

Internal Resistance Variances in Lithium-Ion Batteries and Implications in Manufacturing

By

Radu Gogoana

S.B. Mechanical Engineering

S.B. Management Science

Massachusetts Institute of Technology, 2010

SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN MECHANICAL ENGINEERING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

MAY 2012

[June 2012]

©2012 Radu Gogoana. All rights reserved.



The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author: _____
Department of Mechanical Engineering
May 23, 2012

Certified by: _____
Sanjay E. Sarma
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by: _____
David E. Hardt
Chairman, Department Committee on Graduate Theses

100

100

100

Internal Resistance Variances in Lithium-Ion Batteries and Implications in Manufacturing

by

Radu Gogoana

Submitted to the Department of Mechanical Engineering on May 23, 2012, in Partial Fulfillment of the Requirements for the Degree of Master of Science in Mechanical Engineering

Abstract

This thesis addresses issues in manufacturing that lead to cell DC internal resistance (DCIR) variance, provides an overview of generally accepted cell degradation mechanisms and modeling techniques associated with IR as a function of cycling, models the cycle life of a simple battery pack of two parallel-connected cells that start off as unbalanced, and validates it with experimental data from cycle life testing of parallel-connected LiFePO₄ cell groups.

Experimental results from samples of commercially available cells show that variance in the thickness of the electrode layer is correlated to differences in capacity and DCIR. In cycle-life testing, parallel-connected cell groups with larger differences in DCIR between the cells experienced faster cycle life degradation. The proper matching of DCIR values within a battery pack, relative to the designed C-rate capability of the pack, is important to ensuring maximum useful life of the battery pack. This is especially important for parallel-connected cell groups, where the current distribution to each cell is typically not monitored in order to reduce battery management system complexity.

Thesis Supervisor: Sanjay E. Sarma

Title: Professor, Mechanical Engineering

Acknowledgements

The network of people who have supported me in my education is too large to properly mention here; you know who you are, and understand that the endeavor would require another master's thesis. I'll stick to the people and organizations who were more closely involved with batteries.

I'd like to thank my advisor, Sanjay Sarma, for giving me an unprecedented amount of freedom in pursuing this project. I came into grad school two years ago saying "I want to work on battery manufacturing" and he's provided all of the resources I ever asked for in making it happen. Along with being a great academic advisor, he's been a great mentor. Thank you.

Next up is everyone who's been involved in the MIT Electric Vehicle Team. I was able to pursue ridiculous projects with very few barriers, learning more in the process than I ever could in a commercial setting or a classroom. Nobody ever really told me "no", at least in ways that I couldn't get around, and the work I did with EVT led me to pursue this thesis. I'd like to thank Steve Banzaert and Sandi Lipnoski from the Edgerton Center for dealing with us, Dan Frey for being our ever supportive advisor, and all of our sponsors for providing us with awesome equipment. To Lennon Rodgers, Erick Fuentes, Mark Jeunnette, Randall Briggs, all of the guys who were part of the eLEven crew, and Irene Berry for starting the whole organization.

To the guys in my lab group for helping me out on data processing, idea support and simply hanging around late into the night – special thanks to Dylan Erb, Austin Oehlerking, Erik Wilhelm, Rahul Bhattacharyya, Sumeet Kumar, Isaac Ehrenberg and Stephen Ho.

To Professor Yet Ming Chiang, for providing additional insight on the electrochemical side and donating tons of batteries (literally) and to Greg Tremelling of A123 Systems, for his continued support over the past four years in all things battery related, taking personal time off of work to help out a bunch of undergrads working on some crazy project (and then me, with this project into grad school). To Quinn Horn of Exponent, for providing expert battery consulting advice as well as making available the high-end battery testing resources at his company's labs. To all of the people in China and Taiwan who showed me around battery factories over the summer of 2011; if you read this, you know who you are and I appreciate your generosity.

To the guys at MITERS, for simply existing. MITERS is the true hard core of MIT geekery and badass projects at their finest. They are an inspiration for anyone wanting to build cool stuff on their own, from musical Tesla coils to overpowered electric shopping go-karts and five-horsepower EV kick-scooters. Big shout out to Shane Colton and Charles Guan.

To MIT SEVT for building EVs before they were cool and to FSAE for continuing the trend.

Lastly, to my parents, for supporting me to where I am now.

TABLE OF CONTENTS

Chapter 1	Introduction and project motivation	9
1.1	Hypothesis summary	9
1.2	Background and motivation	9
Chapter 2	Lithium-Ion Cell Fundamentals	10
2.1	Theory of Operation: The Lithium-Ion Battery	10
2.1.1	Anode	10
2.1.2	Cathode	12
2.1.3	Separator	13
2.1.4	Electrolyte	15
2.2	Cell Design and Mechanical Construction	15
2.2.1	Electrode Thickness vs. Cell Rate Capability	17
2.3	Cell Degradation Mechanisms	20
2.3.1	Degradation Mechanisms of Carbon Based anodes	21
2.3.2	Degradation Mechanisms of Lithium Metal Oxide Cathodes	22
Chapter 3	Cell Modeling	23
3.1	Review of Equivalent-circuit Based Models	24
3.2	Review of Continuum-Based Chemistry Models	25
3.3	Review of Modeling Efforts That Account for Capacity Fade and IR Increase	27
3.4	Modeling the Current Distribution in Unbalanced Parallel-Connected Cells	28
3.4.1	Modeling assumptions	29
3.5	Modeling Results and Discussion:	31
Chapter 4	Cell Manufacturing	33
4.1	Active Material Powder and Slurry Production	33
4.1.1	Impact of Manufacturing Variance in Active Material Production	34
4.2	Electrode Coating	34
4.2.1	Impact of Manufacturing Variance in Electrode Coating	36
4.3	Cell Assembly	38
4.3.1	Cell Form Factors:	38
4.3.2	Cell Assembly Process	41

4.3.3	Impact of Manufacturing Variance in Cell Assembly.....	43
4.4	Formation Cycling and Aging.....	44
4.5	Sorting/Binning/Classification.....	44
Chapter 5	Battery Pack Design.....	45
5.1	Mechanical Considerations.....	45
5.2	Thermal Considerations.....	47
5.3	Pack Electrical Balancing.....	49
5.4	Cell Matching.....	50
5.4.1	Capacity Matching.....	50
5.4.2	Internal Resistance Matching.....	51
5.5	The Correlation Between DCIR and C Rate.....	53
Chapter 6	Experiments.....	59
6.1	Experimental Design.....	60
6.1.1	Selection of Sample Cells and Test Parameters:.....	61
6.2	Validation Experiment.....	64
6.3	Selection of test parameters for mass cell cycling experiment:.....	66
Chapter 7	Experimental Results and Discussion.....	70
7.1	Capacity fade vs. initial DCIR Δ	70
7.1.1	Analysis of current distribution within air-cooled 4C cell groups.....	71
7.2	Test results from cells cycled at 8C.....	77
7.3	Effects of water cooling on cells being cycled at 4C and 8C.....	77
7.4	Experimental validation of modeling results.....	77
Chapter 8	Summary and Future Work.....	78
8.1	Summary of key points:.....	78
8.2	Future work.....	79
Chapter 9	References.....	80
Chapter 10	Appendices.....	81
10.1	Combined capacity fade of all 4C air-cooled cells.....	81
10.2	Cycling data from High DCIR Δ Cell group Sample 1 (initial DCIR Δ = 18.4%).....	82
10.3	Cycling data from High DCIR Δ Cell group Sample 2 (initial DCIR Δ = 20.3%).....	83
10.4	Cycling data from Medium DCIR Δ Cell group Sample 1 (initial DCIR Δ = 9.6%).....	84
10.5	Cycling data from Medium DCIR Δ Cell group Sample 2 (initial DCIR Δ = 9.5%).....	85

10.6	Cycling data from Low DCIR Δ Cell group Sample 1 (initial DCIR Δ = 0.17%).....	86
10.7	Cycling data from Low DCIR Δ Cell group Sample 2 (initial DCIR Δ = 0.0%).....	87

Table of Figures

Figure 2-1: Model of a lithium-ion cell "sandwich" (Source: own depiction)	10
Figure 2-2: Diagram showing the "sandwich" configuration of a prismatic Li-Ion cell.	11
Figure 2-3: Brightfield Light Microscopy image of the layer structure inside an aged Li-Ion battery. The red box encompasses one "sandwich layer"; the blue bracket identifies the separator layer.....	13
Figure 2-4: SEM surface photomicrograph of a Celgard® Monolayer Polyethylene battery separator.....	14
Figure 2-5: Volumetric composition of a power cell vs. an energy cell (LiFePO ₄ 2.2Ah pouch cell on the left, LiCo-based 18650 laptop cell on the right). Data compiled by author from cell cutaway measurements.....	16
Figure 2-6: Pore wall flux of lithium ions as a function of dimensionless distance from the anode. $I = 10 \text{ A/m}^2$. Negative values of j_n are for insertion. [2].....	18
Figure 2-7: Concentration profiles in the cathode at long times; $I = 10 \text{ A/m}^2$ discharge. Dashed line divides the separator and the composite cathode. Initial concentration is 1000 mol/m^3 . [2]	18
Figure 2-8: Temperature rise as a function of electrode thickness [3]	19
Figure 2-9: Electrode thickness vs. specific power [4].....	20
Figure 2-10: Diagram of a Li-Ion cell on discharge [5]	21
Figure 2-11: Degradation mechanisms at the anode/electrolyte interface [6]	21
Figure 2-12: Overview of basic aging mechanisms of cathode materials [6].....	23
Figure 3-1: Lumped internal resistance - voltage source model of a cell.....	24
Figure 3-2: Equivalent-circuit model of a single cell [7]	24
Figure 3-3: Half-cell discharge curves showing various capacity-fade phenomena [5]	25
Figure 3-4: Ragone plot showing the dependence of specific power on electrode thickness [4]	27
Figure 3-5: Dimensionless loss of lithium and total resistance at MCMB negative electrode vs. cycle number (EOCV: 4.2 V, DOD: 0.6) [11]	28
Figure 3-6: Influence of DOD on discharge capacity as a function of cycle life [11]	29
Figure 3-7: Correspondence of cell resistance as a function of SOC and discharge rate [12].....	30
Figure 3-8: Cycling current distribution among parallel-connected cells, as a function of IR increase due to cycling.....	31
Figure 3-9: Capacity fade as a function of cycling	31
Figure 4-1: Cutaways showing the inner workings of ball-mills. Recirculating steel balls grind the desired material into a progressively finer powder as it moves through the mill.	34
Figure 4-2: Process flowchart for a typical Li-Ion electrode coating process.....	35
Figure 4-3: Battery electrode calendering machine (left), Calendered electrode spool (right)...	36
Figure 4-4: Pouch cell total thickness vs. DCIR and capacity	37
Figure 4-5: Pouch cell capacity vs. DCIR	37
Figure 4-6: 2.2 Ah LiFePO ₄ pouch cell teardown and thickness measurement.....	38
Figure 4-7: Exploded view of cylindrical cell construction (left), commercial 18650-size laptop cell (right).....	38

Figure 4-8: Exploded view of wound prismatic cell construction (left), commercial Li-Ion digital camera cell (right).....	39
Figure 4-9: Exploded view of stacked prismatic cell construction (left), commercial 40Ah LiFePO ₄ cell (right).....	40
Figure 4-10: Exploded view of pouch cell construction (left), commercial 20 Ah LiFePO ₄ pouch cell (right).....	40
Figure 4-11: Process flowchart for a typical cell assembly process.....	42
Figure 5-1: Diagram of battery pack terminology	45
Figure 5-2: Ford Fusion Hybrid battery pack module stack (left), disassembled module (right). 46	
Figure 5-3: Ford Fusion Hybrid battery pack module interconnect tabs (left), metallic pack enclosure (right).....	46
Figure 5-4: (from top left, clockwise): CAD of battery module, actual battery module after completion, CFD streamlines of air moving throughout the battery module, CFD with the module housing, exploded view of cell mechanical constraint and tab electrical connections..	48
Figure 5-5: Simplified diagram of cell internal resistance (DCIR)	49
Figure 5-6: Diagram of problems arising from poor capacity matching in series, limiting total pack capacity.....	50
Figure 5-7: Voltage curve dispersion in a series battery pack (Source: Author's data)	51
Figure 5-8: Current distribution between two parallel-connected cells with different DCIR values on discharge.....	52
Figure 5-9: Voltage traces at different C-rates for a high power LiCO _x based cell, on discharge	54
Figure 5-10: Voltage curves for each cell in a 6S battery pack on discharge (1C).....	55
Figure 5-11: Voltage curves for each cell in a 6S battery pack on discharge (5C).....	55
Figure 5-12: Voltage curves for each cell in a 6S battery pack on discharge (10C).....	56
Figure 5-13: Voltage curves for each cell in a 6S battery pack on discharge (17C).....	56
Figure 5-14: Voltage curves for each cell in a 6S battery pack on discharge (30C).....	57
Figure 5-15: Voltage curves for each cell in a 6S battery pack on discharge (52C).....	57
Figure 5-16: Battery pack rate capability as a function of manufacturing tolerances and designed cell rate capability.....	58
Figure 6-1: Hypothetical curve of IR variability vs. battery pack cycle life	59
Figure 6-2: Hypothetical surface of IR variability vs. battery pack cycle life, accounting for C-rate	60
Figure 6-3: Anticipated effect of cycling: convergence of the current distribution profiles.....	60
Figure 6-4: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 4C discharge	61
Figure 6-5: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 6C discharge	62
Figure 6-6: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 8.5C discharge	62
Figure 6-7: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 17C discharge	63
Figure 6-8: Individual cell voltage traces from a DCIR pulse test	64

Figure 6-9: Current distribution profiles in a 2P A123 26650 cell group, at cycles 10-12 (left) and 300-302 (right)	65
Figure 6-10: Current distribution profile at cycle 500-502 (left) and cell-group capacity fade (right)	65
Figure 6-11: Capacity fade of each cell in the 2P A123 26650 cell group (left).....	66
Figure 6-12: Sampling areas for the curve of DCIR vs. cycle life (red line is hypothetical)	67
Figure 6-13: Experimental cell-cycling setup	69
Figure 7-1: Capacity fading for cell test group at 4C	70
Figure 7-2: Initial DCIR values vs. capacity fade of air-cooled cells cycled at 4C.....	71
Figure 7-3: 4C air-cooled high DCIR Δ cell group capacity and average current over time. (Left = Sample 1) (Right = Sample 2)	72
Figure 7-4: 4C air-cooled high DCIR Δ Sample 2 current traces at cycles 20-22 (left) and 1395-1397 (right)	73
Figure 7-5: 4C air-cooled high DCIR Δ Sample 1 current traces at cycles 20-22 (left) and 1395-1397 (right)	73
Figure 7-6: Current distribution convergence (Left = High DCIR Δ Sample 2) (Right = Low. DCIR Δ Sample 1)	74
Figure 7-7: Current distribution divergence (Left = Medium DCIR Δ Sample 1) (Right = Medium DCIR Δ Sample 2).....	74
Figure 7-8: Current distribution inversion (Left = High DCIR Δ Sample 2) (Right = Low DCIR Δ Sample 1)	75
Figure 7-9: Initial DCIR value vs. initial capacity of A123 26650 cells (capacity value is averaged over the first 10 cycles).....	76

Chapter 1 Introduction and project motivation

1.1 Hypothesis summary

I hypothesize that manufacturing variances from one lithium-ion cell to another leads to current distribution imbalances when these cells are assembled into a battery pack. These imbalances cause uneven wearing of the battery pack and reduce useful life. In this report, I intend to answer questions about lithium-ion battery manufacturing tolerances that can lead to:

- A better understanding of *how much* variance between cells leads to *how much* loss of usable cycle life in a battery pack
- The correct repurposing of cells that were previously thought to be out-of-spec, to lower performance applications that may be able to use them (application specific binning)
- The feasibility of smaller-scale manufacturing lines, where the use of high precision equipment can be reduced
- Higher throughput of existing battery lines. Certain manufacturing processes (such as electrode coating, calendaring, etc) are rate-limited because of the need to maintain tolerance windows

1.2 Background and motivation

My previous work in battery pack assembly (from small model airplane to full EV-sized packs) motivated these questions. The performance limits of commercially available batteries are typically well-explored by the radio-controlled hobby community; these small-vehicle applications have the advantage of using small (and therefore cheap) battery packs, which is conducive to experimentation. Here, safety and cycle life issues aren't as pressing as they are for full-sized EVs. With the competitive atmosphere of radio-controlled car and airplane racing, battery packs are usually pushed to their limit.

