

Design and Analysis of a Battery for a Formula Electric Car

by

Samuel Reineman

Submitted to the
Department of Mechanical Engineering
in Partial Fulfillment of the Requirements for the Degree of


Bachelor of Science in Mechanical Engineering

at the

Massachusetts Institute of Technology

June 2013

ARCHIVES



© 2013 Massachusetts Institute of Technology. All rights reserved.

Signature of Author: _____
Department of Mechanical Engineering
May 17, 2013

Certified by: _____
Anette Hosoi
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by: _____
Anette Hosoi
Professor of Mechanical Engineering
Undergraduate Officer

Design and Analysis of a Battery for a Formula Electric Car

by

Samuel Reineman

Submitted to the Department of Mechanical Engineering
on May 17, 2013 in Partial Fulfillment of the
Requirements for the Degree of

Bachelor of Science in Mechanical Engineering

ABSTRACT

The purpose of this paper is to present the philosophy and methodology behind the design of the battery pack for MITs 2013 Formula SAE Electric racecar. Functional requirements are established for the pack. An overview of cell chemistry, pack size and configuration selection process to meet these requirements is given. Next, the mechanical and electrical design and analysis of the major pack components is discussed. Finally, a transient thermal model of the pack is established to guide design choices about cooling.

Thesis Supervisor: Anette Hosoi
Title: Professor of Mechanical Engineering

Table of Contents

1.0	Overview of the FSAE Electric Competition	8
2.0	Introduction to the Tractive System	8
2.1	Functional Requirements	9
2.2	Safety	10
2.3	Power Density vs Energy Density	11
2.3	Track Model.....	13
3.0	Accumulator Design.....	14
3.1	Chemistry Selection.....	14
3.2	Pack Sizing	15
3.3	Tab connections.....	18
3.3.1	Clamping Plate	19
3.3.2	Clamping Wedges.....	20
3.3.3	Clamping Bars.....	21
3.4	Pack architecture	22
3.4.1	Cell Separation	23
3.4.2	Endplate Design	24
3.5	Battery Management System	25
3.5.1	Electrical Layout.....	26
3.5.2	Mechanical Layout.....	27

4.0	Thermal Analysis	28
4.1	Heat Generation.....	29
4.2	Adiabatic Model.....	30
4.3	Cooling Model	32
4.4	Results	34
5.0	Conclusion	38
6.0	Appendix.....	40
6.1	Natural Convection	40
7.0	References	41

List of Figures:

Figure 2.1: Block diagram of the Tractive System 9

Figure 2.2: Power density vs Energy Density 12

Figure 2.3: Simulated current draw over one lap of the endurance race 14

Figure 3.1: A123 20Ah Pouch Cell Discharge Profile at 10A..... 17

Figure 3.2: Tab Clamping Overview 19

Figure 3.3: FEA of the Clamping Plate 20

Figure 3.4: Exaggerated Wedge Displacement Under Load 21

Figure 3.5: FEA of the Clamping Bar Under Load 22

Figure 3.6: Full 28s Module Architecture..... 23

Figure 3.7: Cross Section Showing Cell Separation 24

Figure 3.8: Endplate FEA under 20g Crash Loading Conditions..... 25

Figure 3.9: BMS Electrical Block Diagram..... 27

Figure 3.10: 28s Module with BMS Board 28

Figure 4.1: Basic Cell Model 29

Figure 4.2: Test Setup for Measuring Internal Resistance 30

Figure 4.3: Transient Pack Temperature over an Endurance Race 35

Figure 4.4: Pack Temperature Rise With and Without Heat Spreaders 36

Figure 4.5: Fan Cooling Performance Demands..... 37

1.0 Overview of the FSAE Electric Competition

Formula SAE is an international, intercollegiate design competition focused around the design and fabrication of a small-scale formula style race car. Several different versions of the competition exist, ranging from gas powered ICE cars to hybrid to full electric cars. In all of the competitions, students build cars guided by a set of rules published each year by the Society of Automotive Engineers [1]. The rules primarily focus on safety; the vast majority of design decisions are carried out by the participating teams, resulting in a wide variety of cars at the final competition, which are then raced and judged based on performance, design, and a variety of other metrics.

MIT Motorsports, the student team at MIT which participates in the Formula SAE competition, switched this year from the gas competition (FSAE) to the electric competition (FSAE Electric); this change has introduced a host of new engineering challenges in the design of a safe and reliable fully electric powertrain.

2.0 Introduction to the Tractive System

MIT's 2013 car is built around the tractive system. A block diagram overview of the system, created by Brian Sennett (MIT '13), is presented in Figure 2.1. In general, power is provided by the batteries (more generically referred to as an accumulator) to two motor controllers which take inputs from the vehicle controller to torque-control two independent motors, each powering one of the rear wheels of the car. A series of safety devices compose a shutdown circuit, which has the ability to open the Accumulator Isolation Relays (AIRs) separating the batteries from the car in case of an emergency. The accumulator itself is

composed of A123 20Ah pouch cells arranged in custom built modules held inside a kevlar composite accumulator housing, along with a variety of safety and monitoring circuits.

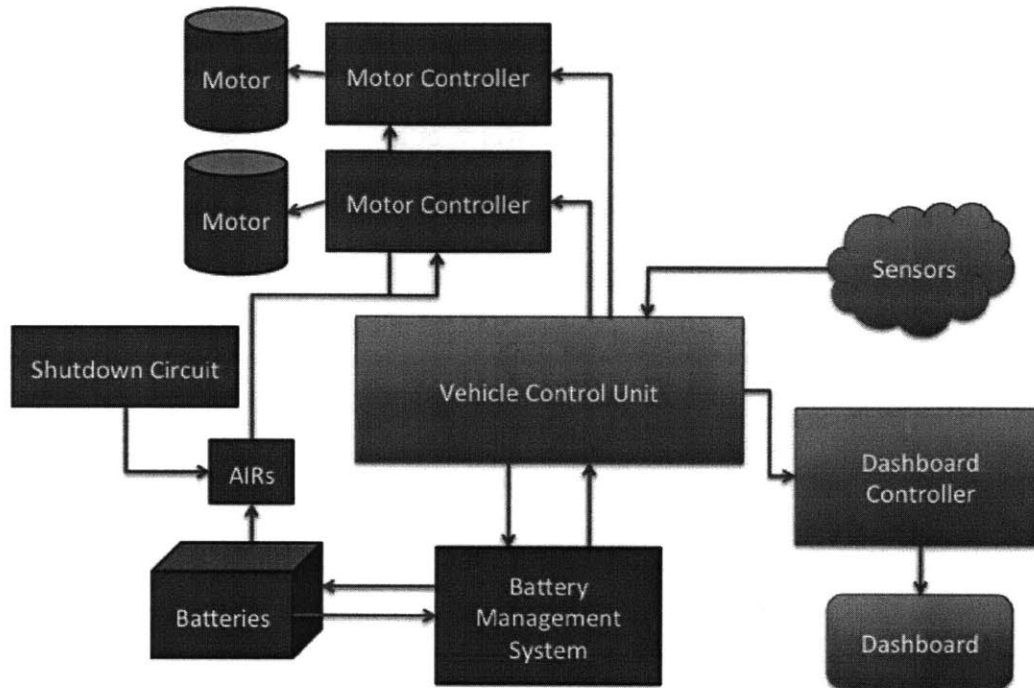


Figure 2.1: Block diagram of the Tractive System [2]

