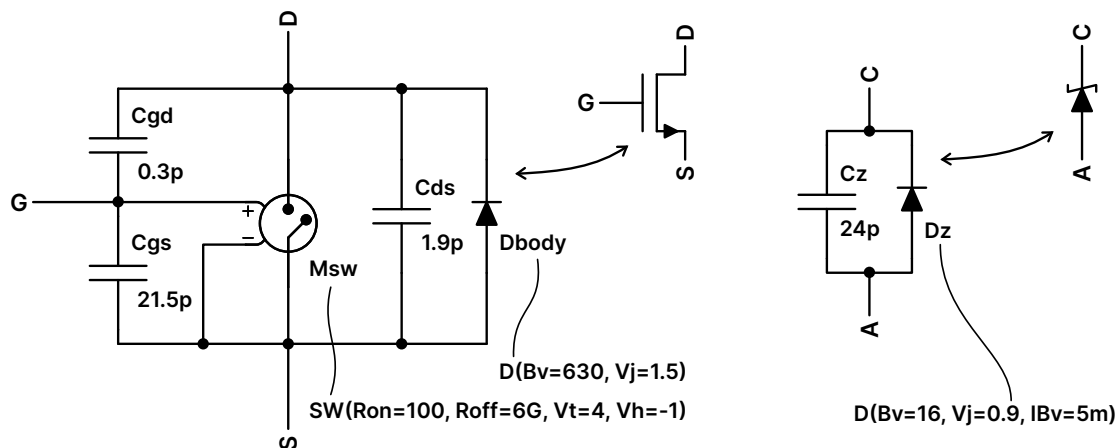


Figure 5.19: Equivalent models of MOSFETs and Zener diodes used for simulations in LTSpice.

MOSFETs are modeled as a voltage controlled switches. In effect, these are variable resistors

that smoothly (linearly) transition from $R_{off} = 6 \text{ G}\Omega$ at or below a gate voltage of 3 V to $R_{on} = 100 \Omega$ at or above a gate voltage of 5 V . The MOSFET model also includes parasitic capacitances between each pin. Zener diodes are modeled as ideal diodes with a parallel parasitic capacitance. Parameter values are taken from the respective datasheets of the components used in the physical implementation (BSS127, Diodes Inc.) and (BZT52C16SQ, Diodes Inc.)

The linear switching behavior of the MOSFET model is a fairly strong simplification of how the device actually switches. However, because we will be operating well into the saturation region, such a simplification is reasonable. Likewise, the parasitic inductances within the MOSFET are neglected, however, we are not operating at sufficiently high frequencies for their effects to be substantial.




Figure 5.20: A comparison of the simulated and observed switching behavior of a BSS127 MOSFET.

To validate this model we compare the switching characteristics of the simulation to benchtop test data. We test a single MOSFET (BSS127, Diodes Inc.) as a low-side switch with a $200\text{ k}\Omega$ load and a 300 V rail supplied by a Trek 610E and a 15 nF stabilization capacitor. The MOSFET is directly driven by a function generator (33500B, Keysight) to a nominal gate voltage of 10 V . We measure the voltage across the load resistor with an oscilloscope (SDS1202X, Siglent) and a $10\times$ probe. Simulations include an extra 13 pF of capacitance across the MOSFET drain-source pins to account for the capacitance of the probe. Two tests are performed, and their results are shown in figure 5.20. The first is a 10 kHz square wave with a 50% duty cycle (pulse width of $50\text{ }\mu\text{s}$), and the second is a comparable ramp ($50\text{ }\mu\text{s}$ rise, $50\text{ }\mu\text{s}$ fall).

These results demonstrate a close match between the simulation and the benchtop test, and thus validate our model. It should be noted that an additional 7 pF of capacitance across the MOSFET drain-source pins yields an even closer result. However, this capacitance has no obvious physical manifestation or source, and is not included in the model.

5.4.1 Transient Gate Behavior

Simulation can provide insight into those behaviors which are prohibitively difficult to evaluate using benchtop tests. One such behavior is the change in gate voltage during turn-on and turn-off switching transients: these dynamics occur on the order of 100s of nanoseconds or even less for comparable designs in the literature [49, 50, 93, 97], yet can have a magnitude of several kilovolts, and complete characterization would require 8 or even 12 simultaneous measurements. Nonetheless, these dynamics are crucial to the performance of the switch, and are therefore of particular interest.

Figures 5.21 and 5.22 show the simulated results of the gate voltage for each MOSFET in the stack during turn-off and turn-on, respectively. Each run simulates an $8\times$ stack of BSS127 MOSFETs, modeled as described in section 5.4, operating as a high-side switch for a $200\text{ k}\Omega$ resistive load. The rail is simulated as a 4 kV ideal voltage source. Input pulses to the base MOSFET M_x are simulated as floating, ideal 10 V voltage sources.

Neither the timing of the turn-on, or the timing of the turn-off appear to be substantially affected by the value of the balancing resistance R_b or the coupling capacitance C_c . However, as hypothesized in section 5.3.1, insufficiently small values of C_c will lead to a lower gate-source voltage, especially for MOSFETs higher in the stack. This could lead to only partial activation of some MOSFETs in the stack. Likewise as hypothesized, unnecessarily large values of C_c appear to be mitigated by the zener diodes.

We note that consistent with our expectations, switching is sequential, but extremely fast when compared to all other relevant dynamics. Switches turn-on from the bottom to the top of the stack, starting with the activation of the base MOSFET M_x by the gate driver. When a MOSFET is activated, current flows through the coupling capacitor of the next MOSFET in the stack, and the gate voltage of the next MOSFET rises, causing it to turn on. This cascade continues until the entire stack is on.

Likewise, turn-off occurs in the same order, from the bottom to the top of the stack. In this case, the deactivation of the base MOSFET M_x by the gate driver blocks current flow, in turn redirecting current through the zener diode of the next MOSFET in the stack, causing a drop in gate voltage, and the deactivation of the subsequent MOSFET.

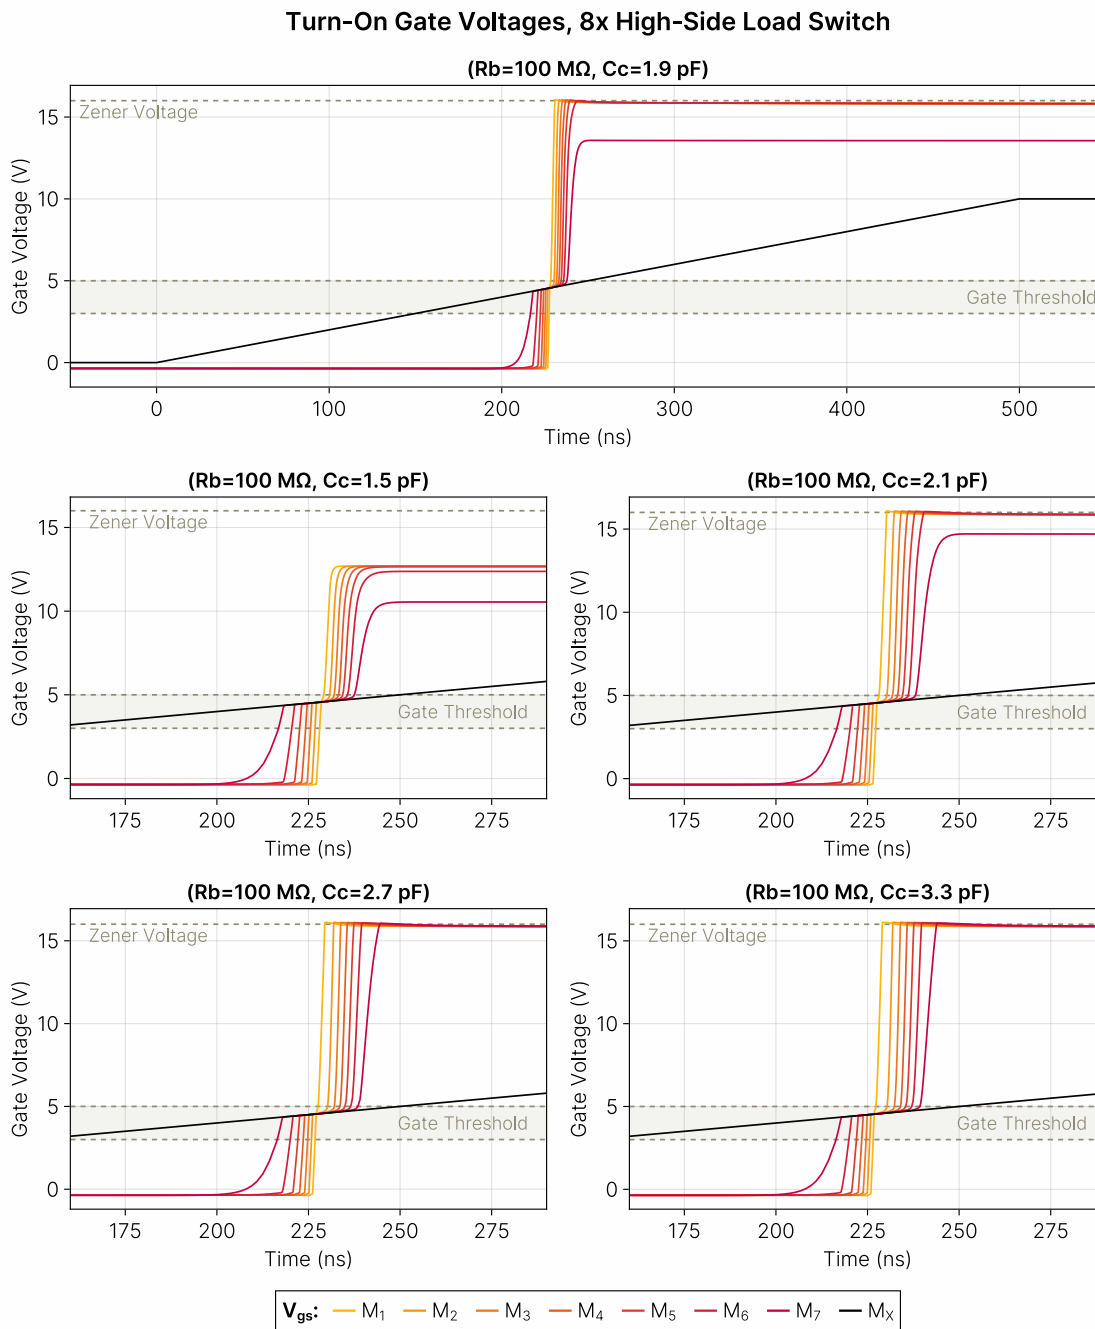


Figure 5.21: Effect of variations in coupling capacitance C_c on gate turn-on timing throughout a capacitively-coupled series MOSFET stack.