For these small, high-performance battery packs, it is especially critical that the cells are well-matched, as manufacturing differences become more evident as the cells are pushed to their performance limits. In the previous era of Nickel-Cadmium (NiCd) and Nickel Metal-Hydrate (NiMH) based batteries, hobby suppliers would sell "race-grade" battery packs at a steep premium due to the fact that their cells were well-matched in terms of internal-resistance and capacity.

Similar issues exist for larger battery packs that aren't pushed to such high charge and discharge rates. High energy density cells will typically have lower power densities and the differences in their discharge-rate performance will become more apparent.

By exploring the relationship between discharge rate capability, manufacturing tolerance and cycle life, it is possible to quantify the largest allowable manufacturing tolerance windows. In turn, this can aid cost reductions in high-volume manufacturing, as well as enabling the use of less sophisticated production equipment in low-volume manufacturing.

Chapter 2 Lithium-Ion Cell Fundamentals

In this chapter I cover basic theory terminology that is used in the rest of the document. The goal is to provide enough background to understand the issues that arise in manufacturing and their impact on cell performance; a comprehensive review of battery chemistry can be found in other resources and is beyond the scope of this paper. Many of the basic insights in the sub-chapters below are covered in books such as *Lithium Batteries: Science and Technology* [1].

2.1 Theory of Operation: The Lithium-Ion Battery

Similar to lead-acid, nickel-cadmium and nickel metal-hydride batteries, the lithium-ion battery consists of two electrodes immersed in an electrolyte fluid. Ions move between the cathode and the anode through this fluid. Electrons pass through the cathode and anode, which causes current to flow. The electrodes (which are electrically charged, nominally at 3.6 volts) are separated by a dielectric film called the “separator”, which allows ions to pass but keeps the electrodes from physically touching and short-circuiting to one another in the electrolyte solution. The lithium ions pass from the anode to the cathode when the battery discharges, and from the cathode to the anode when the battery charges, as can be visualized in Figure 2-1.




Figure 2-1: Model of a lithium-ion cell "sandwich" (Source: own depiction)

2.1.1 Anode

The anode is typically composed of a matrix of carbon particles, into which the lithium ions intercalate during the charging process. During discharge, the ions leave this matrix and diffuse over to the cathode side of the cell.

While carbon powder alone wouldn't be structurally stable enough for practical use in a battery cell, this powder is mixed with a polymer binder to hold the matrix together; polyvinylidene fluoride (PVDF) is typically used for lithium-ion cells. The combined powder matrix of PVDF and carbon particles form the active material of the anode (the part that contributes to storing energy).

This powder (PVDF and carbon particles) is coated onto a thin copper foil, which is used for distributing current into the active material. The active material matrix isn't as good of an electrical conductor as copper, so thin foils of copper are distributed throughout a lithium-ion cell to aid in this distribution. Copper is used for the anode side as it is at the same electrochemical potential as carbon in this application (use of another metal, such as aluminum, could cause a "micro-battery" to form between the anode active material and its current collector).

The use of many distributed current-collecting foils effectively lowers the internal resistance of the cell, otherwise the I^2R ohmic losses of current passing through the active material would heat up the cell. This would limit how fast the cell could be discharged, as too high of an operating temperature can cause premature degradation (more on this in 2.2.1).

Another important observation is the ion concentration shown by Figure 2-1; as the cell is cycled, the concentration of ions inside the electrode layer tends to increase when closer to the interface with the separator. The charge or discharge rating of the cell is limited by the concentration polarization mechanism, which is a practical limitation on mass transport. The reaction taking place happens faster than the ions can diffuse into the porous electrode. If the voltage at the terminals is allowed to rise or fall below certain voltage limits on charge or discharge (due to high C-rates, not because of SOC window limits) lithium plating can occur.

Prismatic Cell Electrode Stack




Figure 2-2: Diagram showing the "sandwich" configuration of a prismatic Li-Ion cell. Multiple current collectors distributed throughout the cell allow for thinner electrode layers and more current pathways, reducing internal resistance. (Image source: mpoweruk / Woodbank Communications Ltd)

On this point, it is important to note that the thickness of the electrode active material layer is an important parameter in designing the cell. The thicker this portion is, the more resistance it

has, and the more it limits the rate at which the cell can be discharged. (More on cell design in chapter 2.2.1)

The SEI (solid-electrolyte-interface) layer is formed at the initial charge/discharge cycling of the cell. Because the carbon anode is thermodynamically unstable in relation to the electrolyte, a thin passivating layer builds up at the boundary where they meet. This layer moderates the charge rate and restricts current. This adds to the internal resistance of the cell because it limits the rate of ion transfer; more discussion about the SEI layer is in chapter 2.3.1.

Most commercial lithium-ion batteries today use carbon anodes, although there have been recent developments in titanate and silicon nanowire based anodes. Some of these alternative anode materials also have the advantage of not depending on an SEI layer for stability, which can lead to higher charge/discharge rates (and potentially greater cycle life, as buildup of the SEI layer is typically a method of cell capacity degradation, discussed further in chapter 2.3).

2.1.2 Cathode

The cathode active material is also a three-dimensional porous structure. It is typically comprised of a lithium metal oxide or phosphate powder (Cobalt-Oxide, Manganese-Oxide, Iron-Phosphate, etc), a polymer binder for mechanical stability (typically also PVDF), and conductive additives (such as carbon powder) to aid in the electronic conductivity throughout the cathode active material matrix.

It is important to note that there is no solid-electrolyte interface layer between the cathode active material and the electrolyte, as the thermodynamic equilibrium issues present at the anode do not arise here.

The first generation lithium-ion cells that were commercialized in the early 1990s used lithium cobalt-oxide based cathodes. There have since been more developments in the state-of-the-art; a comparison of different cathode active material compounds is in Table 2-1.

Table 2-1: Lithium-Ion cathode chemistry comparison (used with carbon anodes)
 (Source: mpoweruk / Woodbank Communications Ltd)

Cathode Material	Typical Voltage (V)	Energy Density		Thermal Stability
		Gravimetric (Wh/Kg)	Volumetric (Wh/L)	
Cobalt Oxide	3.7	195	560	Poor
Nickel Cobalt Aluminum Oxide (NCA)	3.6	220	600	Fair
Nickel Cobalt Manganese Oxide (NCM)	3.6	205	580	Fair
Manganese Oxide (Spinel)	3.9	150	420	Good
Iron Phosphate (LFP)	3.2	90-130	333	Very Good

2.1.3 Separator

Because volumetric space inside of a battery is at a premium (the goal of designing a battery is usually to store energy in as small of a space as possible), the separator is in intimate contact with the cathode and anode active material layers. Also, because it does not contribute to charge storage within the battery, a typical cell design goal is to make this layer be as thin as possible.




Figure 2-3: Brightfield Light Microscopy image of the layer structure inside an aged Li-Ion battery. The red box encompasses one "sandwich layer"; the blue bracket identifies the separator layer.
 (Image source: Azom.com / Carl Zeiss)

The layer outlined by the blue bracket in the Figure 2-3 shows the separator; it is bound on both sides by anode and cathode active material (the SEI layer on the anode side is not very

visible in this image). The beige-colored stripe is the copper anode current collector; the light gray stripe is the cathode current collector. The space outlined by the red box is one 3.6V cell “sandwich”; the rest of this pattern repeats as more electrode layers are added in parallel. Current collectors are coated with active material on both sides.

The construction of this separator is typically a flexible plastic film or a fiberglass cloth; some advanced work is being done on ceramic separators although at the time of this writing, they are not usually found in commercial cells.

The functional requirements of the separator are to:

- Mechanically keep the cathode and anode layers from touching one another, providing electrical insulation
- Permit ions to pass through the electrolyte solution in which it is immersed
- Resist penetration by burrs or dendrite growths between the electrodes (which can cause high self-discharge rates)
- Maintain physical integrity at high operating temperatures within the cell

An SEM surface photomicrograph of a polyethylene based battery separator can be seen in Figure 2-4; uniform porosity is important in these applications.




Figure 2-4: SEM surface photomicrograph of a Celgard® Monolayer Polyethylene battery separator (Image source: Celgard, LLC)

Specialty separators also exist for high temperature battery applications (where they won’t lose their physical properties) and safety-critical applications (where the separator’s pores will close when the battery gets too hot, impeding ion transfer and “shutting down” the cell internally. Much development work is targeted toward reducing separator thickness while maintaining its physical properties, as thinner separators lead to higher total energy densities.

2.1.4 Electrolyte

The electrolyte of a lithium-ion cell is typically composed of a lithium salt (such as LiPF_6) dissolved in an organic solvent (such as ethylene carbonate). Design requirements include:

- Stability over a broad voltage range (2-4.2 V)
- Consistent performance over a broad temperature range (0 °C – 100 °C)
- High ionic conductivity

The thermodynamic instability between the electrolyte and the anode causes the formation of the SEI layer; as the cell ages, the electrolyte tends to decompose and add to the thickness of the SEI layer. Proprietary additives have been developed by cell manufacturers to reduce this effect in order to improve cycle life.

A factor that limits the energy density of current lithium-ion batteries is the instability of currently available electrolytes at high voltages. In a manner similar to that of a car lead-acid (PbA) battery “bubbling off” hydrogen gas from its sulfuric acid/water electrolyte when it is over-charged, most lithium-ion electrolytes will also release gas when charged past 4.3 volts; most lithium-ion cells have a maximum charge level of 4.2 volts. While PbA batteries can usually vent off this gas and more distilled water can be added to the electrolyte, lithium batteries are sealed. Once they vent, they are considered damaged beyond repair.

2.2 Cell Design and Mechanical Construction

When designing a lithium-ion cell, there are tradeoffs to be made between energy density, power density and cycle life. In comparing an automotive electric vehicle (EV) cell to hybrid-electric vehicle (HEV) cell to a laptop cell, Table 2-2 shows a rough ballpark of design characteristics:

Table 2-2: Comparison of Li-Ion battery design parameters

Application	Energy Density	Power Density	Temperature Range	Required Cycle Life
EV	Med	Med	-20C to 50C	Med (>2,000)
HEV	Low	High	0C to 40C	High (>10,000)
Laptop	High	Low	0C to 50C	Low (>300)

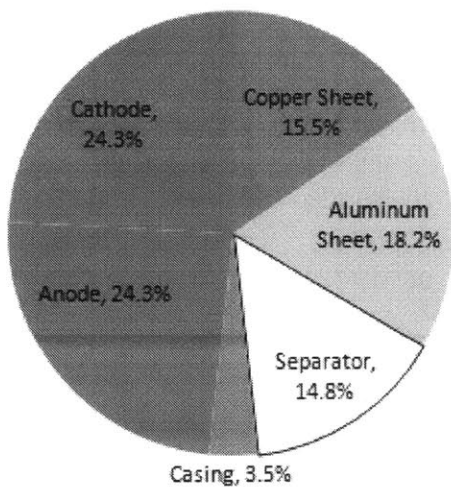
To illustrate the example further, I will compare the specifications of a high-power HEV cell to a high-energy laptop cell.

A laptop battery is typically charged over the span of 1-2 hours (0.5C – 1C), discharged over the span of 2-5 hours (.2C-.5C) and is extremely space constrained (in modern portable electronics, smaller size and lighter weight is usually a premium), so energy density (both volumetric and gravimetric) is important. The industry standard design life is 300 cycles to 70% of the original rated capacity (roughly one year’s worth of use).

An HEV automotive cell is typically charged over the span of a few seconds on deceleration (although they are typically only cycled over a small fraction of their capacity, so overall charge currents are between 3C and 5C) and are discharged in a similar timeframe on acceleration. Here, the ability for a small battery to absorb and release as much power as possible is important, whereas total energy storage capability is less important. The design lifetime of the battery is for the life of the vehicle, which is typically 10 years / 150,000 miles. This means that there are tens of thousands of stop/acceleration cycles, although they don't span the full capacity of the battery (and are thus referred to as "micro-cycles"). The depth-of-discharge window in which a battery is operated is important to determining its cycle life and for a hybrid is typically much smaller than for an EV (Hybrid cells are typically operated in a window of 40-60% of total charge; EV cells can be stressed as far as 0-100%). More on this in Chapter 3.3.

Within the volumetric confines of a cell, it is important to note how much of the volume is actually used for energy storage (anode carbon powder, cathode metal oxide powder and electrolyte). Much of the cell contents (copper current collector, aluminum current collector, PVDF binder, separator, carbon conductivity enhancing powder) do not partake in the chemical reaction that contributes toward energy storage. As can be seen in Figure 2-5 (numbers taken from the teardown of a high-power lithium iron-phosphate cell), less than half of the volume of the cell is composed of cathode and anode powder layers. These powder layers also contain polymer binders and conductive additives, so the actual energy storing materials are present in even lower proportions.

**Power LiFePO4 Cell
Volumetric Composition**



**Energy Laptop 18650 Cell
Volumetric Composition**




Figure 2-5: Volumetric composition of a power cell vs. an energy cell (LiFePO4 2.2Ah pouch cell on the left, LiCo-based 18650 laptop cell on the right). Data compiled by author from cell cutaway measurements.

By comparison, a high-energy cell doesn't need as many current collectors distributed throughout it in order to reduce internal resistance, so there is more room for active material powders.

As mentioned in Chapter 2.2.1, the thickness of the electrode active material layer is a key design parameter when manufacturing a cell. Thicker electrodes have higher electrical resistance (due to the longer distance of the electrical pathway in the poorly conducting powder matrix) and lower ion diffusivity (due to greater distance that ions must travel for intercalation). Combined, they lower the discharge and charge rate capability of the cell.

KEY POINT:

From a manufacturing standpoint, it is important to precisely control the thickness of this layer; if it varies from the design specification, the cell will end up performing unexpectedly (i.e. a high-power cell that has too thick of an electrode active material layer may behave like a high-energy cell and therefore have unacceptably high internal resistance).