2.1 Functional Requirements

Functional requirements for the tractive system were established based on the events the car participates in during competition. These events are divided into two groups: static events which cover the design of the car, and dynamic events, which cover the performance. Our functional requirements were primarily drawn from desire to perform well over the whole range of dynamic events, combined with limitations placed on the car by the rules. The dynamic events involve an acceleration event (fastest time to travel 100m), a skid-pad event (car's ability to corner in a figure-8), autocross (fastest time for a 1km long racecourse, tests acceleration, handling, and driving skill), and endurance (20 laps of the 1km long autocross course). Of these events, the endurance race and the acceleration both test the performance of

the accumulator (skid-pad and autocross are more concerned with the general handling of the car, which while it of course involves accumulator performance, is greatly influenced by a large number of other factors). In order to perform optimally in acceleration, the car must be able to draw the maximum amount of power allowed by the rules from the accumulator (85kW). In order to perform optimally in endurance (or perform at all, for that matter), the accumulator must contain enough energy to power the car over the 20km course. Additional functional requirements are derived by looking towards the design competition, rules, and ease of assembly of the pack. To summarize the requirements we developed, in order of importance:

Functional Requirements of Accumulator:

- 1.** Supply the maximum allowable power at the desired voltage to the motors
- 2.** Store enough energy to power the car for the entire endurance event
- 3.** Comply with all the safety regulations laid out in the rules
- 4.** Not overheat during any event under reasonable conditions
- 5.** Not add significant weight above that of the bare cells
- 6.** Able to be disassembled quickly and non-destructively

After establishing these functional requirements, all design work on the accumulator was undertaken to address specific items in this list.

2.2 Safety

As the battery pack operates at 300 Volts and can output over 1000 Amps during a short circuit condition, resulting in instantaneous power dumps on the order of 0.3 MW, pack safety is of utmost importance. A series of safety devices and sensors, designed in accordance with the rules, were integrated into the pack. The primary safety device is the Shutdown Circuit, designed by Brian Sennett (MIT '13). This takes input from a variety of sensors, and if activated, will open two Kilovac EV200 relays mounted at both poles of the accumulator. These

contactors are capable of breaking 2000 Amps at 600V; they are the main panic switch for the car in case of an emergency. The positive pole of the pack is fitted with a fast-acting semiconductor fuse that would blow instantaneously in the case of a full pack short. The fuse is sized so that it is the first of any element to blow during an overcurrent event. The fuse is just a backup, however. The Battery Management System constantly monitors the state of the accumulator, including the current, and can tell the shutdown circuit to open the relays in case of an emergency. These devices, along with a variety of others described in the rules, make sure the accumulator is only operating under safe conditions.

In addition to the physical safety devices, the team has put a lot of work into the way we think about high voltage safety. Every conceivable failure mode of the pack (and the entire car, for that matter) has been documented in a Failure Mode Effect Analysis document, as dictated by the rules, along with the proper responses to those failures. In order to work in close proximity to the accumulators or car when the accumulators are installed, team members must first take and pass the same high voltage safety course and test as Tesla powertrain engineers must. Part of this course is a description of the proper safety equipment needed to work on the accumulator – our team has several sets of the correct equipment, and this equipment must be worn while working on the accumulator. It is our hope that a combination of good safety design and good safety consciousness will prevent any hazardous situations long before they become dangerous.

2.3 Power Density vs. Energy Density

In the world of energy storage, there is an inherent tradeoff between specific power and specific energy, or in other words, how much power a cell can output for a given weight, which correlates to an Electric Vehicle's (EV's) acceleration, and how much total energy it can store, which correlates to an EV's range. This trade-off is especially noticeable when looking at

Lithium-Ion battery chemistries. Figure 2.2 shows a diagram of specific power plotted against specific energy of different methods of energy storage. The Lithium-ion chemistries have the largest span; note that any particular make of lithium-ion cell inhabits a point on the graph, and the red area represents the total range of available cell types of various lithium-ion chemistries.

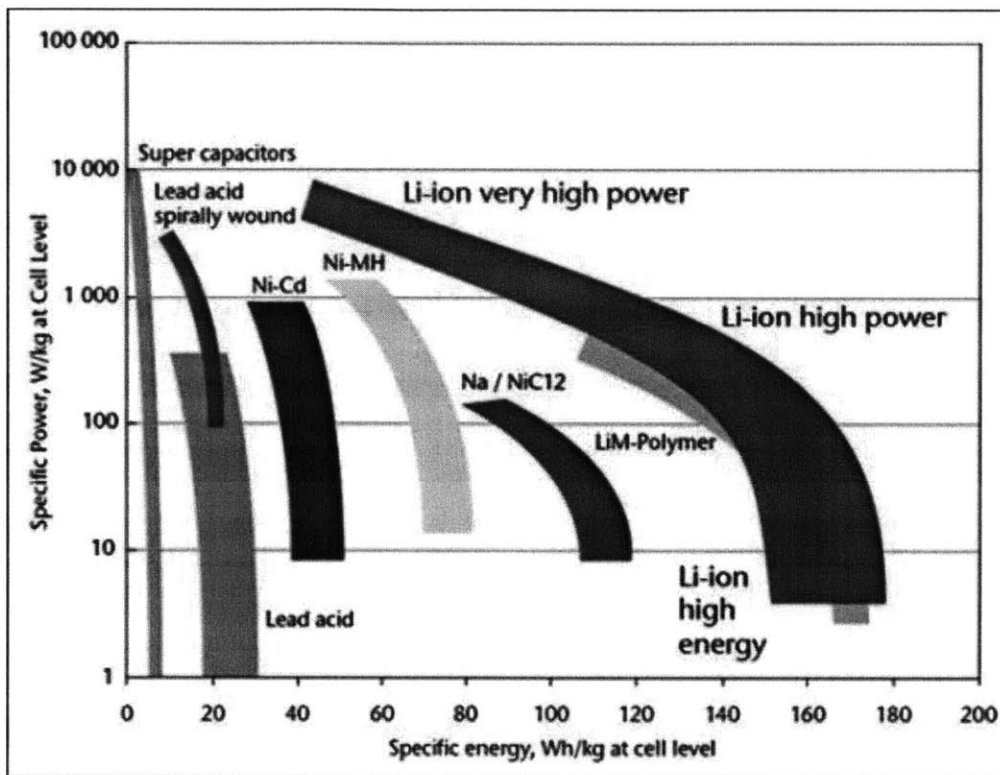


Figure 2.2: Power density vs Energy Density [3]

