

Battery Pack Design and Transient Performance Modeling for High-Power Legged Robots

by

Christopher K. Evagora

B.S. in Electrical Engineering and Computer Science, MIT, 2023

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

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND
COMPUTER SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2025

© 2025 Christopher K. Evagora. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Christopher K. Evagora
Department of Electrical Engineering and Computer Science
Jan 24, 2025

Certified by: Sangbae Kim
Professor of Mechanical Engineering, Thesis Supervisor

Accepted by: Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Battery Pack Design and Transient Performance Modeling for High-Power Legged Robots

by

Christopher K. Evagora

Submitted to the Department of Electrical Engineering and Computer Science
on Jan 24, 2025 in partial fulfillment of the requirements for the degree of

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND
COMPUTER SCIENCE

ABSTRACT

Legged robotics has recently shifted toward advanced optimization-based control methods, such as Model Predictive Control (MPC), to generate agile and energy-efficient locomotion. By casting the control problem as an optimization task, robotic systems can account for complex robot dynamics and operational constraints, including joint limits and actuator capabilities. However, high-performance maneuvers also demand rigorous consideration of onboard battery constraints. This work presents an empirically derived lithium-ion battery model that captures transient voltage sag and time-dependent internal battery state, enabling more accurate prediction of feasible power delivery. Additionally, a custom high-power battery pack was designed to meet the power demands of the MIT Humanoid, emphasizing power density, safety, and maintainability. Although the work presented in this thesis does not integrate the battery model into a trajectory optimization framework, it establishes the foundation for future research that aims to couple battery and robot dynamics in robot control. Ultimately, this approach will facilitate safer and more capable legged robots by ensuring that planned trajectories respect both physical and electrochemical constraints.

Thesis supervisor: Sangbae Kim

Title: Professor of Mechanical Engineering

Acknowledgments

I would like to thank:

Sangbae Kim for allowing me to conduct research and learn more about robotics and batteries in an environment that offers great freedom to do so.

Elijah Stanger-Jones for offering continuous guidance and advice with my thesis, research, and engineering endeavors.

Andrew Saloutos for always being willing to give advice, whether it be my engineering problems or robotics in general.

Hongmin Kim for advising me on mechanical design and how to properly perform GD&T for mechanical drawings.

Matthew Chignoli for helping me integrate battery models into his trajectory optimization code.

Charles Khazoom for always lifting my spirits.

Miguel Talamantez for helping me lay the groundwork for battery data collection and with battery component validation.

Miguel Flores-Acton for spending long hours writing data collection code and perfecting our data collection methods and battery models.

The MIT Biomimetics Lab for pursuing academic excellence, possessing general prowess in robotics, and an eagerness in each individual to share their expertise. I have been influenced greatly by you all.

Contents

<i>List of Figures</i>	9
1 Introduction	13
1.1 Project Motivation	13
1.2 Battery Pack Modeling	16
1.3 Battery Pack Design	16
1.4 Related Work	16
2 Battery Pack Design	19
2.1 Cell Selection and Battery Sizing	19
2.1.1 Cell Selection	19
2.1.2 Battery Sizing	20
2.1.3 Thermal Considerations	28
2.2 Mechanical Design	28
2.2.1 Mechanical Structure	31
2.2.2 Electrical Interconnect	39
2.3 Electrical Design	46
2.3.1 The Carrier PCB	47
2.3.2 The Power Board	47
2.3.3 The Segment	48
2.3.4 The BOB	49
2.3.5 Electrical Assembly	51

3	Modeling Batteries	53
3.1	Modeling OCV vs SOC	54
3.2	Modeling Cell Impedance	57
3.2.1	The Ideal Voltage Source	57
3.2.2	The Simple IR Model	57
3.2.3	The Randles Model	59
3.3	Electrochemical Impedance Spectroscopy	61
3.3.1	Performing EIS	62
3.3.2	Data Collection Validation	66
3.3.3	Fitting the Extended Randles Model	67
3.3.4	Proof of Mapping ERM from Battery Cell to Battery Pack	71
4	Simulating Batteries	75
4.1	Forward Simulation of the Extended Randles Model	75
4.2	The Issue of Imaginary Voltage	78
4.2.1	Linear Extrapolation of the Voltage Curve	79
4.3	Constraining The Optimization	81
5	Conclusion and Future Work	83
A	ERM Equivalent Circuit Model Parameters for P45B	85
	<i>References</i>	89

List of Figures

1.1	MIT Humanoid Robot and typical power drawn, easily exceeding 3kW in the case of a simple jump.	14
2.1	Cross Section of Humanoid Torso with 48 2170 battery cells	21
2.2	Simple IR model of a cell	22
2.3	Predicted Battery Voltage vs Output Power, graphed with Desmos[3]	25
2.4	Predicted Battery Voltage vs Output Power, graphed with Desmos[3]	27
2.5	Final CAD of battery design	29
2.6	Final battery assembly	30
2.7	Battery pack inserted into Humanoid Torso	30
2.8	Double braced polycarbonate structure that is bonded to cell surfaces to form main battery structure	32
2.9	Aluminum spacer with male tapered dovetail	33
2.10	Aluminum piece with female tapered dovetail	33
2.11	Flexure bar highlighted in blue	34
2.12	Compliant interface mating with female tapered dovetail through large surface contact.	35
2.13	Compliant interface highlighted in blue mating with internal side strut of the humanoid using a large boss feature and bolts for retention.	36
2.14	Compliant interface with female tapered dovetail and flexure retention mechanism installed in humanoid torso.	36

2.15	The flexure mechanism is unable to open due to the installation of the rear battery panel.	37
2.16	The BOB highlighted in blue makes all electrical connections to the battery when it is inserted.	38
2.17	Example nickel interconnect connecting six cells into a 2S3P configuration.	39
2.18	Final routing solution for pack.	40
2.19	Unique nickel interconnects excluding piece that wraps back overtop.	41
2.20	Carrier PCB installed on module.	42
2.21	Nickel interconnect placed and soldered on carrier PCB	42
2.22	Nickel warts that are soldered to the carrier PCB.	43
2.23	Each interconnect is spot welded by hand.	44
2.24	A layer of Kapton insulates the spot welded interconnects.	45
2.25	The final interconnect that folds over is soldered to the exposed interconnects and is insulated over.	45
2.26	The three main PCBs in the battery pack.	47
2.27	The power board.	48
2.28	The segment.	49
2.29	The BOB highlighted in blue makes all electrical connections to the battery when it is inserted.	50
2.30	The BOB fully assembled.	50
2.31	The power board and segment installed in a test battery.	51
2.32	The electrical assembly is covered by a 3D printed cover.	51
3.1	Cell testing setup consisting of a DC load, DC supply, and water cooling loop to maintain temperature.	55
3.2	Cell jig which includes a kelvin connection as well as two water cooling blocks.	55
3.3	Discharge curve showing 1A constant discharge as well as a DCIR compensated OCV vs SOC curve.	56
3.4	The battery voltage decays during a current pulse which lasts 5 seconds.	59

3.5	The Randles Model with Warburg Impedance	60
3.6	The Extended Randles Model with a chain of n RC circuits.	60
3.7	Test bench for conducting electrochemical impedance spectroscopy on lithium ion battery cells.	64
3.8	Collected EIS data on the P45B cells. The scaling of the real and imaginary axes are the same.	65
3.9	The constructed fake cell.	66
3.10	The two control impedances and their corresponding Nyquist plots. The scaling of the real and imaginary axes are the same.	67
3.11	Fitted ERM with 1 RC circuit	68
3.12	Fitted ERM with 2 RC circuits	69
3.13	Fitted ERM with 3 RC circuits	70
3.14	Fitted ERM with 4 RC circuits	71
3.15	Cell level ERM on the left and pack level ERM on the right with only one RC circuit.	72
4.1	ERM with single RC circuit with nodes, currents, and charges annotated. . .	76
4.2	Battery voltage curve which transitions into a linear function, graphed with Desmos[3].	80

Chapter 1

Introduction

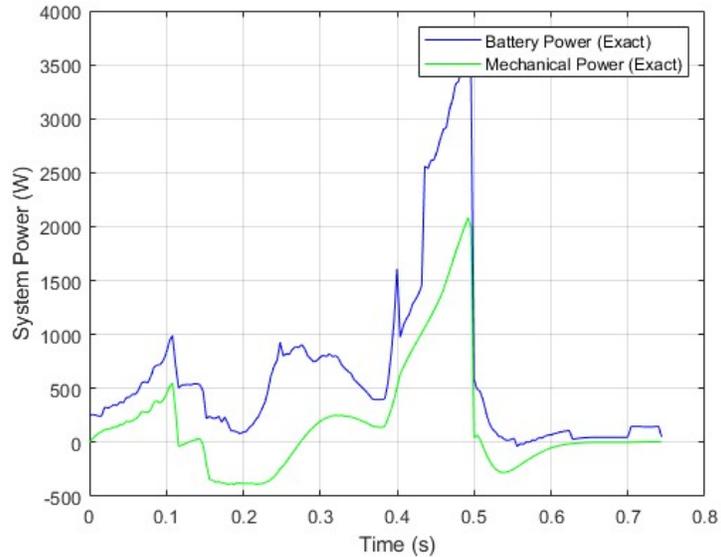
1.1 Project Motivation

The field of legged robotics has seen a recent shift, moving away from conventional and conservative approaches to locomotion planning. In particular, researchers have increasingly embraced the use of sophisticated control algorithms capable of addressing the inherent nonlinearities found in the dynamics of legged robots. Among these advanced control methods, model predictive control (MPC) has emerged as a powerful tool, enabling the generation of energy-efficient trajectories that closely resemble the locomotion patterns observed in biological systems. At the core of many of these complex control algorithms lies the idea of formulating the control problem for highly dynamic systems as an optimization problem. This paradigm shift has opened up new possibilities for achieving agile and adaptable legged robot locomotion.

Control algorithms based on formulating control as an optimization problem offer a promising solution to address the limitations of traditional methods. By casting the control problem as an optimization problem, researchers can leverage powerful nonlinear optimization solvers to generate optimal or near-optimal trajectories for both the system's state variables and control inputs, while accounting for various constraints.



(a) MIT Humanoid



(b) Example power profile for MIT Humanoid

Figure 1.1: MIT Humanoid Robot and typical power drawn, easily exceeding 3kW in the case of a simple jump.

By leveraging the optimization framework, researchers can take advantage of the rich mathematical tools available in this domain. They can define objective functions that capture specific performance criteria, such as energy minimization, stability maximization, or task completion time reduction. Additionally, constraints can be incorporated to ensure the robot adheres to physical limitations, such as joint limits, contact forces, or actuator capabilities.

Traditionally, the development of control strategies for legged robots relied on heuristic approaches, often based on the concept of the Zero Moment Point (ZMP). While effective to some extent, these conservative methods failed to fully exploit the capabilities of legged robots and often struggled to cope with the complex dynamics and interactions encountered during locomotion. Such methods of locomotion focus primarily on position control of the joints of a robot, a scenario where high speed or power carry the possibility of violating the very strict assumptions made by control methods like ZMP about the dynamics of such a system. Modern humanoid robots such as the MIT humanoid [1] juxtapose this paradigm. By using optimal control to more fully realize the physical capabilities of the MIT Humanoid,

output powers are easily capable of exceeding 5kW on the MIT humanoid. This is an order of magnitude above traditional approaches that could get by with low power, high gear ratio actuators.

Through the use of its actuators, a robot is capable of converting chemical energy from an on board battery to mechanical energy. Lithium ion batteries are a very common choice when choosing an energy storage solution due to the chemistry's energy and power density. Despite this, it is important to understand and enforce the output power limits of a battery while solving for a valid trajectory to avoid commanding control inputs that are not possible to track on real hardware. All batteries are fundamentally limited in their maximum output power which is a function of a large number of factors such as the state of charge (SOC), the internal resistance (IR) described in certain models of an electrochemical cell, and other time dependent internal chemical processes. Given that there exists an empirical battery model derived from real world testing and an optimization framework that allows constraints to be applied while solving for valid trajectories, it naturally follows to have the solver accept constraints that enforce the dynamics of the battery alongside the dynamics of the robot itself. This idea was touched upon in [1], though such constraints were checked after a solution was found, not enforced during solving.

It is here we can find the motivation for this research. By constraining a trajectory to obey both the physical dynamics of the robot as well as the battery dynamics, the power limits of the robot's on board battery can be better understood ahead of time before real world deployment. The effects of voltage sag on actuator torque-speed limits can be well understood. The feasibility of a certain task given to the robot can therefore be assessed in simulation without the possibility of failure on the physical hardware, which can lead to potential damage to the robot itself in highly dynamic maneuvers. Having an adequate model of the battery will also minimize the difference between simulated and real world trajectories the robot may take, explaining a difference in performance that otherwise would be labeled unknown error.

1.2 Battery Pack Modeling

A key contribution of this research is the development of a battery model that characterizes how a battery pack will respond to transient loads by modeling hidden time-dependent internal battery states. This model, built upon empirical testing of lithium-ion cells, is suitable for integration into advanced control algorithms such as MPC. By incorporating the battery model directly within the optimization framework, one can enforce more intelligent power limits and account for voltage sag of the onboard battery in simulation. Although the current work did not extend to incorporating this battery model into an actual trajectory optimization, the model itself lays the groundwork for such future research. In principle, coupling this model with trajectory optimization or MPC would allow dynamic maneuvers to be planned with full knowledge of the instantaneous battery power available, ensuring feasible solutions that respect both physical dynamics and battery constraints.

1.3 Battery Pack Design

In addition to modeling battery behavior, this project undertook the design of a physical battery pack tailored to the high-power demands of the MIT Humanoid. The design of the pack ultimately focused on achieving peak electrical power to weight ratio, though the design process was also driven by practical considerations of safety, ease of maintainability, and ease of use.

1.4 Related Work

The application of battery constraints and dynamics in the work that will be done for this thesis will be based on the existing hardware and infrastructure made for the MIT humanoid. Described further by Chignoli et al. [1], the MIT humanoid is a platform for highly dynamic and acrobatic maneuvers. The paper discusses the impact of battery output power limits in

section 2, and highlights the typical constraints that are necessary to protect the battery. This paper will take battery modeling a step further, modeling the time dependent effects that all batteries have.

Battery testing and modeling is a very large and relevant field of study today following the growing importance of lightweight and power dense electronics such as in handheld devices or electric vehicles. The demand for performance in these cutting edge scenarios in turn translates into building better battery models. Muenzel et al. [2] acknowledges this need for better quantification of performance among different manufactures and cells. Included in the paper are several popular power and budget cells undergoing both electrochemical impedance spectroscopy (EIS) testing and long term capacity testing. A Randles model is used and its parameters best solved for using a non linear least squares (NLL) algorithm, allowing the cells and their models to be simulated with any ODE solver of choice, including in the optimization of humanoid robot trajectories. Similar work is done in this thesis, though the cells that were tested for this thesis were the highest power cells readily available for purchase at the time.

Chapter 2

Battery Pack Design

2.1 Cell Selection and Battery Sizing

Before work on how to mechanically house and electrically interconnect battery cells together can begin, it is first necessary to decide which battery cell will be used as well as how many battery cells are needed and in what configuration.

2.1.1 Cell Selection

When designing a battery pack, the designer must first decide on what battery cell the pack will be designed around. The two main form factors high power lithium ion battery cells are typically sold as are the cylindrical and prismatic packages.

Cylindrical lithium ion cells are the most popular form factor for lithium based batteries. They offer many advantages such as having numerous options for consumers due to the maturity of cylindrical cell technology, the low cost to manufacture, and the inherent resistance to internal cell pressures due to the geometry of the thin walled steel can itself.

Prismatic cells, sometimes called pouch cells, offer advantages over the cylindrical format such as improved packing density and drastically lower internal resistance, on the order of one fifth the resistance of a cylindrical can per unit weight. This allows prismatic cells, by

virtue of their geometry, to output much higher power per unit weight or volume compared to cylindrical cells. This makes pouch cells the optimal choice for high output power batteries.

Despite our battery pack design seeking to maximize output power to weight ratio, a cylindrical cell format was chosen over a prismatic cell due to concerns about safety and reliability. Prismatic cells are much harder to package into a battery pack for many reasons, some of which include the absence of a steel backing to hold the battery's shape, their need to be mechanically compressed to reduce the effects of battery swelling, the difficulty in structurally supporting the cells to an outer frame, and so on. Due to the absence of a thin walled steel can, prismatic cells are also much more likely to swell and potentially burst if not properly mechanically constrained and compressed. Cylindrical cells are much safer in this regard due to the outer casing acting as a thin walled pressure vessel which keeps the internals of the battery in a consistent and satisfactory state. Cylindrical cells also can be easily adhered to through the outside surface of the steel can.

In this battery pack design, the highest power cylindrical cell that was commercially available at the time of design was the MOLICEL P45B. Comparing to other cells from other manufacturers, the P45B easily had the lowest DC and AC internal impedance according to the datasheet with a $7m\Omega$ AC impedance and a $15m\Omega$ DC impedance which matched well with the cell testing that was done on them. Therefore, the P45B was chosen to be the cell to be used in this battery pack design.

2.1.2 Battery Sizing

For the MIT Humanoid, it was found that the best number of cells would be the maximum amount of cells that could fit into the cross section of the torso. Ensuring symmetry, it was possible to fit 48 P45B cells within the cross section of the torso of the MIT Humanoid.

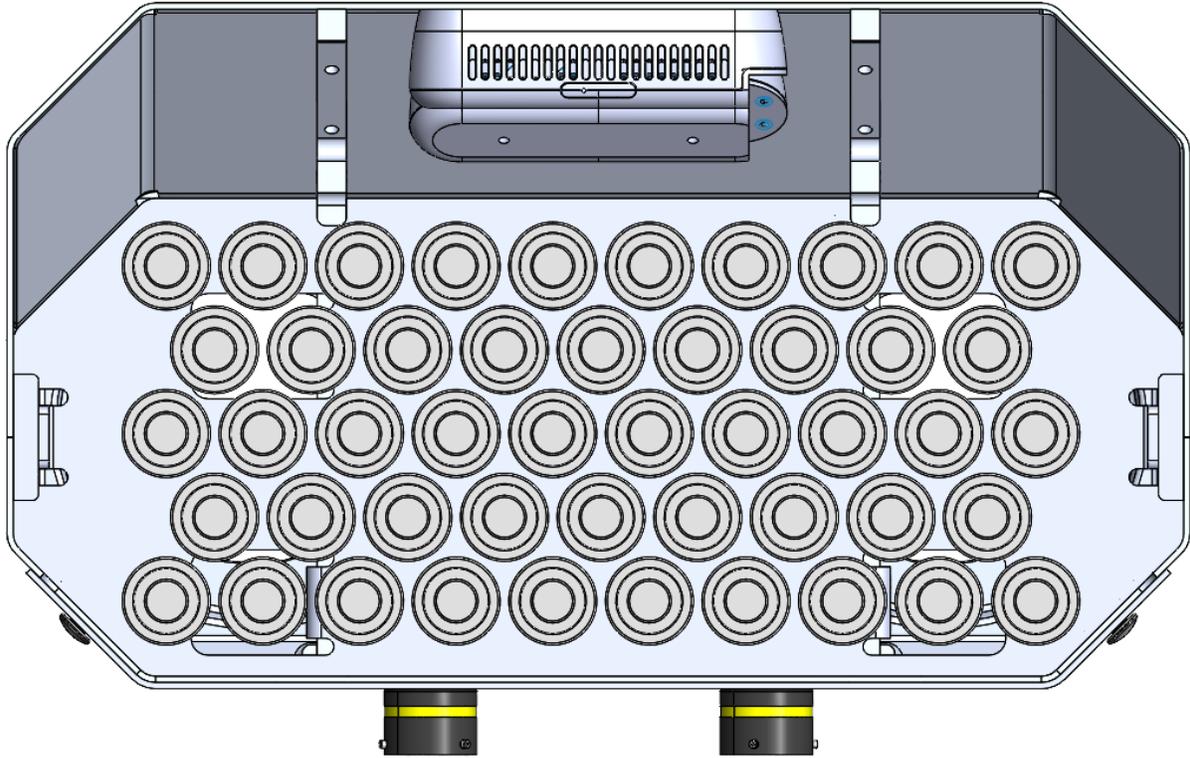


Figure 2.1: Cross Section of Humanoid Torso with 48 2170 battery cells

In order to make sure that these 48 cells will be able to produce enough power to allow the humanoid to perform high power motions like jumps and backflips, a simple battery model was made to ensure that the cell voltages would remain within their minimum and maximum voltages at different power levels. In this case, each cell is modeled using an ideal voltage source in series with a resistor which represents the DCIR (direct current internal resistance), or the apparent resistance of the cell when subjected to a DC load. The ideal voltage source is equal to the voltage the output terminals of the battery cell would be at if no current from the cell is being drawn, known as the OCV or the open circuit voltage. The output voltage of the cell may differ from the internal ideal voltage source due to losses inside the cell while current is flowing. This can be modeled by a simple series resistor representing the IR or the internal resistance of the cell causing a voltage drop in response to a current flowing either into or out of the cell. The output voltage of the cell is therefore known as the

CCV, or the closed circuit voltage. When no current is flowing, it is true that $OCV = CCV$. This model is referred to as the simple IR model.