2.2.1 Electrode Thickness vs. Cell Rate Capability

The way that internal resistance is defined in this paper (DC voltage sag under current loads) is closely related to rate capability. At some point, the electrochemical kinetics can't keep up with the current demand and the voltage at the cell terminals sags.

Thicker electrodes require diffusion through a porous solid across a longer distance before reaching the current collector, slowing the rate of reaction. Doyle, *et al* [2], proposed the following parameter for modeling diffusion in the electrode layer:

$$S_c = \frac{R_s^2 I}{D_s F (1 - \epsilon) c_T \delta_c}$$

Equation 2-1

Where S_c is the dimensionless ratio representing diffusion within the electrode as a function of the following parameters: R_s is the radius of the cathode active material particles, I is superficial current density (in A/m^2), D_s is the diffusion coefficient of lithium in the solid matrix (m^2/s), ϵ is the porosity of the electrode, c_T is the maximum concentration in the solid (mol/m^3) and (δ_c) is the thickness of the electrode.

Thus, this diffusion parameter should vary linearly with electrode material thickness. However, overall cell rate capability is not dependent on this factor alone. The authors' modeling of pore wall flux as a function of distance from the anode shows what happens over time when the cell receives a current step input (Figure 2-6).




Figure 2-6: Pore wall flux of lithium ions as a function of dimensionless distance from the anode. $I = 10 \text{ A/m}^2$. Negative values of j_n are for insertion. [2]

Similarly, for the cathode, concentration profiles also change as a function of time (Figure 2-7).




Figure 2-7: Concentration profiles in the cathode at long times; $I = 10 \text{ A/m}^2$ discharge. Dashed line divides the separator and the composite cathode. Initial concentration is 1000 mol/m^3 . [2]

It is difficult to say that the diffusion coefficient based on the thickness of the electrode is linearly related to rate capability. A pulse discharge of a cell will lead to a brief concentration overpotential at the electrode interface surface, which takes time to equalize through the electrode.

Botte, *et al* [3] modeled the thermal rise of lithium-ion cells during discharge as a function of electrode thickness. This is related to rate capability, as cells with lower “lumped internal resistance” (described in this paper as DCIR) internally heat less for higher discharge rates. In cell construction, the volume of the active material doesn’t scale with the surface area available for charge transfer. In keeping the discharge rate constant (at 1C), the effective applied current-per-surface-area increases. Thus, this increase in temperature can not be attributed only to electrode thickness, as the resistive SEI layer also needs to pass more current.

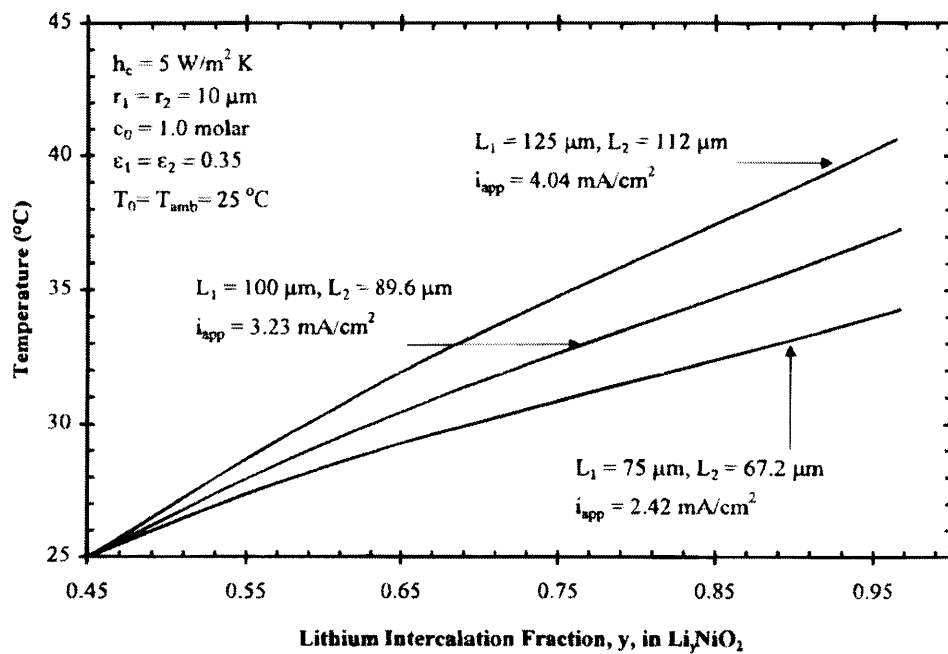


Figure 2-8: Temperature rise as a function of electrode thickness [3]

Tarascon, *et al* [4] experimentally determined the effect of cathode material thickness on specific power capability of a cell. As introduced in the Ragone plot in Figure 3-4, if electrode thickness is graphed vs. rate capability, the relationship is clearly non-linear. For cathode thicknesses of fewer than 100 microns, a 10% difference in thickness will lead to a more than 10% increase in power capability.

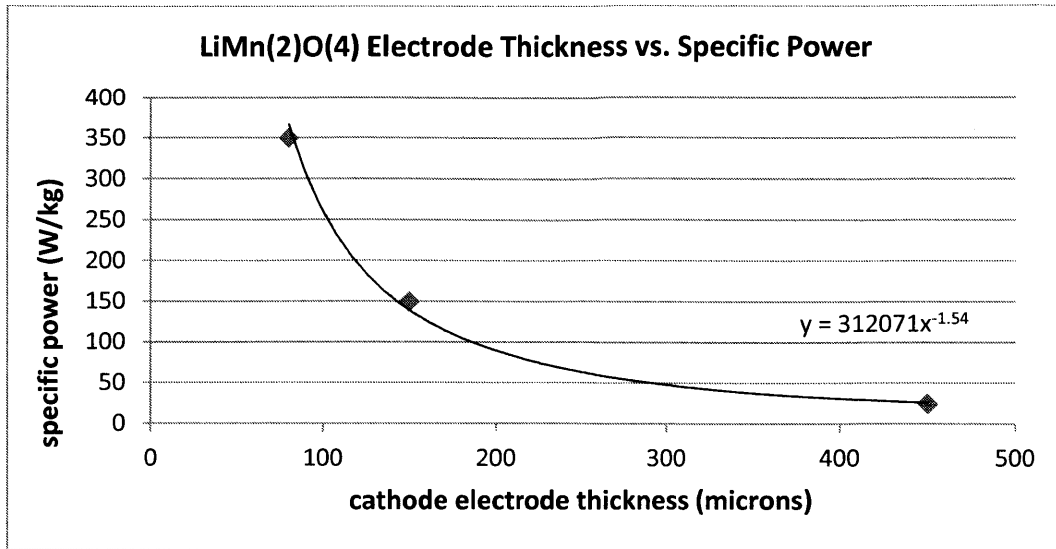


Figure 2-9: Electrode thickness vs. specific power [4]

KEY POINT:

While the exact relationship between rate capability and electrode thickness is chemistry-dependent and difficult to generalize for all electrodes, previous work has shown that this relationship is not necessarily linear. Furthermore, it appears that this relationship is not favorable to tolerance differences in practical batteries with thin electrodes, as small differences in electrode thickness amplify differences in rate capability. Most commercial Li-Ion batteries have electrode thicknesses between 40 and 100 microns, which is in the steepest portion of the curve depicted in Figure 2-9.

2.3 Cell Degradation Mechanisms

This topic has been studied for over two decades since the commercialization of the lithium-ion battery by Sony in the early 1990s. Given the breadth of information available on the topic, in this chapter I will briefly review the generally accepted theories of battery degradation. I will only discuss the degradation mechanisms associated with normal cycling operation, as most large battery packs are equipped with battery management systems keep the cells within reasonable operating windows and prevent conditions of over-charge, over-discharge, over-current and operation during temperature extremes.

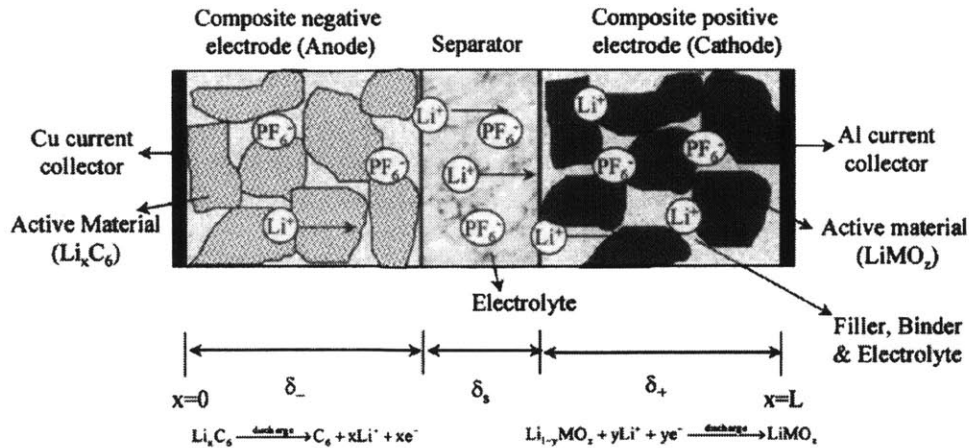


Figure 2-10: Diagram of a Li-Ion cell on discharge [5]

2.3.1 Degradation Mechanisms of Carbon Based anodes

The anode electrode of lithium-ion cells is typically constructed of particulate carbon (either graphite or petroleum coke based powder), mixed in slurry with binding agents and thinly cast onto a copper current collector foil. Because battery anodes operate at voltages that are outside of the electrochemical stability window of electrolyte components[6], reductive electrolyte decomposition occurs on the anode surface, forming a passivating layer. Its growth impedes the transfer of ions into the anode (increasing resistivity) and consumes active lithium from the electrolyte (decreasing the amount of lithium available for intercalation). The formation and stability of this layer is the topic of much research, as it is a primary mechanism for capacity degradation. In this chapter I discuss the formation and stability of this layer, as well as other side-reactions that can occur on the compound electrode (current collector + carbon powder + PVDF binder).




Figure 2-11: Degradation mechanisms at the anode/electrolyte interface [6]

The solid-electrolyte-interface layer is formed during the first few cycles of a lithium-ion battery. After commercial lithium-ion cells are assembled, they are placed on cycling machines

which carefully charge and discharge the cells (known as “formation cycling”, described further in chapter 4.4) to ensure the proper formation of this layer before the cell can be used. The smoothness with which this layer is formed is important, as a uniform SEI layer can help to reduce abnormal future growth of the layer during the cell’s lifetime. This has a direct impact on cycle-life and the charging profile used during formation is the topic of much intellectual property among battery manufacturers.

The decomposition of the electrolyte into the SEI layer is accelerated by higher temperatures. This can happen even when the cell is not being cycled and it is a primary factor for the “calendar life” rating of the cell.

During cycling, portions of the SEI layer can dissolve into the electrolyte and then re-precipitate onto the original layer, causing instability and opening fresh sites for further electrolyte deposition.

On a long time scale, the SEI layer penetrates into the pores of the electrode and, in addition, may also penetrate into the pores of the separator. In some cases it can rupture through the separator and touch the cathode side of the cell, creating a current path and increasing the self-discharge rate of the cell. In extreme cases, if this internal short is of low enough resistance it can lead to thermal runaway.

Lithium plating can occur if the cell is discharged too quickly (limited by the diffusion of lithium ions through the SEI) or if the SEI or electrode thicknesses are uneven (causing inhomogenous current distribution). This can lead to further undesired side-reactions.

2.3.1.1 Degradation mechanisms in the bulk electrode:

Volume changes during intercalation are typically small on the anode side, but over time these mechanical stresses can lead to cracking within the material. This loss of contact leads to higher electrical resistance and can occur at the interfaces between the carbon particles, between the current collector and the carbon particles, between the current collector and binder, and between the binder and the carbon particles.

2.3.1.2 Degradation mechanisms at the current collector:

Volume changes at the interface of the electrode and the current collector can lead to cracking at the surface of the current collector, increasing electrical resistance. These volume changes also place a mechanical strain on the current collector itself, leading to work-hardening of the copper foil. In some cases this can lead to cracking of the current collector, also increasing electrical resistance [5].

2.3.2 Degradation Mechanisms of Lithium Metal Oxide Cathodes

While degradation mechanisms for cathode materials are specific to the type of compounds used (commonplace commercial materials include cobalt, nickel and manganese oxides – and more recently, iron phosphate), there are commonalities. This chapter will cover some most widely accepted mechanisms that lead to cell degradation




Figure 2-12: Overview of basic aging mechanisms of cathode materials [6]

2.3.2.1 Degradation mechanisms in the bulk electrode:

Phase transitions during cycling: for example, LiCoO_2 will transition from single hexagonal to two hexagonal phases during its normal cycling window; large changes in lattice parameters can lead to volume expansions that cause cracking within the compound electrode. As the cathode is a mixture of binder, carbon particles, metal oxide particles and aluminum foil (current collector), cracking can occur in between any of these interfaces, leading to loss of conduction within the electrode (higher internal resistance). Recent advancements in doping and alloying (combinations of nickel, manganese, cobalt and aluminum in particular) have led to stabilization in these phase transitions and consequently, improved cycle life.

Dissolution of active material into the electrolyte: This is most pronounced with manganese-spinel based cathodes; some of these dissolved products can end up on the anode side in the SEI layer. With less active material available for intercalation, capacity fading occurs. This is most pronounced at high or low states of charge; operating the cell in the middle of its cyclable range tends to mitigate this effect and can lead to a higher total Wh-throughput for the life of the cell.

2.3.2.2 Degradation mechanisms at the electrode surface:

Catalytic decomposition of organic solvents in the electrolyte can lead to a surface layer that blocks the passage of ions and increases cell internal resistance. Although this is assumed to happen, its effects are not believed to be as significant as those of the SEI layer on the anode.

Chapter 3 Cell Modeling

In this chapter I briefly address the most generally accepted battery modeling techniques, with additional emphasis on modeling related to cell degradation and capacity fade.

3.1 Review of Equivalent-circuit Based Models





Figure 3-1: Lumped internal resistance - voltage source model of a cell

In a very basic form, a battery can be modeled as a voltage source in series with a resistor. While it is useful to visualize what is going on in a cell, this is a crude model because it hides the underlying physics. Many of the values that are assigned to the electrical component within the model change as a function of charge, temperature, current within the cell and other parameters. Figure 3-2 expands the equivalent-circuit concept further.



E1, E2	single electrode potential	W1, W2	diffusional (Warburg) impedance
L1, L2	current collector lead inductance	R3	resistance of solution and separator
C1, C2	double layer capacitance	C3	inter-electrode capacitance
R1, R2	charge transfer resistance		

Figure 3-2: Equivalent-circuit model of a single cell [7]

As a starting point for further analysis:

- E1, E2: this is the potential of each electrode on its own, as measured with a reference electrode placed within the electrolyte. For reference, see Figure 3-3.
- L1, L2: this is the inductance of the leads to the battery and the current collectors themselves. Typically this is very small (though for some measurements like electrochemical impedance spectroscopy, its effects need to be accounted for)
- C1, C2: the capacitance due to the electrical double layer that occurs on the interface between the electrode surface and the electrolyte

- W1, W2: the Warburg impedance attempts to capture the electrical effects of the reaction kinetics that occur at each electrode. This is a lumped-model approximation.
- R3: the resistance of the electrolyte and the separator; this is usually more complex and there can be a Warburg diffusion coefficient assigned to this part of the cell-sandwich as well
- C3: the capacitance of the current-collector foil electrodes; if all components were removed from a lithium-ion cell except for the current collectors, it would have the construction of a capacitor. This effect still exists when the lithium-ion cell is assembled.