From this, it becomes obvious that optimal performance comes from choosing a cell chemistry with a balance of power density and energy density best suited to the needs of the car. Adding to the maximum power output of an accumulator adds weight, and adding to the maximum energy storage of an accumulator adds weight. A pack composed of the ideal cell chemistry for any specific application reaches its desired power output at the same weight it reaches its desired energy storage. Although in many EV applications it can be tricky to find this balance, the strict usage cases of a Formula SAE Electric car (ie, just the well-characterized

dynamic events) make this easy to find. The process used to choose our chemistry and pack size is described in Section 3.1

2.4 Track Model

In order to properly design the battery pack, the power and energy demands of the car during a race must be well understood. As the car was not yet built, this characterization was done with a full track model designed and implemented by Sammy Khalifa (MIT '12). The track model utilizes an iterative model of the car starting with a velocity of zero. The torque-speed curve of the motor is used to determine the torque being applied at the wheels at a given time. That torque is used to solve for the next velocity of the car using Euler's method. The maximum velocity that the car is allowed to travel at a given point is defined by the instantaneous radius of curvature of the simulated track and what is known as the traction circle. The traction circle is a plot of the maximum longitudinal traction that the tires can apply to the ground for a given lateral acceleration. If the velocity of the wheel is greater than the maximum allowed velocity the wheel "slips" and the car's velocity remains at the maximum allowed velocity. Energy and power changes are calculated directly from the output velocities, taking into account the appropriate efficiencies.

Figure 2.3 shows the power draw from the accumulator of the car over one simulated lap of the endurance race. Note that the current frequently becomes negative – this is because during braking, our motors and motor controllers switch into a regenerative braking mode to recover power, modeled as a negative power loss at the pack. The amount of power fed back into the accumulators is controlled by the motor controller. In this simulation, only 10kW of power was input back into the motors.




Figure 2.3: Simulated current draw over one lap of the endurance race

By integrating the power over 20 laps of this virtual track, we can calculate total energy used by the car over the endurance race. Additionally, we know the instantaneous current draw over the entire race, which is very useful in creating thermal and electrical models of the pack and tractive system components

3.0 Accumulator Design

3.1 Chemistry Selection

In order to meet our first 2 functional requirements, which set performance goals in terms of power and energy storage, the appropriate size and chemistry had to be chosen for the pack. To meet our desired performance goals, we knew that the pack had to output 85kW

and operate as close to 300V as possible (both numbers are the upper limits as allowed by the rules). Running our track model and integrating the power results over 20 simulated laps, we found that we needed approximately 5.1kWh of energy storage. These three numbers, 85kW, 300V, and 5.1kWh, provide enough information to size the pack.

Our approach was to make two simulated packs out of each different cell chemistry. The first pack was sized to be able to output 85kW (by dividing 85kW by the specific power of the chemistry), and the second pack was sized to be able to store 5.1kWh (by dividing 5.1kWh by the specific energy of the chemistry). Whichever of these two packs was the biggest was the minimum pack size needed to meet both the energy and power requirements. By comparing this minimum pack size across a range of cell technologies initially chosen for their availability, we found that packs composed of the Lithium Iron Phosphate (LiFeP04) pouch cells produced by A123 were the lightest, primarily due to their extremely high power density. Upon further investigation, a number of other advantages of the A123 LiFeP04 cells became apparent. These included very low internal resistance (meaning the cells produce little heat during charging and discharging), safer operation due to inherent resistance to thermal runaway (a dangerous phenomenon associated with many lithium chemistries), laxer safety regulations mandated by the FSAE Electric rules, and close company ties to MIT, which we hoped would aid in the procurement of cells or pre-built packs free of charge.

3.2 Pack Sizing

Cells can be added to a pack either in series or in parallel. In either case, the total energy capacity of the pack increases with every cell added. However, adding a cell in series increases the voltage by the voltage of the cell; adding a cell in parallel does not. Groups of cells in series all experience the same current, whereas groups of cells in parallel all experience the same voltage. This means that parallel cell groups will passively balance each other to the

same voltage, which aids in monitoring, as the voltage of the entire parallel cell block is representative of all the cells in it. This same effect, however, is potentially dangerous – if a cell in a parallel block is damaged, the cells in parallel with it can potentially dump all of their energy into it in an attempt to balance its voltage, which can result in catastrophic pack failure (fire, explosions, etc). As a result, packs with large parallel blocks need to be carefully designed so that, in the event of a faulty cell, fuses protect the parallel bank from auto-discharging into the bad cell.

The general methodology of sizing a pack is quite simple: 1) Increase the series count of the pack until the desired voltage is reached. 2) Increase the parallel count of the pack until the desired energy capacity is reached. This math is trivially easy; the tricky part is finding cells with the right capacity so that, after the above process is complete, you land close to your desired capacity. To demonstrate the problem, consider the example of trying to build a 7kWh hour, 300V pack out of 3.3V, 20Ah cells. Following the above steps, one would first add cells in series until we reached 300V, which would be at 90 cells. Next, calculate the capacity of 90 cells, which is 6kWh – less than our 7kWh goal. In order to reach the desired 7kWh capacity, one would now have to add a cell in parallel to get a 90s2p (shorthand for 90 cells in series, 2 cells in parallel) configuration, ending up with a total pack size of 12kWh – nearly twice as much battery (and weight) as is actually needed.

In the end, the sizing of our pack was effectively chosen for us by a generous donation of A123 20Ah M1 pouch cells to MIT Motorsports by Professor Ian Hunter. Going through the pack sizing process with the minimum pack size of 5.1kWh calculated by our track model and the 3.65V max charge of the 20Ah A123 cells, we found that a pack configuration of 82s1p provided us with 300V and 5.3kWh of energy. Noting that the track model was never fully validated, the proximity of 5.3kWh to our absolute calculated minimum of 5.1kWh was alarming. To ease our margins, we took advantage of the fact that, especially in LiFeP04 chemistries, very little of the total energy of the cell is stored at the maximum voltage. This can be seen in Figure 3.1, a discharge profile of one of our typically performing A123 20Ah cells, by noticing that the voltage falls below 3.55V almost immediately. By dividing the area under the curve between 3.65 and 3.55V by the total area under of the curve, we find that only about 0.5% of the total cell energy is contained within this voltage range.




Figure 3.1: A123 20Ah Pouch Cell Discharge Profile at 10A

This means that from a capacity standpoint, it is worth it to artificially limit the voltage our cells can charge to at 3.55V, rather than going all the way to the 3.65V they can be charged to. The new 3.55V max allows a pack configuration of 84S1P, a capacity gain of 129 Wh, while only losing 27Wh from the decrease in voltage of every cell. There are also significant advantages to a cell count evenly divisible by 3: due to FSAE Electric rules, any given "cell stack" cannot exceed 120V. Our 84S1P cells can now be divided into 3 identical modules connected in series by maintenance plugs, in accordance with the rules. This 28s1p module architecture is the backbone of our accumulator design.