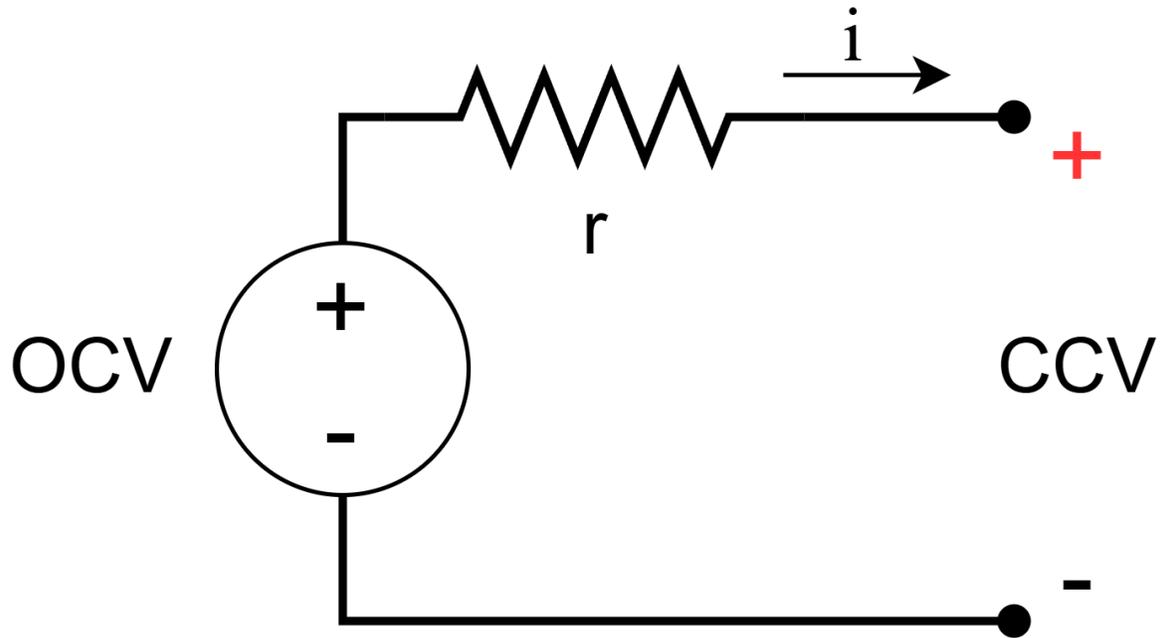


Figure 2.2: Simple IR model of a cell

Using this model, we can write an equation that represents what the output voltage of the cell will be given the electrical power either sourced or sunk into the cell. In this case, a positive power represents the cell outputting power. Given internal resistance r and open circuit voltage OCV , we can write an expression for the output voltage CCV given an input current i .

$$CCV = OCV - i \cdot r \quad (2.1)$$

Solving for i gives the following equation.

$$i = \frac{OCV - CCV}{r} \quad (2.2)$$

Knowing that output power of the cell can be written as the product of CCV and i , we

can obtain an equation that returns the battery voltage given an instantaneous power drawn from the cell.

$$P = CCV \cdot i = \frac{CCV \cdot OCV - CCV^2}{r} \quad (2.3)$$

Solving for CCV gives the final equation relating output power and battery voltage. The larger of the roots of the quadratic is chosen since it corresponds to the solution that requires less current to achieve the same output power, though the other solution is also technically valid.

$$CCV = \frac{OCV + \sqrt{OCV^2 - 4 \cdot P \cdot r}}{2} \quad (2.4)$$

This method of predicting battery voltage is a conservative estimate since any internal battery transients would only serve to uphold the battery voltage during high output powers and reduce the battery voltage during high negative output powers, or during times of high regen.

Using this equation, it is possible to predict what the voltage of a battery pack will be given expected loads or electrical power that will be drawn. Instead of populating values for OCV and r with that of a single cell, equivalent values can be used that represent an entire pack depending on that pack's cell configuration.

Battery cells can be wired in series or parallel, and a battery pack's configuration is given by the number of the number of cells in parallel, the p-count, and the number of those parallel lines of cells in series, the s-count. Thus a battery pack with three pairs of cells in parallel all assembled in series could be described as a 3S2P configuration, with the total number of cells given by the product of S and P, which is 6.

Given a battery pack configuration with a specific S and P count, equivalent values for OCV and r can be calculated that allow the simple IR model to represent an entire battery pack of identical cells. In this case, the OCV of the pack is simply the OCV of a battery

cell times how many cells are in series. The r of the pack is simply the r of a battery cell times the number of cells in series divided by the number of cells in parallel. This is because resistance increases in series and decreases in parallel.

$$OCV_{pack} = OCV_{cell} \cdot S \quad (2.5)$$

$$r_{pack} = r_{cell} \cdot S/P \quad (2.6)$$

For our battery pack model, we can use a nominal OCV of 3.6V per cell and use the DCIR from the P45B datasheet of $15m\Omega$ to build our pack model. Given that we have 48 cells, a configuration of 16S3P was chosen since it has a reasonable output voltage over the P45B's voltage limits. Mapping the P45B cell's voltage limits of [2.5V, 4.2V] to the pack, the pack would have voltage limits from [40V, 67.2V].

We can plot the function of CCV vs power drawn on a graph to see the range of possible powers that the battery pack is capable of. In this case, an extra $30m\Omega$ of resistance was added to the pack to estimate the resistance of the nickel bus bars that interconnect the cells.

In the following graph, we can see the curve representing the battery's output voltage vs output power in blue. Another curve is shown in red which estimates the bus voltage required for the MIT Humanoid to perform a motion that requires a certain peak power. This curve was derived from a worst case trajectory optimization predicting a peak power of 8kW given a bus voltage of 60V which was derived from one of the spin-jump trajectories generated in [1]. It is assumed that for smaller bus voltages, the humanoid is capable of drawing proportionally less power. This was used as a rough estimate to see if the battery would have enough voltage at higher powers to satisfy potential trajectories.

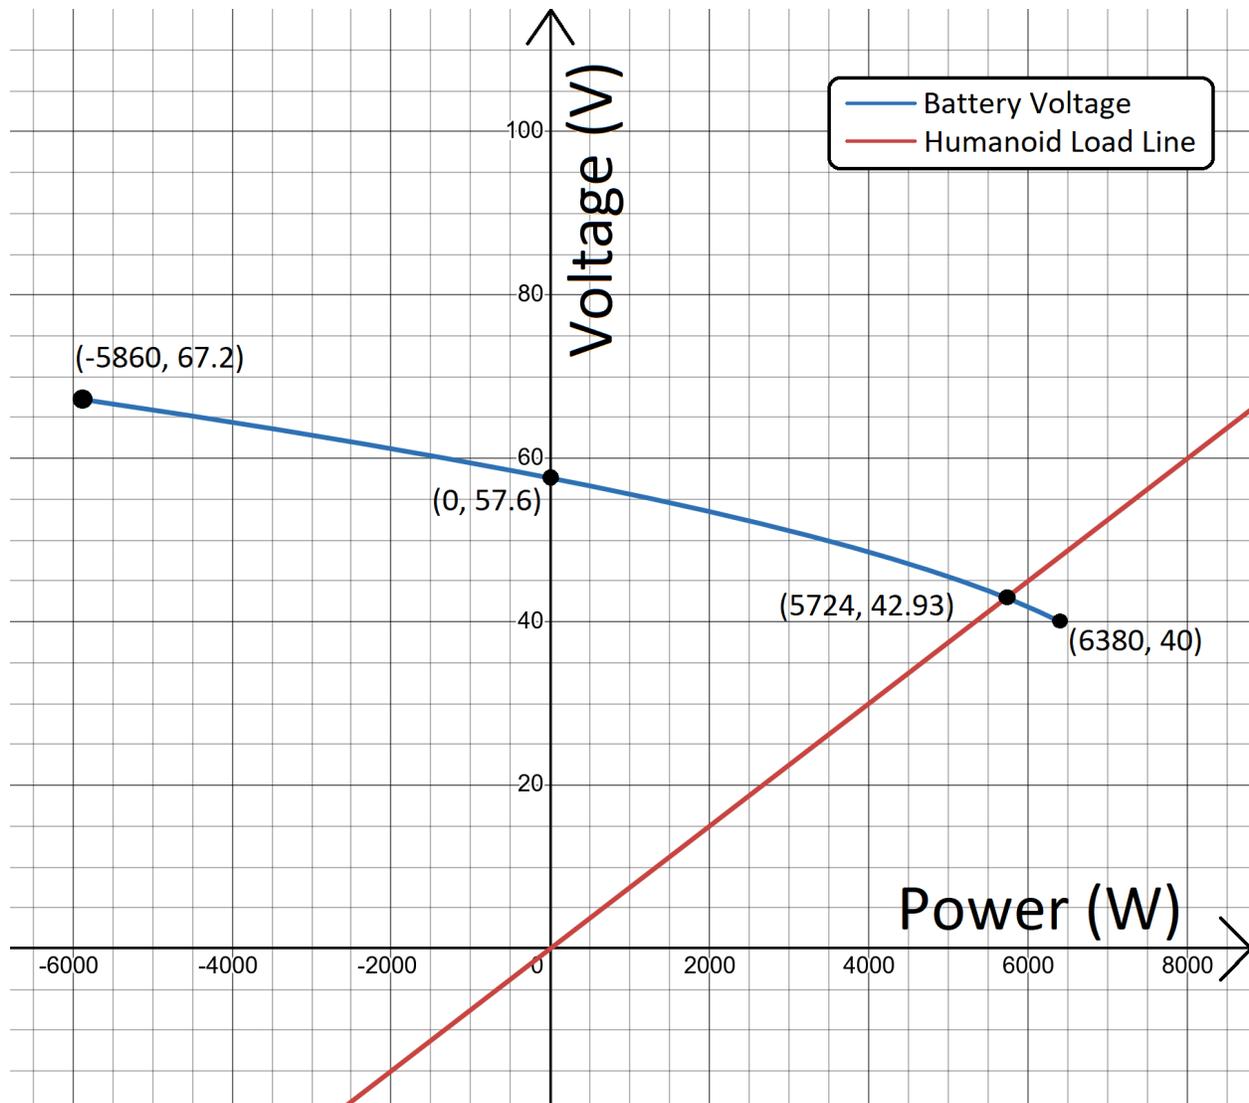


Figure 2.3: Predicted Battery Voltage vs Output Power, graphed with Desmos[3]

(a) Assuming 16S3P and cell OCV = 3.6V and cell DCIR = 15mΩ and pack resistance of 30mΩ

As you can see from the graph of the battery whose cells are charged to 3.6V, the blue curve representing the battery voltage spends most of its valid operating area from 40V to 67.2V above the red curve. For that region, it can be said that the MIT Humanoid is capable of tracking any reasonable trajectory and the actuators will not be voltage limited up until an output power of 5.7kW. To the right of this region where the red curve exceeds the blue implies that the worst case trajectory would require more voltage than the battery is capable of providing, causing actuators to potentially become torque limited at speed. This

is assuming the trajectory is generated without a battery model constraining it which can be avoided by simply including a battery model and constraining the trajectory to obey its voltage limits. At the far right of the battery voltage curve, it can be seen that the peak power capable of being sourced from the battery is 6.38kW. At 3.6V per cell, it is also clear that this battery pack will have no problems absorbing power for regen, capable of -5.8kW of regen. This is a large benefit to powering a legged robot with a battery instead of a wall connected power supply: power is free to flow back into a battery with minimal voltage rise while a power supply typically has no place to sink such large power transients and will instead cause the bus voltage to rise dramatically, potentially destroying everything connected to it. On the other hand, robotic power systems never typically see regen powers that are this large since robotic actuators can be quite inefficient in regen, especially at lower speeds.

At a mid-level charge of 3.6V per cell, this battery pack boasts a respectable peak output power of 6.38kW. This should be sufficient for most trajectories that involve highly dynamic motions such as jumps and even backflips, however if more power is required it is possible to charge the battery to a higher voltage to compensate. The downside of charging to a higher voltage is that the battery pack will not be able to absorb such high levels of power in the case of regen since the voltage is already closer to the upper limit. Due to the general inefficiency of actuators in regen and the ability of the battery to absorb -5.8kW of regen at 3.6V per cell, it is acceptable in this case to charge to a higher voltage.

Below is the same graph but with each cell charged to 4.0V per cell in the battery pack.

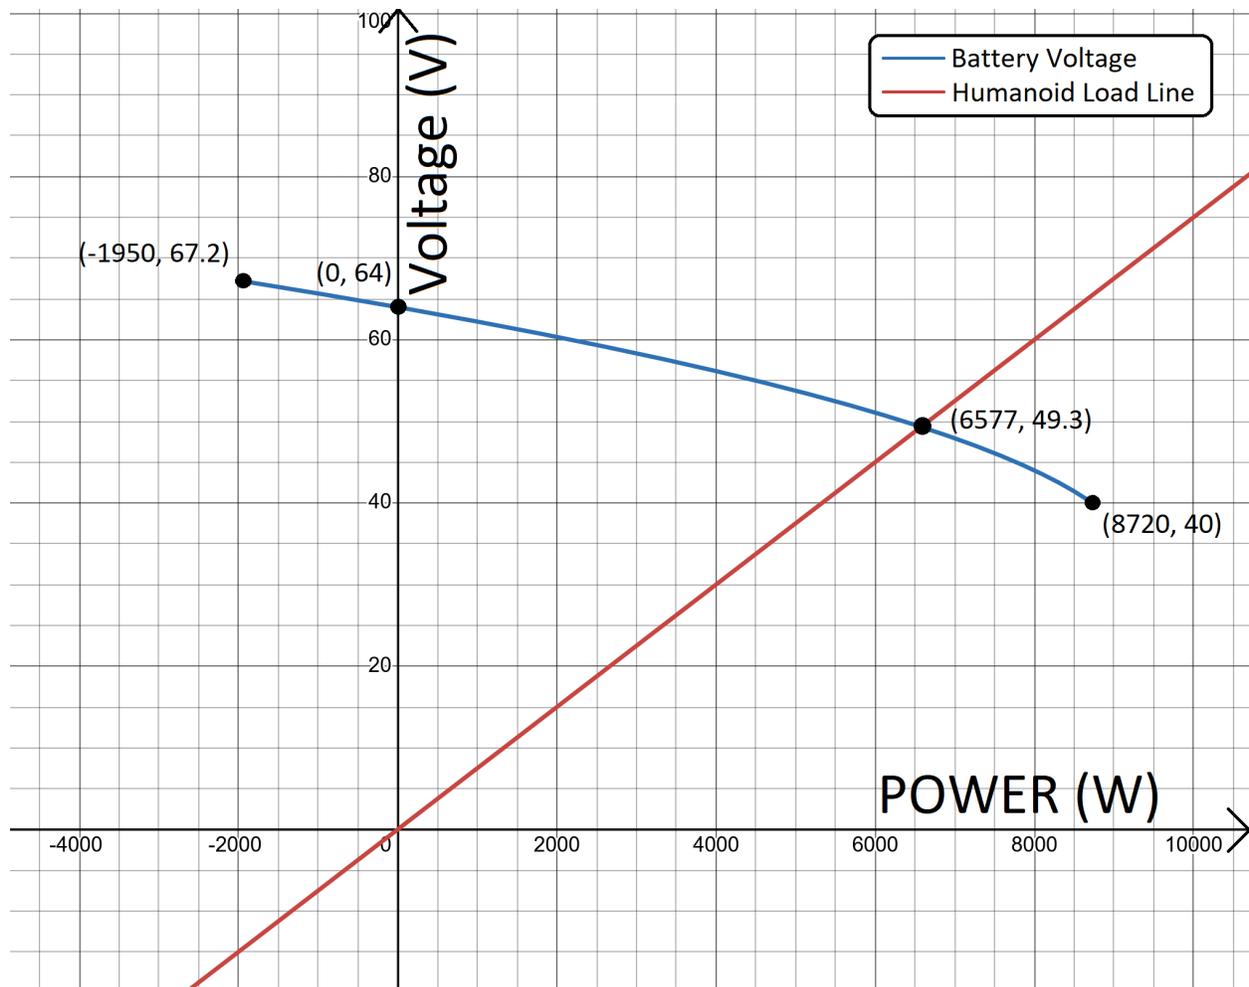


Figure 2.4: Predicted Battery Voltage vs Output Power, graphed with Desmos[3]

(a) Assuming 16S3P and cell OCV = 4.0V and cell DCIR = 15mΩ and pack resistance of 30mΩ

In this case where the cells are charged to 4.0V, the battery pack is capable of outputting 6.5kW before troubles with voltage limiting arise and can output a peak output power of 8.7kW. This is definitely more than the MIT Humanoid will need to perform highly powerful motions like backflips or spin-jumps. At 4.0V per cell, the battery is capable of around -2kW of regen which should be sufficient for any scenario given that a simple jump does not regen any power at all according to Figure 1.1b.

Given that 48 P45B cells would be capable of performing the motions that we expect to do with the MIT Humanoid over a fairly large voltage range and because 48 cells fit nicely into the cross section of the humanoid torso, it was decided that 48 P45B cells in a 16S3P

would be the best configuration for the battery pack for the MIT Humanoid.

2.1.3 Thermal Considerations

In many scenarios where a battery is being designed for an application, one of the main areas of focus is the thermal performance of a battery pack. In applications like electric vehicles, batteries are meant to be kept at an ideal golden temperature, with mechanisms to warm the battery in very cold climates and mechanisms to cool the battery in highly demanding scenarios. The application of a battery pack in the MIT Humanoid is quite different to that of electric vehicles in that assumptions can be made about the environment in which the robot will operate and the power profile that will be demanded of the battery. The MIT Humanoid will not operate in environments that are far from room temperature. The MIT Humanoid will also not draw large amounts of power for prolonged periods of time. It is true that for small periods of time, usually on the order of a few hundred milliseconds, very large amounts of power will be drawn from the battery. Due to the relatively large amount of thermal mass of each cell however, the temperature of the battery pack will not rise a significant amount following an experiment. Arbitrarily large amounts of time between experiments can also be implemented if the battery temperature does begin to rise. Internal temperature sensing will ensure that the battery can shut down or alert the researchers that temperatures are reaching unacceptable levels. It is for these reasons that this battery pack was designed without any cooling mechanisms.

2.2 Mechanical Design

A large aspect of this battery project for the MIT Humanoid was the mechanical of the battery pack itself. The ultimate goal of this design work was to maximize the peak electrical power to weight ratio, however secondary goals such as designing for good maintainability and ease of modification can be easily overlooked when designing prototype hardware in a

research setting. The necessity of "irreversible" assembly steps in the assembly process of battery packs such as the application of structural adhesives or the welding of bus bars to cells make maintenance and modification to existing battery packs very difficult. Due to the lengthy and difficult assembly process of assembling a battery pack and the need for assembly to be done in house, there is great value in designing a battery pack that can still be repaired or modified in its fully assembled state. Design decisions were made to prioritize maintainability by offloading as much complexity as possible into components that remain removable even after irreversible assembly steps are complete.

Figure 2.5 shows a final rendering of the battery pack assembly alongside Figure 2.6 which shows the final battery assembly in real life. The weight of the 48 cells comes in at 3.36kg while the weight of the battery pack itself was measured to be 3.9kg, implying that the cells themselves account for 86% of the weight of the battery pack.

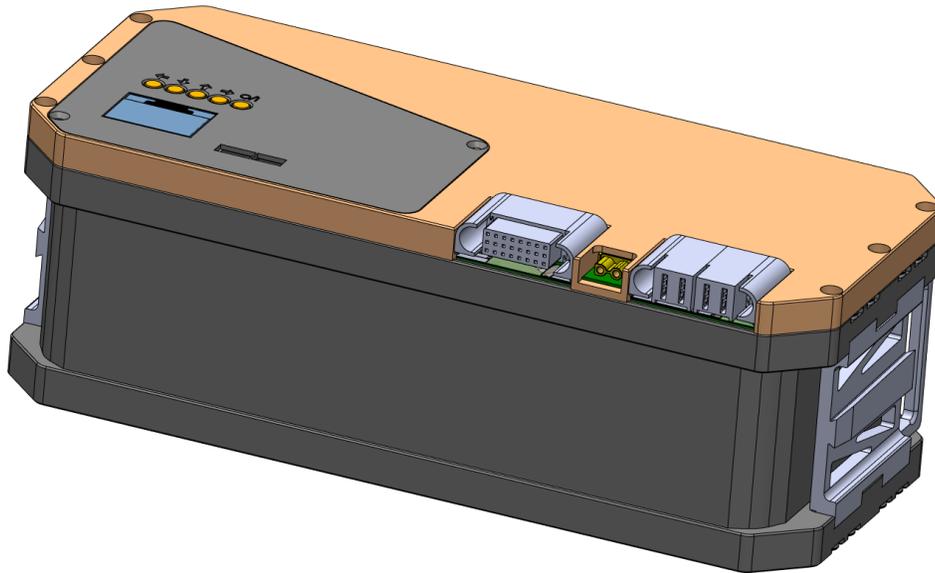


Figure 2.5: Final CAD of battery design



Figure 2.6: Final battery assembly

Figure 2.7 shows how the battery pack is inserted into the MIT Humanoid's torso from the rear.