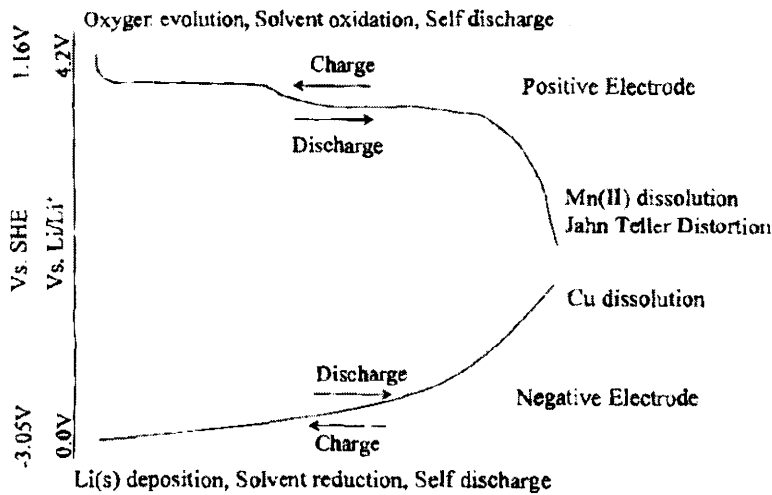


Figure 3-3: Half-cell discharge curves showing various capacity-fade phenomena [5]

These models can be used for first-order approximation and can be refined with empirical data from cell testing. They also have the advantage of not being as computationally intensive as chemical models.

The definition of DCIR (DC Internal Resistance) used in this paper follows this simple circuit based model, and is described below.

$$\frac{(\text{nominal cell voltage}) - (\text{voltage under load})}{\text{load current (A)}} = DCIR$$

Equation 3-1: Calculation of DCIR

Such models are not very useful when predicting the behavior of new cell materials or empirically tested cells that are subjected to drastically different operating conditions. For this, more complex models that directly address the underlying physics are necessary.

3.2 Review of Continuum-Based Chemistry Models

This topic has been explored since the development early secondary cells; the principles that underlie the electrochemical cell-sandwich models that have been developed for NiCd, NiMH

and PbA batteries form the foundation for today's Li-Ion mathematical models. This is a research area on its own [2, 3, 8-10], seeking to further expand the parameters outlined in chapter 3.1.

Because the focus of this paper lies in manufacturing and practical battery pack longevity, in this chapter I will focus on mathematical models for predicting internal resistance as a function of electrode thickness, and predicting capacity fade and internal resistance increase as a function of cycling. In manufacturing lithium batteries, maintaining an even electrode thickness is a challenge for even the most sophisticated automated machinery. Of all factors that contribute to internal resistance variance, consistent electrode thickness is one of the most difficult to control.

KEY POINT:

Variations in the thickness of the electrode within a cell can lead to localized hot-spots (thicker areas that have higher diffusion resistance). Variations in electrode thickness from cell to cell can lead to an entire cell taking more of the current when placed in parallel with other cells in a battery pack.

The difference in power density that is accompanied by thinner electrode coatings is shown in Fig. 8.




Fig. 27. Ragone plots of $\text{LiMn}_2\text{O}_4/\text{carbon}$ (petroleum) cells of three different thicknesses of the composite electrodes, in EC + DEE (50:50) + 1M LiClO_4 electrolyte, extrapolated to practice hardware with a hypothetical realization factor of 33%, as compared to presently commercialized secondary lead acid (portable sealed battery of 61.7 cm^3) and D-size (55 cm^3) nickel-cadmium cells (13). The pellet thicknesses used for the $\text{LiMn}_2\text{O}_4/\text{carbon}$ cells are the following. 1: $e_+ = 0.085 \text{ mm}$ and $e_- = 0.08 \text{ mm}$. 2: $e_+ = 0.17 \text{ mm}$ and $e_- = 0.15 \text{ mm}$. 3: $e_+ = 0.51 \text{ mm}$ and $e_- = 0.45 \text{ mm}$.

Figure 3-4: Ragone plot showing the dependence of specific power on electrode thickness [4]

3.3 Review of Modeling Efforts That Account for Capacity Fade and IR Increase

A well-studied effect of lithium cell cycling is the fade of capacity, which is accompanied by a decrease of rate capability. A cause for the decrease in rate capability is the growth of the SEI layer, which increases the effective DCIR. Because the deposits that form the SEI layer also consume available lithium from the electrolyte, capacity also fades.

Christensen, *et al* [8], developed a mathematical model for the lithium-ion negative electrode solid electrolyte interface layer, modeling its thickness as a function of time. They worked to capture its effects on decreasing the available lithium content for the reaction of the cell (capacity fade) and its contribution to the resistance of the system, leading to power loss.

Ning, *et al* [11], developed a mathematical model to predict capacity fade and SEI resistance based on the loss of active lithium ions due to solvent reduction and the precipitation of insoluble products onto the anode surface. They captured the effects of SOC voltage windows and the impact of the depth-of-discharge on a Li_xCoO_2 / mesocarbon microbead (MCMB) electrode. Their model was developed with empirical cell-cycling data.

Dubarry, *et al* [12], outlined a general purpose battery model based on equivalent-circuit modeling and empirical cell testing data, describing the limits to which a non-chemistry-based model can be useful for practical simulation.

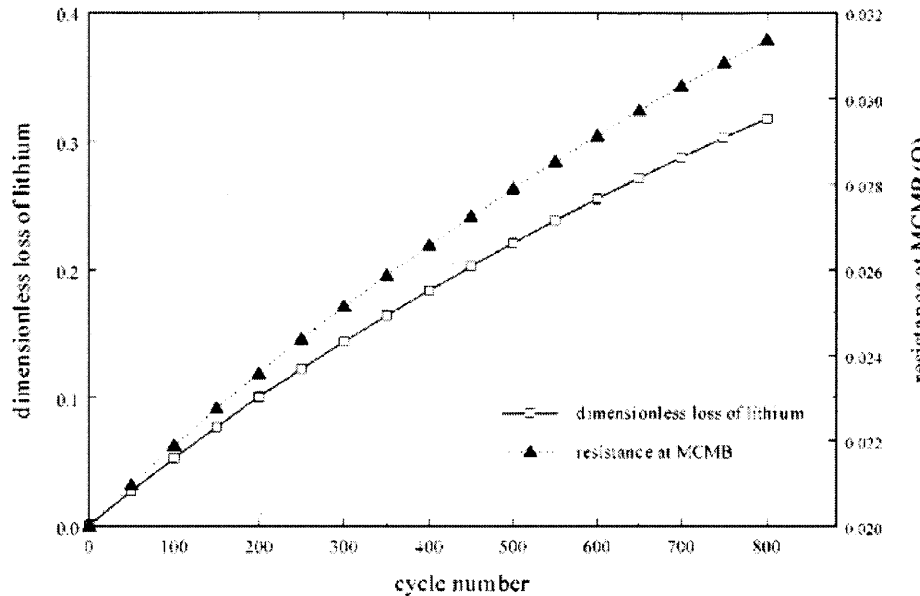


Figure 3-5: Dimensionless loss of lithium and total resistance at MCMB negative electrode vs. cycle number (EOCV: 4.2 V, DOD: 0.6) [11]

3.4 Modeling the Current Distribution in Unbalanced Parallel-Connected Cells

In this chapter I will quantify the effect described in 0, where two cells with different internal resistances are connected in parallel. By knowing the IR rise of a cell as a function of cycle number and depth of discharge, it is possible to predict when two unbalanced parallel-connected cells will eventually wear to the same point and share the current load evenly.

The following model builds upon the work done by Ning, *et al* in modeling depth-of-discharge and cycling effects on IR increase.




Figure 3-6: Influence of DOD on discharge capacity as a function of cycle life [11]

It is expected that higher depth-of-discharge (DOD) will lead to more capacity fade and faster IR increase as cycling progresses. In the case of two parallel-connected cells where one cell has a lower IR, it will receive more of the current load and thus experience a higher DOD. This will cause its IR to rise faster than that of its neighbor, to the point where it will eventually match the IR of its neighbor and the current will be distributed evenly.

3.4.1 Modeling assumptions

Extrapolating from the modeling data provided by Ning, *et al*, I made the following assumptions for the parallel-connected-cell current distribution model:

- The difference in the rate of capacity fade between the 40% DOD cell and the 60% DOD cell can be extrapolated to the difference in IR rate increase. That is, this model assumes that capacity fade and IR increase are proportional. This is not always true, depending on the cathode chemistry used and electrode thickness
- The resistance of the SEI layer is representative of the system resistance of the entire cell; that is, the resistance of the cell system increases proportionately with SEI layer resistance.
- That the cells are being operated at a low enough C-rate where the DOD is the dominant factor for capacity fade and IR increase. There was no incorporation of the effect of C rate on cycle life.
- The IR of the cell remains constant throughout its discharge range.

Note: efforts by Dubarry, *et al*[12] on LiCoO₂ / Graphite cells shows that for low C-rates, IR does not vary significantly over the discharge range. See Figure 3-7.

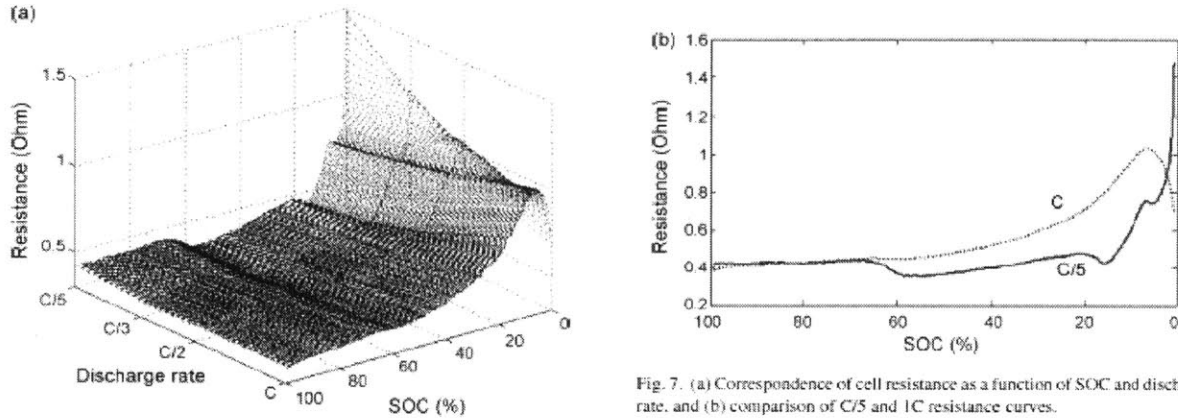


Fig. 7. (a) Correspondence of cell resistance as a function of SOC and discharge rate, and (b) comparison of C/5 and 1C resistance curves.

Figure 3-7: Correspondence of cell resistance as a function of SOC and discharge rate [12]

These are rough assumptions and do not create a perfect model; this is an approximation for when the point occurs that both parallel cells reach the same IR.

After each discharge cycle, the IR of the stronger cell will increase faster than the weaker cell. This means that on the next cycle, the IR difference between the two cells will become smaller, which leads to the DOD window of the higher-IR cell getting larger (compared to the previous cycle) and the DOD window of the lower-IR cell getting smaller (compared to the previous cycle). As the cells become more equalized, the rate at which they converge becomes slower.

The purpose of the experimentation (starting with Chapter 6) is to verify these modeling results, which are based on lumped-model assumptions. This is required to verify what happens in real-life systems that include more complex side-reactions.

3.5 Modeling Results and Discussion:




Figure 3-8: Cycling current distribution among parallel-connected cells, as a function of IR increase due to cycling

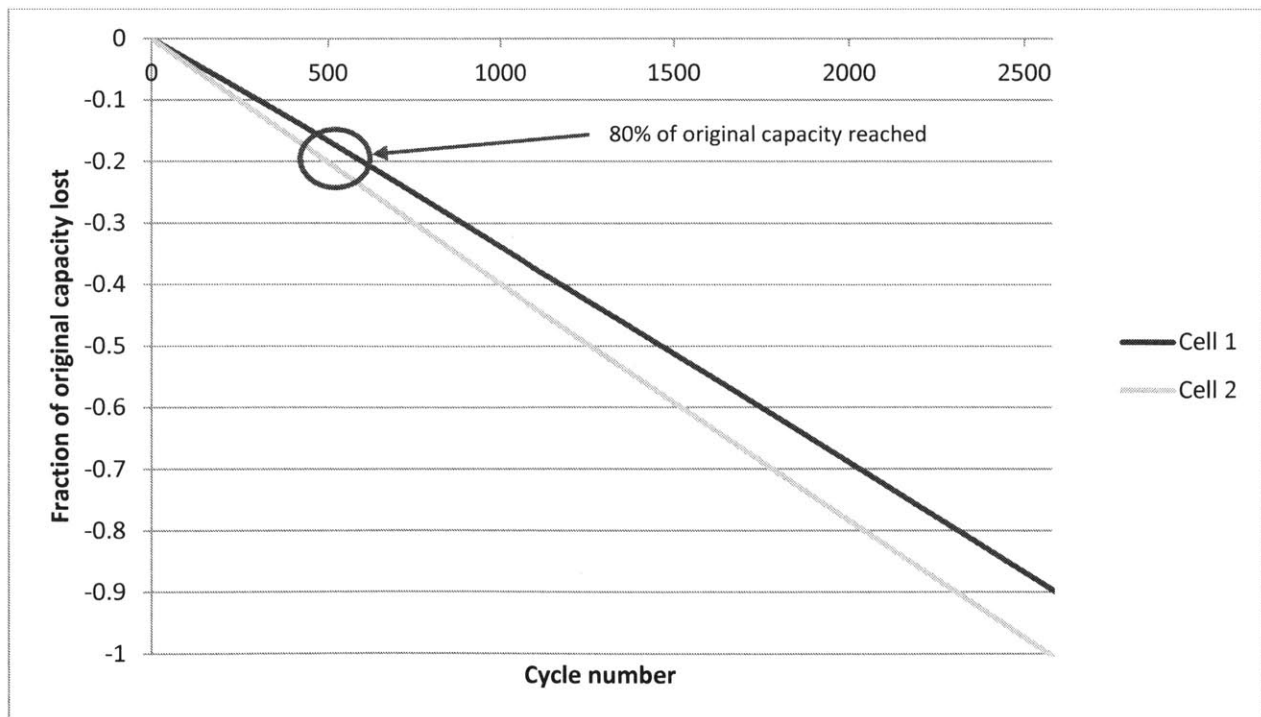


Figure 3-9: Capacity fade as a function of cycling

Based on the rate at which the IR increases as a function of DOD and cycle number in Ning, *et al*'s model, the eventual convergence of two cells that start off at drastically different IR values (.03 Ohm and .02 Ohm) is seen in Figure 3-8 and Figure 3-9.

Experimental results of commercially assembled lithium-ion cells typically do not degrade so linearly all the way to zero capacity. My model doesn't capture the effects of other side reactions, which usually lead to nonlinearities along the degradation curve. However, degradation to 80% remaining capacity can be approximated as linear.