3.3 Tab connections

To link the pack in series, the positive tab of a pouch cell is connected to the negative tab of the next cell. Most production-scale packs have their cell tabs laser or spot welded together directly in the desired arrangement. The primary advantages of this process are scalability and low weight; the disadvantages, however, are massive initial cost and setup, and inability to non-destructively disassemble the cells. As a primary goal of the design of this battery pack was ease of disassembly, tab welding was not considered as an option for cell connection.

The alternative we considered was mechanically clamping the tabs together. As the tab spacing in our packs is very tight, traditional bolted approaches were difficult to implement. To rectify this problem, we designed a tab clamping system that presses the two tabs together between an aluminum wedge and a Hydlar Z (kevlar reinforced nylon) clamping plate with the same angle on one face. This general approach can be seen in Figure 3.2:




Figure 3.2: Tab Clamping Overview

By clamping the wedge down to the Hydlar Z piece with a clamping bar, we can achieve a mechanical clamping advantage of $1/(\sin\Phi)$, where Φ is the angle of the wedge. After some experimentation, 15 degrees was chosen as a good balance between high mechanical advantage and large angle that the thickness tolerances don't need to be extreme to prevent the wedge from slipping through. At 15 degrees, there is a 3.86X mechanical advantage – a 1000N force on the wedge from the clamping bars is transformed into 3860N of force clamping the tabs together. Additionally, the way the aluminum is constrained lets it distribute this force more evenly over the surface than it would using a conventional bolt system.

3.3.1 Clamping Plate

The clamping plate is the critical structural piece of the clamping system. FEA was used to make sure the piece could withstand the loads placed on it by the wedge and clamp bar, and an iterative design/FEA process was used to minimize weight while still achieving the required strength. The FEA results can be seen in Figure 3.3.

Model name: Tab_clamp
Study name: Study 1
Plot type: Factor of Safety Factor of Safety1
Criterion: Automatic
Factor of safety distribution: Min FOS = 1.1




Figure 3.3: FEA of the Clamping Plate

There is some concern with the Factor of safety of 1.1 – however, after careful analysis of the areas at 1.1, it seems to be odd meshing effects around areas where contours meet, and not reflective of true stress concentrations. As the part has to be an insulator while still withstanding high loads, Hydlar Z was chosen – it is a kevlar reinforced plastic with a yield strength of 17000 psi – nearly twice as high as normal nylon. The plastic is also known for its dimensional stability, excellent thermal properties, and relative ease of machining.

3.3.2 Clamping Wedges

The clamping wedges are designed to evenly distribute pressure over the cell tabs. They are made of aluminum for stiffness to accomplish this and to add thermal mass to the tabs to stabilize their temperatures (ie. avoid high temperature spikes) during the rapid pulse loading typical of our power usage profiles from the track model. FEA was used to determine the optimum position along the clamping wedge that the load from the clamping bars should be

applied to attain the most uniform clamping force; the final deformation (highly exaggerated) is shown in Figure 3.4:




Figure 3.4: Exaggerated Wedge Displacement Under Load

The deformed surface clamping the tabs has a flatness of about 0.0015mm – definitely within the tolerances of machining the wedge or clamping plate, and small enough to be totally absorbed by elastic deformation of the clamping plate.

3.3.3 Clamping Bars

The clamping bars were designed to transfer the force from threaded inserts in the clamping plate to the top of the aluminum clamping wedges. By controlling the torque put on the bolts holding the clamping bar down, the downward force on the clamping wedges can be fairly tightly controlled. The bars were designed iteratively with FEA to minimize weight while maintaining strength and stiffness. The results of the FEA can be seen in Figure 3.5.

Model name: boltstrip
Study name: Study 2
Plot type: Factor of Safety Factor of Safety1
Criterion : Automatic
Factor of safety distribution: Min FOS = 1.7




Figure 3.5: FEA of the Clamping Bar Under Load

Hydlar Z was chosen for the clamping bars for the same reason as the clamping plates – strength and thermal stability. The clamping bars are designed to be waterjet to reduce manufacturing time and cost. This also allows us to have several different versions of the clamping bars easily produced, which is important as each clamping bar has a slightly different pattern to which the BMS PCB, discussed in Section 3.5, is bolted.

3.4 Pack architecture

Our overall module architecture is guided by the design of A123's own prebuilt modules. This manifested itself in two ways: the first way is in the compression of the cells. A123 recommends placing their cells under approximately 10 psi of pressure to reduce swelling, which can damage the internal resistance and capacity of the cells. The clamping system we designed can be seen in Figure 3.6. The other way is in the separation of the cells within the 28s module.




Figure 3.6: Full 28s Module Architecture

3.4.1 Cell Separation

Following A123's lead, each cell is contact with a thin layer of fire-retardant neoprene on one side, and a high purity aluminum heat spreading plate on the other. The neoprene absorbs small imperfections in cell flatness to make sure pressure is evenly distributed across the entire cell surface, and aids in the vibration damping characteristics of the modules. The aluminum plates conduct heat out from the center of the module to the edges; the plates are bent 90 degrees at the edges to provide mechanical protection for the sides of the cells. A cross section of the pack can be seen in Figure 3.7: the alternating cell-aluminum-cell-neoprene pattern can be seen, as well as the protective effect of bending the heat spreader inwards.




Figure 3.7: Cross Section Showing Cell Separation

3.4.2 Endplate Design

Endplates were designed with two goals in mind. They need to be able to withstand the forces from the compressive load with minimal deflection to make sure pressure is evenly distributed on the cells, and they need to constrain the module in the event of a crash, which the rules define as a 20g lateral and 10g vertical impulse [1]. To meet these goals in the lightest possible manner, the endplate is split into two components, a rigid, flat composite plate made out of a kevlar-honeycomb composite and designed to spread the compressive loading force evenly over the cells, and a waterjet aluminum plate designed to transfer the compressive load to the composite plate and to bolt the modules down strongly enough to withstand the

required impact loading. FEA was run on the aluminum endplates to minimize their weight while maintaining their ability to withstand the high loads placed on them. These results can be seen in Figure 3.8.




Figure 3.8: Endplate FEA under 20g Crash Loading Conditions

3.5 Battery Management System

The Battery Management System (BMS) is a critical part of the accumulator, and one of the most important safety devices in the entire car. Its main purpose is to constantly monitor the inputs and the outputs of the accumulator in order to have a clear picture of the current state of the accumulator. It has sensors that measure the voltage and temperature of every cell in the pack, and the current flow through the entire pack. It is hardwired into the shutdown circuit – if any parameters exceed the operating bounds, it opens the AIRs and forces the car to shut down. The other important function of the BMS is cell balancing – if any cell goes above or below the operating voltage ranges, permanent damage can occur. As a result, the entire pack

is shutdown during a discharge when the lowest cell is at the low voltage cutoff (LVC) and during a charge when the highest cell reaches the high voltage cutoff (HVC). Consequently, any imbalance of the pack can drastically and artificially reduce the capacity of the whole pack. Our BMS uses dissipative balancing during charging cycles – in other words, it discharges cells with comparatively high voltage through built in shunt resistors.