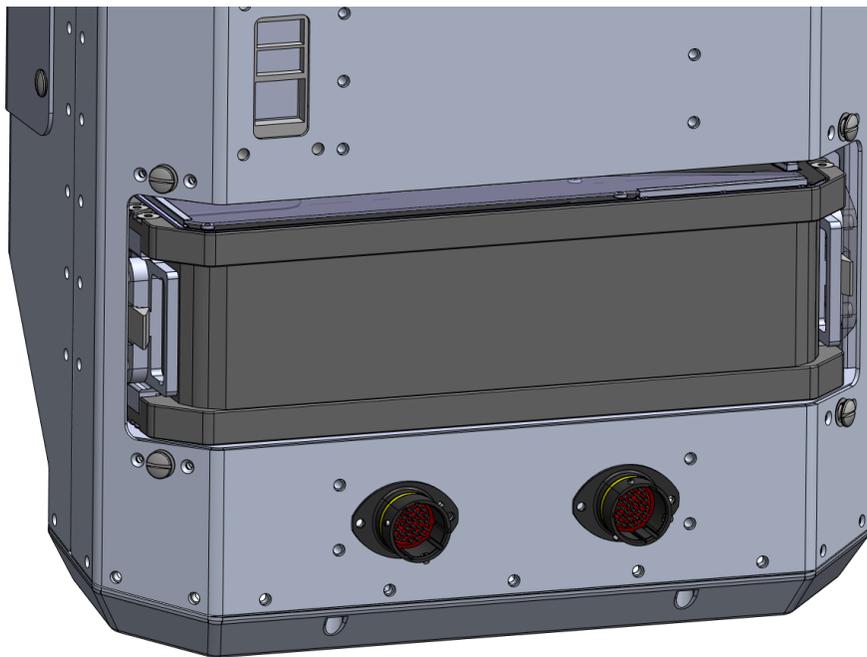
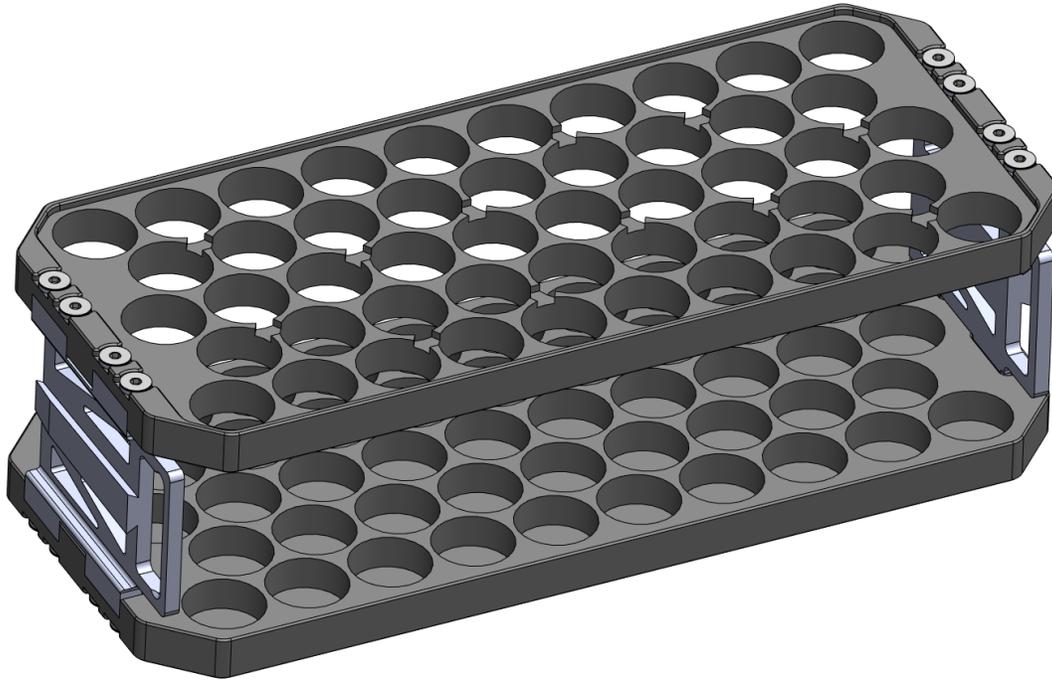


Figure 2.7: Battery pack inserted into Humanoid Torso

2.2.1 Mechanical Structure

The battery pack packages all 48 P45B cells in a vertical orientation where the positive terminal always faces up. Prior to assembly, every cell has its insulation stripped away and has its outer can cleaned with rubbing alcohol. This is to ensure that the structural epoxy used to hold the cells together makes an excellent bond to the double braced polycarbonate structure which maintains electrical isolation between the cells but also offers great strength and impact resistance. Although the adhesive ensures a very strong and rigid structure, this is the first example of an irreversible assembly step. Any mistakes made during design or manufacturing of the final battery pack that prevent this cell block from functioning correctly necessitates the disposal of all 48 cells. Figure 2.8 shows this double braced polycarbonate structure with and without the cells inserted. Aluminum spacers which mate with the ends of the polycarbonate structure maintain spacing while the adhesive cures as well as bearing all the inertial loads of the battery during acceleration.



(a) Without cells



(b) With cells inserted

Figure 2.8: Double braced polycarbonate structure that is bonded to cell surfaces to form main battery structure

Each side of the polycarbonate structure uses a spacer made of aluminum to minimize weight. Each spacer has a small handle for removal as well as a male tapered dovetail to help

guide the battery pack while being inserted into the torso while also reacting all the inertial loads of the pack during accelerations. The spacer makes contact with the polycarbonate sheets on precisely machined surfaces to ensure that loads are transferred through the contact of these large surfaces.

Each spacer mates with an aluminum piece with a corresponding female tapered dovetail. Both pieces are anodized to prevent galling of the surfaces when in contact. Not depicted in the cad model of the female tapered dovetail is a backstop to ensure that an inserted battery does not get stuck when inserted with significant momentum.

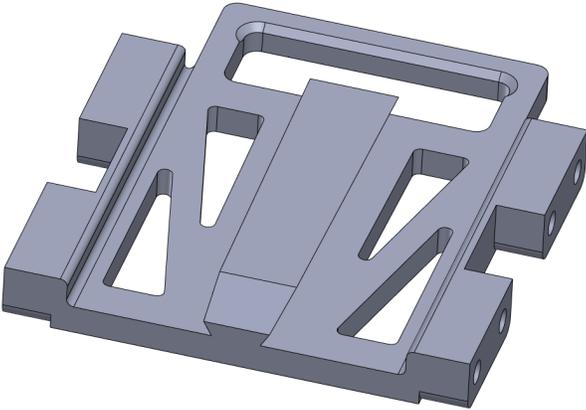


Figure 2.9: Aluminum spacer with male tapered dovetail

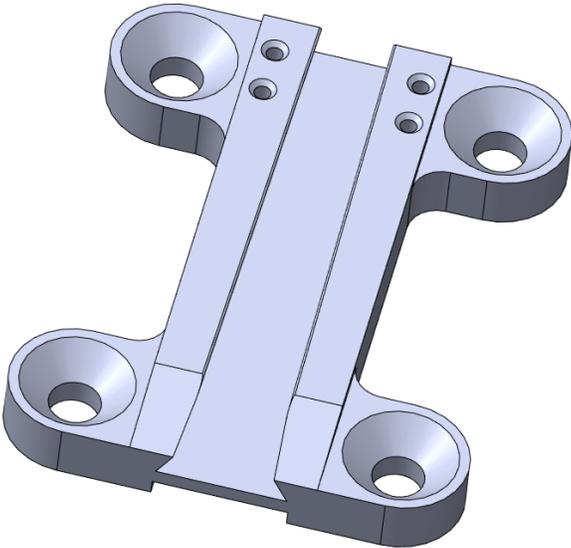


Figure 2.10: Aluminum piece with female tapered dovetail

The female tapered dovetail piece also has four tapped bolt holes for mounting a flexure retention mechanism that keeps the battery pack dovetails in full contact. The female tapered dovetail piece has a recession machined in the backside to ensure that the flexure remains flush when not in bending. The front tapered edge of the male tapered dovetail pushes the flexure open when being inserted. Manual flexing of the flexure is necessary for battery pack removal.

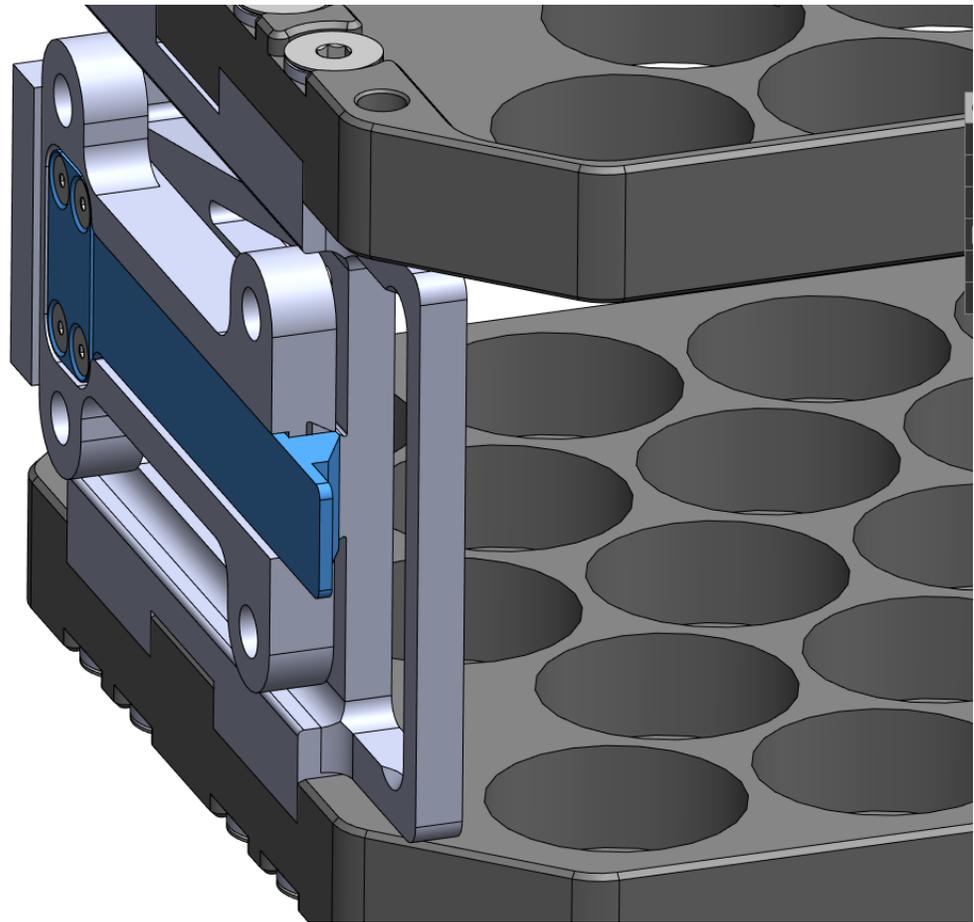


Figure 2.11: Flexure bar highlighted in blue

The female dovetail piece connects to each side of the robot using a compliant interface. The idea behind the compliant interface is that very large accelerations of the torso will cause a small deformation of the compliant interface to avoid shock loading the battery pack structure. Shock loading can very easily occur during a collision of the robot torso with objects or the ground and could potentially cause the polycarbonate sheets of the main

structure to crack. Because the magnitude of these loads is unknown, an estimate of a worst case acceleration was set to be 50G. Under this acceleration, it was chosen somewhat arbitrarily that the compliant interface in pure shear should experience a shear strain of 0.1% to ensure proper energy absorption. An ABS material was chosen for this piece given the desired load, strain, and geometry.

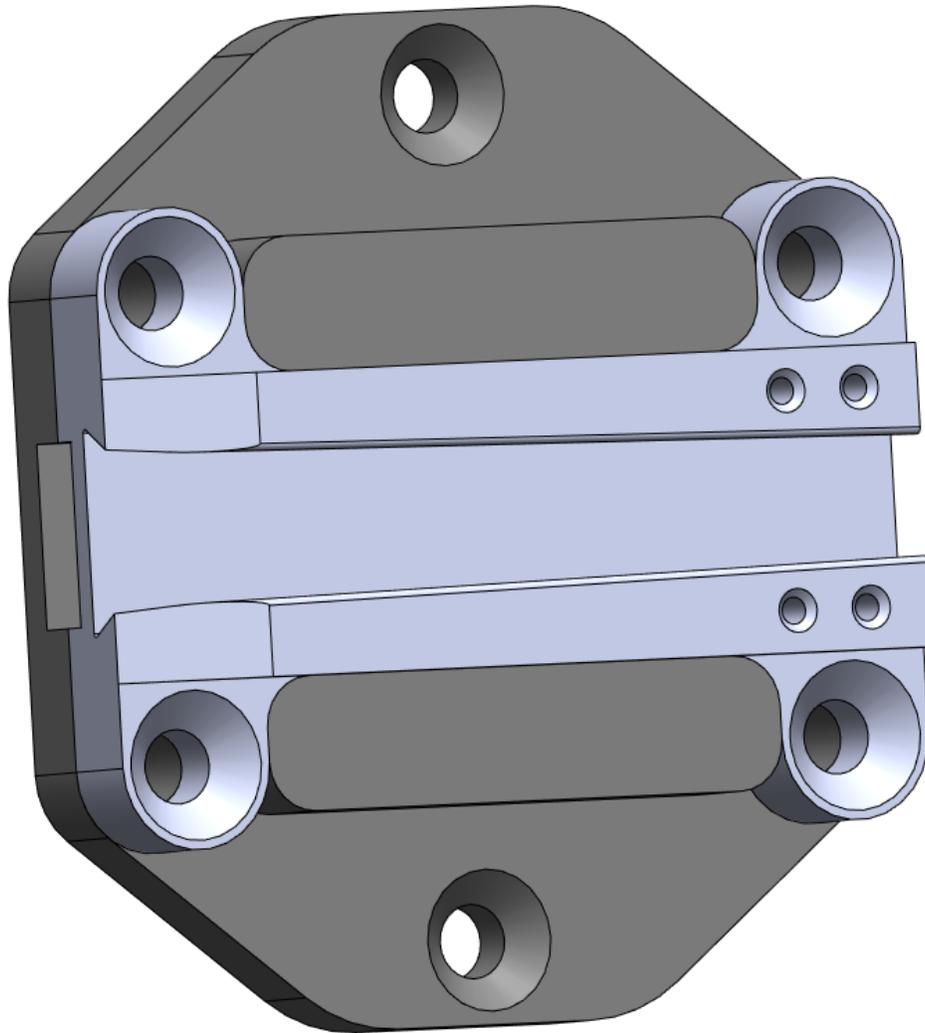


Figure 2.12: Compliant interface mating with female tapered dovetail through large surface contact.

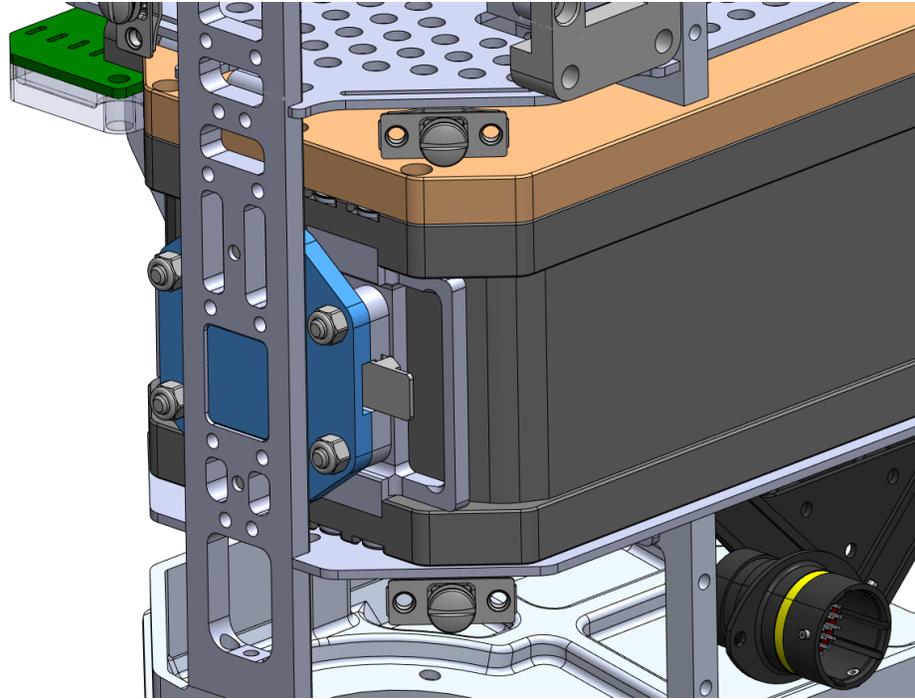


Figure 2.13: Compliant interface highlighted in blue mating with internal side strut of the humanoid using a large boss feature and bolts for retention.

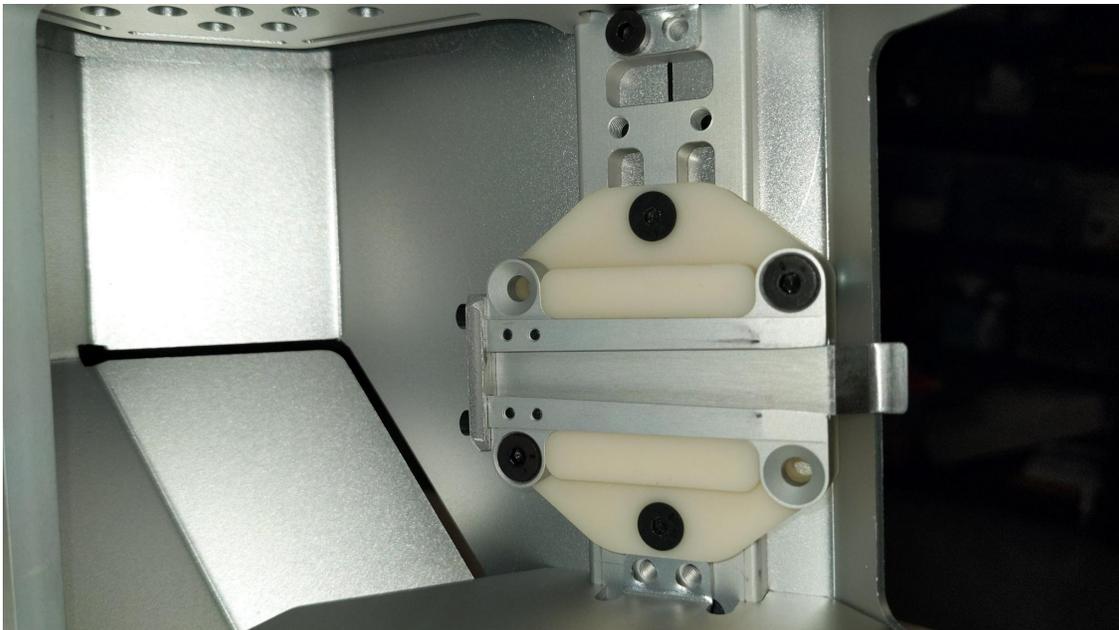


Figure 2.14: Compliant interface with female tapered dovetail and flexure retention mechanism installed in humanoid torso.

When the rear battery panel which covers the battery is installed on the MIT Humanoid,

the flexure mechanism is prevented from opening due to it interfering with the rear panel. The rear panel is installed using captive quarter turn wing nuts to ensure that the entire process of installing and removing the battery pack requires no tooling. This ensures the battery remains positively locked while installed inside the humanoid.

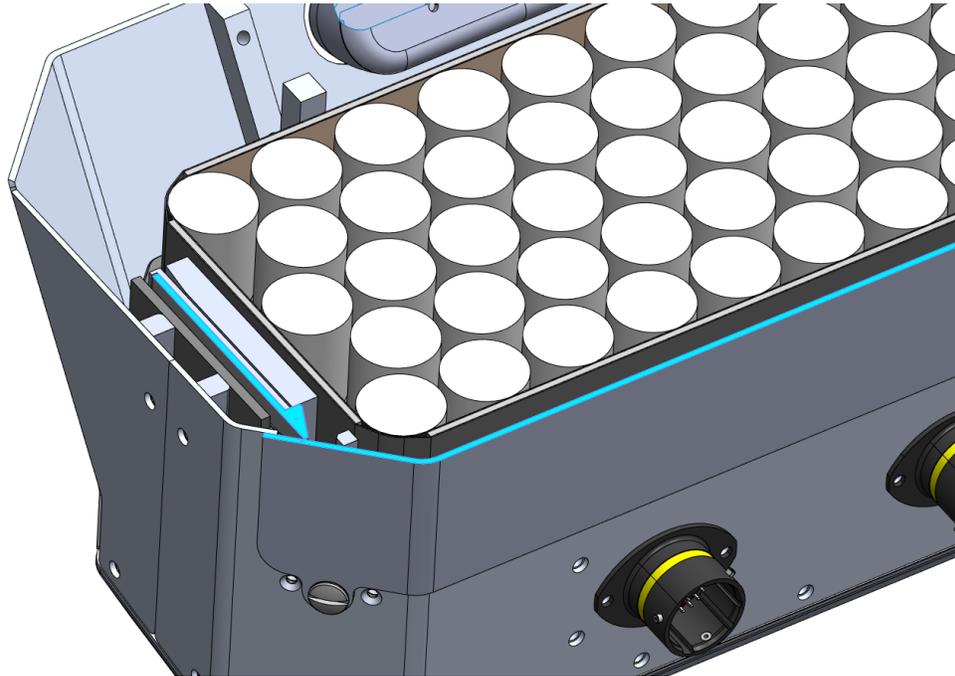


Figure 2.15: The flexure mechanism is unable to open due to the installation of the rear battery panel.

Finally, the inserted battery makes all relevant electrical connections when inserted due to the presence of the BOB, or the break out board PCB. The BOB breaks out both the high power motor output as well as all other signals like computer power, wakeup lines, and isolated communications.

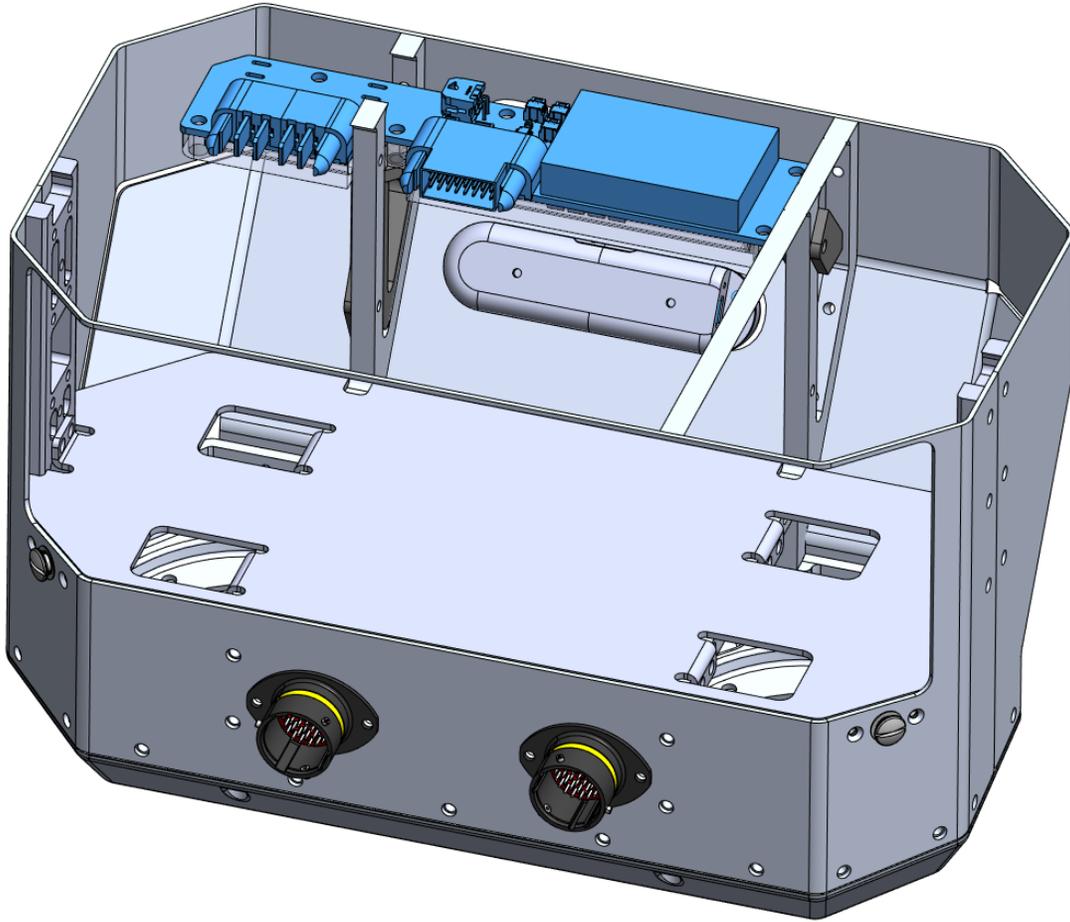


Figure 2.16: The BOB highlighted in blue makes all electrical connections to the battery when it is inserted.