A note about the "end-of-life" of a cell being defined as 80% of its initial capacity: this is a common industry standard. Some laptop cells are even rated to as low as 70% of their initial capacity. These are chosen arbitrarily, mostly driven by what a consumer would find acceptable. Some industrial applications may find batteries to be useful at much lower remaining capacities, such as end-of-life grid storage solutions that employ used EV battery packs. For this paper, I will use the 80% end-of-life standard, as it is comparable to what many other cells are rated to.

Some cell chemistries degrade linearly with cycle life down to as far as 50% of their initial capacity. For the imbalanced cells in the model, Cell 2 hit that point after 497 cycles. At this point, Cell 1 still had 83.5% of its capacity left. Averaged among the two cells, they hit the 80% mark of the original 2-cell-pack's capacity at 544 cycles.

As the "stronger" cell gets worn faster due to higher cycling loads, its DCIR value increases as its capacity decreases. So while DCIR values converge, the capacities of the two cells diverge as they age.

If the cells were perfectly balanced to begin with and started off at identical IRs of 0.025 Ohms, they would have degraded identically and the 80% capacity remaining point would have been hit after also 544 cycles.

KEY POINT:

If the degradation curve is non-linear (as a function of DOD and C-rate), there will be optimization points that make it possible to plot battery pack cycle life as a function of initial variance in cell internal resistances. This can lead to further optimization in industrial planning, where the cost to achieve those tolerances is known. Thus, cycle life can be plotted as a function of incremental manufacturing cost.

Chapter 4 Cell Manufacturing

In this chapter, I provide a brief overview of the processes used in manufacturing lithium-ion cells, to provide background for the subsequent experimental design chapter. The impact of manufacturing variance at each step of the process is highlighted.

In this chapter I will not cover the manufacturing of aluminum and copper foils (used as current collectors), electrolyte, or separator materials. The sheet-metal industry is relatively mature and these processes are well documented elsewhere. The high-purity chemical supply industry is also mature and the chemicals used for electrolytes are commodity materials. The techniques used for separator manufacture are typically proprietary (not documented in academic literature) and I didn't visit any separator manufacturers so I am unable to add further knowledge on the topic.

I did however visit seven cell assembly plants over the span of this research and I will add my commentary below based on this experience.

4.1 Active Material Powder and Slurry Production

The anode's active material matrix is composed of fine carbon powder and a PVDF binder (typically dissolved in an organic solvent). The powder and binder are mixed together to form a slurry, which is spread onto the copper current collector foil (details of which are described in the following chapter on "electrode coating").

A typical anode will contain 95% (by weight) graphite powder and 5% PVDF, dissolved in an organic solvent such as NMP (*N*-Methyl-2-pyrrolidone), which is evaporated in the coating process.

Reducing the particle size of the graphite powder has the advantage of increasing the effective surface area of the anode, providing slightly higher capacity, but more importantly, higher rate capability (the ability to charge and discharge quickly).

From a manufacturing standpoint, the same issues of granularity apply to the powders used in the cathode active material matrix, composed of lithium-metal-oxide powder, PVDF binder and conductivity-enhancing carbon particles.

These powders are either mechanically milled or grown through chemical processes. Some of the more recently developed chemical processes claim to be able to produce particles of smaller size and higher uniformity, but many of these developments are proprietary and the technical details are not discussed here [13].

Precision ball mills are typically used for the mechanical powder formation process, with screens at the end of them that filter out particle size.




Figure 4-1: Cutaways showing the inner workings of ball-mills. Recirculating steel balls grind the desired material into a progressively finer powder as it moves through the mill.
 (Image sources: mine-engineer.com, nsiequipments.com)

4.1.1 *Impact of Manufacturing Variance in Active Material Production*

The uniform composition of these powders is critical in cell manufacturing; particles that are too large can provide less surface area than anticipated, reducing capacity and rate capability. Sophisticated manufacturing plants have scanning electron microscopes to not only verify the average particle size of the powder that they receive from suppliers, but also their shape and uniformity from particle-to-particle.

4.2 **Electrode Coating**

Electrode coating is one of the most closely guarded processes of the battery assembly industry, outside of active material powder production. I visited seven cell assembly plants over the course of this research; none of these factories allowed visitors into the electrode coating line. The reason for this is twofold:

- Maintaining a consistent electrode thickness is relatively difficult because the dimensions are so small (see Chapter 2.2.1 for more discussion on this topic). While nanometer-level manufacturing is common in the semiconductor industry, mechanical layer deposition is typically less precise than the features that can be achieved by processes such as chemical vapor deposition (CVD). These slurries must be deposited mechanically, with target thicknesses ranging between 40 and 100 microns, depending on whether the cells are designed for energy or for power density. These tolerances become even more important for high power cells that have thin electrodes.
- Electrode thickness is critical to cell performance. When these tolerances are out of control, the manufacturer may end up with a high energy cell when they think they're making a high power cell and vice-versa. Furthermore, for cylindrical cells, the fit of the electrode spiral within the can is related to the thickness of the electrode coating. Too thin of a coating leads to a spiral electrode stack that won't fill the entire inner diameter of the can, which leads to mechanical constraint issues (and other problems regarding electrode-to-electrode pressure; see chapter 4.3.1.1).

Typically, the electrode coating is applied to the current collector foil through a sputtering or blade-casting process. Blade-casting involves using a precision knife-blade to deposit the active material slurry onto the current collector. The distance between the knife blade and the foil determines the thickness of the active material layer, before it is compressed through calendaring. The coating processes are almost identical for both anode and cathode. Purity is key in this step; these processes are done in clean rooms, and the anode is usually coated in a separate room from the cathode.

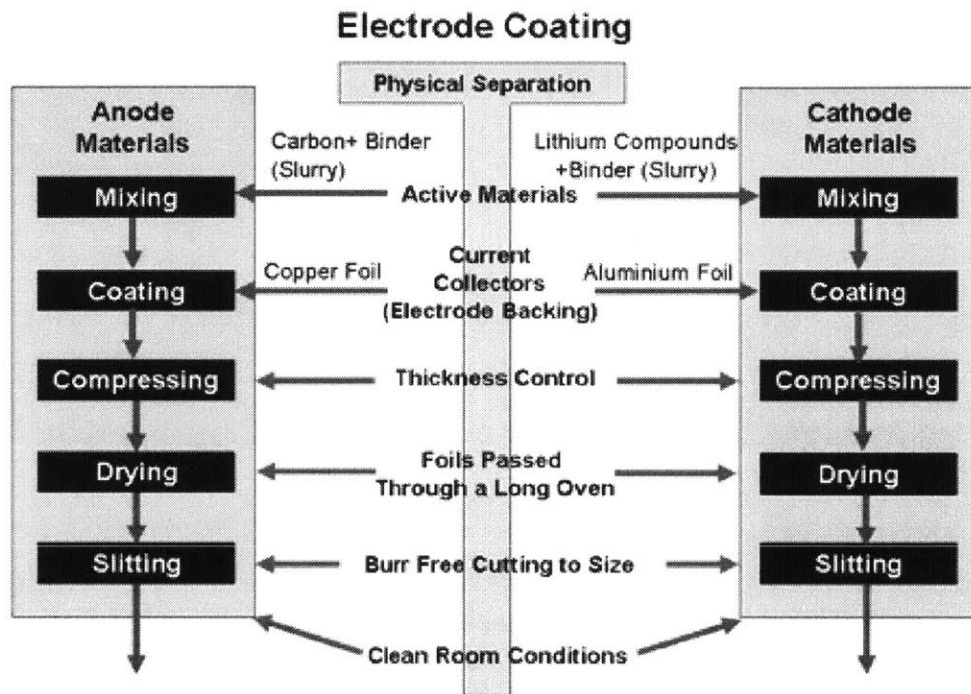


Figure 4-2: Process flowchart for a typical Li-Ion electrode coating process
 (Image source: mpoweruk / Woodbank Communications Ltd)

The calendaring process involves passing the electrode sheet through a set of two precision rollers, where the gap between them is tightly controlled. Controlling this gap and maintaining a uniform electrode thickness is difficult. For reference, cathode electrode thicknesses are in the range of 40 to 100 microns, depending on whether the cell is designed for a high-power or a high-energy-density application. Thus, a +/- 10% deviation in thickness for a high-power cell that has a target electrode thickness of 40 microns, means maintaining a 4 micron tolerance on the machinery that controls this gap thickness. Today's sophisticated CNC machining centers (~\$100,000 for a 3 axis HAAS CNC mill) can maintain tolerances down to about 5 microns. Extremely high-end micro-mills can achieve sub-micron accuracy, however, keeping these tolerances is not an easy task, especially on larger machinery. The rollers of a calendaring machine on lithium-ion battery production line are in the range of about 1 meter in diameter and in the range of .5 meters to 2 meters long. Machining the surfaces of these cylinders to be straight within a few microns (the width of a human hair is 60 microns) is a challenge in itself.





Figure 4-3: Battery electrode calendaring machine (left), Calendared electrode spool (right)
(Left image source: Xingtai Spiry Trade Co., Ltd.) (Right image source: JCI-Saft)

After the calendaring process (of which there may be more than one step; the same sheet can pass through multiple machines), the coated electrode will enter a long continuous oven to fully dry the electrode layer by evaporating any remaining organic solvent. It is then slit and spooled back up again into large rolls, awaiting final assembly into a complete cell sandwich (cathode + separator + anode).

4.2.1 Impact of Manufacturing Variance in Electrode Coating

As mentioned previously in this chapter, the thickness of the electrode layer has a significant impact on cell performance, in terms of both capacity and rate capability. In some cells where the volumetric energy storage capacity of the anode active material is different from that of the cathode (i.e. for the cell's storage capacity to be balanced internally, the anode and cathode layers are designed to be of different thicknesses), maintaining proper electrode thickness becomes even more important.

Below are graphs of cell thickness vs. capacity and DC internal resistance, as measured from a batch of fifteen 2.2Ah LiFePO₄ pouch cells. These are high power cells (designed for hobby model airplane use), whose volumetric breakdown is presented earlier in Figure 2-5.




Figure 4-4: Pouch cell total thickness vs. DCIR and capacity

The thickest cell in this batch was just 5.1% thicker than the thinnest cell, but the difference in capacity (between highest and lowest) was 17%. The difference in internal resistance was even more striking, with the worst cell having almost twice the internal resistance of the best cell (measured with a 15 second 12C pulse, at 50% SOC). The correlation between internal resistance and cell thickness aligns with the theoretical explanations earlier, in chapter 2.2.1.

However, when graphing the sample measurements against each other (capacity vs. DCIR), the correlation is not so high (Figure 4-5).

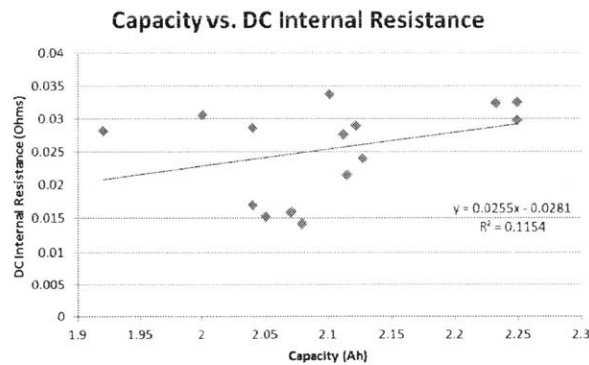


Figure 4-5: Pouch cell capacity vs. DCIR

KEY POINT:

Although the slope of this correlation is positive (higher capacity cells have higher DCIR values, which makes sense from a theoretical standpoint), the correlation of this trend is relatively low. Potential reasons for this include variance in DCIR measurements (cell rest periods are an important factor, as can be seen in Figure 2-7), as well as variance within the active materials for each cell (it is unknown whether they came from the same production batch)




Figure 4-6: 2.2 Ah LiFePO₄ pouch cell teardown and thickness measurement

The method used to determine internal resistance was with a 15 second 26A pulse (12C), at 50% SOC. DCIR was determined using Equation 3-1.

4.3 Cell Assembly

Although the cell's operating principles don't fundamentally change, different physical form factors need their own considerations in the manufacturing process.

4.3.1 Cell Form Factors:

I describe the four most popular lithium-ion cell form factors in commercial use today, in the sub-chapters below.

4.3.1.1 Cylindrical




Figure 4-7: Exploded view of cylindrical cell construction (left), commercial 18650-size laptop cell (right)
 (Left image source: mpoweruk / Woodbank Communications Ltd) (Right image source: Yoko Energy)

Cylindrical cells have a tightly wound spiral electrode sandwich that is inserted into a cylindrical can and then crimp-sealed at the top (another method is to assemble an extruded cylindrical housing with two end-caps via laser-welding). Due to the slippery nature of the electrolyte, this

spiral electrode sandwich is colloquially known as a “jellyroll” in the industry. The tension with which the jellyroll is wound is important, as well as its fit within the can.

As lithium ions intercalate into the anode on charging, they cause the anode to swell slightly. In order to maintain a reasonable cycle life, it is important to minimize the strain that is incurred in the anode side of the cell, so proper mechanical constraint of this swelling is important. In a pouch cell (see the chapter 4.3.1.4, describing that construction), this is done by applying compression plates to both ends of the cell. In a spiral-wound cylindrical, this force is applied by the level of tension with which the spiral is wound, initially. After the first few charge/discharge cycles, the electrode spiral expands radially towards the walls of the metal can, which then becomes the mechanical constraint. For these reasons, the tension applied to winding the electrode spiral, as well as its final diameter, are critical.

Current take-off tabs are welded to the positive and negative current collectors in the electrode spiral, and then welded to the cell can and positive tab. This is typically done through friction (vibration/ultrasonic) welding. For higher current cells, there are typically more current take-off tabs spread throughout the current collector to distribute current more evenly and lower cell internal resistance (sometimes as many as 5 tabs for each anode and cathode).

4.3.1.2 Wound Prismatic




Figure 4-8: Exploded view of wound prismatic cell construction (left), commercial Li-Ion digital camera cell (right) (Left image source: mpoweruk / Woodbank Communications Ltd) (Right image source: digitalpowerpro.com)

The construction of these cells is similar to the cylindrical cell, except that the electrode spiral is wound around a flat mandrel to create an “ovalized” spiral. On larger cells requiring compression, this is typically provided by mechanical constraint added to the battery pack.

4.3.1.3 Stacked Prismatic




Figure 4-9: Exploded view of stacked prismatic cell construction (left), commercial 40Ah LiFePO₄ cell (right)
(Left image source: midwestenergynews.com) (Right image source: CALB)

In this configuration, the positive and negative electrodes are cut into squares and stacked on top of each other along with the separator. Each electrode has its own current take-off tab; current distribution is relatively even and ohmic losses within the cell are low. This design also has the advantage of high volumetric space efficiency (almost all of the volume inside of the case is used) and its prismatic shape allows it to be efficiently packaged into packs, as opposed to the space lost in cylindrical designs.