There are two major design philosophies used in BMSs: centralized and distributed. In a centralized system, voltage sense wires from every cell run to a single box, where the balancing and computation is done. A distributed system does the balancing and measuring on small circuit boards distributed through the pack; typically each circuit board handles only one cell, and a single module or series of modules monitors the all cell boards. We chose to avoid the centralized BMS for two reasons. The first is the hassle of routing so many wires – as every cell needs two sense lines running to it, the numbers can get out of hand quickly. The second is the inherent danger of running two small wires with the full potential of the pack right next to each other. A small nick in the protective coating of the wire or a slip of the hand could cause an extremely dangerous arc flash. An MIT student team had this exact scenario happen several years ago, and the resulting arc caused serious injury to a student. With all of this in mind, we chose a distributed BMS system for the 2013 car.

3.5.1 Electrical Layout

Our accumulator employs a fully distributed BMS system called EMUS, manufactured by Elektromotus. There is a circuit board, called a cell board, for each of the 84 cells in the pack. Each cell board contains voltage sense lines to monitor the cell, a shunt resistor for balancing, and a thermistor to measure cell temperature. These boards are linked in series to each other and communicate at each end of the stack, through opto-isolators, to a module which runs all of the computations and talks to the rest of the car via CAN communication. Additionally, to

prevent the loss of an entire cell board in the case of a short circuit, fuses are placed on the board every 7 cell boards apart. These fuses will protect the vast majority of the BMS in case of a short on the communication line, which would have the potential to ruin all 84 cell boards without them. Figure 3.9 shows a block diagram layout of the BMS cell boards.

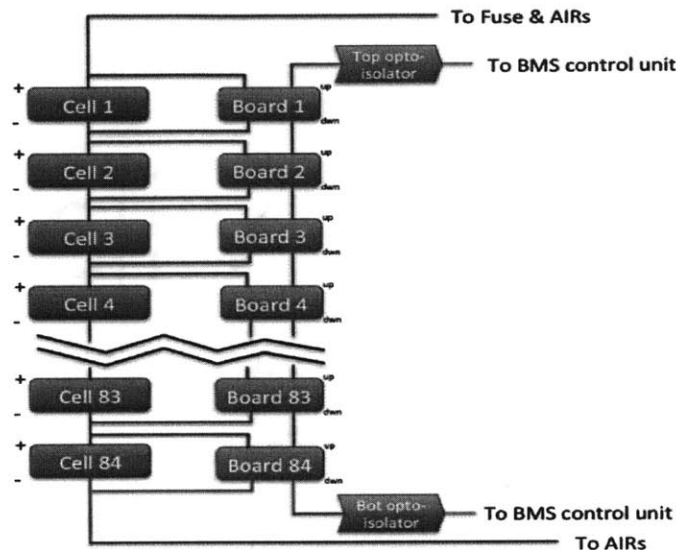


Figure 3.9: BMS Electrical Block Diagram [2]

3.5.2 Mechanical Layout

In order to place the cells as close as possible to the tabs without significantly modifying our tab clamping method, 28 cell boards are mounted on a custom designed PCB, which sits on top of the 28s modules. This PCB can be seen in green sitting on top of a 28s module in Figure 3.10.




Figure 3.10: 28s Module with BMS Board

Working towards the goal of quick and easy disassembly of the pack, the PCB is held to the pack mechanically with only 6 screws. When screwed down, small spring loaded pins (not pictured in the CAD) extend down and contact the surface of the aluminum clamping bar: the spring loaded contacts make the electrical connection between the voltage sense wires of the cell board and the tabs of the cells. This makes assembly and disassembly times virtually non-existent, critical for an application we expect to require a large amount of debugging time.

4.0 Thermal Analysis

The thermal characteristics of the accumulator are critical to its safety, performance, and longevity. Quantifying the temperature rise of the pack during competition allows us to answer two major questions about the pack architecture. The first is whether or not the aluminum heat

plates are necessary to sufficiently cool the pack. Although they appear at first glance to be a primary component of the packs thermal dissipation, they add significant weight to the overall pack architecture; if we can maintain proper thermal performance of the pack without them, they are definitely worth removing. The second question is what mass flow of air we need to blow through the accumulator container in order to maintain the pack at temperature levels acceptable for the cell chemistry - under 55°C for discharging [4].

4.1 Heat Generation

For the purposes of quantifying heat generation of a battery pack, cells can be modeled as a resistor in series with a voltage source, as seen in Figure 4.1.




Figure 4.1: Basic Cell Model [5]

As a current I flows through the cell, the internal resistance ($R_{internal}$) causes ohmic heating of the cell following the equation:

$$\dot{q}_{cell} = I^2 R_{internal}$$

The primary cause of heat generation in a battery pack is the ohmic heating of the cells caused by their internal resistance, and to a much smaller extent, the internal and contact resistance of the wiring and connections. In order to quantify the internal resistance of our cells, as well as their capacity and tab clamping resistance of the wedge tab-clamping method, a 6s1p test pack was designed and built around the same principles the large 28s1p pack uses. Figure 4.2 shows

the test setup: a PowerLab 6 charger monitors the voltage and internal resistance (among a host of other metrics) of every cell in the pack during charging and discharging, and balances every cell during the charge cycles.




Figure 4.2: Test Setup for Measuring Internal Resistance

Julia Kimmerly (MIT '13) used this setup to run full charge/discharge cycles on all 90 of our cells. The result was an average internal resistance of 1.67 mOhms with a standard deviation of 0.595 mOhms. This number includes both the internal resistance of the cells and the resistance of their contact to the aluminum block, which should be roughly equal to the contact resistance of the tabs to each other. This should give us a more complete picture of heat generation in the clamped tab- cell system than just the internal resistance of the cell.

4.2 Adiabatic Model

Lumped parameter models have been proven highly accurate at modeling battery cells; as the cells are made up of thin layers of highly thermally conductive material pressed tightly together, the fact that this model works well make sense. If we take the cell as a lumped

thermal mass with mass m and specific heat C_p , we find that the energy stored in the cell,

q_{stored} , can be found with:

$$q_{stored} = mC_p\Delta T$$

A first order estimate of the thermal behavior of the pack can be obtained by making a no cooling assumption in which all of the ohmic heat generation is absorbed by the cells [5]:

$$q_{stored} = \int \dot{q}_{cell} dt = \int_0^{t_{final}} I(t)^2 R_{pack} dt$$

And $I(t)$ is the current draw at the pack given to us by the track model. Performing this integration numerically, we find the total ΔT to be 75°C over an endurance race.