2.2.2 Electrical Interconnect

The question of how to interconnect the cells into a 16S3P configuration is non trivial in the case of this battery pack. With all cells standing vertically and in the same orientation, connections must be made to the anode in the center of the cell and the cathode along the outer edge of the can. Typically, the negative edge of one cell's can can be electrically connected to the positive terminal of the cell adjacent to it in a repeating pattern until the entire pack is electrically connected together. Due to the pack having a geometric width of 5 cells instead of 3 along the length of the cell block, a more creative way than just accumulating voltage along a single direction is needed. There are many ways to solve the problem of how to accumulate voltage in a pack that is not regularly shaped, but the solution that minimizes the number of unique metal interconnect pieces needed should be preferred to keep part count and overall complexity low. Figure 2.17 shows an example of how cells are interconnected with a thin sheet of nickel sheet metal. Each connection to a cell's anode can also be individually fused if the width of the connection is properly tuned.

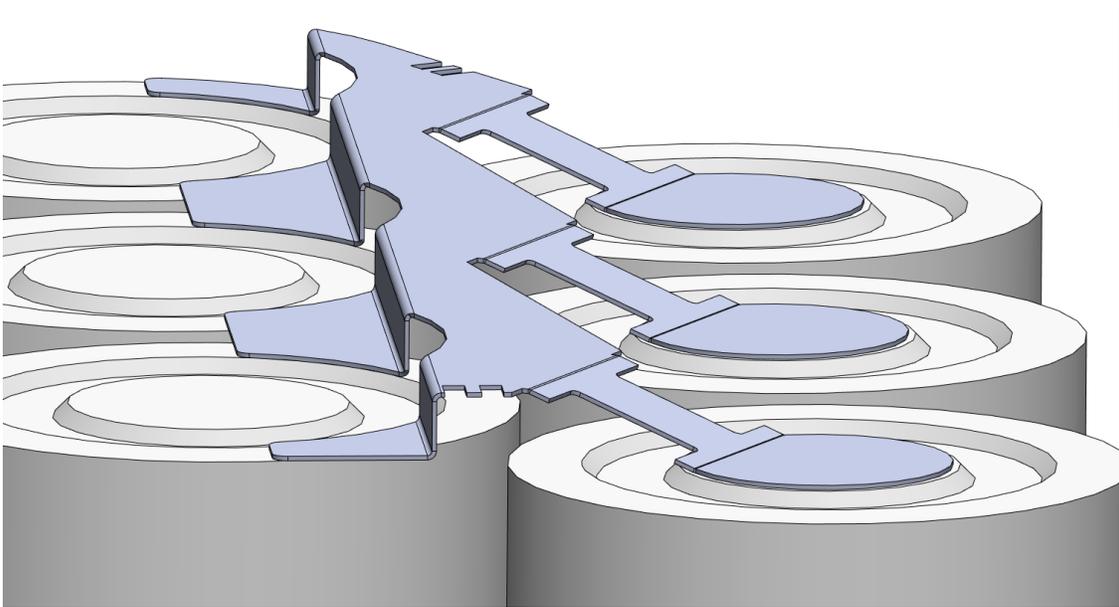


Figure 2.17: Example nickel interconnect connecting six cells into a 2S3P configuration.

For the MIT Humanoid’s battery pack, the routing solution is shown below. Annotations show where current concentrations occur and how an interconnect was routed back overtop other connections. Overall, this solution kept the unique part count low at 7 unique parts if mirrored parts are considered redundant. Mirrored parts are ordered as the same piece of sheet metal from our manufacturer but are simply bent in the opposite direction to make a mirror part.

The red circles show where conductor reinforcement with copper bars were necessary. The green outline shows where the current path is folded back over the top of the pack in order to complete routing.

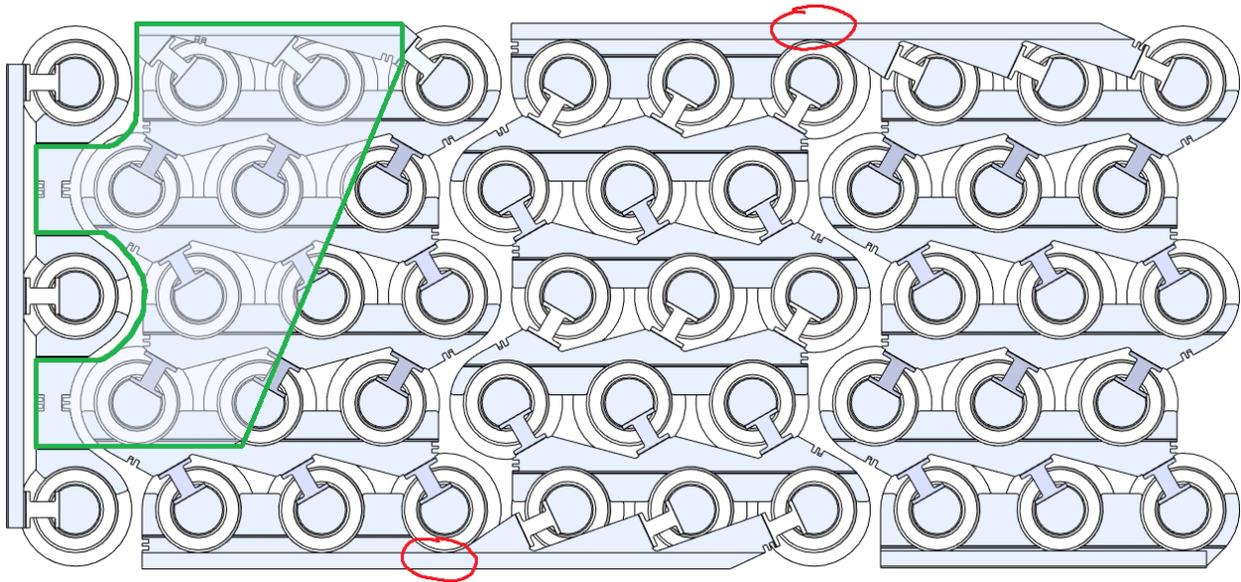


Figure 2.18: Final routing solution for pack.

Below are all the unique nickel sheet metal parts including their mirror versions, excluding the piece that wraps back over top.



Figure 2.19: Unique nickel interconnects excluding piece that wraps back overtop.

Before the nickel interconnect can be inserted, the carrier PCB must first be attached to the top of the cell block. The carrier PCB serves a couple of purposes, one of them being the alignment of these nickel interconnects overtop the cells. The carrier PCB has many cutouts precisely beneath where the nickel interconnects bend down and make connections to the cells beneath. This ensures that the nickel interconnects make electrical connection only in the areas in which they were designed to touch. The carrier PCB is fixed to the tops of the cells using a sheet of laser cut VHB double sided tape which also has the same cutouts as the carrier PCB to allow the nickel interconnects to bend down and make connections to the cells beneath. The carrier also provides voltage sensing by allowing each nickel interconnect to be soldered to pads on the PCB as well as providing temperature sensing by distributing temperature sensors on the underside.



Figure 2.20: Carrier PCB installed on module.



Figure 2.21: Nickel interconnect placed and soldered on carrier PCB

Each nickel interconnect is soldered down to the PCB using features called "warts". These small tabs are meant to be able to deform under vibration or thermal loading to ensure that

the pads that they are soldered to are not ripped off of the PCB.

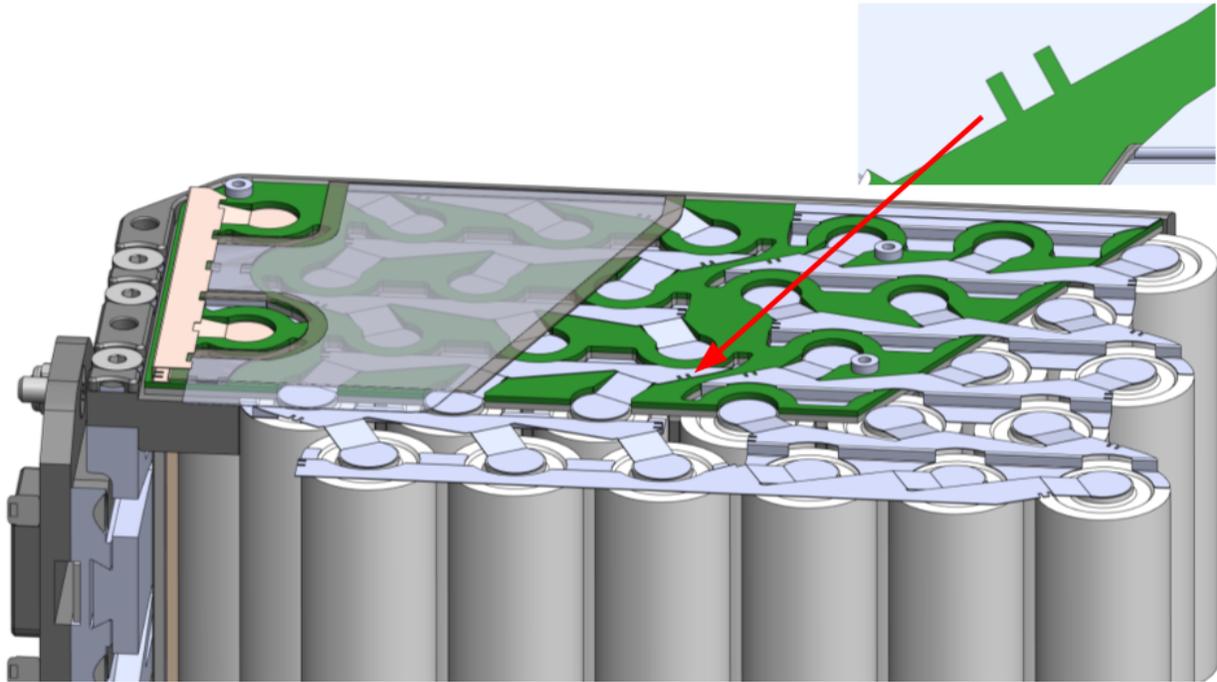


Figure 2.22: Nickel warts that are soldered to the carrier PCB.

The next step in the assembly process is to spot weld every nickel interconnect to the battery cells below. This is done manually using a spot welding station. Great care must be made to not damage the battery pack during this operation, as it is the most likely step to cause permanent damage to the pack. Spot welding is also another example of an irreversible assembly step in this process, as each thin nickel sheet holds down the carrier PCB to the cells while also being impossible to remove without scarring the surfaces of the cells where the interconnect is meant to be welded to.



Figure 2.23: Each interconnect is spot welded by hand.

After spot welding is completed, a protective layer of Kapton sheet is applied over the nickel interconnects to ensure nothing is capable of shorting the battery internally. This also prevents the piece that folds over from shorting to other interconnects internally.



Figure 2.24: A layer of Kapton insulates the spot welded interconnects.



Figure 2.25: The final interconnect that folds over is soldered to the exposed interconnects and is insulated over.

2.3 Electrical Design

In this battery pack design, there are four PCBS in total that make up the electrical system of the battery. Each PCB serves a different purpose in the monitoring and protection of the battery. The electrical system is capable of monitoring individual cell voltages, pack voltage, pack current, and cell temperatures. The electrical system is capable of individually disabling the motors but not the main computer, disabling all power including to the main computer, and blowing an integrated fuse inside the battery if certain programmable conditions are met to permanently disable the battery. The electrical system can also be put into a sleep mode that draws only $30\mu A$ which would drain a fully charged battery in 51 years if the cells themselves could not self discharge.

The battery pack itself is home to three of four total PCBS, the carrier PCB, the power board, and the segment. A fourth PCB called the BOB, or the breakout board, simply exists inside the humanoid torso and waits for a connection to be made. The power board and segment are bolted to the surface of the carrier using standoffs. The segment is a rigid-flex PCB allowing the power board and the segment to make many high speed and low impedance connections to allow them to function as a single lumped circuit.

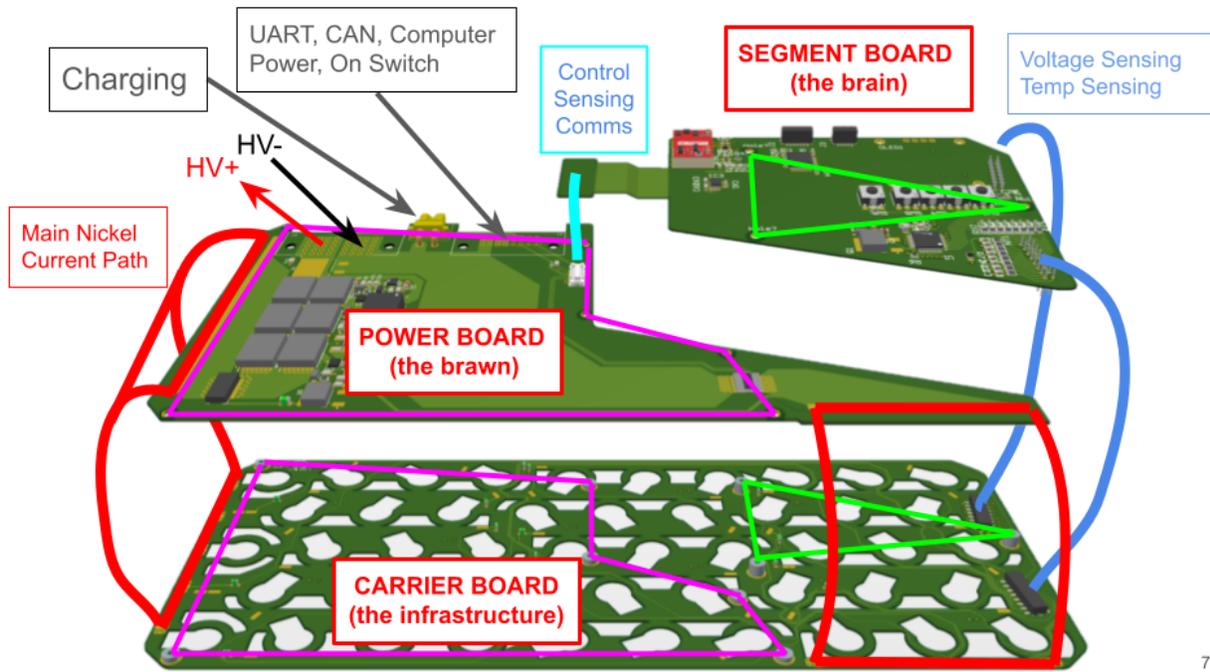


Figure 2.26: The three main PCBs in the battery pack.

2.3.1 The Carrier PCB

The carrier PCB is integral in the assembly process when interconnecting the cells. The carrier PCB also makes electrical connections to every nickel interconnect to pass back to another PCB, the segment, for cell monitoring and balancing. The carrier PCB also has 12 distributed temperature sensors on its underside, allowing most of the cells to be in contact with a temperature sensor. These temperature signals are also fed back to the segment for monitoring.

2.3.2 The Power Board

The power board is home to the high current path of the battery pack. The two main nickel tabs that protrude from the battery interconnect beneath fold over and are soldered to the power board. The power board has the main fuse, all switches for enabling and disabling power to the motors and computer, a free wheeling diode to allow for fast turn off and

give any inductive energy a place to dissipate, precharge for both the computer and motors, localized high speed gate drive for the motor switch, and a current sensing shunt resistor. The power board also has footprints for both the connectors that mate with the BOB as well as a third xt30 connector for external charging. All other complexity was offset to the segment due to the ease of removal. Any broken components on the power board would be painful to fix because the main nickel path would have to be unsoldered and folded back for the power board to be removed.

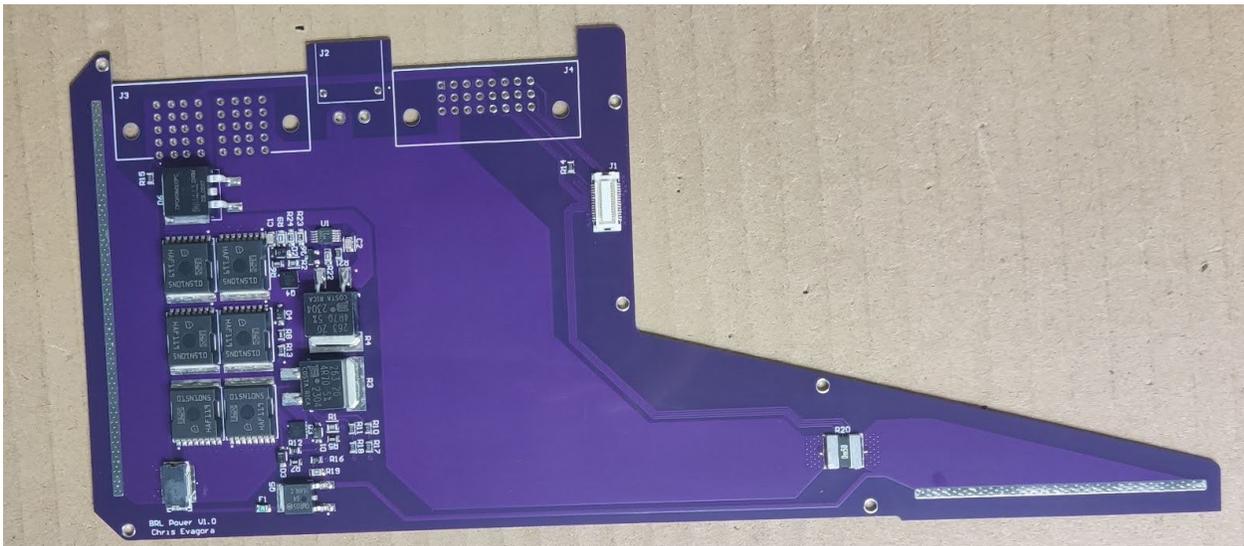


Figure 2.27: The power board.

2.3.3 The Segment

The segment is the main PCB in the electrical system. It is the easiest of all the PCBs to remove and so most of the complexity was offset to this board. The segment has the main battery management system IC, or BMS IC, as well as a microcontroller that stores the configuration for the BMS IC. The IC handles cell monitoring and balancing, current sens amplification, and gate drive for the charge and discharge switches on the power board among other protective features. The segment has an on board OLED screen as well as some buttons for the user to be able to interact with the system and view the state of the battery or any errors present. An on board buzzer is capable of communicating to researchers that a

battery error has occurred if the battery is inserted into the torso and sealed with the battery panel cover and cannot be directly seen.

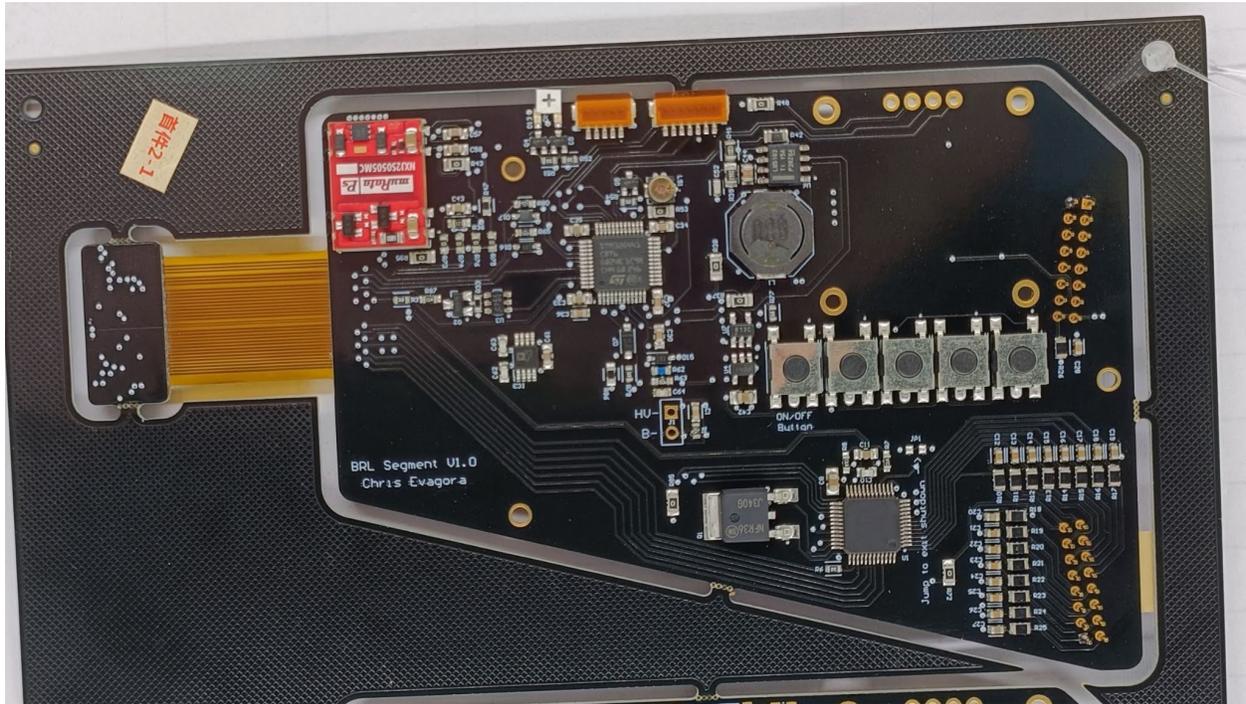


Figure 2.28: The segment.