5.5 Instrumentation

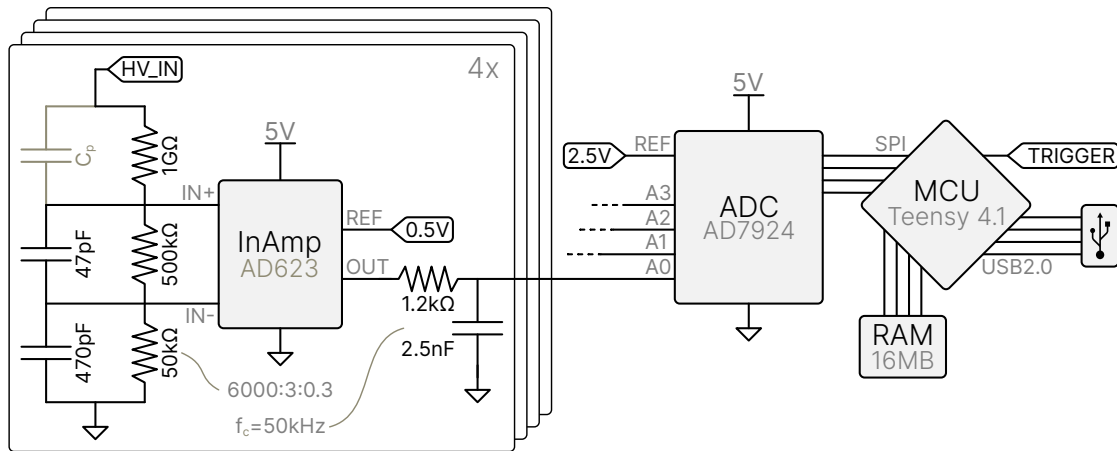


Figure 5.23: High speed, 4-channel high voltage instrumentation system. For additional measurement channels, multiple instruments can be synchronized with a trigger pulse.

Validation and analysis of our design requires the ability to precisely and simultaneously measure multiple currents and voltages in the kilovolt range at moderately high speeds. For this task, the buffered voltage divider shown in figure 3.4 in section 3.2 is insufficient. Instead, we use the improved instrumentation system shown in figure 5.23.

We construct a 6000:3:0.3 voltage divider and take a differential measurement across the middle resistor using an instrumentation amplifier (AD623, Analog Devices). These measurements are then converted using a high-speed ADC (AD7924, Analog Devices). The base of the instrumentation system is a Teensy 4.1 microcontroller (PJRC) augmented with an additional 16 MB PSRAM. It communicates with the ADC over a SPI bus running at 80 MHz. A RC lowpass filter is also added in between the instrumentation amplifier and the ADC. The filter has a cutoff frequency of 50 kHz – the Nyquist frequency given a nominal sample rate of 100 ksp/s per channel.

It should be noted that each of the resistors in the divider has a parasitic capacitance which would normally be negligible. Due to the overall high impedance of the ladder, this capacitance dominates during the fast switching transients we are interested in observing. Because the ratio of

inverses of these parasitic capacitances is roughly equivalent for the three resistors, and certainly not our desired 6000:3:0.3 ratio, we add 47 pF and 470 pF capacitors in parallel to the two lower resistors to counteract the parasitic capacitance of the 1 G Ω resistor, which we estimate to be about 25 fF.

For most tests, multiple instruments are used simultaneously. In these cases, their respective microcontrollers are connected together and synchronized by a trigger pulse. After each test completes, each data set is transferred to a PC, where it is filtered, processed, and plotted using scripts written in Julia.

5.6 Manufacturing




Figure 5.24: Schematic and render of the prototype switch implementation.

We implement an individual switch - consisting of a 12-MOSFET stack and an optoelectronic gate driver - on a single 1 mm thick FR4 PCB (PCBWay), which when assembled, measures 72 \times 18 \times 3.2 mm and weighs 4 g. Prototypes were assembled by hand, but the design is suitable for complete automated assembly using a pick-and-place machine and reflow or wave soldering. All high-voltage traces are routed within a single layer, suggesting the potential for vertically stacking multiple switches to build compact multi-channel systems. Boards were coated with acrylic conformal coating (419D, MG Chemicals) to prevent arcing between components and to protect against moisture and dust. However, this step was found to not be necessary for resilient

switch operation in the lab environment.

Numerous MOSFETs exist which would be suitable for our proposed stack design. We use the BSS127 (Diodes Inc.), a device which is commonly available in a small (SOT23) package and has relatively low parasitic capacitances, thus reducing the demands on the external coupling capacitors, balancing resistors, and gate driver.

Unlike other high-voltage switches in this class, such as power optocoupler or reed relay-based switches, this design can be produced entirely from commonly available components using industry-standard mass-manufacturing methods such as pick-and-place and reflow or wave soldering.

In August 2023, the bill-of-materials cost was \$18.16 for 10 switches, and \$8.31 for 1000 (Mouser Electronics), substantially less than any currently available alternative. This approach reduces the cost to drive an individual HASEL actuator from approximately \$300 using direct drive (A60P-5, XP Power) or approximately \$100 using optocouplers [69,99] to approximately \$20 (when including the cost of PCB manufacturing and assembly). With future optimizations, this cost could be reduced further.

Substantial reductions in footprint size are realistically possible for future iterations of this prototype design. Capacitively-coupled series MOSFET stacks are adaptable to IC technology, and physical implementations have been demonstrated on a silicon-on-insulator (SOI) process (0.35 μm , Honeywell), which were able to switch greater voltages than the rated breakdown voltage of the process used [97]. Thus, the MOSFETs and Zener diodes of this cascade could theoretically be implemented within a single IC.

Flex PCBs are PCBs implemented on a flexible substrate, typically 25 to 50 μm thick polyimide (Kapton). At this thickness, it should be possible to implement gate coupling capacitors as features on the PCB itself instead of using discrete surface mount capacitors, due to the small capacitance values required. These two modifications would consolidate the majority of this design into a single component and some PCB routing features.

5.7 Performance

In this section, we evaluate the performance of our prototype switches from section 5.6, against various metrics.

5.7.1 Leakage Current

To evaluate off-state leakage current, switches were connected directly to a HV-SMPS (610E, Trek) set to supply mode and providing a constant 6 kV output. Off-state leakage current through each switch was determined to be the total load current measured by the HV-SMPS.

Off-state leakage current was measured to be $5\ \mu\text{A}$ on startup, consistent with our expectations given the combined $1.2\ \text{G}\Omega$ impedance of the resistive balancing ladder. After 1 minute, the leakage current had risen to $6 \pm 2\ \mu\text{A}$ and remained in that range for the 3 hr duration of the test.

5.7.2 Slew Rate

We are fundamentally interested in various circuits' ability to drive capacitive loads with arbitrary voltage waveforms. The performance of any drive circuit towards achieving this goal will primarily depend on two factors: the magnitude of the load capacitance, and the peak current supplied through the driver. For switched bridge-style amplifiers, this current inversely correlates to the on-state impedance of the switches from which they're constructed. The maximum slew rate, defined as the maximum rate of voltage change that the drive circuit is capable of achieving for a given load capacitance, provides a useful metric by which to quantify this performance. We measure the slew rate as the average slope of the load voltage from 5% to 65% of the rail voltage value (nominally 6 kV) in response to a step input. We select this range over the other commonly used definitions (eg. 10% to 90% or 5% to 95%) to minimize the effects of rail droop on the larger capacitive loads.

We quantify the maximum slew rate for half-bridge amplifiers constructed using the solid state switches presented in this chapter, as well as using the optocouplers presented in chapter 3 -

using both 60 mA and 600 mA of nominal LED input current. We expect the performance of the feedback controller to have a substantial effect on the maximum slew rate in a complete half-bridge amplifier implementation. We therefore design the test to operate in open-loop, and thus decouple the effects of the controller. Consequently, this test provides an upper bound on the performance of amplifiers constructed using a particular switching technology rather than a direct evaluation of the performance of a particular amplifier. To provide a baseline for comparison, we conduct an analogous test and quantify the maximum slew rate directly driven loads. All tests were conducted using the same power supply (610E, Trek).




Figure 5.25: Experimental setups for measuring the slew rate of direct-drive from an HV-SMPS (left) and of parallelizable switching circuits (right).

To evaluate the maximum slew rate when using bridge-style switching circuits, we use the test setup shown in the right side of figure 5.25. We attempt to decouple the performance of the power supply by approximating an idealized high voltage rail which can source practically unlimited amounts of current (albeit for increasingly short durations of time). This 6 kV rail is implemented with a HV power supply (610E, Trek) set to supply mode and augmented with a 180 nF MLCC-based capacitor bank (FKP 1, WIMA). The reader should note that this can be a particularly dangerous test setup, and is reminded to read [113] and implement appropriate precautions before attempting to replicate this, or any comparable experiment². The switching technology under test is then used to construct a half-bridge amplifier. 100 kΩ resistors are used to measure current

² Future implementations of this test setup could, perhaps counterintuitively, be implemented more safely using a high voltage power supply capable of sourcing larger amounts of current augmented with a smaller capacitor bank.

and simultaneously limit current when testing NMOS-based switches, as required by the design. It should be noted that only the high-side switch is strictly necessary to conduct this test. However, a half bridge is useful, as it provides a means to fully discharge the load capacitor before and after each test.