As a 75°C rise puts us out of the allowable operating temperature range of the cells, it is likely we will have to air cool the entire pack. It is also beneficial to calculate the amount of mass flow required to keep the pack at thermal equilibrium. We can do this using the equation [6]:

$$\dot{m} = \frac{\dot{q}_{pack}}{C_p\Delta T}$$

Where C_p is the specific heat of air, ΔT is the allowable temperature change of the air, derived from the maximum allowable temperature of the pack and the initial air temperature, and \dot{q}_{pack} is the total heat generation of the pack. By dividing q_{stored} by the total time elapsed in the race, we find the average heat power generation of the pack. For all 84 cells, we have about 1.2kW of heat generation. As our simple model assumes the air flow removes all of the generated heat, this quantity equals \dot{q}_{pack} . Plugging everything in, we find the mass flow

required to remove the heat generation and stay within our thermal limits is about 60 grams/s, or about 120 CFM.

4.3 Cooling Model

The results from our initial first pass calculation are instructive, but in need of serious refinement. One major oversimplification is evident: it assumes all of the heat generated in the pack is removed by air. In actuality, only part of the generated heat leaves the pack; the rest is stored in the cell as thermal energy [5].

$$q_{gen} = q_{stored} + \int_0^{t_{final}} \dot{q}_{cooling} dt$$

Heat escapes from the cell in 4 major fashions – through the faces of the cells, through the sides of the cells, through the top and bottom of the cells, and through the tabs. Based on the way the cells are arranged in the 28s modules (with insulating plates on the top and bottom), we can make adiabatic assumptions for one of the faces and for the top and bottom of the cells. This leaves us with only three paths for heat to escape through: through one face of the pack to an aluminum heat spreader to the environment, through the sides of the pack to the environment, and through the tabs into the environment. By modeling each of these heat paths as branches of a resistive circuit in parallel with each other, the total $\dot{q}_{cooling}$ can be calculated.

The first branch of the resistive circuit, representative of heat flow into the heat spreaders and out to the environment, has two thermal resistances in series. The first is from thermal conduction from the cell into the non-protruding body of the heat-spreader. The second is from convective heat loss in the protruding part of the heat spreader, modeled as a convective-tipped fin. As there is air from the cooling fans flowing over the heat spreaders, it is possible that forced convection, rather than natural convection, dominates the heat transfer of

the fin. Determining the dominant form of convection can be done by calculating the Archimedes number, a ratio of the Grashof and Reynolds number [7]. These numbers are calculated with 120CFM and values at the median points of expected pack conditions.

$$Ar = \frac{Gr}{Re^2}$$

If $Ar \gg 1$, then natural convection dominates. If $Ar \ll 1$, then forced convection dominates. Plugging in our median values, we find that for the plates, Ar is around 100, indicating natural convection is dominant.

The second branch represents the heat escaping through the sides of the pack, and is modeled as a single thermal resistance dominated by convection along an idealized, flat, vertical surface the width and height of the cell. Due to the geometry of the pack, there is not significant air flow experienced by the sides of the cells. Consequently, natural convection was assumed to be the dominant mode of heat transfer. As each cell has two sides, a value of 0.5 was factored into the final resistance value.

The third branch of the resistive circuit models the tabs as vertical fins protruding from the pack. Again, due to the geometry of the pack, significant airflow is not expected over the fins. As a result, natural convection was assumed. Additionally, as each cell contains two tabs, the same factor of 0.5 was again multiplied into the final resistance of the branch.

For calculation of the convective heat transfer coefficient, the relevant dimensionless quantities were calculated with static, median values in our expected range. They were subsequently used to determine which equations would be used in the transient solver. A brief overview of the relevant dimensionless numbers and equations can be found in Appendix 6.1

By comparing the resistance values of each branch at our static median pack conditions, we can get a high level sense of what pathways dominate heat transfer from the pack to the

surrounding air. Table 4.1 compares the value of these resistances, and the overall resistance found by adding the three branches in series. Resistance is given in K/W.

Table 4.1: Heat Transfer Path Resistance

Cell Sides	Heat spreaders	Tabs	Total
6.12	1.60	48.1	1.23

These values indicate that the heat spreaders are the dominant mode of heat transfer out of the pack, but also that neither of the other two resistances are large enough to be outright discarded. As a result, all three numbers were calculated and factored in when solving for the transient heat transfer of the pack. The general approach to solving for this transient behavior was to find the heat energy generated by the pack over a discrete time interval; part of this energy was removed through the heat transfer mechanisms already discussed, and the rest went into the heating the pack. A discrete temperature rise is calculated for each time step, and this, along with the amount of energy dissipated from the pack, is recorded. Note that the model assumes that the temperature of the air in the pack does not rise; this necessitates that all of the energy dissipated from the pack into the surrounding air is removed by forced air cooling of the pack enclosure.

4.4 Results

The transient solver was first used to evaluate the pack using the adiabatic model discussed in Section 4.2, and compare it to cooling model established in Section 4.3. Figure 4.3 shows the total temperature rise in the pack over the course of the endurance rise. The fact that the temperature rise seen in the curve for the adiabatic model closely matches that the temperature rise seen in the adiabatic transient solution previously calculated indicates our transient solver is working correctly.