2.3.4 The BOB

The last of the four PCBS is the BOB, or breakout board, which serves to branch out relevant signals like motor power, computer power, and isolated comms from the battery when it is inserted into the torso of the humanoid. This PCB does not exist inside of the battery pack, rather it is mounted rigidly inside the torso of the humanoid and waits for a connection to be made.

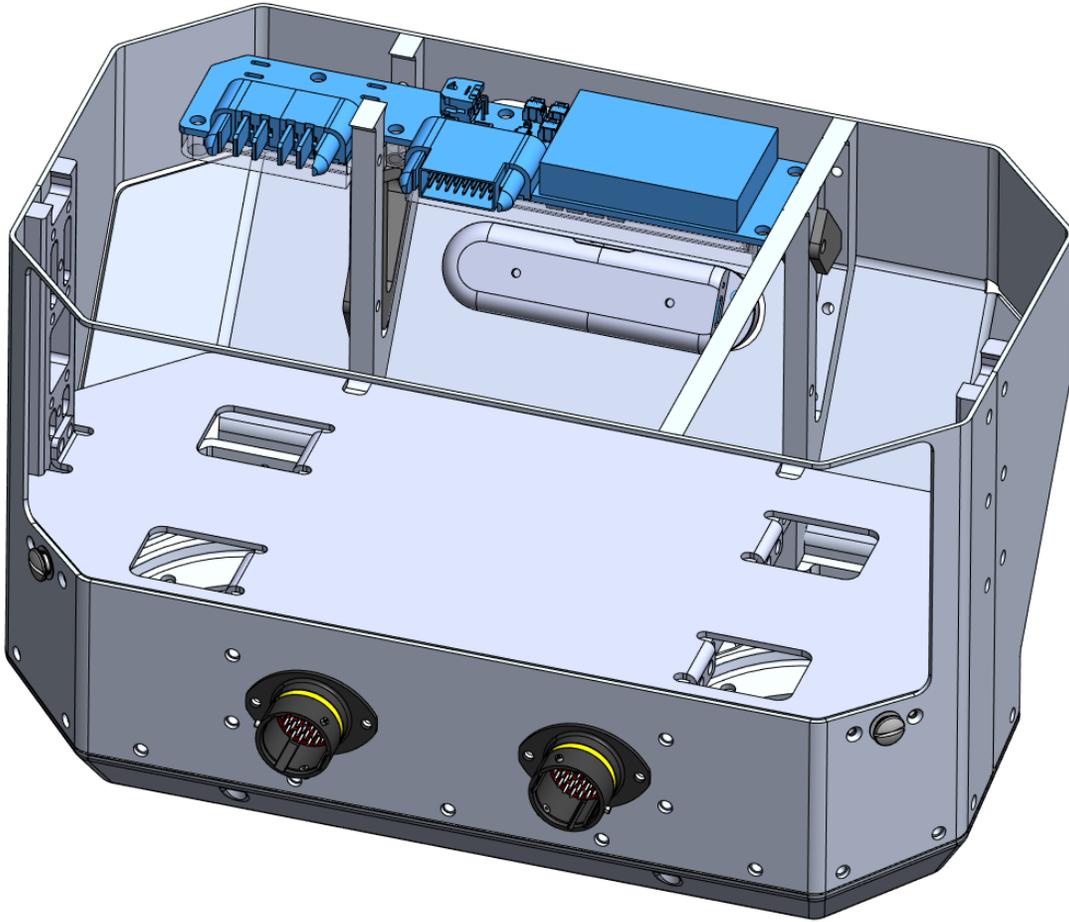


Figure 2.29: The BOB highlighted in blue makes all electrical connections to the battery when it is inserted.

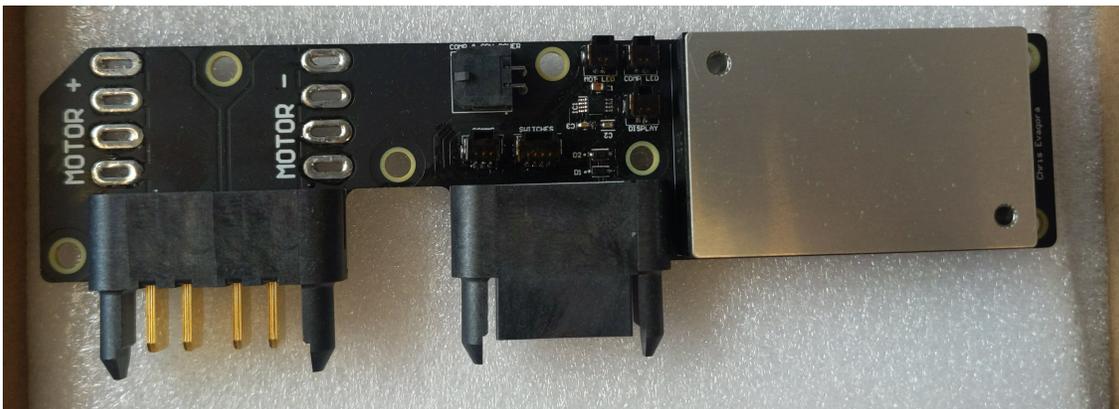


Figure 2.30: The BOB fully assembled.

2.3.5 Electrical Assembly

Below you can see the power board and segment installed inside a test battery. The segment can be debugged remotely over external cabling to ease development of the firmware.

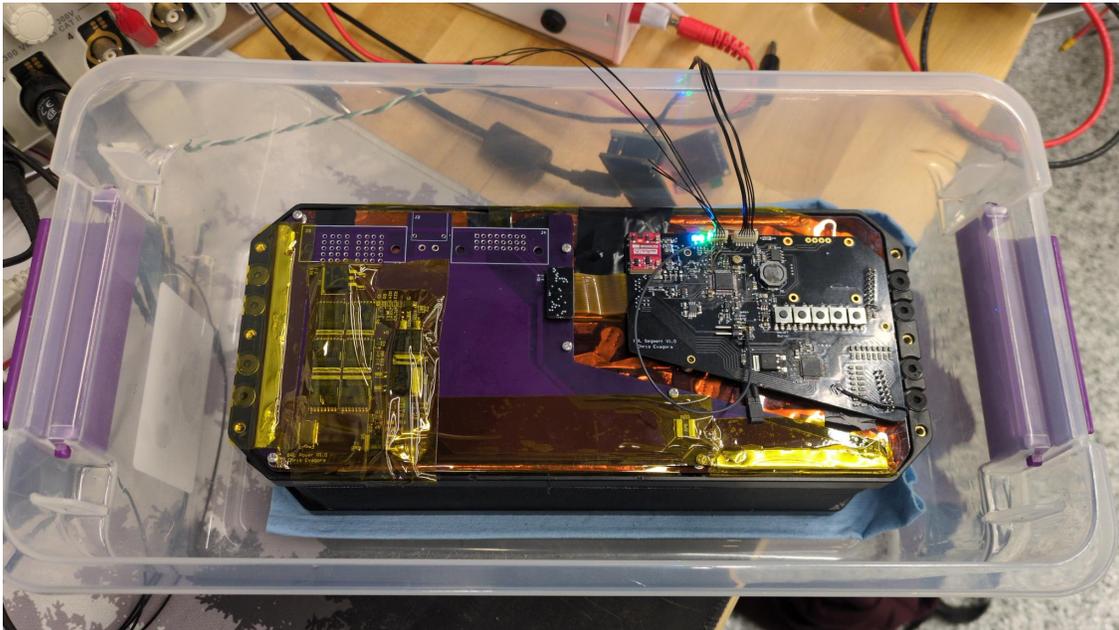


Figure 2.31: The power board and segment installed in a test battery.

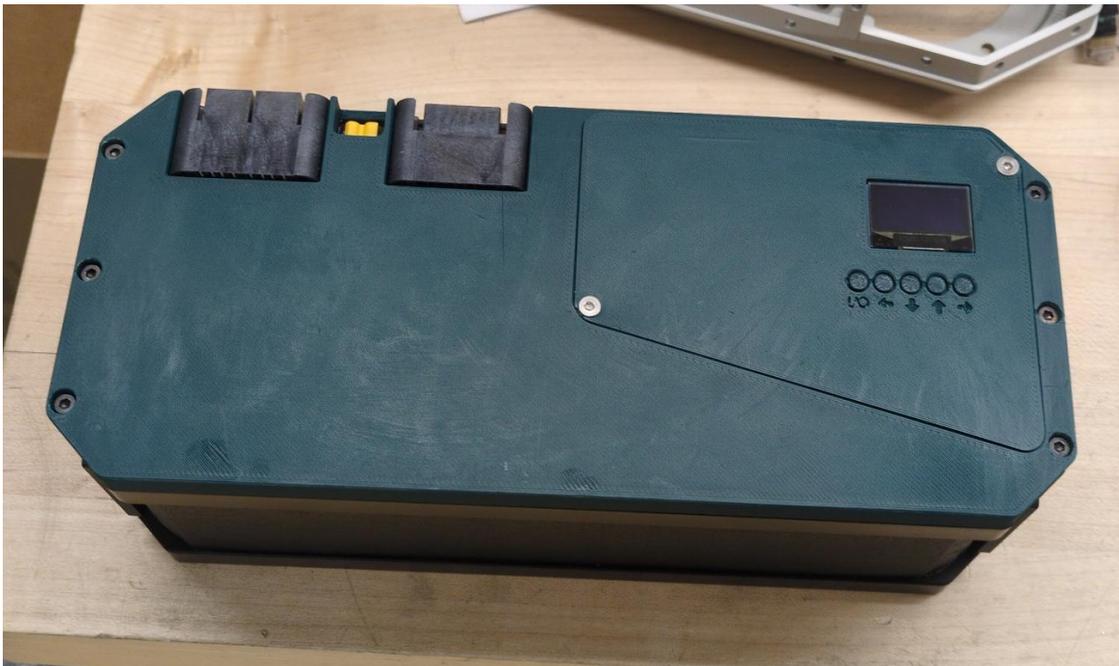


Figure 2.32: The electrical assembly is covered by a 3D printed cover.

Chapter 3

Modeling Batteries

Being able model a battery pack well for use in legged robotics can be very beneficial. Knowing that batteries need to operate within a certain output voltage range in order to prevent significant wear and damage to the internals of the battery cell, great care must be taken to not exceed the output power capabilities of a battery cell. In the case of lithium ion batteries, although each cell is different and will have a safe operating voltage range specified within its datasheet, most lithium ion battery cells adopt a safe operating range of 2.5V to 4.2V by virtue of lithium ion chemistry. It is therefore important that whatever load a battery cell is subjected to does not cause an instantaneous voltage drop or increase that causes its output voltage to leave this operating range. Thus, the only way to ensure that a battery cell has its voltage limits respected is to modulate the load that the battery cell is subjected to. If the battery itself can be modeled well enough that a load profile can be generated such that the battery's output voltage is at or near its limit of 2.5V for the entire load profile, it can be said that the absolute limits of the battery's output power was closely followed for the duration of the load profile. This would maximize energy output during that duration without violating the battery's output voltage limits. The analogy to legged robots is that a trajectory optimization, aware of a battery model and constraints on what the output voltage of that battery are, would be capable of generating a trajectory that both ensures that the

battery does not leave its safe operating range but also maximizes an objective such as jump height or task speed. The objective of this research is therefore to build a battery model that most accurately predicts what the output voltage of a battery will be given a power profile.

3.1 Modeling OCV vs SOC

It is important to be able to predict what the battery pack's voltage will be under no load depending on the battery pack's SOC, or state of charge. The SOC can be represented as a percentage of how much charge is stored in the battery divided by the total amount of charge the battery can store. In this paper, the SOC will simply refer to the amount of charge in coulombs that is stored in the battery. To measure what the OCV, or open circuit voltage, of the battery pack will be depending on its SOC, tests on an individual cell can be done and scaled by the S count of the battery pack.

A test bench was set up to measure the OCV vs SOC curve of the P45B cell. The setup included a DC load to discharge the cell, a DC supply to recharge the cell back to full voltage, and a water cooling loop that maintained the temperature of the outside surface of the battery cell. The water cooling loop used an aquarium chiller to continuously remove energy from the water supply while a sous vide maintained a temperature of 25C within 0.1C.

The cell was inserted into a cell jig which included a kelvin connection to pass high currents through the cell while not causing voltage drops across the voltage measuring probes. The jig could be folded over and locked with consistent pressure using a thermal gap pad to ensure that the cell made a good thermal connection to each of the two water cooling blocks.

The tests were orchestrated using a raspberry pi communicating with the DC load and supply to ensure that test conditions and state transition thresholds were consistent between tests. The cell was charged to 4.2V and would begin a discharge once the charge current reached 100mA.

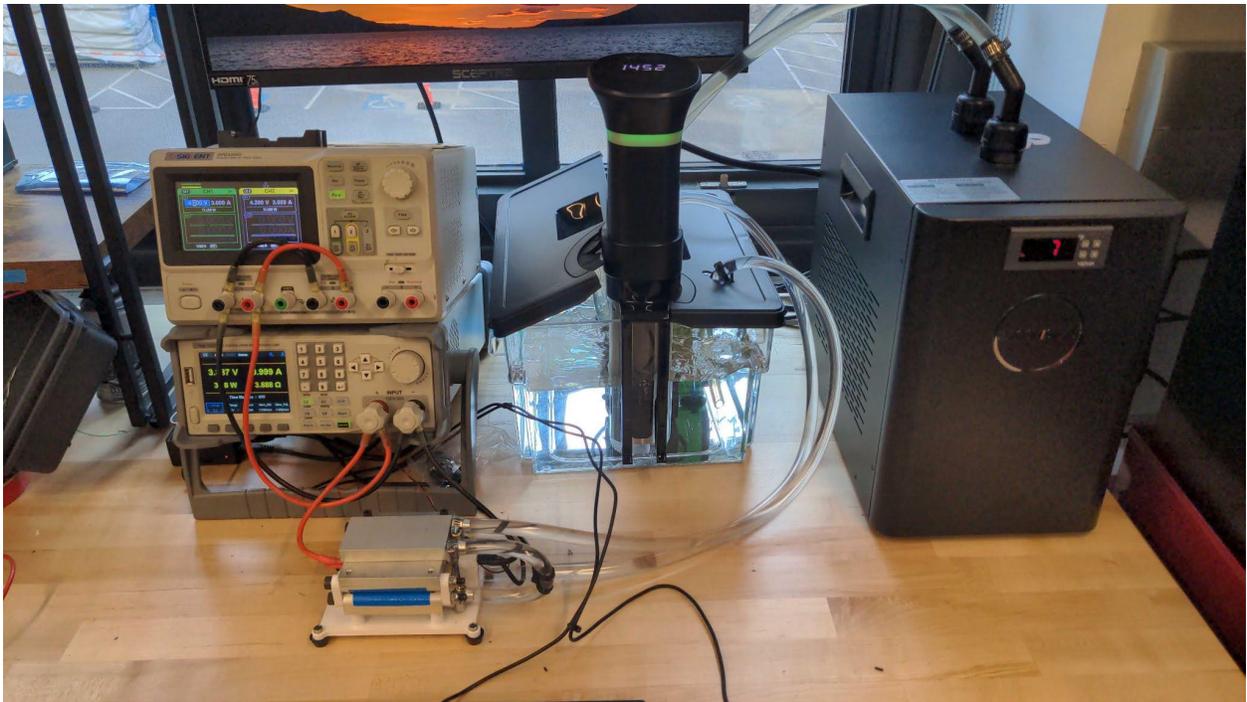


Figure 3.1: Cell testing setup consisting of a DC load, DC supply, and water cooling loop to maintain temperature.

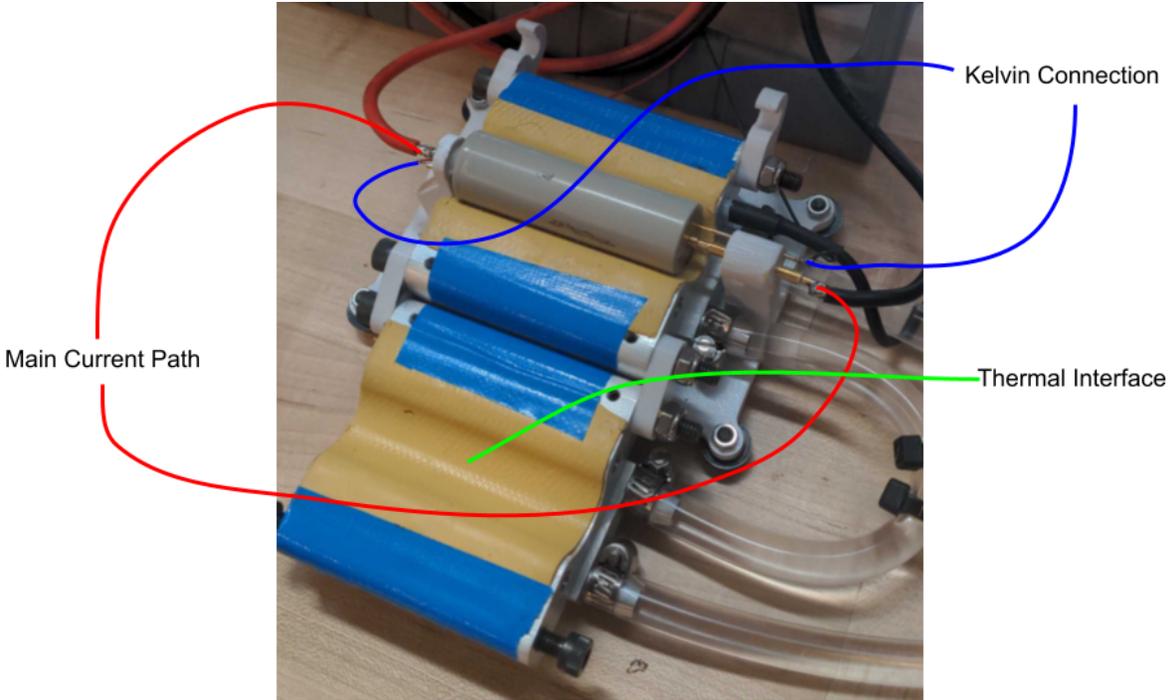


Figure 3.2: Cell jig which includes a kelvin connection as well as two water cooling blocks.

The OSC vs SOC curve of the P45B could be measured using this setup. Because current must be drawn from the battery cell while it is discharging, the voltage measured at the battery terminals is the sum of both the internal OCV as well as the resistive drop due to that current. For the data collected, a relatively small current of 1A was drawn to discharge the cell. That curve is shown in the figure below in blue. To get as close as possible to the true OCV vs SOC, we can compensate for the DCIR, or DC internal resistance, of the cell by adding the cell's DCIR times the discharge current at every point on the graph. Assuming the cell's DCIR is $15m\Omega$, we can produce a more correct OCV vs SOC curve shown in orange.

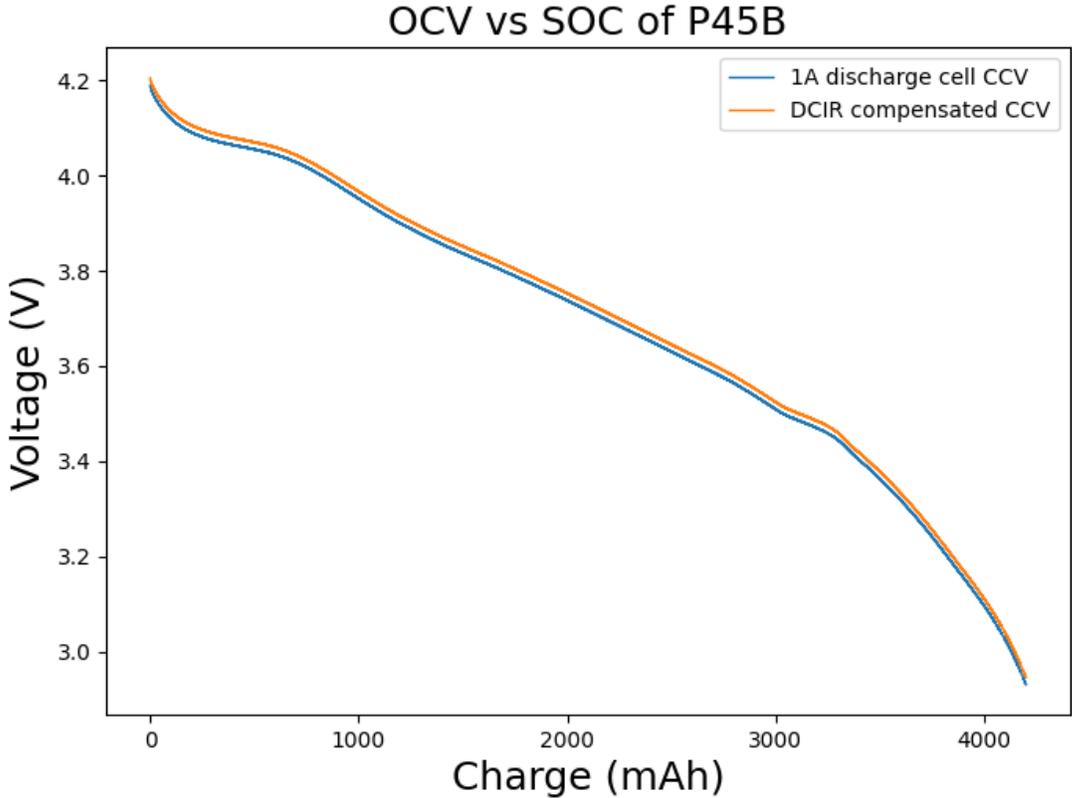


Figure 3.3: Discharge curve showing 1A constant discharge as well as a DCIR compensated OCV vs SOC curve.

Using this curve, the battery voltage can be known by implementing a lookup table. If need be, a polynomial approximation of this curve can also be used if storing the data is a problem or if differentiability is required.

3.2 Modeling Cell Impedance

There are many ways to model battery impedance, the easiest to analyze being those that use an equivalent circuit model that represents the battery's I-V characteristics. An equivalent circuit model is also easy to interpret and simulate.