To provide a realistic baseline for the maximum slew rate we consider the case of direct drive using a commercially available HV-SMPS, specifically using the test setup shown on the left side of figure 5.25. The power supply (610E, Trek) is set to amplifier mode, whereupon it acts as a $1000\times$ amplifier given an analog input. An input signal is generated using a digital output pin on a microcontroller (Teensy 4.0, PJRC) to drive a MOSFET (PMV20EN, Nexperia) which connects a 6 V signal provided by a benchtop power supply (E3630A, Agilent) to the analog input of the Trek 610E. In effect, the setpoint of the Trek's regulated output is driven between 0 and 6 kV depending on the state of the digital pin. A $200\text{ k}\Omega$ resistor is used to measure current while ensuring that the resistance on the load path matches the $200\text{ k}\Omega$ of the other tests. In both test setups, the voltage rise on the load is recorded using the instrumentation described in section 5.5 at a sample rate of 100 kHz.

The load capacitors used to conduct this test were either ceramic disk capacitors (CDCs) ranging from 4.7 pF to 1 nF (assorted dielectrics, Vishay), or multilayered ceramic capacitors (MLCCs) ranging from 235 pF to 68 nF (FKP 1, WIMA). We expect the MLCCs to be more representative of an ideal capacitive load, whereas due to their higher equivalent series resistance (ESR), we expect the CDCs to exhibit lower slew rates than equivalent MLCCs. It should be noted that ESAs typically also exhibit non-trivial series resistances. As such, the maximum slew rate given a MLCC load capacitor can be considered as the upper bound for an ESA with equivalent capacitance.




Figure 5.26: Characteristic trace of a slew rate evaluation test pulse delivered to a 2x12 Peano-HASEL actuator stack by a half-bridge amplifier circuit constructed from stacked series MOSFETs.

The results of a characteristic trial are shown in figure 5.26, specifically a trial characterizing one of the 2x12 Peano-HASEL actuators used throughout this work. This particular actuator was found to have a capacitance of approximately 860 pF, as measured using an LCR meter (Model 891, BK Precision). Five such trial runs were conducted for every permutation of load capacitor and drive topology, and the aggregate results are shown in figure 5.27.




Figure 5.27: Maximum slew rate achieved using different switching technologies as a function of load capacitance.

As evident from the results shown in figure 5.27, this test demonstrates that parallelizable driver topologies for capacitive loads do not necessarily decrease performance: rather, under the right conditions, they may even increase it. These results also highlight the fundamental relationship between on-state impedance and amplifier performance. Under the entire range of load capacitances tested (which span almost 5 orders of magnitude), the NMOS stack-based bridge exceeded the direct-drive performance by a factor of 5-10. Likewise, the optocoupler-based bridge consistently exceeded the direct-drive performance by roughly a factor of 2 when driven with 600 mA of input current. Only in the case where LED input current was reduced to 60 mA did the optocoupler-based bridge underperform direct-drive. Consistent with our expectations, the maximum slew rate of the HV-SMPS approached its datasheet value in the low capacitance limit.

5.7.3 Efficiency and Thermals

We are also fundamentally interested in the efficiency and thermal characteristics of the high voltage switches we use. Our rationale is simple: the higher the on-state impedance of the switches relative to that of the load, the more energy is dissipated (and wasted) across them instead of being delivered to the load.

A known shortcoming of optocouplers is their relatively high on-state impedance, which causes the switch to dissipate substantial energy as heat. As discussed in section 3.1.5, excessive, unmanaged heat can lead to thermal runaway and catastrophic failure of the device. A constant-current LED driver, such as a current mirror, could be used to mitigate thermal runaway. However, any substantial energy waste is nonetheless problematic in a power-limited system.

Consequently, low on-state impedance switches are fundamentally desirable for the construction of efficient, high-performance, high voltage drive circuitry for capacitive loads. Stacked NMOS switches provide this low impedance, and the significance is clearly visible in the thermal images shown in figure 5.28. The stacked NMOS switches (right) not only deliver the majority of the energy to the load, but also deliver more than 4 times the average current. As discussed below, this value could be even higher with a larger power supply.




Figure 5.28: Thermal imagery highlighting power delivery disparities using optocoupler (left) and solid-state (right) load switches to drive a resistive load.

This test was conducted by connecting the switch technology under test as a high-side switch for a $200\text{ k}\Omega$ load resistance under a 6 kV rail voltage provided by a HV-SMPS (610E, Trek) set to supply mode and stabilized with a 15 nF MLCC (FKP 1, WIMA). Average load path current was calculated after measuring the DC voltage drop across the $200\text{ k}\Omega$ load resistor (made from two 1% , $100\text{ k}\Omega$ resistors in series) using a multimeter (115, Fluke) and a high voltage probe (80K-40, Fluke). When testing power optocouplers, the LEDs received 600 mA of input current in the on-state. Conducting this test with the HV rail disconnected showed that the LEDs generate a comparably negligible amount of heat on their own.

Images were taken using a thermal camera (T650sc, FLIR). Note that the absolute values of temperature will highly depend on the duration of the test. Absolute test duration was not tightly controlled: these photos were taken after running for approximately 3-4 minutes. Much more significant is the relative distribution of the temperatures (and therefore energy dissipation) between the switch and the load. Switches were driven with a 10 kHz pulse train at 10% duty cycle. These test conditions were selected both due to the limitations of our power supply, and to reduce the sensitivity of the results to the duration of the test. The reader should note that 2 mA of output current is the upper limit for continuous output from a Trek 610E power supply. Reaching

this current with a $200\text{ k}\Omega$ load means the power supply will lose regulation, enter a current limiting mode, and the rail will sag substantially.

5.7.4 Pulse Width Modulation

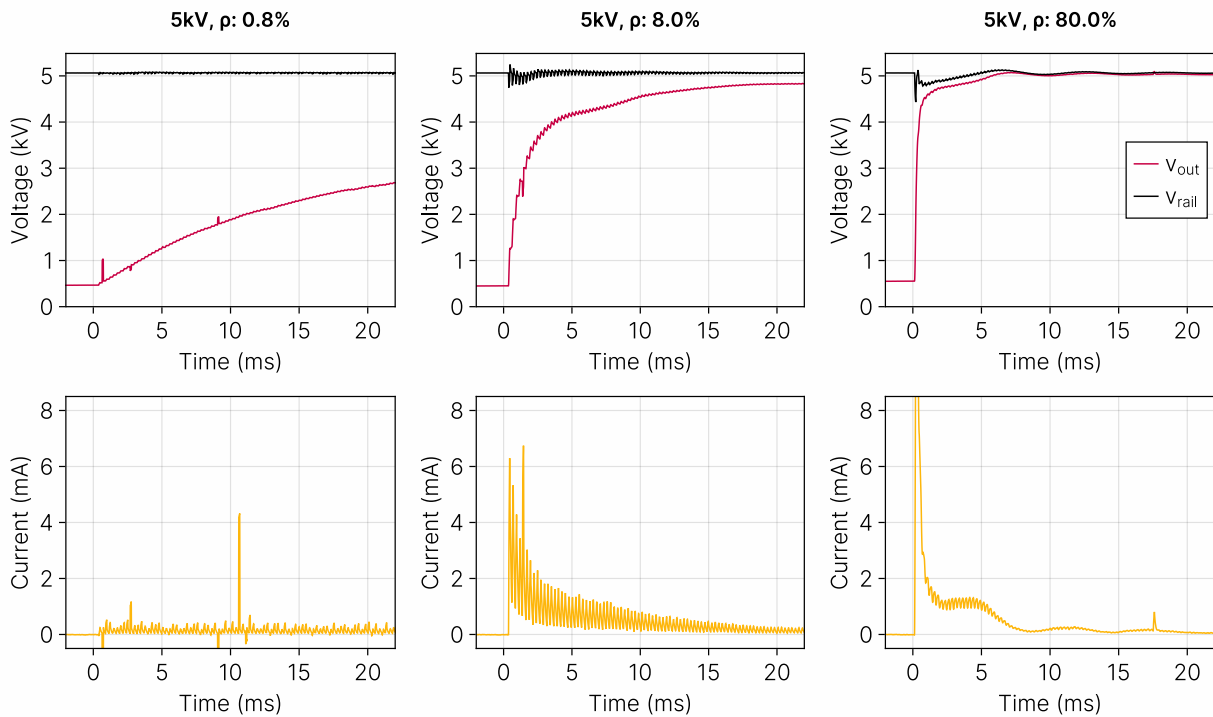


Figure 5.29: Effect of PWM duty cycle variation on the output voltage and current to an ESA driven by a solid-state half-bridge amplifier.

The ability to vary the magnitude of charge delivery through a switch is a fundamental requirement for the construction of switching amplifier circuits with closed loop controls. A common scheme for this is pulse width modulation (PWM), as introduced in section 2.6. In this section, we demonstrate that our proposed switching technology is suitable for use with the typical PWM outputs available to most microcontrollers.

A half-bridge amplifier is constructed from two $12\times$ MOSFET stacks. It connects to a 5 kV rail provided by a Trek 610E set to supply mode and stabilized with a 15 nF MLCC capacitor (FKP 1, WIMA). A 2×12 Peano-HASEL actuator (of the same type as first presented in chapter

2) is connected as a load to the midpoint of the bridge.