4.3.1.4 Pouch




Figure 4-10: Exploded view of pouch cell construction (left), commercial 20 Ah LiFePO₄ pouch cell (right)
(Left image source: Exa Energy) (Right image source: A123 Systems)

Pouch cells (also known in the industry as “Lithium-Polymer” cells, which are not always the same as true lithium polymer cells that have a solid polymer separator/electrolyte) have a construction similar to that seen in stacked prismatic cells. The pouch cell construction is sometimes used for true lithium-polymer solid-electrolyte cells, but is also found with traditional porous separator/electrolyte stacks. The pouch casing is typically made from a plastic-coated aluminum foil. This has the advantage of space-saving, but requires extra care in packaging. Internal pressure from electrolyte side-reactions can also lead to leakage; proper battery pack management and the development of gas-reducing electrolyte additives are also important to the proper use and longevity of pouch cells.

These cells are traditionally compressed for optimum cycle life, but some recent developments in electrolytes that are less gassing have allowed for cells without compression. Uncompressed cells are typically found in consumer electronics applications where cycle life isn't as critical as it is on larger battery systems, such as in automotive applications.

4.3.2 Cell Assembly Process

Regardless of form factor, most cells follow the same procedure from electrode coating to final packaging in a finished cell.

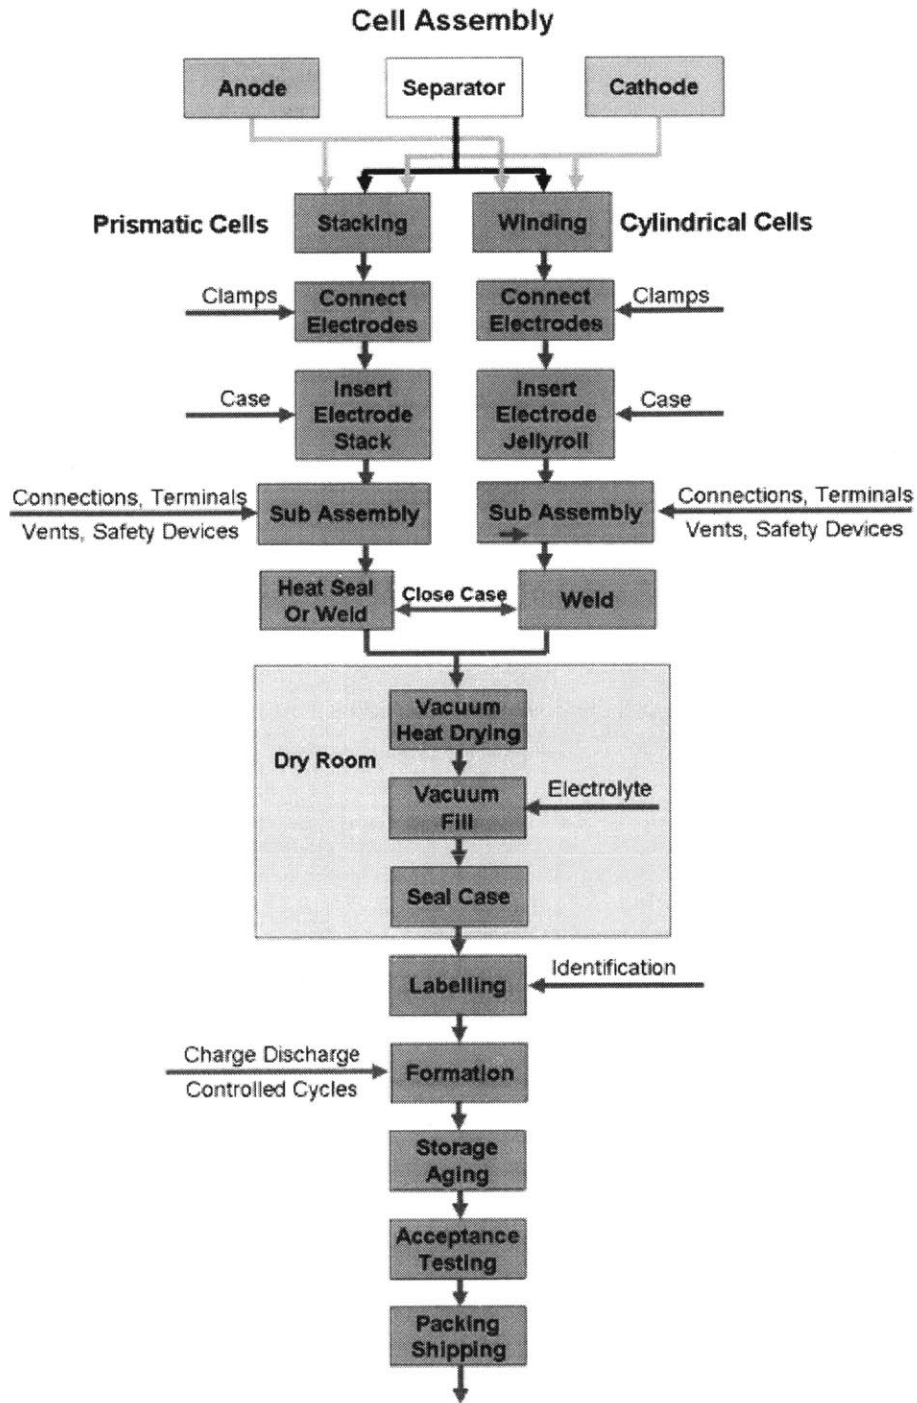


Figure 4-11: Process flowchart for a typical cell assembly process
 (Image source: mpoweruk / Woodbank Communications Ltd)

4.3.3 Impact of Manufacturing Variance in Cell Assembly

Many of the errors that occur in cell assembly are critical defects that lead to completely unusable cells and are caught in quality control during the formation cycling, storage aging and acceptance testing phase of production.

Table 4-1: Easily detectable cell assembly manufacturing errors (Source: author's research)

Defect Cause	Effect	Quality Control
Improper tab welding	High DC internal resistance	Impedance spectroscopy at the end of the aging storage process can determine IR (more sophisticated manufacturers typically have this capability)
Particulate contamination within the electrode stack	Internal short-circuiting of the cell, or abnormally high self-discharge rates	Self-discharge rate is measured during the cell aging process.
Incomplete absorption of the electrolyte within the electrode stack	Low capacity	Capacity is measured during formation cycling
Improper can sealing	Electrolyte leakage, humidity contamination within the cell	Visual leakage inspection, vacuum testing during electrolyte filling

Errors that are more difficult to catch during the quality control process are listed in the table below. These typically impact long-term cycle life, and are difficult to measure in the relatively few cycles done during formation and after aging.

Table 4-2: Difficult-to-detect cell assembly manufacturing errors (Source: author's research)

Defect Cause	Effect	Mitigation
Improper jellyroll winding tension	Reduced cycle life due to improper electrode-to-electrode pressure within the spiral	High quality winding equipment that carefully measures the tension of the jellyroll winding, accounting for differences in force due to the increase in diameter of the jellyroll as it is wound.
Improper Dry Room Humidity Control (too much humidity)	Reduced cycle life due to side-reactions between water and the electrolyte	Carefully designed dry-rooms with stringent humidity control. Not allowing more than a few operators within the dry room at a given time due to humidity from human perspiration/exhalation.

Table 4-1 and Table 4-2 provide a list of defects that are typically watched for in advanced factories. It does not cover completely egregious errors that are seen in hand-assembly of cells in less developed facilities, ranging from uncontrolled dry rooms, hand-cutting of electrode current take-off tabs, etc.

4.4 Formation Cycling and Aging

After cells are assembled, they are placed on an electronic test rack containing hundreds of test channels. Here the cells are carefully charged for the first time, which activates the electrode materials and creates the SEI layer on the anode. The initial formation of this layer is important, as a smooth initial layer promotes its even growth in future cell cycling. The capacity and impedance of these cells is measured before they are charged to a high SOC (state-of-charge) and put into storage for aging.

The purpose of cell aging is to further build up the SEI layer in a controlled manner (which occurs more rapidly when the anode is at a high potential difference from the electrolyte, when the cell is fully charged) and to measure the rate of self-discharge. A high self-discharge rate can be indicative of particle contamination within the jellyroll/electrode-stack, a contaminated electrolyte solution, dendrite growth through the separator during SEI layer formation and potentially other issues that are not mentioned here. High and uneven self-discharge rates lead to cell balance issues once they are assembled into a battery pack (more on this in chapter 5.3)

4.5 Sorting/Binning/Classification

After the aging process, cells are typically sorted based on their capacity. The window for acceptable capacities can be as high as 10 percent, and manufacturers have different grades of cells at the end of their production line based on capacity.

Some manufacturers will also bin a lower grade of cells based on self-discharge rate, although there is usually only a single “B-grade” class for this, as opposed to the multitude of classes available for capacity.

There is usually a tolerance window for AC impedance measurements and cells outside of this range are rejects.

KEY POINT:

Binning by DCIR values hasn't been done at the factories that I visited, including two of the larger manufacturers in China who use highly automated production lines. Chapter 5.4.2 explains more about why I advocate for taking DCIR into consideration when binning.

Chapter 5 Battery Pack Design

Because the voltage of a lithium-ion cell is relatively low (less than 4 volts nominally), higher power applications require cells to be connected in series in order to minimize ohmic “copper losses” incurred from operating high-power loads at low voltages.

For greater ampere-hour capacity, either larger-capacity cells or multiple parallel-connected cells can be used.




Figure 5-1: Diagram of battery pack terminology

When constructing a multi-cell battery pack there are mechanical, thermal and electrical balancing considerations to be made. Because the focus of this thesis is on electrical current distribution and balancing issues, the chapters on mechanical and thermal design will be brief.

5.1 Mechanical Considerations

Mechanical constraint and vibration testing is especially important in portable applications, where stresses over time can lead to the failure of inter-connect tabs, cell enclosures and battery management wires. The fraying of wires and interconnects due to movement can lead to serious safety issues, such as internal battery pack short-circuiting.




Figure 5-2: Ford Fusion Hybrid battery pack module stack (left), disassembled module (right)

Figure 5-2 shows the constraint of cylindrical modules within a hybrid battery pack. The metal cell cans are end-welded to one another with polymer sleeves around the ends. This constrains the cells within the plastic module housings (which also act as cooling shrouds). The plastic modules each contain 8 cells and the modules interlock mechanically with tabs in the injection molded housing. These modules are then stacked 26 deep and are constrained with aluminum bars and cast aluminum endplates. In case anything fails, the cells are covered with 1mm-thick plastic shrouds for electrical insulation. The interconnect tabs and associated covers are also 1mm+ thick injection molded plastic.




Figure 5-3: Ford Fusion Hybrid battery pack module interconnect tabs (left), metallic pack enclosure (right)

High voltage pack assembly is also a concern, where it is desirable to have as few exposed connections as possible during the assembly process. The left photo in Figure 5-3 shows an interconnect tab with an integrated plastic housing (the metallic tab interconnect is hidden below the screws), that interlocks with the previous tab, covering those bolts and interconnections. In this way, the voltage of the pack is brought up gradually as these tabs are added to interconnect the cell modules. With each previous tab being covered, the maximum

exposed voltage between any two tabs in assembly is 20V (the voltage of the entire pack is 275V).


Electromagnetic interference (EMI) is also a concern when designing high voltage battery packs; the high-power cables that run from the motor controller to the battery pack are typically shielded; the battery pack itself is then usually in a metallic housing to prevent further EMI leakage that can interfere with other electronics onboard the vehicle.

5.2 Thermal Considerations

Operation at low temperatures can lead to lithium plating during cycling; operation at high temperatures leads to accelerated growth of the SEI layer and therefore faster cycle life degradation. Hybrid battery packs need to be cooled due to self-heating from high-rate charge/discharge cycling (because of internal resistance losses); they need to be heated to ensure proper operation at low temperatures.

Liquid-cooling has the advantage of allowing for more compact packaging. Air cooling is typically less complex. For both types of cooling, fluid flow simulation is necessary to ensure even heat distribution in as compact and lightweight of a package as possible (at least for mobile applications; large stationary battery packs usually don't have such stringent weight or size requirements).

To use the example of a custom battery pack that I designed with the MIT Electric Vehicle Team, 220 26650-sized LiFePO₄ cells were assembled into a 1.6kWh battery pack for a motorcycle. This was a demonstration project for a rapid-recharge battery pack (10 minute charge from 0 to 90% SOC) so cooling during charging was necessary.






Figure 5-4: (from top left, clockwise): CAD of battery module, actual battery module after completion, CFD streamlines of air moving throughout the battery module, CFD with the module housing, exploded view of cell mechanical constraint and tab electrical connections.

This battery pack was air-cooled during the charging process. Temperature sensors were placed on the cells that were estimated from the CFD analysis to have the most and least airflow (in the CFD streamline diagram, the cells on the top left layer of the pack receive lots of airflow, while those on the bottom left receive the least). The difference in temperature between those two cells was just one degree Celsius.

This type of design and analysis is straightforward from an engineering perspective, but is critical for proper pack operation and cycle life.

5.3 Pack Electrical Balancing

If all cells were created identically, there would be no need for a battery management system that monitors the voltage of each cell group. Lithium-Ion cells have a problem of bursting when they are over-charged, so a monitoring and balancing system is a safety-critical addition.

A battery management system (BMS) monitors the voltage of every parallel-connected cell group (3-4 volts, depending on battery chemistry; see the definition of “cell-group” in Figure 5-1) and balances end-of charge voltages to ensure that all cells are evenly charged, usually by shunting resistors across the cell groups with higher voltages than the rest of the battery pack. This keeps the energy distribution within the battery pack even and brings the “weaker” cells back to a full state-of charge.

Because of manufacturing differences, each cell will have a slightly different self-discharge rate. This means that if a battery pack has been unused for an extended period of time, some cell groups will have lower voltages than others. These voltages (each corresponding to a cell’s state of charge) need to be equalized as the pack approaches full charge. For well-manufactured packs, this becomes problematic on a time-scale of years. If they are not balanced, some cell groups with higher voltages have the potential to be over-charged, leading to safety problems.

What a BMS cannot monitor, however, is the current distribution that occurs within a parallel-connected cell group. Once two cells are connected to the same bus bar, their voltages become identical and to the BMS they both look like one larger cell. However, they don’t necessarily act as one larger cell once connected to a load.




Figure 5-5: Simplified diagram of cell internal resistance (DCIR)

Due to manufacturing differences, some cells will be “stronger” than others and have lower internal resistance (as defined by Equation 3-1). The DCIR value is assigned to the modeled “internal resistors” in Fig. 2.

A simple model of a cell is a voltage source in series with a resistor (where the “resistor” is the lumped model of reaction kinetics, ion diffusion, current collector resistance, etc). When

connected in parallel, cells with lower internal resistances will bear the brunt of the current; modeled as a simple equivalent circuit, more current will flow down the path of less resistance. If R_1 is less than R_2 , the cell on the left will take more of the current, becoming discharged faster and heating more internally.

This is an undesired effect, as it will lead to premature wear of the battery pack and can lead to safety issues in the event of extreme imbalances. This makes it important to quantify how cycle life is affected as a function of internal resistance variance.

5.4 Cell Matching

Cells must be matched in both capacity and internal resistance to extract the maximum usable performance from a battery pack.