3.2.1 The Ideal Voltage Source

It is often easy and convenient to model batteries as an ideal voltage source whose voltage depends on nothing except the SOC, or the state of charge of the battery. This battery model well approximates the behavior of an actual battery cell when the magnitude of the output current of the battery is kept very small. This model will not perform well in the case of highly dynamic and powerful trajectories for legged robots as this model fails to capture any way of describing the finite output power of a battery pack. To achieve the desired output power, more current can simply be drawn from the battery at no expense to the output voltage. This is nonphysical as it is known that a battery cell cannot supply an infinite amount of current, even if its output terminals are shorted with a low impedance wire.

3.2.2 The Simple IR Model

The Simple IR model does a decent job at modeling a battery's output voltage given a power or current drawn from the battery. As discussed in Section 2.1.2, the simple IR model models the cell's output impedance with a resistor representing the IR, or the internal resistance, of the battery cell in series with an ideal voltage source whose voltage depends on the SOC. This model also captures the finite output power capabilities of the cell. When graphing the output voltage of the cell over output power, the shape of the curve follows a \sqrt{x} shape, where for a certain output power the derivative of voltage with respect to power is infinite. The function returns imaginary roots if the output power exceeds the peak power threshold as the cell would be physically incapable of outputting that output power independent of

what current is drawn from the battery.

The simple IR model is also a conservative model. When connecting a battery to a current step load, the simple IR model would predict that the battery's output voltage instantly settles to a voltage which equals the internal ideal voltage source minus the voltage drop accumulated over the IR of the cell due to an external current. This is not what an actual battery cell does in response to a current step load. Instead, the battery cell will decay towards the solution to the Simple IR model with a series of time constants. The Simple IR model therefore returns a voltage that is lower than actuality in response to a load which draws current from the cell, and returns a voltage that is higher than actuality in response to a load which sinks current into the cell. In a sense, the Simple IR model returns a voltage which assumes that the cell always has a worst case internal impedance which equals the DCIR, or the effective series resistance at DC or zero frequency. It is known even from the datasheet of lithium ion cells that at higher frequencies the "AC impedance" of the cell can be much lower, usually around half of the DCIR. This AC impedance usually refers to tests performed with square wave loads at 10kHz which is not the most useful metric, though still highlighting that cells can exhibit apparent lower output resistance when the load is a high frequency signal.

Below is data captured with the test bench setup that shows such time constants existing on the order of 1 - 10 seconds. The output voltage of the battery does not look like a square wave, and instead decays both during the increase in current and the decrease in current.

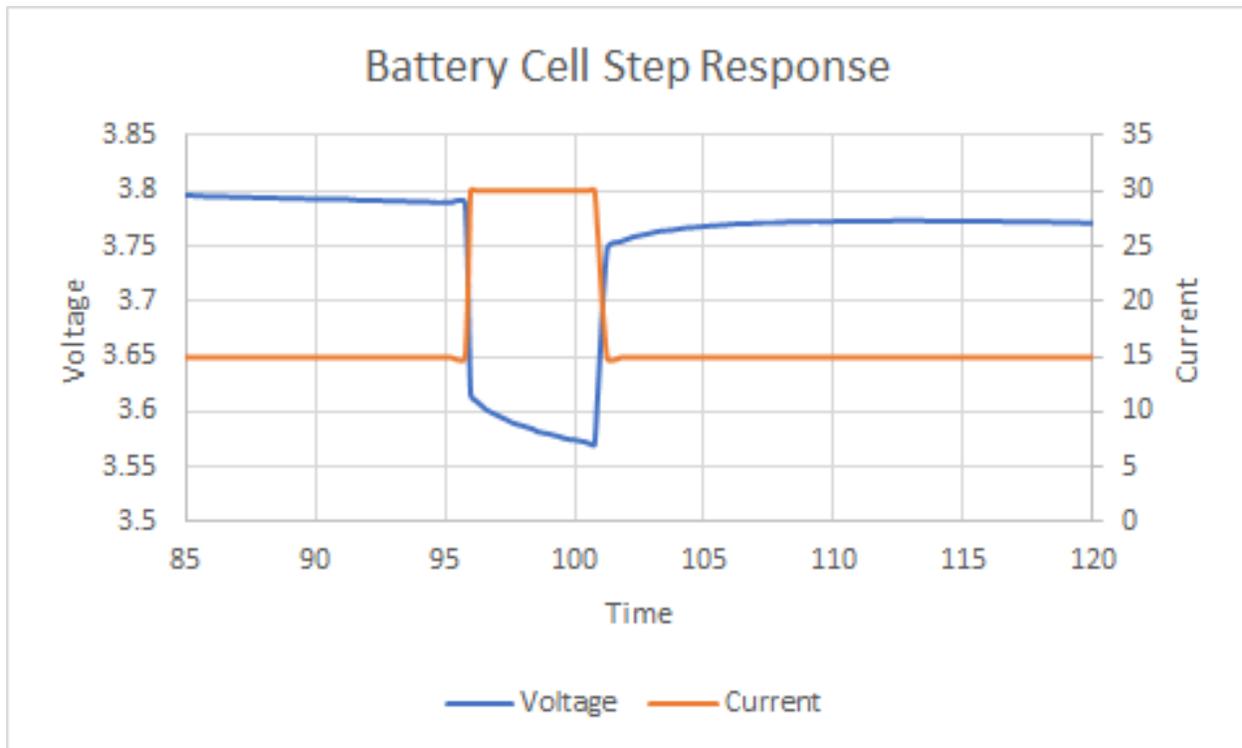


Figure 3.4: The battery voltage decays during a current pulse which lasts 5 seconds.

This figure shows that the Simple IR model, though correct in equilibrium, does not always accurately reflect what the output voltage of a battery cell will be in response to a load profile. The Simple IR model is stateless, returning a voltage given the instantaneous output power drawn from the cell. This means the Simple IR model is incapable of modeling these battery dynamics which cause the battery voltage to decay.

3.2.3 The Randles Model

Battery cells produce electrical current through the diffusion of ions across an ion permeable barrier, with the transport of the ions facilitated by the electrolyte solution. In response to a sudden change in load, the electrochemical processes internal to the cell do not immediately reach a new equilibrium. It is these electrochemical processes that give rise to the internal time dynamics of a battery cell.

The Randles model models the impedance of a battery cell using a network of resistors,

capacitors, and a non LTI Warburg impedance which can model these time-dependent effects.

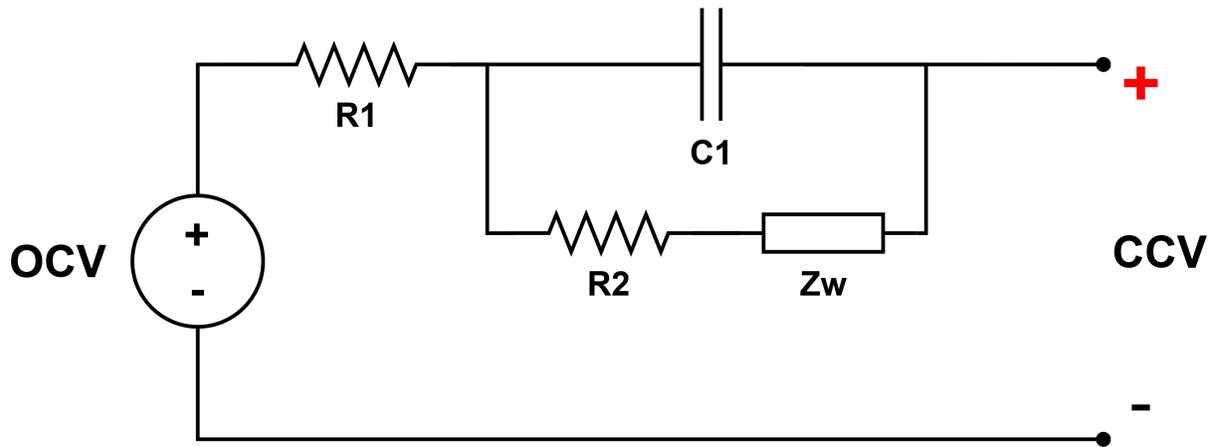


Figure 3.5: The Randles Model with Warburg Impedance

This equivalent circuit model is capable of modeling a battery cell's transient behavior reasonably well if the circuit elements are populated with values that correctly model the cell. The problem with this model is the Warburg impedance which is not an LTI component. The Warburg impedance is meant to be a CPE or constant phase element which implements a constant phase shift of 45 deg regardless of input frequency. Instead of including the Warburg impedance in the equivalent circuit model, a chain of parallel RC components can be used to instead approximate the effects of the Warburg impedance. This model will be referred to as the ERM, or the extended Randles model.

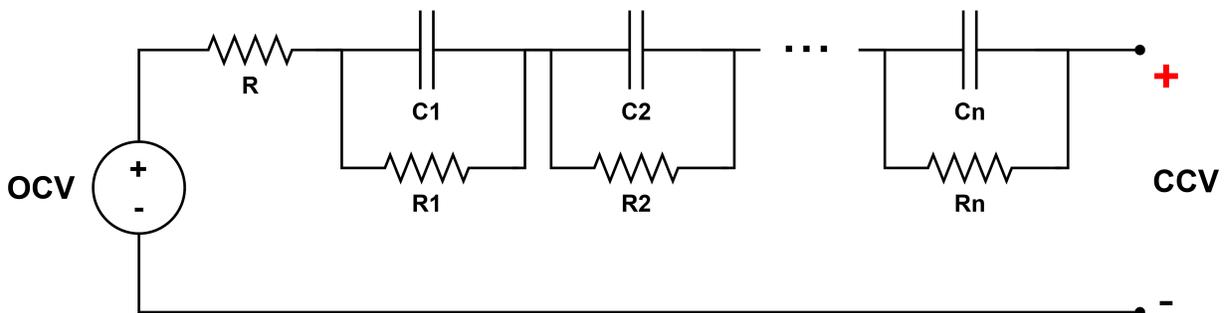


Figure 3.6: The Extended Randles Model with a chain of n RC circuits.

In comparison to the Simple IR model, each of the capacitors included in the ERM will

store an amount of charge that cannot change instantaneously. These charge values represent the internal hidden states of the battery which allow the ERM to predict these voltage transients in response to sudden changes in load.

The ERM is also preferable over the Randles model which includes the Warburg impedance because all components that make up the impedance model are LTI components. This is important because if component values are to be chosen to make the model reflect the behavior of the cell, an analytical function which returns the impedance of the network is needed to allow an optimizer to match the frequency response of the network to the frequency response of real world data collected from the cell. In order to collect this data, a frequency response analysis of an individual P45B cell is needed.

3.3 Electrochemical Impedance Spectroscopy

To have a battery model such as the extended Randles model effectively capture the internal dynamics of a battery cell, a way is needed to stimulate the battery cell to probe these internal parameters which are otherwise unobservable when conducting static tests such as tests which draw constant current or constant power. In the general case, the load a battery cell experiences can be equated to the instantaneous current out of the cell. A case can be made that instantaneous power is a better metric for the load demanded of the cell, but for modeling and simulation purposes thinking of the current drawn from the battery cell as the load will be easier to work with since current itself is a system state variable, while power is a product of system state variables ($V \times I$).

By modulating the input current to the battery cell, a corresponding voltage waveform can be measured. If the shape of the current waveform is a sinusoid and the frequency is swept across a wide range, a corresponding voltage amplitude and phase shift can be observed for every frequency along the sweep. This method of perturbing the cell using sinusoidal currents is called electrochemical impedance spectroscopy, or EIS.

Due to the small impedance of high power battery cells, the voltage waveform produced by a current waveform will be very small in comparison. Whatever driver is used to perturb the battery cell will need to be able to inject a non insignificant amount of current into the battery cell in order to get a decent voltage signal from the battery cell. Another problem with sourcing a driver to inject a current sinusoid into the battery cell is that one half of the sinusoid will have to apply positive power and the other half will have to apply negative power. This means the driver circuit will need to be able to both sink and source power. This difficulty is only exacerbated when going to lower frequencies as the amount of power sourced or sunk in each half cycle increases.

3.3.1 Performing EIS

The following experiments were conducted with a large bi-polar power supply capable of tracking an analog input signal with a bandwidth of around 5-10kHz with a peak current of $\pm 20\text{A}$. The analog input signal was generated using a signal generator.

A four channel 12-bit oscilloscope was used to measure both the battery cell voltage and battery cell current. All ground clips of the oscilloscope probes were connected to earth ground. A pair of probes measured voltage at the anode and cathode of the battery cell, and their difference was used to determine the battery voltage. This pair of probes made a kelvin connection to the battery cell using the cell jig mentioned in Figure 3.2. The other pair of probes measured voltage across a $10\text{m}\Omega$ shunt in order to measure current. Their difference was proportional to the current.

Both the signal generator and oscilloscope were coordinated using a central laptop computer. The laptop would command the signal generator to generate an input voltage waveform to the bipolar supply, causing a current waveform of a specific amplitude and frequency to be sent into the cell. The computer would then enable capture on the oscilloscope and wait at least 11 periods of the input signal plus one second before stopping capture of the channels on the oscilloscope. All four waveforms were then downloaded from the oscilloscope

to the computer where sinusoids are fit to the difference between the voltage measuring channels and the current measuring channels. The magnitude and phase of the voltage and current sinusoids are used to calculate the complex impedance of the cell. The impedance values are saved for every frequency. The current amplitude was set to one amp and the frequency was swept logarithmically from 0.1Hz to 300Hz with 20 data points collected during each test. For the entire test, the cells were kept at 25C using the water cooling loop. The tests were performed with the cell jig in a sand bucket, with another sand bucket ready to be dumped overtop if any battery fire were to occur.

The test was performed five times on five different cells each for a total of 25 experiments. Because the input current waveform is not on average perfectly zero, battery cells would either slightly charge or discharge over the duration of the test. In between experiments, cells were reset to a known SOC using a DC load and supply. The central laptop first discharged the cell below the desired voltage setpoint before charging the cell back up. As soon as the cell was charged to 3.8V, the charging stopped. Each cell's output voltage relaxed back to a consistent 3.72V following the end of charging and was considered ready to be frequency swept again.

A limitation of this setup was the inability to exceed about 300Hz before the data collected began to return erroneous results, with excessive phase shift being measured from simple test resistors. This is likely due to the inductance of the large loops in the harnessing coupled with the fact that the impedances being measured are relatively small. Despite this, mechanical dynamics of robots will typically be much slower than 300Hz. This means that frequencies above this threshold don't really need to be modeled.

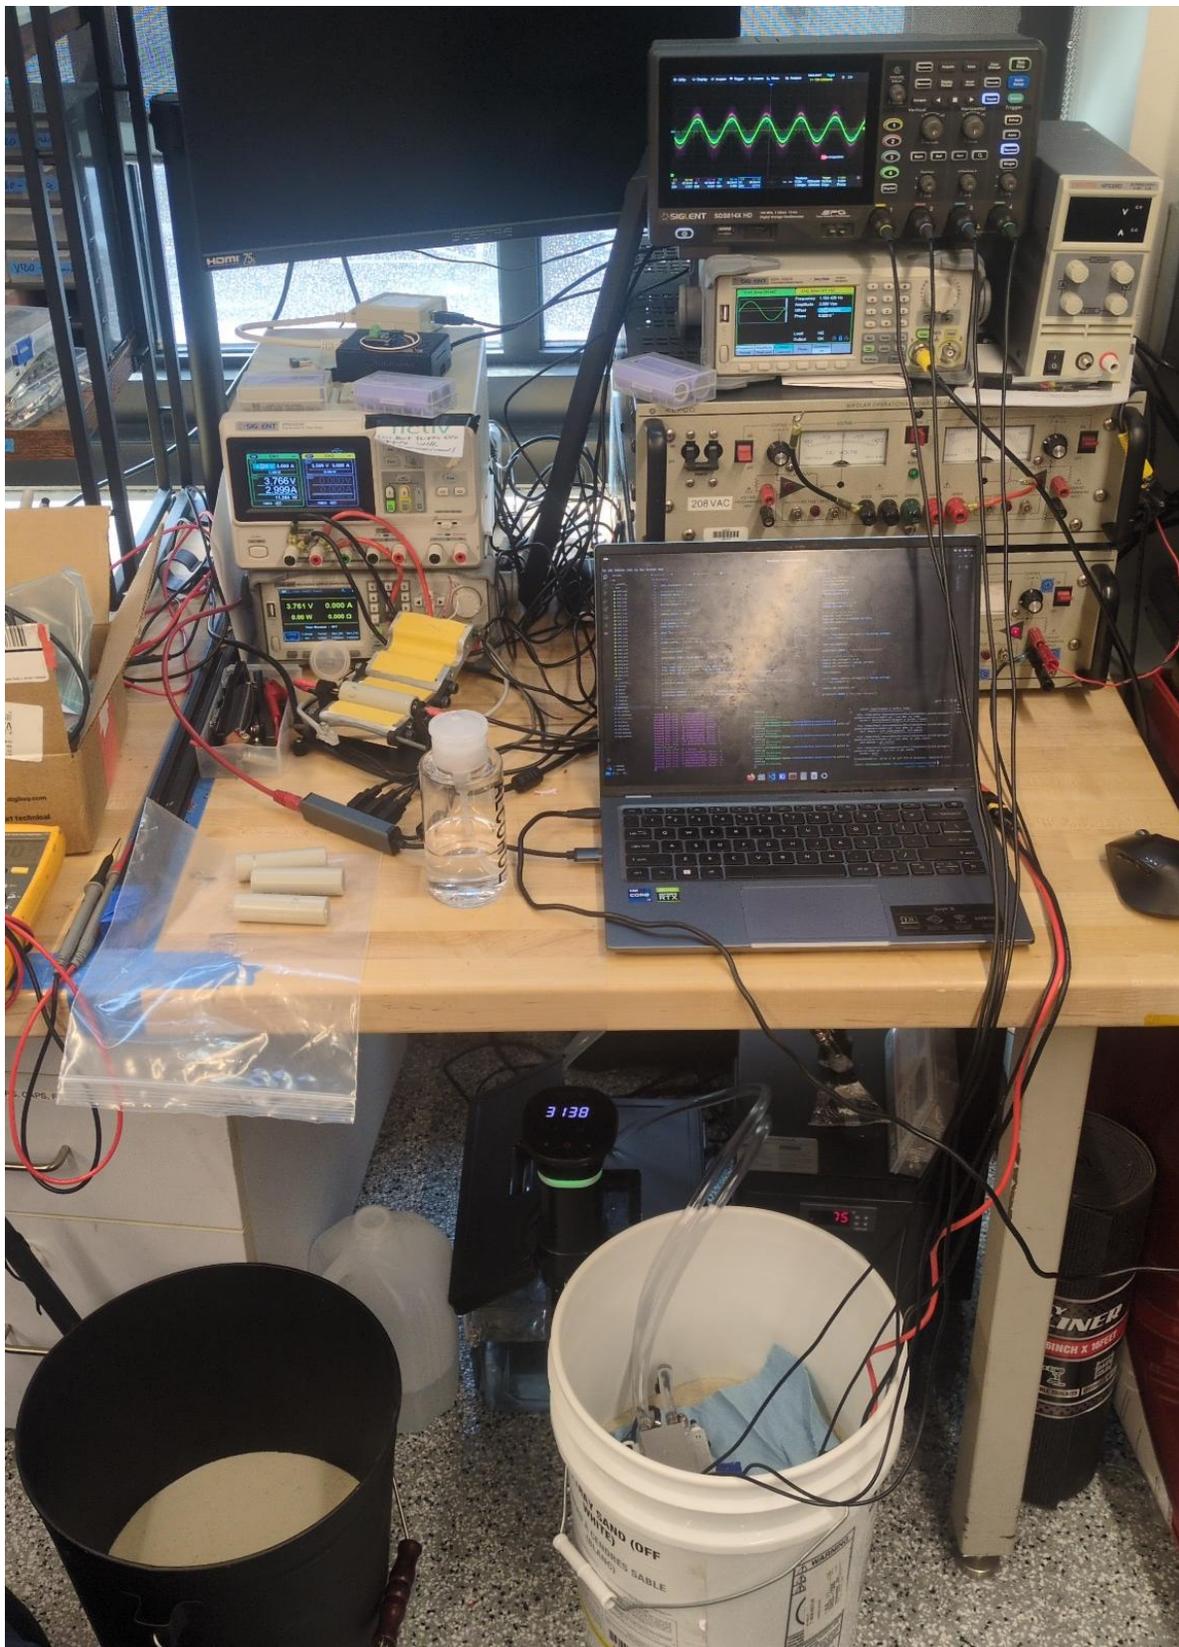


Figure 3.7: Test bench for conducting electrochemical impedance spectroscopy on lithium ion battery cells.

The result of these 25 tests is the following graph showing the impedance of each cell during each test as frequency is swept logarithmically from 0.1Hz to 300Hz using 20 steps.