At rest, this actuator has a capacitance of approximately 860 pF, as measured using an LCR meter (Model 891, BK Precision) across a range of frequencies up to 10 kHz.

however, its capacitance will increase along with actuation. The low side of the bridge is grounded. 100 k Ω resistors are added in series with each MOSFET stack. These resistors act both as current limiting resistors (see section 5.3.4), and as shunt resistors for current measurement. Voltages are measured using the instrumentation described in section 5.5, and currents are calculated from these voltage measurements and resistance values.

Load current is measured as the difference between the currents through the the high- and low-side switches, (charge and drain, respectively).

$$i_{load} = i_{chg} - i_{drn} \quad (5.13)$$

Switches are driven by PWM signals from a microcontroller (Teensy 4.0, PJRC). We use a PWM frequency of 3.921 kHz. While seemingly arbitrary, at this frequency every increment of the 8-bit PWM value conveniently corresponds to an additional 1 μ s of on-state pulse width. The traces shown in figure 5.29 are generated using 0.8 %, 8 %, and 80 % duty cycles and correspond to 2, 20, and 200 μ s-wide pulse trains, respectively, from left to right.

As evidenced by the results shown in figure 5.29, the proposed switching technology can modulate the charge rate, (and consequently, rise time) of a capacitive load by several orders of magnitude using the standard 3.3 V PWM output provided by most commonly-available microcontrollers. These results suggest that the proposed switching technology is well suited for the construction of closed-loop voltage control systems for ESAs.

5.7.5 ESA Step Response

Ultimately, the primary objective of our work is to develop novel circuits which are more capable of driving ESAs. The following test demonstrates this capability. It follows an identical procedure to the tests described in section 5.7.4 above, however, the displacement of the actuator

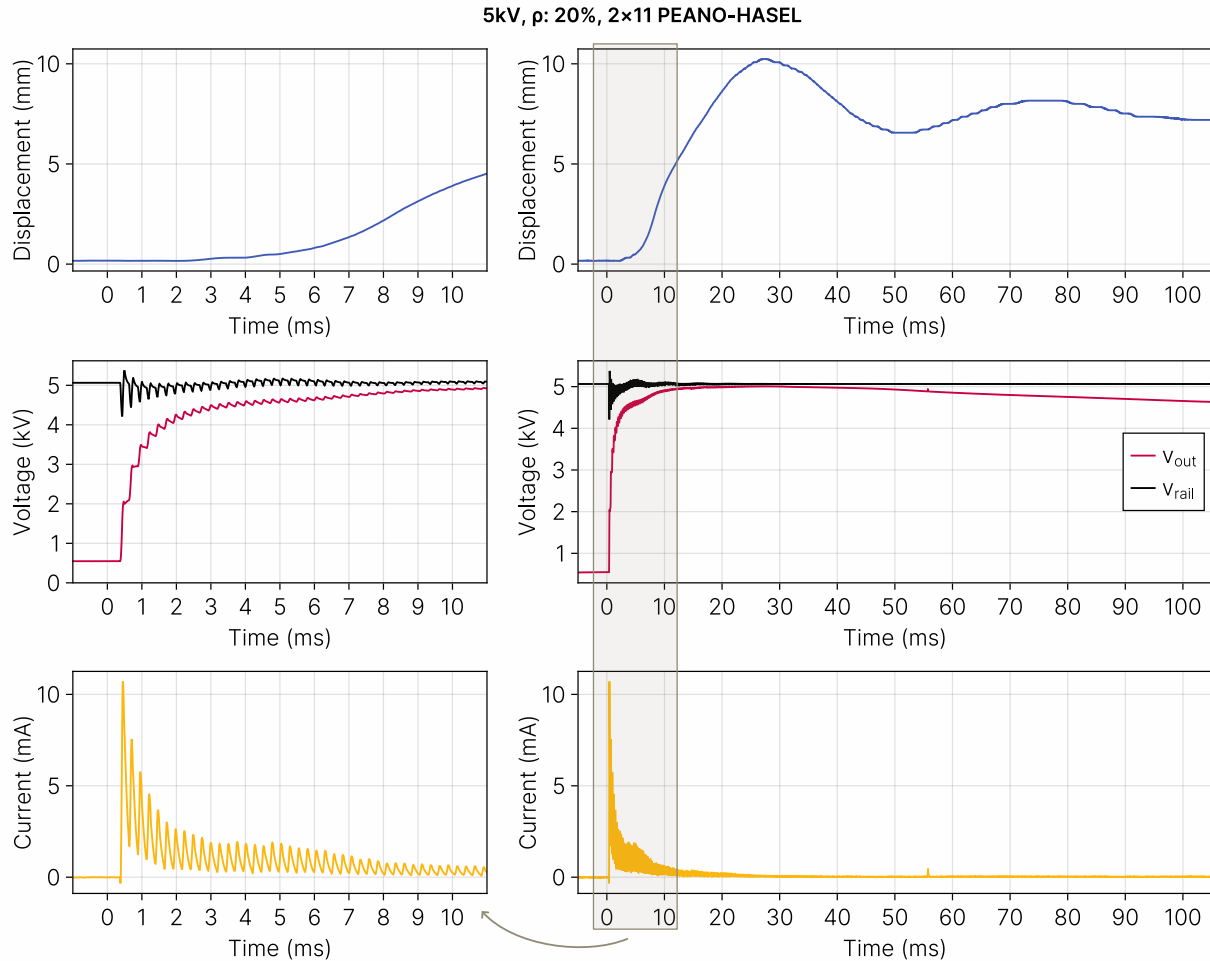


Figure 5.30: Electromechanical step response of an ESA driven by a PWM-controlled solid-state half-bridge amplifier constructed with series MOSFET stacks.

is also measured using a laser displacement sensor (LK-G5000, Keyence) output to, and recorded by, an oscilloscope (SDS1202X, Siglent). The results shown in figure 5.30 were collected using a PWM duty cycle of 20%, corresponding to a pulse width of $50\mu\text{s}$. A 20% duty cycle does not represent the true step response of the system – this value was selected to facilitate the task of fitting different timescales in one figure, and to simultaneously show the low-pass filtering effects of the ESAs displacement dynamics.

This test demonstrates several significant points simultaneously. The output voltage of the bridge begins to droop shortly after reaching the rail voltage - this is because the MOSFET stack

needs a nonzero potential to trigger the switching cascade, as described in section 5.3.1. The voltage response of the actuator begins looking like the quintessential RC response, but becomes nonlinear as the actuator starts to move and its capacitance starts to change, as suggested by section 2.2. The actuator's displacement response to an applied voltage step exhibits the characteristics of a second-order system, consistent with the results from section 4.6. As hypothesized in section 2.6, there exists a substantial discrepancy in the temporal scale of the voltage and displacement dynamics, suggesting they can be decoupled for the purposes of control. In summary, these results demonstrate substantial consistency between our hypotheses and physical reality, and strongly suggest that our switches provide the basis for constructing inherently parallelizable driver circuits for high voltage capacitive loads requiring precise closed-loop voltage regulation.

5.8 Conclusion

This chapter presents a novel design for solid-state switches and switching amplifiers capable of driving ESAs. MOSFETs are stacked in series to achieve high blocking voltages, and capacitively coupled to switch synchronously. Voltage is passively balanced throughout the stack via the use of a load-side resistive ladder. Each stack is controlled by a novel optically-powered, optically-controlled gate driver circuit which floats along with the potential of the switch, and enables the construction of both high- and low-side switches exclusively using higher-performing N-channel MOSFETs.

Like state of the art power optocouplers, this approach is fundamentally parallelizable, and enables the creation of multi-actuator soft robots. However due to drastically reduced on-state impedance, this approach allows the construction of ESA drive circuits which are more compact, efficient, and performant than any existing state of the art technology.

This technology can be used to deliver more current, and thus achieve faster slew rates for larger capacitive loads. Likewise, it is much more efficient: dissipating less heat, delivering more energy to the load, and requiring fewer auxiliary systems for thermal management. The increased inefficiency is furthermore invaluable in the (typical for ESAs) context of power-limited systems.

Perhaps most importantly, these switches can be produced in industrial quantities entirely from commonly available components using only standard mass-manufacturing processes.

This design can be used independently as a load switch, or as a building block for more complex circuits. When combined with a voltage sensor, it forms the basis for constructing bridge-style switching amplifiers for high voltage capacitive loads, and it is demonstrably well-suited for driving ESAs. One of the many branches of the future of robotics involves the coordinated action of hundreds to thousands of artificial electrostatic muscles, controlled by 'synthetic motorneurons' in the form of bridge-style amplifier circuits. The work presented here brings us one step closer to this exiting future of our field.

Bibliography

- [1] E. Acome, S. K. Mitchell, T. G. Morrissey, M. B. Emmett, C. Benjamin, M. King, M. Radakovitz, and C. Keplinger. Hydraulically amplified self-healing electrostatic actuators with muscle-like performance. Science, 359(6371):61–65, January 2018.
- [2] Morteza Amjadi, Ki Uk Kyung, Inkyu Park, and Metin Sitti. Stretchable, Skin-Mountable, and Wearable Strain Sensors and Their Potential Applications: A Review. Advanced Functional Materials, 26(11):1678–1698, 2016.
- [3] Costanza Armanini, Frédéric Boyer, Anup Teejo Mathew, Christian Duriez, and Federico Renda. Soft Robots Modeling: A Structured Overview. IEEE Transactions on Robotics, 39(3):1728–1748, June 2023.
- [4] Hedan Bai, Young Seong Kim, and Robert F. Shepherd. Autonomous self-healing optical sensors for damage intelligent soft-bodied systems. Science Advances, 8(49):eabq2104, December 2022.
- [5] R. J. Baker and B. P. Johnson. Stacking power MOSFETs for use in high speed instrumentation. Review of Scientific Instruments, 63(12):5799–5801, December 1992.
- [6] R.J. Baker and S.T. Ward. Designing nanosecond high voltage pulse generators using power MOSFETs. Electronics Letters, 30(20):1634–1635, September 1994.
- [7] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. SIAM review, 59(1):65–98, 2017.
- [8] M.C. Birch, R.D. Quinn, G. Hahm, S.M. Phillips, B. Drennan, R.D. Beer, Xinyu Yu, S.L. Garverick, S. Laksanacharoen, A.J. Pollack, and R.E. Ritzmann. A miniature hybrid robot propelled by legs. In Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180), volume 2, pages 845–851 vol.2, October 2001.
- [9] Douglas Blackiston, Emma Lederer, Sam Kriegman, Simon Garnier, Joshua Bongard, and Michael Levin. A cellular platform for the development of synthetic living machines. Science Robotics, 6(52):eabf1571, March 2021.
- [10] Richard Bowden-Reid. An Experimental Study of Gridded and Virtual Cathode Inertial Electrostatic Confinement Fusion Systems. Thesis, June 2019.