5.4.1 Capacity Matching

The battery pack's total discharge capacity is limited by the capacity of the weakest cell (or parallel cell-group). As can be seen in the figure below, the total usable capacity of this battery pack of 4 cells in series is 2.0Ah, not 2.2Ah, due to the fact that the lowest-capacity cell will hit its low-voltage-cutout threshold first. The BMS senses this and will disconnect the pack from the load at this point. The capacity of the entire battery pack is reduced by 10% because of this one weak link.




Figure 5-6: Diagram of problems arising from poor capacity matching in series, limiting total pack capacity

5.4.2 Internal Resistance Matching

It is important to match the internal resistance of a battery pack both in parallel (within a cell-group) and in series (from one cell group to another).

5.4.2.1 Series DCIR Matching

In high discharge-rate applications (such as hybrid electric vehicles, power tools, r/c aircraft, etc), the voltage curve of each cell/parallel cell group will be dependent upon their respective DCIR values.




Figure 5-7: Voltage curve dispersion in a series battery pack (Source: Author's data)¹

Figure 5-7 shows a power-cell battery pack being discharged at a high rate, at constant current. The clear outliers in this pack are cells 1 and 6 (the dark blue and orange traces), whose voltage sag is far more than the rest of the pack. Because of its higher voltage sag (due to a higher DCIR value), cell 1 triggered the low-voltage cutoff on this discharge cycle. If we were to follow voltage curves of cells 2,3,4 and 5, the pack could have provided a few more Ah of capacity before cutting out. Again, the weakest link limits the capacity of the entire battery pack.

¹ The experimental setup for this chart: six LiNMC-based cells of 350 mAh capacity were connected in series and balanced to 4.2V per cell. All cells were brought to room temperature (25 °C) before discharge. They were discharged with the current profile displayed on the chart to generate this plot.

5.4.2.2 Parallel DCIR Matching

If the strongest and weakest cells from the battery pack in Figure 5-7 (in terms of DCIR, cell 2 is the strongest and cell 1 is the weakest) were connected in parallel, a current distribution curve similar to the one in Figure 5-8 would be measurable. The cell with the lower DCIR provides more current at the beginning of the discharge cycle. Assuming that both cells have similar capacity values, the lower DCIR cell will be depleted first (because of its higher average current load), after which the higher DCIR cell receives the burden of the entire current load.

Experiment setup: To demonstrate this effect, I conducted an experiment with two nearly identical, fully charged cells of the same type and model, connected in parallel. Current sensors were placed to each cell and they were discharged together. However, the temperature of one of the cells was kept at 65C, while the other was at 25. Because DCIR is correlated to temperature (higher temperatures allow for faster chemical reaction kinetics), the hotter cell had significantly lower DCIR than its neighbor. However, the capacity of both remains largely unchanged, as this value is not as strongly coupled to temperature gradients.

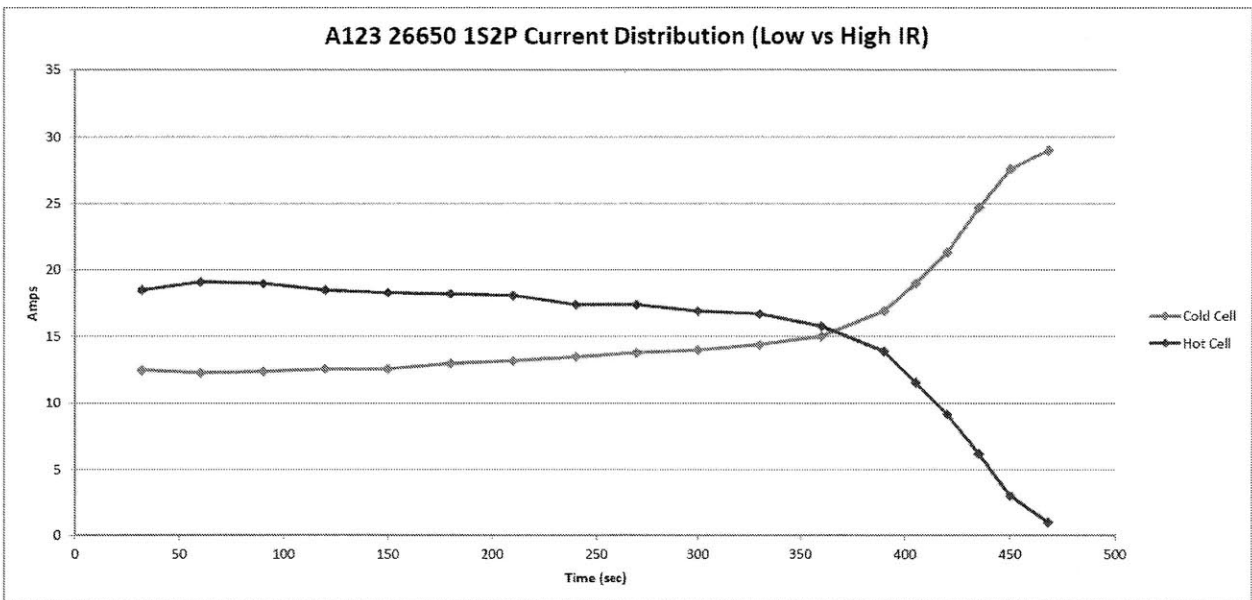


Figure 5-8: Current distribution between two parallel-connected cells with different DCIR values on discharge

KEY POINT:

The cold cell (with higher relative DCIR) still has energy left at around the 400 second mark, whereas the hot cell (with lower DCIR) has already been mostly depleted due to its higher average current load. Although the average current to each cell was supposed to be 15A (30A was applied to the 2-parallel cell group), at the end of the discharge cycle, the cell with the higher DCIR ended up taking almost twice the expected current load. This can lead to premature aging due to lithium plating if the cell is not designed to handle this abnormally high current.

Such uneven current loads can lead to premature wearing of the battery pack, especially because these current distributions are typically not monitored by the BMS. Placing a current sensor on each parallel-connected cell would be prohibitively expensive. For example, the Tesla Roadster battery pack has 99 cell groups in series and 69 cells in each cell group. Monitoring the current to each cell in the battery pack would require 6,831 current sensors.

Manufacturing differences will lead to this problem, but so will uneven cooling of the battery pack. Not only will the hotter cells degrade faster because they are kept at a higher temperature on average, but they are also being discharged at a higher C-rate (more discussion on this in 2.3).

5.5 The Correlation Between DCIR and C Rate

DCIR varies as a function of C rate. DCIR is an imperfect cell characterization metric as it attempts to lump together many electrochemical concepts into one value (more on this in chapter 3.1). However, differences in voltage sag are a serious issue in battery pack design. Although this method of cell characterization is rather crude from an electrochemical standpoint, the DCIR value is a good predictor of current distribution and voltage sag issues that arise in a battery pack.

As with many other mass-manufactured products, the variances from product-to-product tend to show up most when products are tested near the limits of their capabilities. To visualize this, the voltage profiles of each cell in the six-cell pack tested in Figure 5-9 have been graphed at different C-rates (Figure 5-10 to Figure 5-15).

The battery pack being tested is high-power model airplane battery pack that is rated at 65C continuous discharge (rather impressively, more than 80% of the total energy that was delivered at 1C discharge was also delivered at 52C).

The data in Figure 5-10 to Figure 5-15 was generated using the following test procedure:

- Six cells were connected in series, with the voltage of each cell being monitored
- All cells were balanced to 4.2V at the end of the charge cycle before being discharged
- The pack was brought to room temperature after the charging cycle, before the next discharge cycle
- Charging and voltage logging were done using an FMA Direct PowerLab 8 charger




Figure 5-9: Voltage traces at different C-rates for a high power LiCO_x based cell, on discharge

The high rate capability of this cell chemistry can be seen in the voltage sag of the higher current discharges (Figure 5-9), and the fact that it still delivers its rated capacity at discharge rates near 50C. However, the voltage curves in this figure are for only one cell selected from the 6 cell battery pack. Due to manufacturing variances, not all cells follow the same voltage profile as the cell shown in Figure 5-9. These differences are shown in Figure 5-10 to Figure 5-15.

Experimental setup: to generate these discharge profiles, six LiNMC-based cells of 350 mAh capacity were connected in series and balanced to 4.2V per cell. All cells were brought to room temperature (25 °C) before discharge. They were discharged with the current profiles displayed in the charts below.




Figure 5-10: Voltage curves for each cell in a 6S battery pack on discharge (1C)




Figure 5-11: Voltage curves for each cell in a 6S battery pack on discharge (5C)




Figure 5-12: Voltage curves for each cell in a 6S battery pack on discharge (10C)




Figure 5-13: Voltage curves for each cell in a 6S battery pack on discharge (17C)




Figure 5-14: Voltage curves for each cell in a 6S battery pack on discharge (30C)




Figure 5-15: Voltage curves for each cell in a 6S battery pack on discharge (52C)

Although the equipment used to cycle this battery pack had some software issues in maintaining a constant-current profile on the cells (especially at the 17C average rates and above), the cell-to-cell differences in voltage sag that is seen at the higher C rates is clear.

This pack is rated to “65C continuous discharge”, but IR balance issues are clearly seen at 30C average discharge rates and above. Even though the pack’s performance is impressive among any lithium-ion chemistry currently on the market (especially given its 130Wh/kg energy density at 1C), its performance starts being limited by cell imbalances at the 30C discharge level.

KEY POINT:

The takeaway from these graphs is the following: even though these cells may not be manufactured evenly enough to deliver optimal pack performance at 52C, at 10C discharge rates and below, their voltage curves (and corresponding DCIR values) are virtually indistinguishable. The discharge rate capability of a battery pack (when it comes to even current distribution) is limited by a combination of manufacturing variance and chemistry c-rate capability; this is summarized in Figure 5-16: Battery pack rate capability as a function of manufacturing tolerances and designed cell rate capability. Figure 5-16:

		Manufacturing Tolerances	
		Wide	Tight
Designed cell power capability (a function of active material composition and designed electrode layer thickness)	Low	Low C-rate packs (limited by current distribution and series voltage sag issues)	Medium C-rate packs (limited by cell design)
	High	Medium C-rate packs (limited by current distribution and series voltage sag issues)	High C-rate battery packs for high-performance applications

Figure 5-16: Battery pack rate capability as a function of manufacturing tolerances and designed cell rate capability

To add more commentary to Figure 5-16, the term “manufacturing tolerances” refers to the cell-to-cell variability within the cells of a battery pack. These target goals can either be achieved through tight manufacturing processes or through the sorting and binning of cells.

The benefit of having a tightly controlled manufacturing process is that more of the battery packs at the end of the production line end up working as expected. This issue can be approached from a few angles, as manufacturers can either:

- Improve their production processes in order to raise the overall C-rating of the battery packs that come of the production line
- Avoid the investments in production equipment and make the best of what they've got through sorting and binning
- Select better active material powders that can support C-rates that are far beyond what the battery pack application requires (allowing for wider manufacturing tolerances, but at the expense of lower energy density).

Chapter 6 Experiments

The purpose of this experimentation is to verify the modeling results from chapter 0, which are based on lumped-model assumptions. This is required to verify what happens in real-life systems that include more complex side-reactions.

The purpose of these experiments (and in a broader sense, the scope this thesis project) is to provide insight on the shape of the following curve:




Figure 6-1: Hypothetical curve of IR variability vs. battery pack cycle life

The shape of this curve could answer questions about how much variance in DCIR is allowable until there is a significant impact on the cycle life of a battery pack. Knowing the “knee” or inflection point of this curve is useful for both battery manufacturers and pack assemblers. However, because even cells with high IR differences will not exhibit a significant difference in IR distribution until they reach a certain C-rate (which is specific to the cell’s designed power capability), a third axis to the curve in Figure 6-1 is needed. The graph of interest now more resembles the following three-dimensional surface:

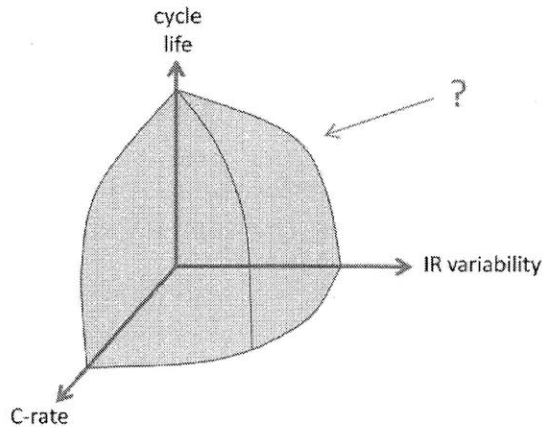


Figure 6-2: Hypothetical surface of IR variability vs. battery pack cycle life, accounting for C-rate

As mentioned in chapter 3.4, cells with lower DCIR will be more heavily loaded than other cells within the cell group. The heavier load should cause them to wear faster, increasing their DCIR values to the point where they eventually match the DCIR values of their neighbors. This is the hypothesis going into the experiment; that the curves of current distribution should converge as the cells age from cycling stress. Drawing from the current distribution shown in the high/low IR cell experiment in chapter 5.4.2.2, the progression of the current distribution profiles should resemble Figure 6-3:




Figure 6-3: Anticipated effect of cycling: convergence of the current distribution profiles

6.1 Experimental Design

The simplest way to begin studying the effects of current distribution is to start off with cell groups of two. The end goal is to predict what the current distribution looks like in a parallel-connected cell group of multiple cells. Because of the need for control samples, multiple tests of the same parameters to ensure repeatability, and the fact that cycle-life tests require hundreds of cycles (which can take months to complete, depending on C-rate), picking appropriate sample points is important to running an efficient experiment. Finding pairs of cell groups with the desired DCIR differences within a batch of manufactured cells also requires testing many more cells than will be used in the experiment, making the test parameter selection even more vital.

6.1.1 Selection of Sample Cells and Test Parameters:

The samples used for the experiment were cylindrical A123 LiFePO₄ 26650-sized 2.2Ah cells, for the following reasons:

- Their relatively small size (low Ah capacity) allowed for smaller and cheaper power electronics for the cycling tests
- The cells came from a relatively well-managed cell assembly facility, which I visited, so the cell-to-cell variance is expected to be from typical manufacturing tolerance issues (and not from wild errors that are found in some of the less reputable manufacturers which rely heavily on manual labor)

To determine the C-rate at which these cells would exhibit significant differences in DCIR values, a randomly selected sample of 8 cells was taken from a batch of 200, which were manufactured in 2006 (the shelf life of these cells is rated to 8 years). The dispersion in cell voltages as a function of C rate is shown in Figure 6-4 through Figure 6-7.

Experimental setup: For the following charts, eight A123 26650 LiFePO₄ cells were connected in series and balance charged to 3.65V per cell. All cells were brought to room temperature (25 °C) before discharge. They were discharged with the current profile displayed on the charts below to generate these plots.




Figure 6-4: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 4C discharge




Figure 6-5: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 6C discharge




Figure 6-6: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 8.5C discharge




Figure 6-7: Voltages profiles of individual cells in an 8S pack of A123 26650 cells, at 17C discharge

As can be seen in these graphs, the greatest spread in DCIR values was seen at 17C discharge rates. This is well within the cell's continuous discharge rating of 32C and its 10-second peak discharge rating of 55C. In order to sort these cells based on DCIR values, a C-rate of 17 was chosen in order to provide an easily measurable difference in voltage (lower C-rates would require the use of an op-amp or very precise voltage measurement with resolution below 1 mV).