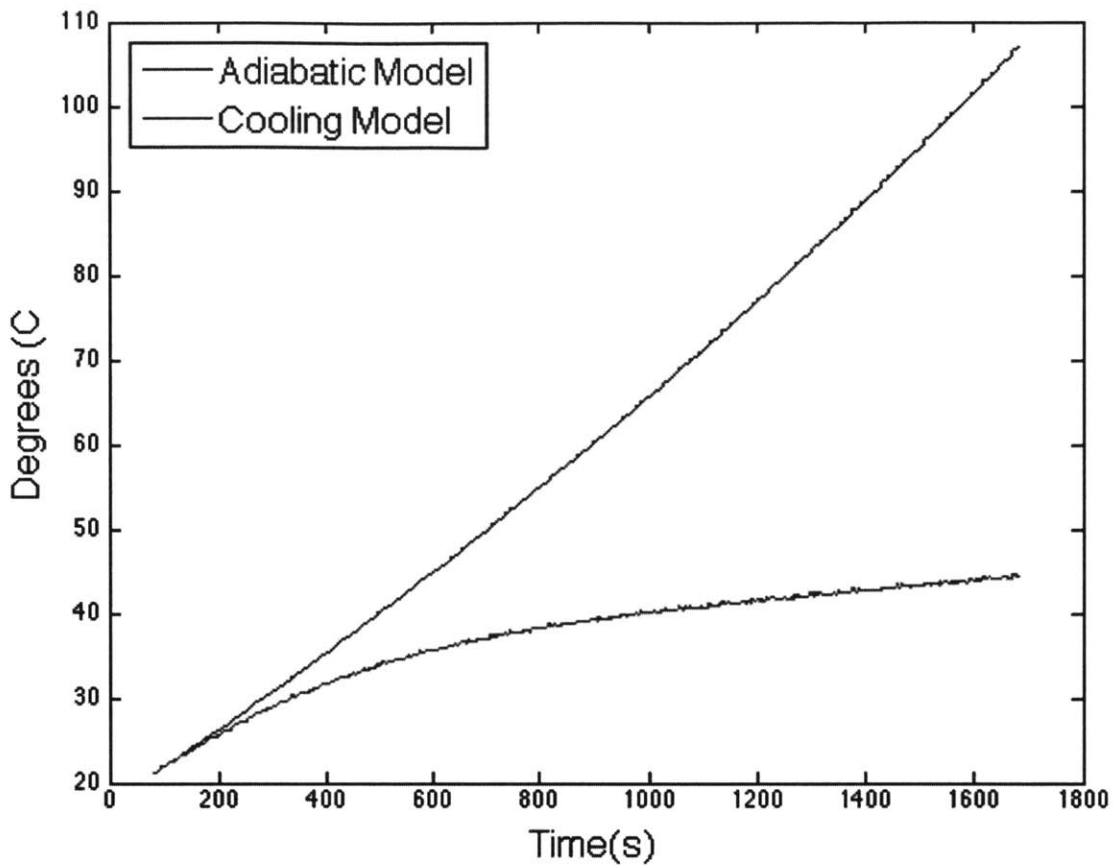


Figure 4.3: Transient Pack Temperature over an Endurance Race

The major takeaway from this graph is that, with reasonable quantities of air cooling, we can keep the pack within the operating temperature parameters of the cells. One of the main questions we wanted to answer was whether or not we can do this without the aluminum heat spreaders. Figure 4.4 shows plots of the cooling model with and without the aluminum heat spreaders included in the heat transfer pathways.

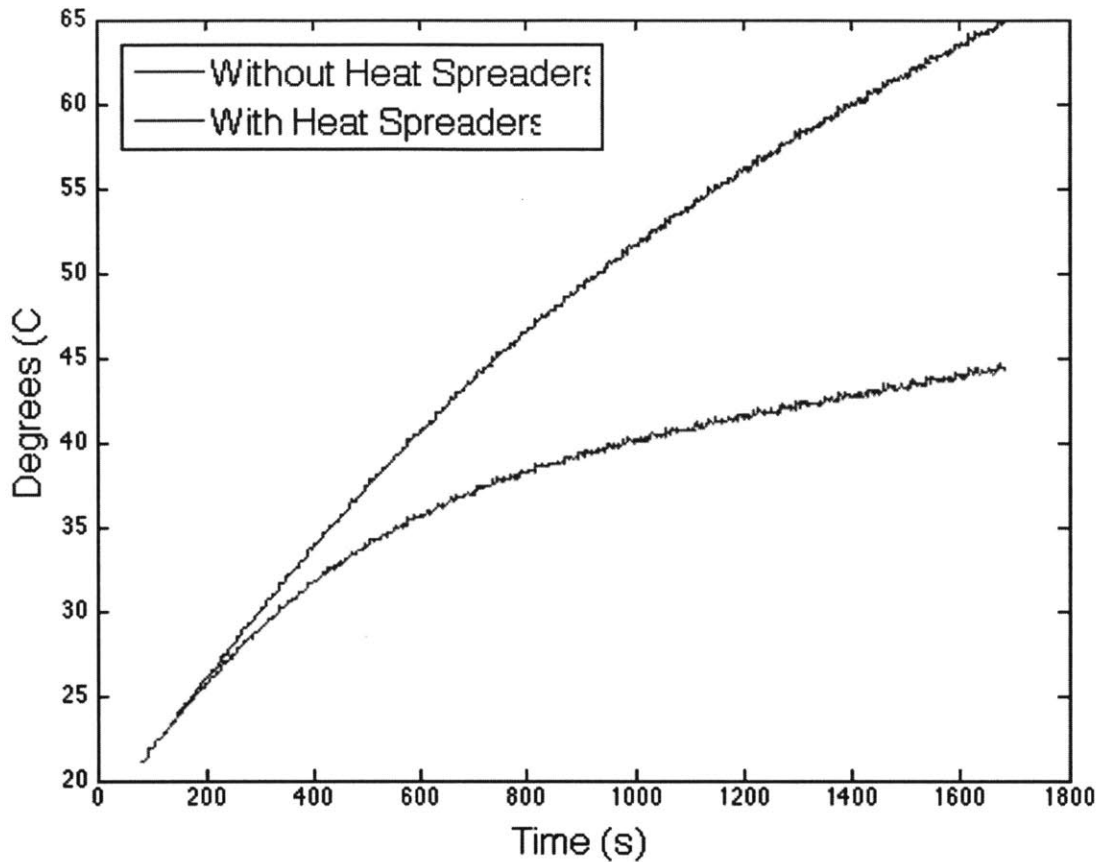


Figure 4.4: Pack Temperature Rise With and Without Heat Spreader

This plot indicates that the heat spreaders are necessary in keeping the pack in its operating temperature range, under 55°C. Overall, a 20°C temperature rise was noted when the heat spreaders were removed - a significant hit to thermal performance. This answers our first major question: we can conclude from Figure 4.4 that the heat spreaders are a necessary part of our pack design.

The other question we want to answer with the thermal model is how much heat we're going to have to remove from the pack with forced air cooling. As this heat removal is done with fans, a more relevant question to the design of the pack is how many CFM (cubic feet per minute) of airflow we need through the pack. Our initial estimates indicated around 120 CFM would be needed to cool the pack. With our transient solver and cooling model, however, we can refine this number and remove the best-case temperature difference assumption the first model forced us to use. Figure 4.5 shows a plot of the total heat transfer to the pack from the air. We can use this heat transfer, combined with the transient temperature of the pack, to calculate the required airflow in CFM, also shown in Figure 4.5.