- [11] Sean R. Bradburn and Herbert L. Hess. An integrated high-voltage buck converter realized with a low-voltage cmos process. In 2010 53rd IEEE International Midwest Symposium on Circuits and Systems, pages 1021–1024, August 2010.
- [12] Daniel Bruder, Xun Fu, R. Brent Gillespie, C. David Remy, and Ram Vasudevan. Data-Driven Control of Soft Robots Using Koopman Operator Theory. IEEE Transactions on Robotics, 37(3):948–961, June 2021.
- [13] Daniel Bruder, Brent Gillespie, C. David Remy, and Ram Vasudevan. Modeling and Control of Soft Robots Using the Koopman Operator and Model Predictive Control, July 2019.
- [14] Daniel Bruder, C. David Remy, and Ram Vasudevan. Nonlinear System Identification of Soft Robot Dynamics Using Koopman Operator Theory. In 2019 International Conference on Robotics and Automation (ICRA), pages 6244–6250, May 2019.
- [15] Jin Cai, Tak Ning, Philip Oldiges, Anthony Chou, Arvind Kumar, Werner Rausch, Wilfried Haensch, and Ghavam Shahidi. SOI series MOSFET for embedded high voltage applications and soft-error immunity. In 2008 IEEE International SOI Conference, pages 21–22, October 2008.
- [16] Fredrik Bagge Carlson. ControlSystemIdentification.jl, October 2023.
- [17] P H Chappell and K J Campden. Switching performance of power MOSFETs with capacitive loads at high frequency and high voltage for square wave generators. Measurement Science and Technology, 3(4):356–361, April 1992.
- [18] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. ACM SIGEVOlution, 7(1):11–23, 2014.
- [19] R.W. Colbrunn, G.M. Nelson, and R.D. Quinn. Design and control of a robotic leg with braided pneumatic actuators. In Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180), volume 2, pages 992–998 vol.2, October 2001.
- [20] Jim Crist, Mattias Fält, and Fredrik Bagge Carlson. ControlSystems.jl. JuliaControl, November 2023.
- [21] F. Daerden and D. Lefeber. Pneumatic artificial muscles: Actuators for robotics and automation. European journal of mechanical and environmental engineering, 2002.
- [22] S.K. De and N.R. Aluru. Full-Lagrangian schemes for dynamic analysis of electrostatic MEMS. Journal of Microelectromechanical Systems, 13(5):737–758, October 2004.
- [23] Cosimo Della Santina, Antonio Bicchi, and Daniela Rus. On an Improved State Parametrization for Soft Robots With Piecewise Constant Curvature and Its Use in Model Based Control. IEEE Robotics and Automation Letters, 5(2):1001–1008, April 2020.
- [24] Cosimo Della Santina, Christian Duriez, and Daniela Rus. Model-Based Control of Soft Robots: A Survey of the State of the Art and Open Challenges. IEEE Control Systems Magazine, 43(3):30–65, June 2023.

- [25] Cosimo Della Santina, Robert K Katzschmann, Antonio Bicchi, and Daniela Rus. Model-based dynamic feedback control of a planar soft robot: Trajectory tracking and interaction with the environment. The International Journal of Robotics Research, 39(4):490–513, March 2020.
- [26] Cosimo Della Santina, Robert K. Katzschmann, Antonio Biechi, and Daniela Rus. Dynamic control of soft robots interacting with the environment. In 2018 IEEE International Conference on Soft Robotics (RoboSoft), pages 46–53, April 2018.
- [27] Anna D. Dikina, Hannah A. Strobel, Bradley P. Lai, Marsha W. Rolle, and Eben Alsberg. Engineered cartilaginous tubes for tracheal tissue replacement via self-assembly and fusion of human mesenchymal stem cell constructs. Biomaterials, 52:452–462, June 2015.
- [28] Mihai Duduta, David R. Clarke, and Robert J. Wood. A high speed soft robot based on dielectric elastomer actuators. In 2017 IEEE International Conference on Robotics and Automation (ICRA), pages 4346–4351, May 2017.
- [29] Anany Dwivedi, Anand Ramakrishnan, Aniketh Reddy, Kunal Patel, Selim Ozel, and Cagdas D. Onal. Design, Modeling, and Validation of a Soft Magnetic 3-D Force Sensor. IEEE Sensors Journal, 18(9):3852–3863, May 2018.
- [30] Tamar Flash and Binyamin Hochner. Motor primitives in vertebrates and invertebrates. Current Opinion in Neurobiology, 15(6):660–666, December 2005.
- [31] Tamar Flash and Terrence J Sejnowski. Computational approaches to motor control. Current Opinion in Neurobiology, 11(6):655–662, December 2001.
- [32] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. inFORM: Dynamic physical affordances and constraints through shape and object actuation. In Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology, pages 417–426, St. Andrews Scotland, United Kingdom, October 2013. ACM.
- [33] Akeem Francis, Andrea Martinez, Kyle Thompson, Kevin Burke, and Jennifer Zirnheld. Space charge accumulation as a contributor to partial discharge activity in dielectric elastomer actuators under high voltage DC. In 2016 IEEE International Power Modulator and High Voltage Conference (IPMHVC), pages 330–335, July 2016.
- [34] M. Frigo and S.G. Johnson. The Design and Implementation of FFTW3. Proceedings of the IEEE, 93(2):216–231, February 2005.
- [35] C. Gerster. Fast high-power/high-voltage switch using series-connected IGBTs with active gate-controlled voltage-balancing. In Proceedings of 1994 IEEE Applied Power Electronics Conference and Exposition - ASPEC'94, pages 469–472 vol.1, February 1994.
- [36] Lee Gill, Luciano A. Garcia Rodriguez, Jacob Mueller, and Jason Neely. A Comparative Study of SiC JFET Super-Cascode Topologies. In 2021 IEEE Energy Conversion Congress and Exposition (ECCE), pages 1741–1748, October 2021.
- [37] Stephan-Daniel Gravert, Elia Varini, Amirhossein Kazemipour, Mike Y. Michelis, Thomas Buchner, Ronan Hinchet, and Robert K. Katzschmann. Low-voltage electrohydraulic actuators for untethered robotics. Science Advances, 10(1):eadi9319, January 2024.

- [38] Guo-Ying Gu, Jian Zhu, Li-Min Zhu, and Xiangyang Zhu. A survey on dielectric elastomer actuators for soft robots. Bioinspiration & Biomimetics, 12(1):011003, January 2017.
- [39] R. Guidini, D. Chatroux, Y. Guyon, and D. Lafore. Semiconductor power MOSFETs devices in series. In 1993 Fifth European Conference on Power Electronics and Applications, pages 425–430 vol.2, September 1993.
- [40] Yaguang Guo, Liwu Liu, Yanju Liu, and Jinsong Leng. Review of Dielectric Elastomer Actuators and Their Applications in Soft Robots. Advanced Intelligent Systems, 3(10):2000282, 2021.
- [41] Ujjaval Gupta, Lei Qin, Yuzhe Wang, Hareesh Godaba, and Jian Zhu. Soft robots based on dielectric elastomer actuators: A review. Smart Materials and Structures, 28(10):103002, September 2019.
- [42] Yoram Gutfreund, Tamar Flash, Graziano Fiorito, and Binyamin Hochner. Patterns of arm muscle activation involved in octopus reaching movements. Journal of Neuroscience, 18(15):5976–5987, August 1998.
- [43] Yoram Gutfreund, Tamar Flash, Yosef Yarom, Graziano Fiorito, Idan Segev, and Binyamin Hochner. Organization of octopus arm movements: A model system for studying the control of flexible arms. Journal of Neuroscience, 16(22):7297–7307, November 1996.
- [44] Yoram Gutfreund, Henry Matzner, Tamar Flash, and Binyamin Hochner. Patterns of motor activity in the isolated nerve cord of the octopus arm. The Biological Bulletin, 211(3):212–222, December 2006.
- [45] David A. Haggerty, Michael J. Banks, Patrick C. Curtis, Igor Mezić, and Elliot W. Hawkes. Modeling, Reduction, and Control of a Helically Actuated Inertial Soft Robotic Arm via the Koopman Operator, November 2020.
- [46] Carter S. Haines, Márcio D. Lima, Na Li, Geoffrey M. Spinks, Javad Foroughi, John D. W. Madden, Shi Hyeong Kim, Shaoli Fang, Mônica Jung de Andrade, Fatma Göktepe, Özer Göktepe, Seyed M. Mirvakili, Sina Naficy, Xavier Lepró, Jiyoung Oh, Mikhail E. Kozlov, Seon Jeong Kim, Xiuru Xu, Benjamin J. Swedlove, Gordon G. Wallace, and Ray H. Baughman. Artificial Muscles from Fishing Line and Sewing Thread. Science, 343(6173):868–872, February 2014.
- [47] Travis Hainsworth, Lawrence Smith, Sebastian Alexander, and Robert MacCurdy. A Fabrication Free, 3D Printed, Multi-Material, Self-Sensing Soft Actuator. IEEE Robotics and Automation Letters, 5(3):4118–4125, July 2020.
- [48] Tim Helps, Christian Romero, Majid Taghavi, Andrew T. Conn, and Jonathan Rossiter. Liquid-amplified zipping actuators for micro-air vehicles with transmission-free flapping. Science Robotics, 7(63):eabi8189, February 2022.
- [49] H.L. Hess and R.J. Baker. Transformerless capacitive coupling of gate signals for series operation of power MOS devices. In IEEE International Electric Machines and Drives Conference. IEMDC'99. Proceedings (Cat. No.99EX272), pages 673–675, May 1999.