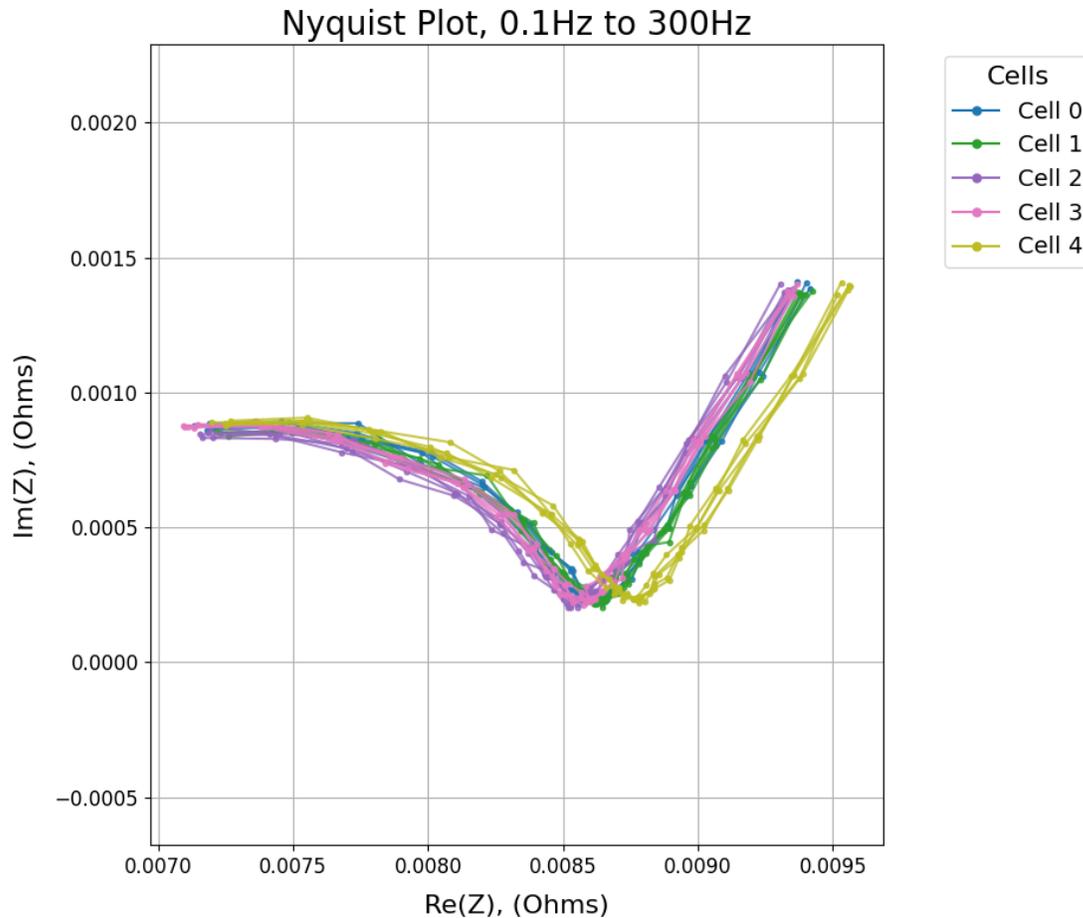


Figure 3.8: Collected EIS data on the P45B cells. The scaling of the real and imaginary axes are the same.

In this graph, the points furthest to the top right of the graph represent the low frequencies, and the points closest to the left represent the highest frequencies. This makes sense since it is expected that the cell's impedance will decrease with increasing frequency. As evident from the collected data, the impedance of the cell does change rather significantly as frequency is swept. From near DC to 300Hz, the impedance changes from a magnitude of about $9.6m\Omega$ to $7.2m\Omega$, a 25% change in the cell's impedance! This lowered impedance could be very useful a battery pack's application requires very high powers for very short amounts of time, which so happens to be the case with the MIT Humanoid.

3.3.2 Data Collection Validation

This data collecting test setup was validated by connecting two control impedances. First, a fake cell was constructed by 3D printing a cylinder of the same dimensions as the P45B and gluing two copper sheets to either side of the cell, connecting them with a wire. Secondly, the voltage measuring probes were removed from the cell jig and attached to the current measuring shunt used to measure the current in this setup. This shunt was measured with a DMM to be $9.96m\Omega$, offering a precision impedance to validate against.



Figure 3.9: The constructed fake cell.

The same frequency sweep that was performed in each of the 25 tests were also performed on these two control impedances. The resulting data is shown in the graph below.

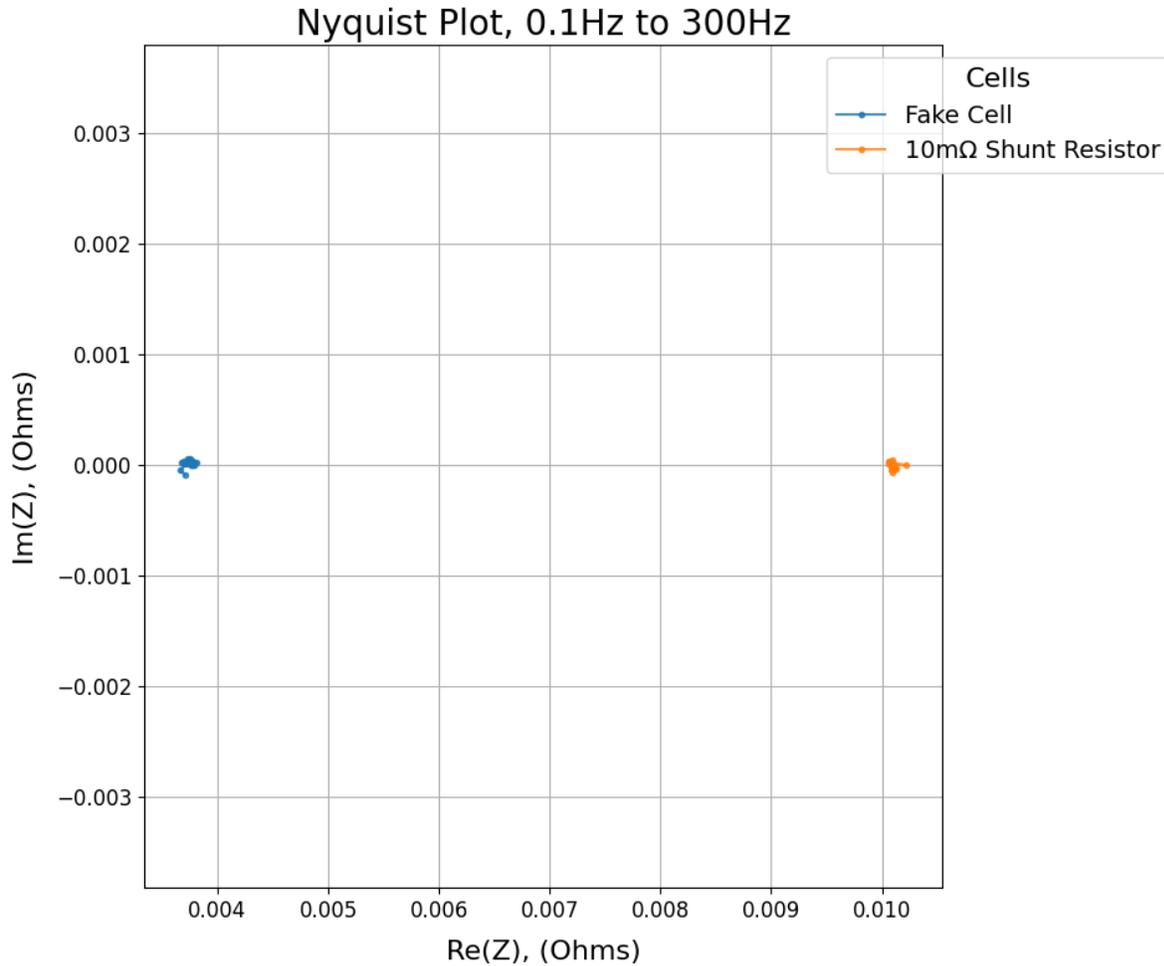


Figure 3.10: The two control impedances and their corresponding Nyquist plots. The scaling of the real and imaginary axes are the same.

As you can see, the impedance of each of the resistors did not move around in the Nyquist plot much as frequency was swept from 0.1Hz to 300Hz. This is to be expected of a resistor since a resistor's impedance is real valued and does not change with frequency. The cluster of points which represent the 10mΩ shunt grouped around a resistance value of around 0.01015Ω, which equates to about a 1.5% impedance magnitude measuring error of the testing setup.

3.3.3 Fitting the Extended Randles Model

Given data from EIS testing, it is now possible to fit an extended Randles model to best match the data. To do this, an optimizer will be used to minimize the difference in impedance

magnitude between the collected data and the ERM for every frequency. The optimizer will find the values of the resistors and capacitors such that the impedance of the model best reflects the real world data. The optimizer chosen to do this was scipy's differential evolution which does not require gradient information and can handle nonlinear objective functions. The number of RC circuits in the ERM can be chosen, and the optimizer will do its best to match each equivalent circuit model to the collected EIS data.

The following graphs will show the best fit ERM for differing numbers of series RC circuits. The values for the equivalent circuit components are written in the appendix to this thesis.

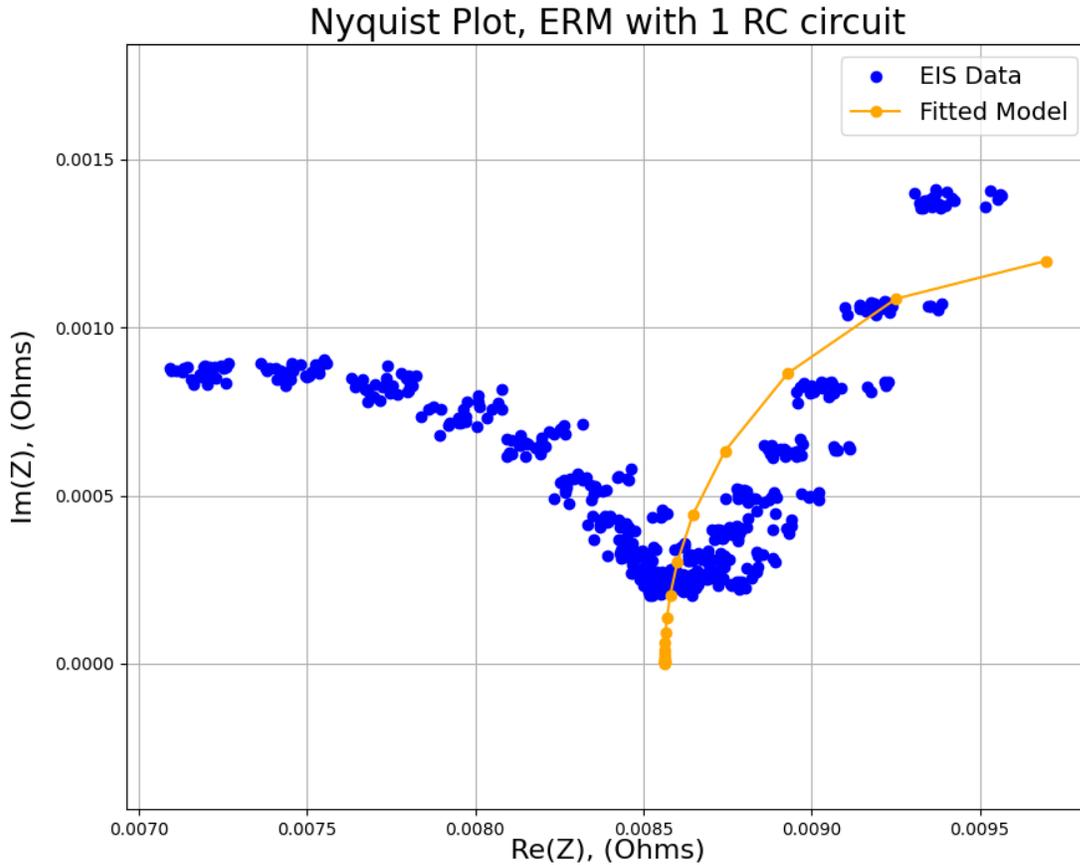


Figure 3.11: Fitted ERM with 1 RC circuit

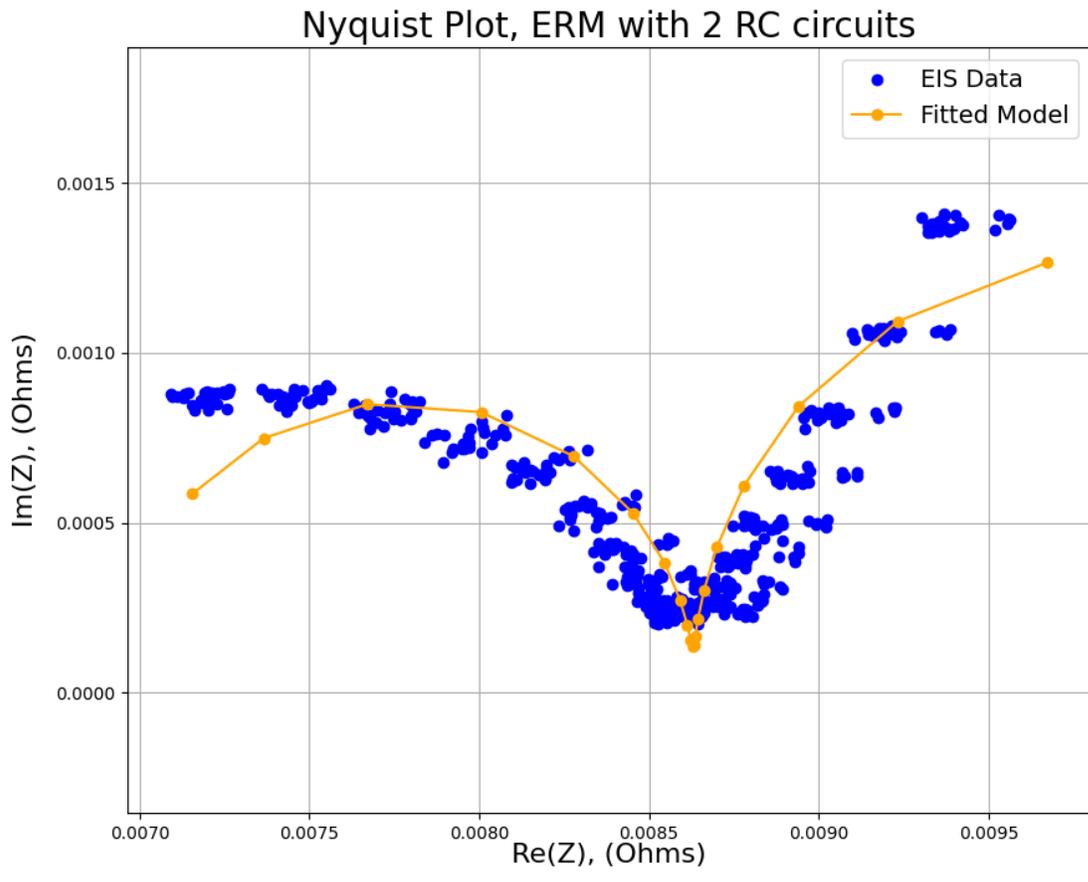


Figure 3.12: Fitted ERM with 2 RC circuits

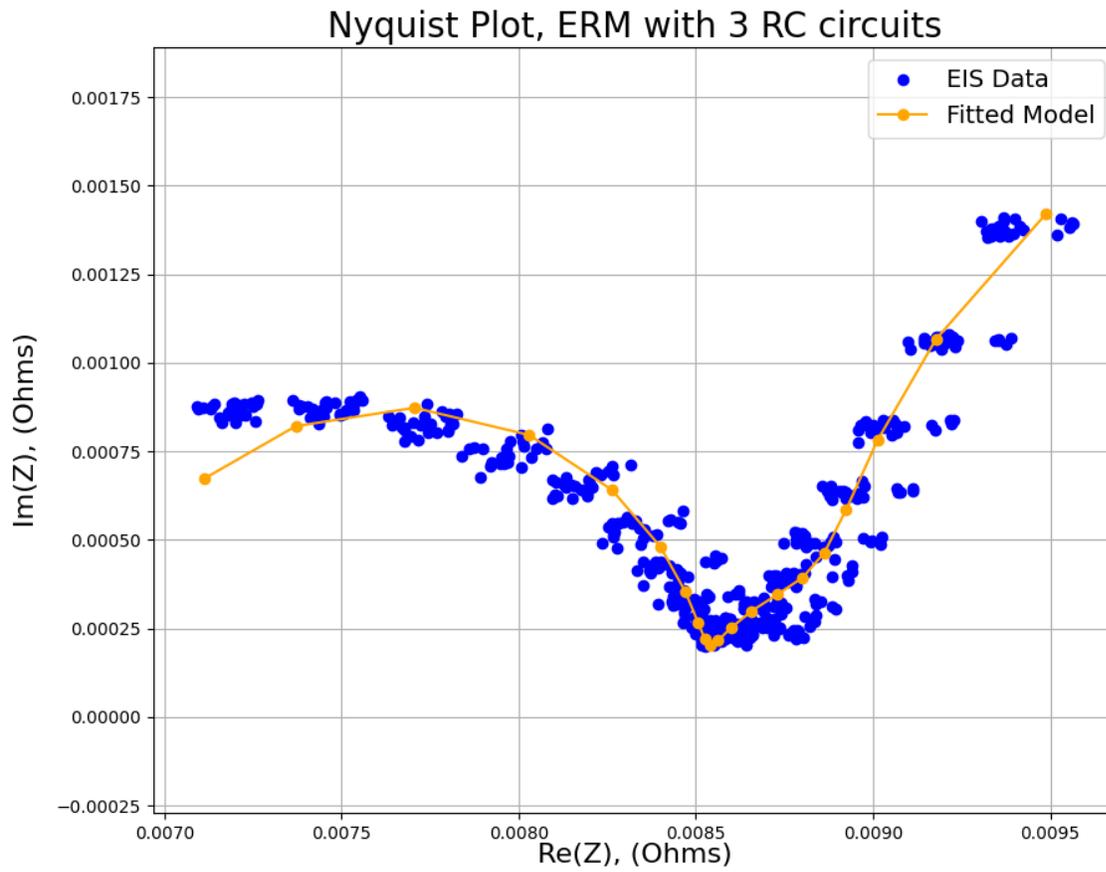


Figure 3.13: Fitted ERM with 3 RC circuits

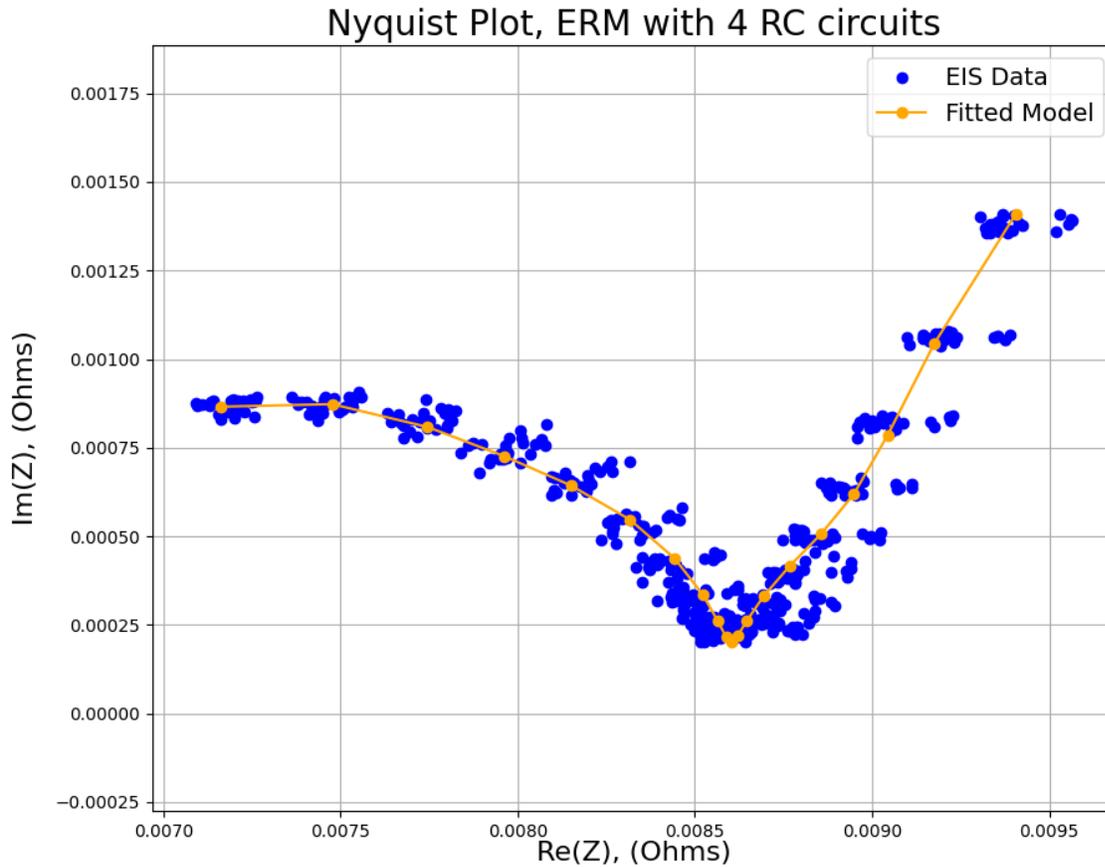


Figure 3.14: Fitted ERM with 4 RC circuits

It is clear that the fourth order ERM with four RC circuits fits the data the best. If a simpler model is needed, the second order ERM with two RC circuits could suffice well.

3.3.4 Proof of Mapping ERM from Battery Cell to Battery Pack

Once an extended Randles model for a cell has been appropriately fitted based on empirical data, it is not immediately obvious how this can be used to model a battery pack which has many cells in series and parallel. If the assumption that all battery cells in a battery pack are identical is made, then it can be said that each cell will experience an identical current profile in response to a power profile demanded of the entire pack. That would mean the internal

state of every battery cell would be identical. Thus, the task of modeling the entire battery pack reduces down to modeling an ERM that represents the pack itself, not an individual cell. How the component values of the pack level model relate to the cell level model is also not immediately obvious. The immediate answer to this dilemma is that resistors scale by the ratio of the S count to the P count, while capacitors scale by the ratio of the P count to the S count. This claim will be proven in the case of an ERM with only one RC circuit, but a more general proof can be made to prove the general case of n series RC circuits.

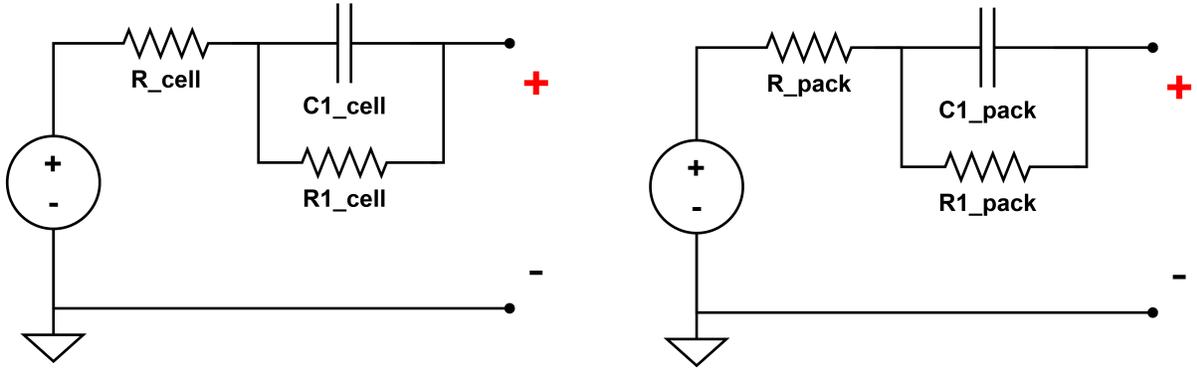


Figure 3.15: Cell level ERM on the left and pack level ERM on the right with only one RC circuit.