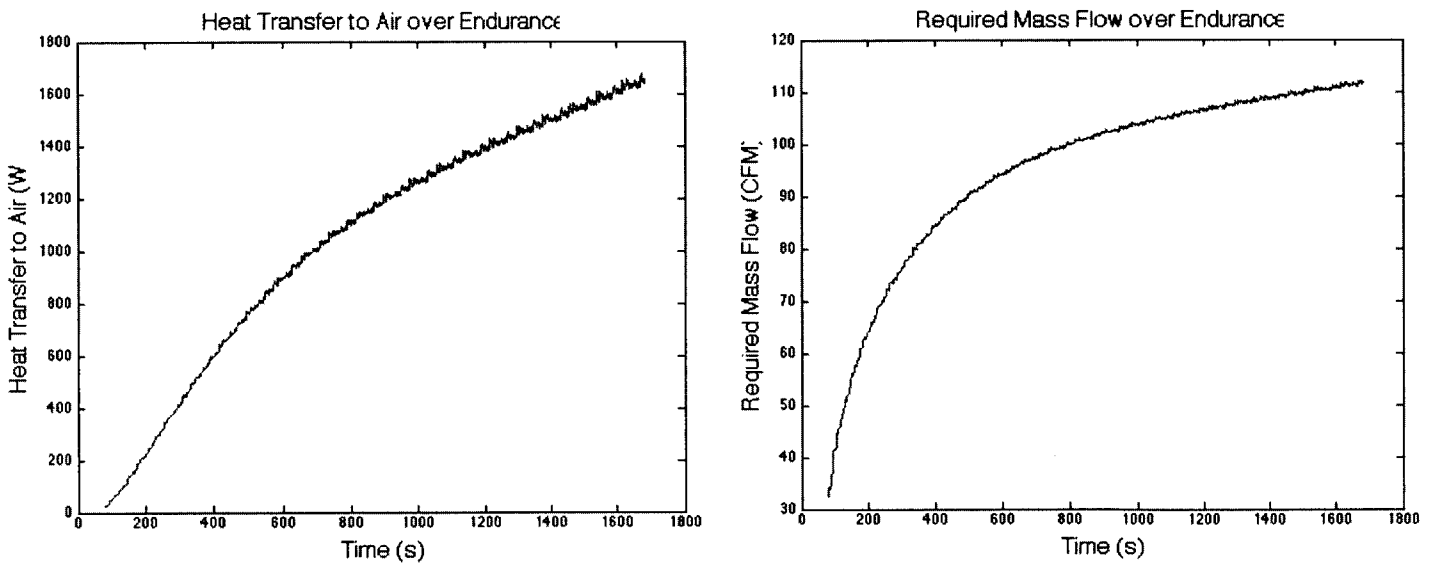


Figure 4.5: Fan Cooling Performance Demands

This data will be used to calculate fan speeds, which correlate to energy demands on the low voltage supply of the racecar. Accurately modeling these loads is important in determining the size of the low voltage battery needed to power the fans and electronics.

5.0 Conclusion

The accumulator was designed to meet a number of functional requirements informed by the rules and our desired car performance. They were:

1. Supply the maximum allowable power at the desired voltage to the motors
2. Store enough energy to power the car for the entire endurance event
3. Comply with all the safety regulations laid out in the rules
4. Not overheat during any event under reasonable conditions
5. Not add significant weight onto that of the bare cells
6. Able to be disassembled quickly and non-destructively

The first two functional requirements were addressed using the results of a track model to determine the light weight combination of cell chemistry and accumulator size and configuration that would meet the requirements.

Pack safety, while certainly not fully solved, was addressed with the inclusion of a multitude of safety devices into the accumulator, as well as a good deal of thought about pack failure modes and safety procedures to follow when building, maintaining, and using the pack.

A fairly extensive thermal model of the pack was developed to quantify the thermal performance of the accumulator and ensure that it was consistent with the requirement of never leaving the operating range of the cells during a race. This model, combined with extensive FEA on the major components of the pack, helped to make it clear where we could and could not remove weight, resulting in a pack with minimal unnecessary weight gain compared to the bare cells.

Finally, the development of novel attachment schemes for cell tabs and for the BMS board makes the final pack design rapidly disassemblable with no permanent modifications to the cells.

Significant work still needs to be done on the assembly and testing of these packs to verify their performance matches the models used. The goal of this thesis, however, was to provide future members of MIT Motorsports an overview of the process I went through in designing the accumulator for MITs 2013 car. I stress that this is not a comprehensive report, but rather an overview of the major thought processes that were used, as well as some of the implementation of some of these thought processes into a design. The final design, while carefully considered, certainly has room for improvements.

6.0 Appendix

6.1 Natural Convection

Relevant Dimensionless Numbers [8]:

	Reynolds (Re)	Grashof (Gr)	Prandtl (Pr)	Nusselt (Nu)
“Local”	$Re_x = \frac{U_\infty x}{\nu_f}$	$Gr_x = \frac{g\beta x^3 (T_0 - T_\infty)}{\nu_f^2}$	$Pr = \frac{\nu_f}{\alpha_f}$	$Nu_x = \frac{h_x x}{k_f}$
“Average”	$Re_L = \frac{U_\infty L}{\nu_f}$	$Gr_x = \frac{g\beta L^3 (T_0 - T_\infty)}{\nu_f^2}$		$Nu_L = \frac{hL}{k_f}$

Relevant Equations for the Nu_L , the “average” Nusselt number:

$10^4 < Gr Pr < 10^9$ $0.00835 < Pr < 1000$ but if $0.6 < Pr < 10$ use (8.18) also, check special cases (7.45b,c)	$10^9 < Gr Pr < 10^{12}$ works for most Pr but if $0.6 < Pr < 10$ use (8.21)
$\frac{Nu_L}{\sqrt[4]{0.25 Gr_L}} = \frac{0.902 Pr^{0.5}}{(0.861 + Pr)^{0.25}}$	$Nu_L = 0.246 Gr_L^{2/5} Pr^{7/15} (1 + 0.494 Pr^{2/3})^{-2/5}$

7.0 References

- [1] "2013 Formula SAE Rules," [Online document], n.d., From url <http://www.fsaeonline.com/content/2013%20FSAE%20Rules.pdf>
- [2] "MIT Motorsports Electrical Systems Form," B. Sennett, S. Reineman, et al, n.d. Not published, available at request from: fsae@mit.edu
- [3] International Energy Agency, *Technology Roadmaps: Electric and Plug-in Hybrid Electric Vehicles*, 2009, p. 12.
- [4] A123 Systems datasheet for Lithium Ion Prismatic Pouch Cell AMP20M1HD-A, 2011.
- [5] "Rapidly Charging Battery Systems," L. Rodgers, P. Karplus, R. Gogoana, M. Narwot, 2010, ASME 2010 International Design Engineering Technical Conference. From url: https://www2.lirimm.fr/lirimm/interne/BIBLI/CDROM/ROB/2010/DETC_2010/data/pdfs/trk-6/DETC2010-29226.pdf
- [6] "Heat Transfer Principles in Electronics Cooling," Class notes from Caird University class MPE 685, n.d, pp. 54-71 From url: <http://www.pathways.cu.edu.eg/ec/Text-PDF/Part%20B-8.pdf>
- [7] *Fundamentals of Heat and Mass Transfer*, F. Incropera and D. DeWitt, John Wiley & Sons; 6th edition, March 10, 2006, p. 565.
- [8] "Forced Convection and Natural Convection Equations," A. Powell, Class notes from MIT class 3.185: Transport Phenomena in Materials Engineering, 2002, from url: <http://dspace.mit.edu/bitstream/handle/1721.1/35739/2/3-185Fall-2002/NR/rdonlyres/Materials-Science-and-Engineering/3-185Transport-Phenomena-in-Materials-EngineeringFall2002/424CAB6E-0B2A-47BA-8FF4-EBD0F84E2AEE/0/NUSSELT3185.pdf>