- [50] H.L. Hess and R.J. Baker. Transformerless capacitive coupling of gate signals for series operation of power MOS devices. IEEE Transactions on Power Electronics, 15(5):923–930, September 2000.
- [51] Ronan Hinchet and Herbert Shea. High Force Density Textile Electrostatic Clutch. Advanced Materials Technologies, 5(4):1900895, 2020.
- [52] Lindsey Hines, Kirstin Petersen, Guo Zhan Lum, and Metin Sitti. Soft Actuators for Small-Scale Robotics. Advanced Materials, 29(13):1603483, 2017.
- [53] Binyamin Hochner. An embodied view of octopus neurobiology. Current Biology, 22(20):R887–R892, October 2012.
- [54] Binyamin Hochner. How nervous systems evolve in relation to their embodiment: What we can learn from octopuses and other molluscs. Brain, Behavior and Evolution, 82(1):19–30, 2013.
- [55] N. Hogan, E. Bizzi, F. A. Mussa-Ivaldi, and T. Flash. Controlling multijoint motor behavior. Exercise and sport sciences reviews, 15:153–190, 1987.
- [56] Neville Hogan and Tamar Flash. Moving gracefully: Quantitative theories of motor coordination. Trends in Neurosciences, 10(4):170–174, April 1987.
- [57] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable Programming for Physical Simulation. ICLR, 2020.
- [58] Yuanming Hu, Tzu-Mao Li, Luke Anderson, Jonathan Ragan-Kelley, and Frédo Durand. Taichi: A language for high-performance computation on spatially sparse data structures. ACM Transactions on Graphics (TOG), 38(6):201, 2019.
- [59] Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. ChainQueen: A Real-Time Differentiable Physical Simulator for Soft Robotics. Proceedings of IEEE International Conference on Robotics and Automation (ICRA), 2019.
- [60] Lina Huang, Zhe Zhang, and Michael A. E. Andersen. A review of high voltage drive amplifiers for capacitive actuators. In 2012 47th International Universities Power Engineering Conference (UPEC), pages 1–6, September 2012.
- [61] J. S. Humbert and A. M. Hyslop. Bioinspired Visuomotor Convergence. IEEE Transactions on Robotics, 26(1):121–130, February 2010.
- [62] J. Sean Humbert, Richard M. Murray, and Michael H. Dickinson. Sensorimotor Convergence in Visual Navigation and Flight Control Systems. IFAC Proceedings Volumes, 38(1):253–258, January 2005.
- [63] Cong Truc Huynh, Minh Khanh Nguyen, Mantas Naris, Gulen Yesilbag Tonga, Vincent M Rotello, and Eben Alsberg. Light-triggered RNA release and induction of hMSC osteogenesis via photodegradable, dual-crosslinked hydrogels. Nanomedicine, 11(12):1535–1550, June 2016.

- [64] J George Hwang, Markus Zahn, and Leif AA Pettersson. Mechanisms behind positive streamers and their distinct propagation modes in transformer oil. IEEE Transactions on Dielectrics and Electrical Insulation, 19(1), 2012.
- [65] Andrew Hyslop, Holger G. Krapp, and J. Sean Humbert. Control theoretic interpretation of directional motion preferences in optic flow processing interneurons. Biological Cybernetics, 103(5):353–364, November 2010.
- [66] Sadaf Ijaz, Hao Li, Manh Cuong Hoang, Chang-Sei Kim, Doyeon Bang, Eunpyo Choi, and Jong-Oh Park. Magnetically actuated miniature walking soft robot based on chained magnetic microparticles-embedded elastomer. Sensors and Actuators A: Physical, 301:111707, January 2020.
- [67] Dimitri Jeltsema and Jacquelen Scherpen. Multidomain modeling of nonlinear networks and systems: Energy- and power-based perspectives. Control Systems, IEEE, 29:28–59, September 2009.
- [68] Xiaobin Ji, Xinchang Liu, Vito Cacucciolo, Matthias Imboden, Yoan Civet, Alae El Haitami, Sophie Cantin, Yves Perriard, and Herbert Shea. An autonomous untethered fast soft robotic insect driven by low-voltage dielectric elastomer actuators. Science Robotics, 4(37):eaaz6451, December 2019.
- [69] B. K. Johnson*, M. Naris*, V. Sundaram, A. Volchko, K. Ly, S. K. Mitchell, E. Acome, N. Kellaris, C. Keplinger, N. Correll, J. S. Humbert, and M. E. Rentschler. A multifunctional soft robotic shape display with high-speed actuation, sensing, and control. Nature Communications, 14(1):4516, July 2023.
- [70] Brian K. Johnson, Vani Sundaram, Mantas Naris, Eric Acome, Khoi Ly, Nikolaus Correll, Christoph Keplinger, James Sean Humbert, and Mark E. Rentschler. Identification and Control of a Nonlinear Soft Actuator and Sensor System. IEEE Robotics and Automation Letters, 5(3):3783–3790, July 2020.
- [71] E. Kamenar, N. Črnjarić-Žic, D. Haggerty, S. Zelenika, E. W. Hawkes, and I. Mezić. Prediction of the behavior of a pneumatic soft robot based on Koopman operator theory. In 2020 43rd International Convention on Information, Communication and Electronic Technology (MIPRO), pages 1169–1173, September 2020.
- [72] 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, April 2012.
- [73] Nicholas Kellaris, Philipp Rothemund, Yi Zeng, Shane K. Mitchell, Garrett M. Smith, Kaushik Jayaram, and Christoph Keplinger. Spider-Inspired Electrohydraulic Actuators for Fast, Soft-Actuated Joints. Advanced Science, 8(14):2100916, 2021.
- [74] Nicholas Kellaris, Vidyacharan Gopaluni Venkata, Philipp Rothemund, and Christoph Keplinger. An analytical model for the design of Peano-HASEL actuators with drastically improved performance. Extreme Mechanics Letters, 29:100449, May 2019.
- [75] Nicholas Kellaris, Vidyacharan Gopaluni Venkata, Garrett M. Smith, Shane K. Mitchell, and Christoph Keplinger. Peano-HASEL actuators: Muscle-mimetic, electrohydraulic transducers that linearly contract on activation. Science Robotics, 3(14):eaar3276, January 2018.

- [76] Christoph Keplinger, Martin Kaltenbrunner, Nikita Arnold, and Siegfried Bauer. Capacitive extensometry for transient strain analysis of dielectric elastomer actuators. Applied Physics Letters, 92(19):1–4, 2008.
- [77] William M. Kier. The musculature of coleoid cephalopod arms and tentacles. Frontiers in Cell and Developmental Biology, 4, 2016.
- [78] William M. Kier and Kathleen K. Smith. Tongues, tentacles and trunks: The biomechanics of movement in muscular-hydrostats. Zoological Journal of the Linnean Society, 83(4):307–324, April 1985.
- [79] William M. Kier and Michael P. Stella. The arrangement and function of octopus arm musculature and connective tissue. Journal of Morphology, 268(10):831–843, October 2007.
- [80] Sangbae Kim, Cecilia Laschi, and Barry Trimmer. Soft robotics: A bioinspired evolution in robotics. Trends in Biotechnology, 31(5):287–294, May 2013.
- [81] Bernard O Koopman. Hamiltonian systems and transformation in Hilbert space. Proceedings of the National Academy of Sciences, 17(5):315–318, 1931.
- [82] Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. A scalable pipeline for designing reconfigurable organisms. Proceedings of the National Academy of Sciences of the United States of America, 117(4):1853–1859, January 2020.
- [83] Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. Kinematic self-replication in reconfigurable organisms. Proceedings of the National Academy of Sciences of the United States of America, 118(49):e2112672118, December 2021.
- [84] Sanjay Lall, Jerrold E Marsden, and Sonja Glavǎski. A subspace approach to balanced truncation for model reduction of nonlinear control systems. International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal, 12(6):519–535, 2002.
- [85] Cecilia Laschi and Matteo Cianchetti. Soft Robotics: New Perspectives for Robot Bodyware and Control. Frontiers in Bioengineering and Biotechnology, 2, 2014.
- [86] Cecilia Laschi, Barbara Mazzolai, and Matteo Cianchetti. Soft robotics: Technologies and systems pushing the boundaries of robot abilities. Science Robotics, 1(1):eaah3690, December 2016.
- [87] Cecilia Laschi, Thomas George Thuruthel, Fumiya Lida, Rochdi Merzouki, and Egidio Falotico. Learning-Based Control Strategies for Soft Robots: Theory, Achievements, and Future Challenges. IEEE Control Systems Magazine, 43(3):100–113, June 2023.
- [88] R Lector. Series connection of MOSFET, bipolar and IGBT devices. In SGS-Thomson Designers’ Guide to Power Products, pages 759–769. SGS Thompson, New York, 1994.
- [89] Yunzhu Li, Hao He, Jiajun Wu, Dina Katabi, and Antonio Torralba. Learning Compositional Koopman Operators for Model-Based Control, April 2020.