The impedance of both the cell and pack level ERM can be found by shorting all internal sources and measuring the impedance from the output terminals. The impedance of the cell and pack level ERM are therefore:

$$Z_{cell} = (R_{cell} + \frac{R_{1,cell}}{1 + \omega^2 R_{1,cell}^2 C_{1,cell}^2}) + j(\frac{-\omega R_{1,cell} C_{1,cell}}{1 + \omega^2 R_{1,cell}^2 C_{1,cell}^2}) \quad (3.1)$$

$$Z_{pack} = (R_{pack} + \frac{R_{1,pack}}{1 + \omega^2 R_{1,pack}^2 C_{1,pack}^2}) + j(\frac{-\omega R_{1,pack} C_{1,pack}}{1 + \omega^2 R_{1,pack}^2 C_{1,pack}^2}) \quad (3.2)$$

Due to the linearity of LTI components, the relationship between the impedance of a cell

and the pack of a battery with a 16S3P configuration is as follows:

$$Z_{pack} = \frac{S}{P} \cdot Z_{cell} \quad (3.3)$$

By looking at special conditions on frequency and matching terms, it is possible to determine the relationship between the resistors and capacitors in the cell level and pack level ERM.

By looking at the real component of each impedance as the input frequency approaches infinity, we find the following to be true:

$$R_{pack} = \frac{S}{P} \cdot R_{cell} \quad (3.4)$$

By looking at the real component of each impedance as the input frequency approaches zero, we find the following to be true:

$$R_{pack} + R_{1,pack} = \frac{S}{P} \cdot R_{cell} + \frac{S}{P} \cdot R_{1,cell} \quad (3.5)$$

Because of equation 3.4, this equation can be simplified to:

$$R_{1,pack} = \frac{S}{P} \cdot R_{1,cell} \quad (3.6)$$

Finally, to find the relationship between the two capacitors the relationship between each of the resistor pairs can be substituted into the real parts of equation 3.3.

$$\left(R_{pack} + \frac{R_{1,pack}}{1 + \omega^2 R_{1,pack}^2 C_{1,pack}^2} \right) = \frac{S}{P} \left(R_{cell} + \frac{R_{1,cell}}{1 + \omega^2 R_{1,cell}^2 C_{1,cell}^2} \right) \quad (3.7)$$

Substituting leaves the following equation:

$$\frac{1}{1 + \omega^2 R_{1,pack}^2 C_{1,pack}^2} = \frac{1}{1 + \omega^2 R_{1,cell}^2 C_{1,cell}^2} \quad (3.8)$$

Simplifying further gives the following equation:

$$\frac{R_{1,pack}}{R_{1,cell}} = \frac{C_{1,cell}}{C_{1,pack}} \quad (3.9)$$

Substituting in equation 3.4 reveals the final relationship between the capacitors:

$$C_{1,pack} = \frac{P}{S} \cdot C_{1,cell} \quad (3.10)$$

Equations 3.4, 3.6, and 3.10 prove that the mapping between component values in the ERM from the cell level to the pack level cause resistors to scale by the ratio of S count to P count while capacitors scale by the ration of P count to S count. Using this information, it is possible to build an ERM that reflects the battery pack itself using only the ERM which models the individual cells of that battery pack and the configuration of the battery pack, the S and P count.

Chapter 4

Simulating Batteries

Now that a battery model such as the extended Randles model has been made and fit according to real world data, forward simulation and utilization in an optimization are yet to be explored.

4.1 Forward Simulation of the Extended Randles Model

Once a pack level ERM has been made, the ability to predict battery voltage and update the charge on the internal capacitors given an external power profile is still to be explored. In the following derivations the assumption that the ERM only has one RC circuit is made, however following the general structure of the derivations can reveal how to scale to ERMs that have an arbitrary number of RC circuits in series.

In order to properly forward simulate an ERM and produce updates for the charge stored in the internal capacitors for use in model validating simulations or trajectory optimization, a three step sequence can be followed. First, the output voltage given a desired instantaneous output power is calculated. This can be done just as easily as the Simple IR model since the charge on the internal capacitors cannot change instantaneously and will remain constant for the small time interval of the integrator being used. Second, the current in the capacitor given the charge in the capacitor and the instantaneous output power is calculated. Finally,

the charge in the capacitor is updated using the calculated capacitor current and time interval of integration.

The following derivation will reveal the equations which represent each step in forward simulating an ERM with one RC circuit.

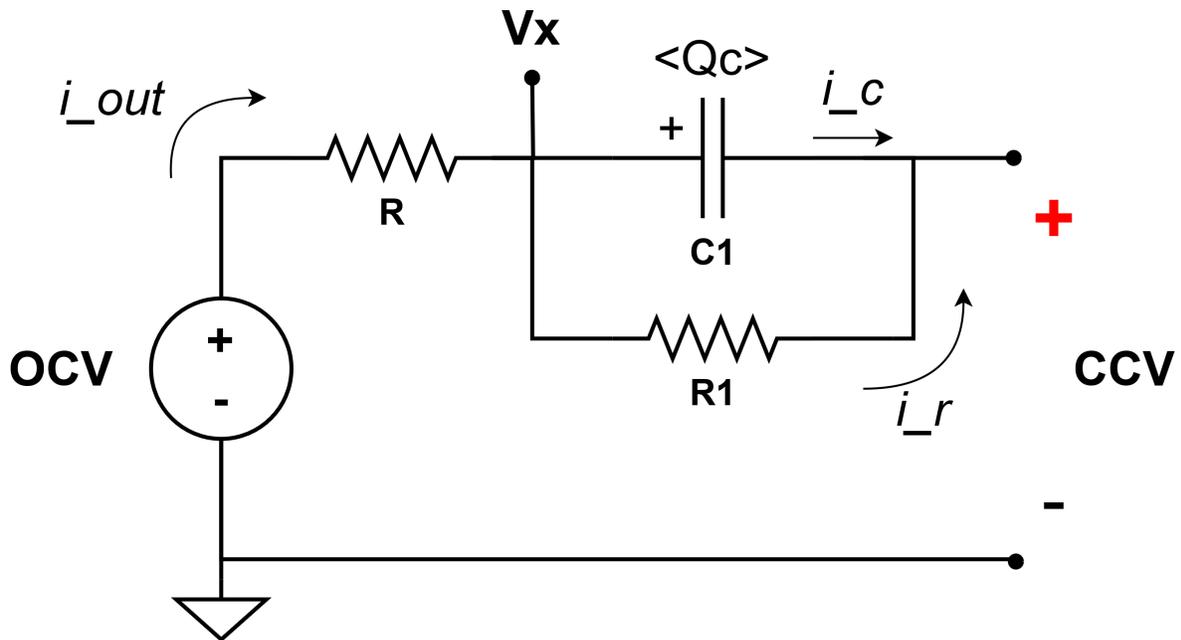


Figure 4.1: ERM with single RC circuit with nodes, currents, and charges annotated.

Given a desired output power P and charge stored in internal capacitor $\langle Q_c \rangle$, the output voltage of the ERM with one RC circuit can be calculated. $\langle Q_c \rangle$ is shown in angle brackets to denote that it is a system state variable.

$$P = CCV \cdot i_{out} \quad (4.1)$$

Using KVL, we find the following equation:

$$OCV - i_{out} \cdot R - \frac{\langle Q_c \rangle}{C1} = CCV \quad (4.2)$$

Substituting Equation 4.1 into Equation 4.2 gives:

$$OCV - \frac{P}{CCV} \cdot R - \frac{\langle Qc \rangle}{C1} = CCV \quad (4.3)$$

Multiplying by CCV and rearranging, we find the quadratic:

$$CCV^2 + \left(\frac{\langle Qc \rangle}{C1} - OCV\right) \cdot CCV + P \cdot R = 0 \quad (4.4)$$

Solving the quadratic and choosing the root with the larger value which corresponds to the solution that requires less current to achieve the same output power, we find the following equation for CCV, or the output voltage of the battery pack given desired power P and internal charge $\langle Qc \rangle$.

$$CCV = \frac{OCV - \frac{\langle Qc \rangle}{C1} + \sqrt{\left(\frac{\langle Qc \rangle}{C1} - OCV\right)^2 - 4 \cdot P \cdot R}}{2} \quad (4.5)$$

This concludes step one. For step two, we must find the current in the capacitor. To do this, we will first use the solution for the output voltage, CCV, and the desired output power, P, to find the current in the cell.

$$i_{out} = \frac{P}{CCV} \quad (4.6)$$

Next, due to KCL, we find the following equation:

$$i_{out} = i_c + i_r \quad (4.7)$$

Next, solving for i_r is as simple as dividing the voltage of the capacitor by resistor R1.

$$i_r = \frac{\frac{\langle Qc \rangle}{C1}}{R1} \quad (4.8)$$

Finally, we can use the previous three equations to solve for the current in the capacitor.

$$i_c = \frac{P}{CCV} - \frac{\langle Qc \rangle}{C1R1} \quad (4.9)$$

This concludes step 2. The final step is to update the charge in the capacitor. To do this, the integration time step ΔT must be known.

Using the calculated value for the current in the capacitor and the integration time step ΔT , a change in charge of the capacitor can be found as follows:

$$\Delta Q = i_c \cdot \Delta T \quad (4.10)$$

This delta is added to the value of charge in the current time step to find the charge in the next time step. That concludes the three step process to forward simulate and update the internal state of the ERM with one series RC circuit.

4.2 The Issue of Imaginary Voltage

If all of the previous steps are followed to populate a battery model and implement the forward simulation and update equations for a trajectory optimization, one may find difficulty in having a solver find solutions reliably due to the intrinsic shape of the battery voltage curve. Whether the Simple IR model or the ERM is used, the underlying function that returns the battery voltage given power and potentially capacitor charge values is a \sqrt{x} function. Recalling the function which returns the battery voltage given the ERM with one RC circuit:

$$CCV = \frac{OCV - \frac{\langle Qc \rangle}{C1} + \sqrt{(\frac{\langle Qc \rangle}{C1} - OCV)^2 - 4 \cdot P \cdot R}}{2} \quad (4.11)$$

If the radicand in this equation were to be negative, the resulting output voltage of the battery pack would be imaginary. When tasking a trajectory optimization with achieving a certain jump height that is infeasible to achieve given the capabilities of the battery pack in that system, the optimizer will still attempt to find a solution. The optimizer will increase

joint torques at high joint speeds to achieve the goal, but in doing so will demand more and more power from the battery model. That manifests are larger values of P inside the radicand which causes the radicand to become more negative. For trajectories that are not feasible for deployment on real robot hardware due to power limits, the optimizer will push the radicand to be more negative which will inevitably return an imaginary solution during an iteration of the optimizer. In order to avoid imaginary numbers, a trick can be used to remedy the situation.

4.2.1 Linear Extrapolation of the Voltage Curve

One method to avoid imaginary numbers is to switch the battery voltage curve to a line function at a power level slightly before the critical power level, P^* , where the square root function has an infinite magnitude derivative.

By defining an indicator function, $i(P)$, a plain expression can be formed for the shape of this curve. P^* will be defined as the point along the x-axis where the battery voltage function has a double root, the analog being the peak power of the battery curve. By transitioning the battery voltage curve to a tangent line slightly before this critical point P^* , the overall function can be real valued for all values of P , or power. The point along the power axis where the transition from the square root function $V(P)$ to the linear function $L(P)$ is P_0 .

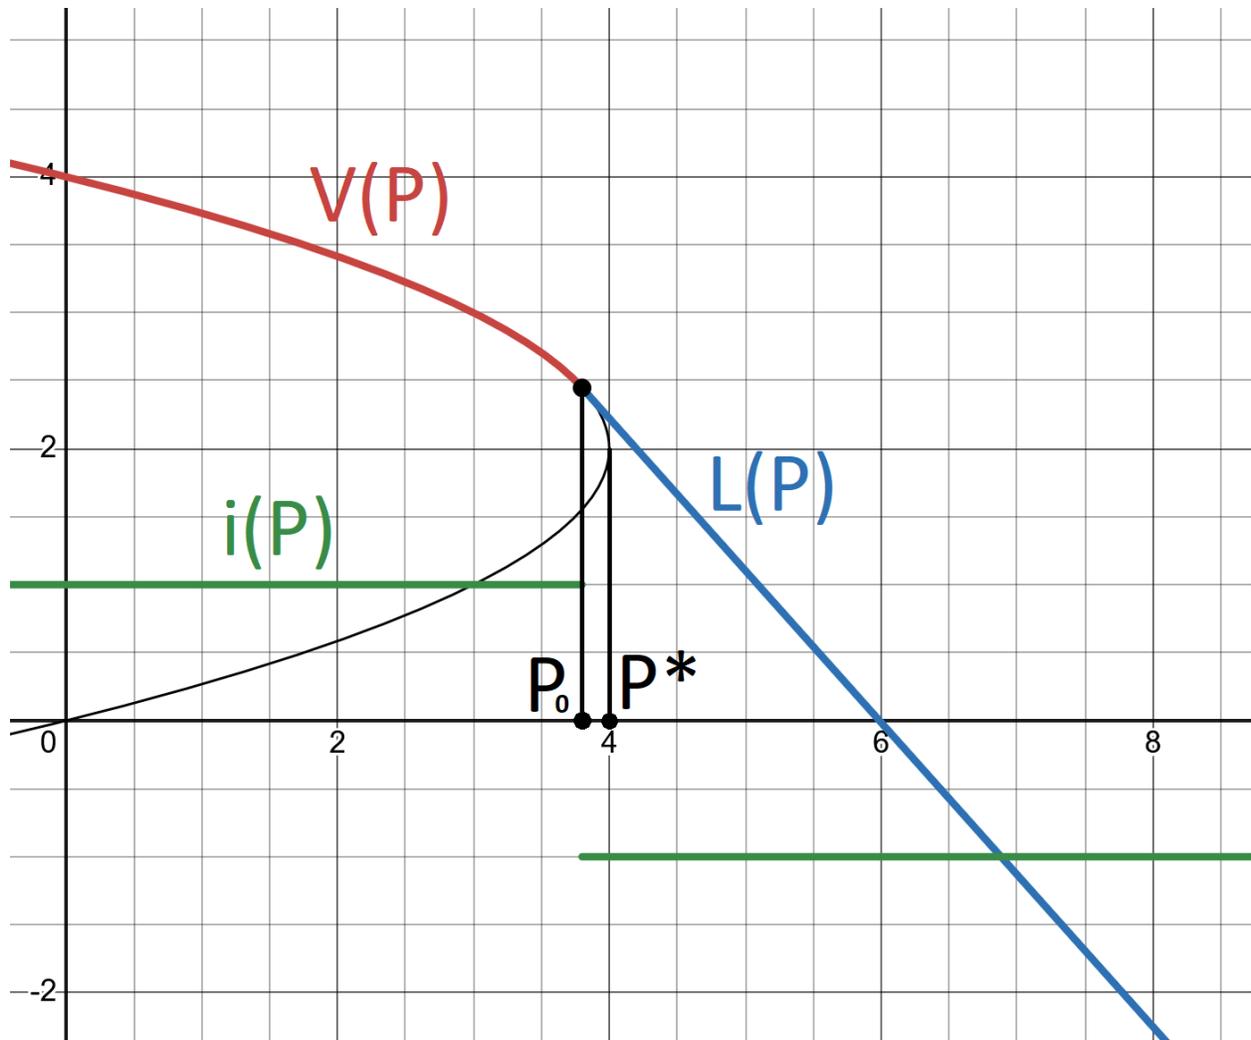


Figure 4.2: Battery voltage curve which transitions into a linear function, graphed with Desmos[3].

The indicator function is defined as:

$$i(P) = \frac{P_0 - P}{|P_0 - P|} \quad (4.12)$$

The linear function is defined as:

$$L(P) = V'(P_0) \cdot (P - P_0) + V(P_0) \quad (4.13)$$

Where $V'(P)$ is the analytical derivative of the battery curve function.

The final function which represents the battery voltage curve with the linear function tacked on is as follows:

$$CCV(P) = V(P)\left(\frac{1}{2} + \frac{1}{2}i(P)\right) + L(P)\left(\frac{1}{2} - \frac{1}{2}i(P)\right) \quad (4.14)$$

As P_0 approaches P^* , the slope of the linear function becomes more and more negative which can help prevent the optimizer from being able to source high powers from the battery pack while still maintaining the battery pack's output voltage. The graph shows an extreme example of the linear function; for better performance it would probably be best to have the slope of the linear function be quite steep.

This method can also benefit from a constraint on power drawn from the battery. By constraining the total power of the robot to be less than P_0 for every time step, a valid trajectory will only be subject to power limits that are reflected from the actual battery curve $V(P)$. If this constraint is violated during solve time the battery voltage function will still not return any imaginary numbers which will help convergence of the optimization.

4.3 Constraining The Optimization

Now that a battery model has been made and can fairly accurately predict battery voltage as a function of power demanded, relevant constraints must be imposed on the solutions found by a the solver in order to reap the benefits of such labor. When it comes to protecting batteries and respecting their limits, the most important thing that can be done is to keep each battery cell's output voltage within the limits specified by its datasheet. Other considerations like keeping the cells within a certain temperature range are also important, but for the application of research into legged robotics this constraint is unlikely to be broken. Arbitrary limits such as current limits often noted in advertising materials or datasheets are based off of thermal limits of the cell, operating at those elevated current levels continuously or semi-continuously for enough time for the cell's internal temperature to reach dangerous

levels. Due to the nature of power profiles used in high power legged robots, it would be highly unlikely that the battery pack is pushed to such current levels for a long enough period of time for any problem to arise. High power trajectories like jumps last only hundreds of milliseconds, and any energy dissipated in the cell is gladly subdued by the thermal mass of the cells themselves. It is therefore imperative that the battery pack's output voltage remain within the operating voltage range of the cell which makes up the pack, times the S count. This would manifest as two inequality constraints on the output voltage, or the CCV of the battery pack. For every node in the trajectory, the CCV for that node must be constrained to be above $2.5V$ and below $4.2V$.

Chapter 5

Conclusion and Future Work

This thesis has laid the ground work for how to design and model a high power battery pack for use in legged robots. By leveraging the internal transient effects inherent to all electrochemical cells, a less conservative approach to modeling battery limits can be used to push the limits of robotic power systems. It was shown that at higher frequencies, state of the art high power battery cells can have their impedance reduced by up to 25%. A more accurate estimate for battery output voltage also allows for more accurate actuator torque-speed limits. Incorporation of the extended Randles model in a trajectory optimization framework will help close the sim to real gap which plagues robotics, especially in scenarios that require large actuator effort. This method of modeling batteries is promising as it allows existing battery systems to potentially output more power safely and allows future battery designs to be smaller if the desired peak output power is limited in duration.

The original intention of this thesis project was to also incorporate the transient battery model into a trajectory optimization and prove that trajectories generated without the model would either be infeasible or too conservative. Such engineering efforts would require a bit more time to implement. This would be the next step in continuing this work as the ERM model has been fit and concerns about solver feasibility have already been considered. Hardware experiments can also be conducted since the battery cell that was modeled is also

in the battery pack designed, and the MIT Humanoid is a great and well modeled research platform to perform such experiments. The future holds great potential.

Appendix A

ERM Equivalent Circuit Model

Parameters for P45B

1 RC circuit

Fitted Model Parameters:

$$R = 0.008562245725282797$$

$$R1 = 0.0023999509812431047$$

$$C1 = 700.4376758029904$$

Time constants:

$$T1 = 1.6810160873430264$$

$$\text{DC resistance} = 0.0109621967065259$$

2 RC circuits

Fitted Model Parameters:

$$R = 0.006923690842352759$$

$$R1 = 0.0025743852849546894$$

$$C1 = 749.9550113627948$$

$$R2 = 0.0017064870192161953$$

$$C2 = 1.171459782843088$$

Time constants:

$$T1 = 1.9306731456304058$$

$$T2 = 0.0019990809129555526$$

$$\text{DC resistance} = 0.011204563146523645$$

3 RC circuits

Fitted Model Parameters:

$$R = 0.006795703160648186$$

$$R1 = 0.0017348660657844812$$

$$C1 = 0.9668640102276936$$

$$R2 = 0.0037919180765289486$$

$$C2 = 963.3638336436927$$

$$R3 = 0.00035396038802013247$$

$$C3 = 371.8916297644393$$

Time constants:

$$T1 = 0.0016773795615723252$$

$$T2 = 3.6529967350677452$$

$$T3 = 0.13163490557286037$$

$$\text{DC resistance} = 0.012676447690981748$$

4 RC circuits

Fitted Model Parameters:

$$R = 0.006428631596419841$$

$$R1 = 0.0015092416058207665$$

$$C1 = 0.5528122375316784$$

$$R2 = 0.0004040655126863733$$

$$C2 = 663.8165189932115$$

$$R3 = 0.004774060802661011$$

$$C3 = 1084.2430675670112$$

$$R4 = 0.0006627965074183888$$

$$C4 = 6.9798491396501285$$

Time constants:

$$T1 = 0.0008343272290896813$$

$$T2 = 0.26822536207667563$$

$$T3 = 5.176242329428602$$

$$T4 = 0.004626219632067351$$

$$\text{DC resistance} = 0.01377879602500638$$

References

- [1] M. Chignoli et al. “The MIT humanoid robot: Design, motion planning, and control for acrobatic behaviors”. In: *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*. IEEE. 2021.
- [2] V. Muenzel, A. F. Hollenkamp, A. I. Bhatt, J. de Hoog, M. Brazil, D. A. Thomas, and I. Mareels. “A Comparative Testing Study of Commercial 18650-Format Lithium-Ion Battery Cells”. In: *Journal of The Electrochemical Society* 162.8 (2015). DOI: [10.1149/2.0721508jes](https://doi.org/10.1149/2.0721508jes). URL: <https://doi.org/10.1149/2.0721508jes>.
- [3] I. Desmos. *Desmos Online Graphing Calculator*. <https://www.desmos.com>. 2025.