- [90] Xiang Lin, Lakshmi Ravi, Dong Dong, and Rolando Burgos. Hybrid Voltage Balancing Approach for Series-Connected 10 kV SiC MOSFETs for DC-AC Medium-Voltage Power Conversion Applications. In 2020 IEEE Energy Conversion Congress and Exposition (ECCE), pages 3769–3775, October 2020.
- [91] Darren J. Lipomi, Michael Vosgueritchian, Benjamin C-K. Tee, Sondra L. Hellstrom, Jennifer A. Lee, Courtney H. Fox, and Zhenan Bao. Skin-like pressure and strain sensors based on transparent elastic films of carbon nanotubes. Nature Nanotechnology, 6:788–792, 2011.
- [92] Y. Z. Liu, Z. W. Hao, J. X. Yu, X. R. Zhou, Pooi See Lee, Y. Sun, Z. C. Mu, and F. L. Zeng. A high-performance soft actuator based on a poly(vinylidene fluoride) piezoelectric bimorph. Smart Materials and Structures, 28(5):055011, April 2019.
- [93] Tianjun Long, Lei Pang, Geqi Li, Chenhui Zhou, Mingtian Ye, Xuanyu Chen, and Qiaogen Zhang. An 8kV Series-Connected MOSFETs Module that Requires One Single Gate Driver. In 2018 IEEE International Power Modulator and High Voltage Conference (IPMHVC), pages 383–386, June 2018.
- [94] Khoi Ly, Nicholas Kellaris, Dade McMorris, Brian K. Johnson, Eric Acome, Vani Sundaram, Mantas Naris, J. Sean Humbert, Mark E. Rentschler, Christoph Keplinger, and Nikolaus Correll. Miniaturized Circuitry for Capacitive Self-Sensing and Closed-Loop Control of Soft Electrostatic Transducers. Soft Robotics, 8(6):673–686, December 2021.
- [95] Emily A. Margolis, Nicole E. Friend, Marsha W. Rolle, Eben Alsberg, and Andrew J. Putnam. Manufacturing the multiscale vascular hierarchy: Progress toward solving the grand challenge of tissue engineering. Trends in Biotechnology, 0(0), May 2023.
- [96] Michael A Meller, Matthew Bryant, and Ephraim Garcia. Reconsidering the McKibben muscle: Energetics, operating fluid, and bladder material. Journal of Intelligent Material Systems and Structures, 25(18):2276–2293, December 2014.
- [97] Erik J. Mentze, Herbert L. Hess, Kevin Matthew Buck, and Tracey G. Windley. A Scalable High-Voltage Output Driver for Low-Voltage CMOS Technologies. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 14(12):1347–1353, December 2006.
- [98] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. Nonlinear Dynamics, 41(1-3):309–325, 2005.
- [99] Shane K. Mitchell, Trent Martin, and Christoph Keplinger. A Pocket-Sized Ten-Channel High Voltage Power Supply for Soft Electrostatic Actuators. Advanced Materials Technologies, 7(8):2101469, 2022.
- [100] Shane K. Mitchell, Xingrui Wang, Eric Acome, Trent Martin, Khoi Ly, Nicholas Kellaris, Vidyacharan Gopaluni Venkata, and Christoph Keplinger. An Easy-to-Implement Toolkit to Create Versatile and High-Performance HASEL Actuators for Untethered Soft Robots. Advanced Science, 6(14):1900178, 2019.
- [101] Concepción A. Monje, Bastian Deutschmann, Jorge Muñoz, Christian Ott, and Carlos Balaguer. Fractional Order Control of Continuum Soft Robots: Combining Decoupled/Reduced-Dynamics Models and Robust Fractional Order Controllers for Complex Soft Robot Motions. IEEE Control Systems Magazine, 43(3):66–99, June 2023.

- [102] Bobak Mosadegh, Panagiotis Polygerinos, Christoph Keplinger, Sophia Wennstedt, Robert F. Shepherd, Unmukt Gupta, Jongmin Shim, Katia Bertoldi, Conor J. Walsh, and George M. Whitesides. Pneumatic Networks for Soft Robotics that Actuate Rapidly. Advanced Functional Materials, 24(15):2163–2170, 2014.
- [103] Mantas Naris, Nicholas S. Szczecinski, and Roger D. Quinn. A neuromechanical model exploring the role of the common inhibitor motor neuron in insect locomotion. Biological Cybernetics, 114(1):23–41, February 2020.
- [104] Janna C. Nawroth, Hyungsuk Lee, Adam W. Feinberg, Crystal M. Ripplinger, Megan L. McCain, Anna Grosberg, John O. Dabiri, and Kevin Kit Parker. A tissue-engineered jellyfish with biomimetic propulsion. Nature Biotechnology, 30(8):792–797, August 2012.
- [105] B. J. Nel and S. Perinpanayagam. A Brief Overview of SiC MOSFET Failure Modes and Design Reliability. Procedia CIRP, 59:280–285, January 2017.
- [106] Di Ni, Ronald Heisser, Benyamin Davaji, Landon Ivy, Robert Shepherd, and Amit Lal. Polymer interdigitated pillar electrostatic (PIPE) actuators. Microsystems & Nanoengineering, 8(1):1–8, January 2022.
- [107] Ailish O’Halloran, Fergal O’Malley, and Peter McHugh. A review on dielectric elastomer actuators, technology, applications, and challenges. Journal of Applied Physics, 104(7):071101, October 2008.
- [108] Stacey Page, Andrew Wajda, and Herbert Hess. High voltage tolerant stacked MOSFET in a Buck converter application. In 2012 IEEE Workshop on Microelectronics and Electron Devices, pages 1–4, April 2012.
- [109] Sanghyeon Park, Aaron Goldin, and Juan Rivas-Davila. Miniature High-Voltage DC-DC Power Converters for Space and Micro-Robotic Applications. In 2019 IEEE Energy Conversion Congress and Exposition (ECCE), pages 2007–2014, September 2019.
- [110] Sangtak Park, Yanhui Bai, and John T. W. Yeow. Design and analysis of resonant drive circuit for electrostatic actuators. In 2010 International Symposium on Optomechatronic Technologies, pages 1–6, October 2010.
- [111] Ron Pelrine, Roy Kornbluh, Qibing Pei, and Jose Joseph. High-speed electrically actuated elastomers with strain greater than 100%. Science, 287(5454):836–839, 2000.
- [112] Rolf Pfeifer, Max Lungarella, and Fumiya Iida. Self-Organization, Embodiment, and Biologically Inspired Robotics. Science, 318(5853):1088–1093, November 2007.
- [113] S. Pourazadi, A. Shagerdmootaab, H. Chan, M. Moallem, and C. Menon. On the electrical safety of dielectric elastomer actuators in proximity to the human body. Smart Materials and Structures, 26(11):115007, October 2017.
- [114] André Preumont. Mechatronics: Dynamics of Electromechanical and Piezoelectric Systems, volume 136 of Solid Mechanics and Its Applications. Springer, 2006.
- [115] J. Proctor, S. Brunton, and J. Kutz. Dynamic Mode Decomposition with Control. SIAM Journal on Applied Dynamical Systems, 15(1):142–161, January 2016.

- [116] Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Generalizing Koopman theory to allow for inputs and control. SIAM Journal on Applied Dynamical Systems, 17(1):909–930, 2018.
- [117] Wei Qian, Dong Cao, Jorge G. Cintron-Rivera, Matthew Gebben, Dennis Wey, and Fang Zheng Peng. A Switched-Capacitor DC–DC Converter With High Voltage Gain and Reduced Component Rating and Count. IEEE Transactions on Industry Applications, 48(4):1397–1406, July 2012.
- [118] Christopher Rackauckas and Qing Nie. DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in julia. Journal of Open Research Software, 5(1):15, 2017.
- [119] Giuseppe Rega and Hans Troger. Dimension reduction of dynamical systems: Methods, models, applications. Nonlinear Dynamics, 41(1-3):1–15, 2005.
- [120] Luquan Ren, Bingqian Li, Guowu Wei, Kunyang Wang, Zhengyi Song, Yuyang Wei, Lei Ren, and Qingping Liu. Biology and bioinspiration of soft robotics: Actuation, sensing, and system integration. iScience, 24(9):103075, September 2021.
- [121] Kang Rongjie, Guglielmino Emanuele, Zullo Letizia, David T. Branson, Godage Isuru, and Darwin G. Caldwell. Embodiment design of soft continuum robots. Advances in Mechanical Engineering; London, 8(4), April 2016.
- [122] Jonathan Rossiter and Helmut Hauser. Soft Robotics - The Next Industrial Revolution? IEEE Robotics & Automation Magazine, 23(3):17–20, September 2016.
- [123] Jonathan Rossiter, Peter Walters, and Boyko Stoimenov. Printing 3D dielectric elastomer actuators for soft robotics. In Electroactive Polymer Actuators and Devices (EAPAD) 2009, volume 7287, pages 149–158. SPIE, April 2009.
- [124] Philipp Rothmund, Nicholas Kellaris, Shane K. Mitchell, Eric Acome, and Christoph Keplinger. HASEL Artificial Muscles for a New Generation of Lifelike Robots—Recent Progress and Future Opportunities. Advanced Materials, 33(19):2003375, 2021.
- [125] Philipp Rothmund, Yoonho Kim, Ronald H. Heisser, Xuanhe Zhao, Robert F. Shepherd, and Christoph Keplinger. Shaping the future of robotics through materials innovation. Nature Materials, 20(12):1582–1587, December 2021.
- [126] Philipp Rothmund, Sophie Kirkman, and Christoph Keplinger. Dynamics of electrohydraulic soft actuators. Proceedings of the National Academy of Sciences, 117(28):16207–16213, July 2020.
- [127] Clarence W Rowley, Igor Mezić, Shervin Bagheri, Philipp Schlatter, and Dan S Henningson. Spectral analysis of nonlinear flows. Journal of fluid mechanics, 641:115–127, 2009.
- [128] Scott Rubeo, Nicholas Szczecinski, and Roger Quinn. A Synthetic Nervous System Controls a Simulated Cockroach. Applied Sciences, 8(1):6, 2017.
- [129] Daniela Rus and Michael T. Tolley. Design, fabrication and control of soft robots. Nature, 521(7553):467–475, May 2015.

- [130] Samuel Shian, Katia Bertoldi, and David R. Clarke. Dielectric Elastomer Based “Grippers” for Soft Robotics. *Advanced Materials*, 27(43):6814–6819, 2015.
- [131] Harshal Arun Sonar and Jamie Paik. Soft Pneumatic Actuator Skin with Piezoelectric Sensors for Vibrotactile Feedback. *Frontiers in Robotics and AI*, 2, 2016.
- [132] Wilfried Stoll, Elias Knubben, Johannes Stoll, Florian Mangold, Martin Timmer, Simon Nikolai Fröhlich, Felix Müller-Graf, Alexander Hildebrandt, Steffen Hülsmann, and Martin Thomas Ehrle. *WaveHandling*, 2013.
- [133] Hannah A. Strobel, Anna D. Dikina, Karen Levi, Loran D. Solorio, Eben Alsberg, and Marsha W. Rolle. Cellular Self-Assembly with Microsphere Incorporation for Growth Factor Delivery Within Engineered Vascular Tissue Rings. *Tissue Engineering Part A*, 23(3-4):143–155, February 2017.
- [134] German Sumbre, Yoram Gutfreund, Graziano Fiorito, Tamar Flash, and Binyamin Hochner. Control of octopus arm extension by a peripheral motor program. *Science*, 293(5536):1845–1848, September 2001.
- [135] Germán Sumbre, Graziano Fiorito, Tamar Flash, and Binyamin Hochner. Neurobiology: Motor control of flexible octopus arms. *Nature*, 433(7026):595–596, February 2005.
- [136] Germán Sumbre, Graziano Fiorito, Tamar Flash, and Binyamin Hochner. Octopuses use a human-like strategy to control precise point-to-point arm movements. *Current Biology*, 16(8):767–772, April 2006.
- [137] Keyao Sun, Emma Raszmann, Jun Wang, Xiang Lin, Rolando Burgos, Dong Dong, and Dushan Boroyevich. Modeling, Design, and Evaluation of Active dv/dt Balancing for Series-Connected SiC MOSFETs. *IEEE Transactions on Power Electronics*, 37(1):534–546, January 2022.
- [138] Vani Sundaram, Khoi Ly, Brian K. Johnson, Mantas Naris, Maxwell P. Anderson, James Sean Humbert, Nikolaus Correll, and Mark Rentschler. Embedded Magnetic Sensing for Feedback Control of Soft HASEL Actuators. *IEEE Transactions on Robotics*, 39(1):808–822, February 2023.
- [139] Nicholas S. Szczecinski, Amy E. Brown, John A. Bender, Roger D. Quinn, and Roy E. Ritzmann. A neuromechanical simulation of insect walking and transition to turning of the cockroach *Blaberus discoidalis*. *Biological Cybernetics*, 108(1):1–21, February 2014.
- [140] Nicholas S. Szczecinski, Alexander J. Hunt, and Roger D. Quinn. Design process and tools for dynamic neuromechanical models and robot controllers. *Biological Cybernetics; Heidelberg*, 111(1):105–127, February 2017.
- [141] Nicholas S. Szczecinski, Alexander J. Hunt, and Roger D. Quinn. A Functional Subnetwork Approach to Designing Synthetic Nervous Systems That Control Legged Robot Locomotion. *Frontiers in Neurorobotics*, 11, 2017.
- [142] Nicholas S Szczecinski and Roger D Quinn. MantisBot Changes Stepping Speed by Entraining CPGs to Positive Velocity Feedback. In *Lecture Notes in Artificial Intelligence 10384*, pages 440–52, 2017.

- [143] Nicholas S Szczecinski and Roger D Quinn. Template for the Neural Control of Directed Walking Generalized to All Legs of MantisBot. Bioinspiration and Biomimetics, 2017.
- [144] Ion-Dan Sîrbu, David Preninger, Doris Danninger, Lukas Penkner, Reinhard Schwödiauer, Giacomo Moretti, Nikita Arnold, Marco Fontana, and Martin Kaltenbrunner. Electrostatic actuators with constant force at low power loss using matched dielectrics. Nature Electronics, 6(11):888–899, November 2023.
- [145] K Tai and TH Chee. Design of structures and compliant mechanisms by evolutionary optimization of morphological representations of topology. J. Mech. Des., 122(4):560–566, 1998.
- [146] Yichao Tang and Alireza Khaligh. Bidirectional Resonant DC–DC Step-Up Converters for Driving High-Voltage Actuators in Mobile Microrobots. IEEE Transactions on Power Electronics, 31(1):340–352, January 2016.
- [147] Thomas George Thuruthel, Yasmin Ansari, Egidio Falotico, and Cecilia Laschi. Control Strategies for Soft Robotic Manipulators: A Survey. Soft Robotics, 5(2):149–163, January 2018.
- [148] Michael T. Tolley, Robert F. Shepherd, Bobak Mosadegh, Kevin C. Galloway, Michael Wehner, Michael Karpelson, Robert J. Wood, and George M. Whitesides. A Resilient, Untethered Soft Robot. Soft Robotics, 1(3):213–223, September 2014.
- [149] Angella Volchko, Shane K. Mitchell, Timothy G. Morrissey, and J. Sean Humbert. Model-Based Data-Driven System Identification and Controller Synthesis Framework for Precise Control of SISO and MISO HASEL-Powered Robotic Systems. In 2022 IEEE 5th International Conference on Soft Robotics (RoboSoft), pages 209–216, April 2022.
- [150] Ian D. Walker. Continuous Backbone “Continuum” Robot Manipulators. International Scholarly Research Notices, 2013:e726506, July 2013.
- [151] Michael Wehner, Ryan L. Truby, Daniel J. Fitzgerald, Bobak Mosadegh, George M. Whitesides, Jennifer A. Lewis, and Robert J. Wood. An integrated design and fabrication strategy for entirely soft, autonomous robots. Nature, 536(7617):451–455, August 2016.
- [152] Qiang Xiao, Yang Yan, Xinke Wu, Na Ren, and Kuang Sheng. A 10kV/200A SiC MOS-FET module with series-parallel hybrid connection of 1200V/50A dies. In 2015 IEEE 27th International Symposium on Power Semiconductor Devices & IC’s (ISPSD), pages 349–352, May 2015.
- [153] Youcan Yan, Zhe Hu, Zhengbao Yang, Wenzhen Yuan, Chaoyang Song, Jia Pan, and Yajing Shen. Soft magnetic skin for super-resolution tactile sensing with force self-decoupling. Science Robotics, 6(51):eabc8801, February 2021.
- [154] Yoram Yekutieli, Rea Mitelman, Binyamin Hochner, and Tamar Flash. Analyzing octopus movements using three-dimensional reconstruction. Journal of Neurophysiology, 98(3):1775–1790, September 2007.
- [155] Yoram Yekutieli, Roni Sagiv-Zohar, Ranit Aharonov, Yaakov Engel, Binyamin Hochner, and Tamar Flash. Dynamic model of the octopus arm. I. biomechanics of the octopus reaching movement. Journal of Neurophysiology, 94(2):1443–1458, August 2005.

- [156] Yoram Yekutieli, Roni Sagiv-Zohar, Binyamin Hochner, and Tamar Flash. Dynamic model of the octopus arm. II. control of reaching movements. Journal of Neurophysiology, 94(2):1459–1468, August 2005.
- [157] Yoram Yekutieli, German Sumbre, Tamar Flash, and Binyamin Hochner. How to move with no rigid skeleton? page 6, 2002.
- [158] Zachary Yoder, Nicholas Kellaris, Christina Chase-Markopoulou, Devon Ricken, Shane K. Mitchell, Madison B. Emmett, Richard F. ff. Weir, Jacob Segil, and Christoph Keplinger. Design of a High-Speed Prosthetic Finger Driven by Peano-HASEL Actuators. Frontiers in Robotics and AI, 7, 2020.
- [159] Zachary Yoder, Daniela Macari, Gavriel Kleinwaks, Ingemar Schmidt, Eric Acome, and Christoph Keplinger. A Soft, Fast and Versatile Electrohydraulic Gripper with Capacitive Object Size Detection. Advanced Functional Materials, 33(3):2209080, 2023.
- [160] Jung-Hwan Youn, Seung Mo Jeong, Geonwoo Hwang, Hyunwoo Kim, Kyujin Hyeon, Jihwan Park, and Ki-Uk Kyung. Dielectric Elastomer Actuator for Soft Robotics Applications and Challenges. Applied Sciences, 10(2):640, January 2020.
- [161] Ido Zelman, Meirav Galun, Ayelet Akselrod-Ballin, Yoram Yekutieli, Binyamin Hochner, and Tamar Flash. Nearly automatic motion capture system for tracking octopus arm movements in 3D space. Journal of Neuroscience Methods, 182(1):97–109, August 2009.
- [162] Ido Zelman, Myriam Titon, Yoram Yekutieli, Shlomi Hanassy, Binyamin Hochner, and Tamar Flash. Kinematic decomposition and classification of octopus arm movements. Frontiers in Computational Neuroscience, 7, 2013.
- [163] Huichan Zhao, Kevin O’Brien, Shuo Li, and Robert F. Shepherd. Optoelectronically innervated soft prosthetic hand via stretchable optical waveguides. Science Robotics, 1(1):eaai7529, December 2016.