72 cells were tested from this initial batch of 200 cells and their DCIR values were recorded using the following method:

1. Batches of 8 cells were connected in series and balance-charged to 3.6 V per cell using a computerized charger with 8 balance channels. Balance resolution was 5mV.
2. The entire cell group was discharged at 4C and 1 Ah was removed from the pack (coulomb-counted), bringing the pack to 55% SOC.
3. The battery pack was brought to room-temperature at 25 °C.
4. The 8S battery pack was connected to a 0.52 Ohm power resistor and the voltage of each cell was recorded 15 seconds after initial connection. The DCIR value was measured using Equation 3-1.

The reason for the 15 second delay can be seen in Figure 6-8; the longer delay time allows the effects of the mass-transfer limited reaction to show. Longer delay times could lead to significant self-heating of the cell, which can affect the measured DCIR value (this is the reason behind the positive slope of the cell voltages in Figure 6-7 after the 1 Ah mark).




Figure 6-8: Individual cell voltage traces from a DCIR pulse test

The DCIR values of 72 cells were recorded using this 17C pulse test method and sorted. From these values, the maximum difference between the lowest and highest IR cells was 24.7%. The maximum difference in capacity in this same batch of cells (one full discharge cycle at 17C) was 3.6%. For the purposes of this experiment, the differences in initial capacity were assumed to be negligible compared to the differences in DCIR.

6.2 Validation Experiment

Before setting up a mass cell-cycling experiment, I conducted a quick first-pass experiment with one cell-group was set up to ensure that the effects of uneven current distribution due do DCIR variance in manufacturing were significant and worth measuring.

I paired together two A123 26650 cells with a 16.4% difference in DCIR (as measured per the method in Chapter 6.1) and cycled them at an 8C charge and 8C discharge rate. There was a measurable difference in charge and discharge current to each cell. The shape of the distribution profile changed as the cells aged from cycling, as can be seen in Figure 6-9.




Figure 6-9: Current distribution profiles in a 2P A123 26650 cell group, at cycles 10-12 (left) and 300-302 (right)

The average current difference between these two cells on the discharge cycle grew 230% from cycle 11 to cycle 301. The spikes in the current measurement are due to the PWM switching of the DC/DC buck-boost converter inside the charger used for automation. Any current measurement between -17 A and +17 A was regarded as an outlier and those data points were removed (which is why the tapering of the current in the constant-voltage phase of cycling is not visible).

The capacity fade of the cell group can be seen in Figure 6-10 (graph on the right), with the current difference becoming drastic at the end of the cycle life (graph on the left).




Figure 6-10: Current distribution profile at cycle 500-502 (left) and cell-group capacity fade (right)

This initial data set proved to be different from the original hypothesis. The current profiles were initially predicted to converge (due to faster degradation of the more highly-stressed cell). However, when calculating the capacity of each cell individually throughout cycling, it can be seen that one cell's capacity is fading more quickly than the other (Figure 6-11):




Figure 6-11: Capacity fade of each cell in the 2P A123 26650 cell group (left)
 Temperature of each cell throughout the cycling process (0-602 cycles, graph on the right)

The spikes in the capacity measurement were due to the brief moments when the cyclers were shut off and re-started on a “charge” cycle. These measurements are outliers and for the purpose of data analysis can be ignored.

It is important to note that the cells in this test were not actively cooled; they were left standing in free air (bare cylindrical cell cans). The temperature of the cell receiving more of the current increased throughout the cycling process, as it began receiving more and more current.

At this point, it was uncertain as to why one of the two cells was decreasing in capacity much more rapidly than the other; they both started out with virtually the same capacity. In any sense, it appeared that the current distribution to each cell was diverging and having a positive-feedback effect. It was also possible that one of the two cells was simply a manufacturing “dud”. Either way, more testing was deemed necessary to analyze these effects.

6.3 Selection of test parameters for mass cell cycling experiment:

From these 72 cells that were originally characterized for the validation experiment, it was possible to arrange 4 cell groups that each had between 0% and 0.5% differences in DCIR between the cells, 4 cell groups between 9.5% and 10%, and 8 cell groups between 18% and 20%. This allowed sampling of the curve at the points indicated in Figure 6-12 ; the cell cycling parameters are also expressed in

Table 6-1.




Figure 6-12: Sampling areas for the curve of DCIR vs. cycle life (red line is hypothetical)

Each shaded sample area was trialed twice (2 cell groups per shaded area) to add to the repeatability of results. For only the $\Delta 18-20\%$ region, cell group pairs were also tested with liquid cooling. The two tests at 4C and at 8C are identical in all regards other than C-rate.

The 4C and 8C cycling tests were done with the following parameters

- Constant current charging to 3.65V per cell group
- Constant voltage at 3.65V with termination at C/3
- 1 minute rest period between charge and discharge
- Constant current discharge to 2.40V
- 1 minute rest period between discharge and charge

Table 6-1: Test parameters of cycle life testing

Sample		Low DCIR Delta Group (0-1%)			Med. DCIR Delta Group (9-10%)			High DCIR Delta Group (18-20%)		
		Cell 1 DCIR	Cell 2 DCIR	DCIR Ratio	Cell 1 DCIR	Cell 2 DCIR	DCIR Ratio	Cell 1 DCIR	Cell 2 DCIR	DCIR Ratio
Air Cooling, 4C	1	0.01569	0.01569	1.000	0.01977	0.01794	1.102	0.02103	0.01749	1.203
	2	0.01569	0.01566	1.002	0.01961	0.01791	1.095	0.02023	0.01708	1.184
Air Cooling, 8C	1	0.01507	0.01505	1.001	0.01949	0.01778	1.096	0.02085	0.01733	1.203
	2	0.01362	0.01358	1.003	0.01938	0.01758	1.103	0.02023	0.01710	1.183
Liquid Cooling, 4C	1							0.02001	0.01697	1.180
	2							0.02052	0.01729	1.187
Liquid Cooling, 8C	1							0.01997	0.01693	1.179
	2							0.01996	0.01687	1.183

The 8C charge and discharge rate was chosen for half of the cell parameters due to the interesting results achieved in the validation experiment. Although cycle life was diminished by high charge/discharge rates, there are some high power applications for which 400 cycles are sufficient (such as power tools). Given the high stress of this application and the fact that cooling usually improves cycle life (the electrolyte degrades more rapidly at temperatures over 50 degrees C), forced air cooling was added to all cells.

The high differences in DCIR values between cells was of particular interest, so more tests were carried out in the “high DCIR delta group”. Liquid cooling was added to another test group to further reduce the effect of any thermal gradients between the cells.

These cells were tested on cycling equipment consisting of 16 FMA Direct PowerLab 6 chargers (capable of 40A charge and 40A discharge) and an Agilent 34980A data acquisition system (with 40ch and 70ch multiplexer cards). The forced-air cooled cells were placed in a custom-built flame retardant enclosure to partition off each cell group, while the liquid-cooled cells were immersed in distilled water, with a chiller keeping the water temperature between 23 and 25 degrees Celsius.




Figure 6-13: Experimental cell-cycling setup

Chapter 7 Experimental Results and Discussion

In this chapter I will address trends that were found in data from cell-cycling. A list of conclusions that can be drawn from this information is at the end of the chapter. Data from the 4C charge / 4C discharge air-cooled cells is covered first, expanding on the following topics:

- Capacity fade vs. initial cell group DCIR Δ
- Initial capacity vs. initial cell group DCIR Δ
- Convergence/Divergence/Inversion of current distribution between the cells in a cell group
- The influence of rest periods on the current distribution between cells

7.1 Capacity fade vs. initial DCIR Δ




Figure 7-1: Capacity fading for cell test group at 4C

The test results for the cells that were cycled at 4C are displayed in Figure 7-1. The x-axis represents the difference in DCIR between the two cells in each cell group; the y-axis represents the fraction of the cell group's initial capacity that remained after 1,100 cycles.

While there appears to be a correlation between initial DCIR Δ and the average capacity fade of each cell, the variance in these measurements is high; there are only six sample points that form the curve. The difference in cycle life fading is significant between each DCIR Δ bin (Low,

0-1%, Medium, 9-10% and high, 18-20%). More data for this would be nice but it is not easy to obtain, as it involves testing many cells from one batch in order to get the desired DCIR Δ values and then cycling those cell groups for over one month.

It is important to note that on an individual cell level, there is almost no correlation between initial DCIR and capacity fade, as evidenced by the low R^2 value of the linear fit.




Figure 7-2: Initial DCIR values vs. capacity fade of air-cooled cells cycled at 4C

KEY POINT:


The low correlation between initial DCIR value and capacity fade supports the argument that capacity fade is faster within mismatched parallel-connected cell groups, and that this fade is independent of the initial DCIR values of the individual cells.

7.1.1 Analysis of current distribution within air-cooled 4C cell groups

In order to gain more insight as to why there was such a large difference in capacity fading between the two samples in the High DCIR Δ group, the current distribution between the cells in the parallel-connected cell groups was analyzed.

The graphs of the individual cell capacities (within the cell group) over time are shown in Figure 7-3, as well as the average current per discharge cycle to each cell. Graphs of sample 1 are on the left, Sample 2 are on the right.

High DCIR Δ Sample 1



High DCIR Δ Sample 2






Figure 7-3: 4C air-cooled high DCIR Δ cell group capacity and average current over time. (Left = Sample 1) (Right = Sample 2)

The capacity of cell 2 in sample 2 (the red trace) degraded more quickly over time, leading to an inversion in the average current distribution to each cell (a cell with less capacity will take less average current during the discharge period). This effect can also be seen in the graphs of the current to each cell, in Figure 7-4.




Figure 7-4: 4C air-cooled high DCIR Δ Sample 2 current traces at cycles 20-22 (left) and 1395-1397 (right)

Cell 2 started out as the “stronger” cell with lower DCIR and then became the “weaker” cell, starting at about cycle number 400 (see the bottom right graph in Figure 7-3). The current distribution between cells inverted again after about 1900 cycles. This same effect did not occur in sample 1, where the cell that started out strongest at the beginning remained that way throughout the cycle testing (see Figure 7-5).




Figure 7-5: 4C air-cooled high DCIR Δ Sample 1 current traces at cycles 20-22 (left) and 1395-1397 (right)

There are three ways that the plots of average current to each cell can evolve: the traces can converge, diverge, or invert (cross).

7.1.1.1 Current distribution convergence




Figure 7-6: Current distribution convergence (Left = High DCIR Δ Sample 2) (Right = Low DCIR Δ Sample 1)

Of the six samples, two of the current distribution profiles converged; these were from one sample each of the high and low DCIR Δ groups. This is the expected degradation mechanism between two cells, as the cell that takes more of the current is expected to wear out faster and experience a faster rise in DCIR, eventually taking less current.

7.1.1.2 Current distribution divergence

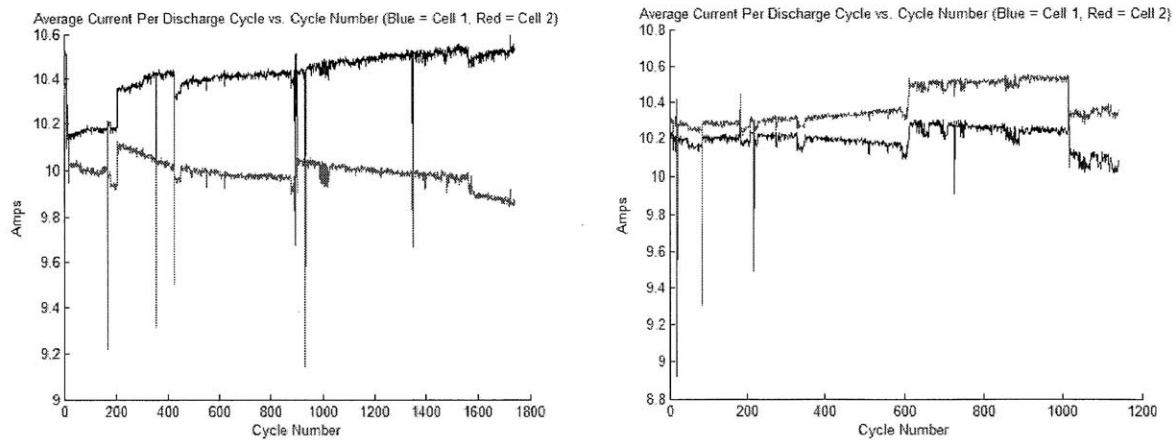


Figure 7-7: Current distribution divergence (Left = Medium DCIR Δ Sample 1) (Right = Medium DCIR Δ Sample 2)

The current distribution profiles of both samples from the medium DCIR Δ group diverged. This is due to the faster capacity degradation of one cell from that of its neighbor. Although these current profiles have been diverging up to this point (1800 cycles for sample 1, 1100 cycles for sample 2), I expect that they will begin to level off once the average current difference gets great enough to begin significantly affecting the life of the other cell. This requires further cycling.

7.1.1.2.1 The effect of rest periods

Another effect to note in Figure 7-7 is the effect of rest-periods. Each of the spikes downward on these graphs represents a period when the cycling equipment was re-started (this disturbed the coulomb counting algorithm and created an outlier). After these rest periods, the difference in current distribution between the two cells is not as large as it was before the rest period, although the curves begin to diverge afterwards again. A possible reason for this is chemical stabilization within the electrode as it can take up to a few hours for the electrode to internally equalize (see Figure 2-7).

The jumps in overall current (for both cells) are due to differences in the charger's current-set programming.

7.1.1.3 Current distribution crossing / convergence




Figure 7-8: Current distribution inversion (Left = High DCIR Δ Sample 2) (Right = Low DCIR Δ Sample 1)

Sample 2 of the high DCIR Δ group and Sample 1 of the low DCIR Δ group displayed an inversion in current profiles. I hypothesize that these current level inversions will keep occurring as the cell group ages. The difference in current distribution to each cell is relatively minor (less than 4%, for example, for the graph on the left), and a certain current differential is required to get a cell to age significantly more rapidly than the other. Why these curves cross, instead of asymptotically approaching one another, is still uncertain.

7.1.1.4 Initial DCIR value vs. Initial Capacity




Figure 7-9: Initial DCIR value vs. initial capacity of A123 26650 cells (capacity value is averaged over the first 10 cycles)

For the A123 26650 LiFePO₄ cylindrical cells, it is impossible to measure electrode thickness in the same way that it was done with the 2.1Ah pouch cells in chapter 4.2.1. The correlation between DCIR value and capacity is similarly vague (see Figure 4-5: Pouch cell capacity vs. DCIR). This demonstrates that there are more factors to determining DCIR and capacity than simply electrode coating thickness, among cells of the same type.

Manufacturing issues that can affect DCIR and capacity values:

- Electrode coating thickness
- Variance in how finely the active material powders are ground
- Variance in the porosity of the separator
- Chemical variances within the electrolyte and active material

The capacity of a cell is also typically limited by the cathode side and not the anode. Internally, manufacturers typically keep the tolerances of the cathode electrode coating tighter for this reason. Situations could arise where a thick anode and a thin cathode create a cell that has both a low capacity and high DCIR.

KEY POINT:

DCIR, as described in this paper, may also not be a very repeatable measurement of a cell's capabilities. Issues such as rest periods (covered in chapter 7.1.1.2.1) of days to weeks can affect a cell's voltage during pulse discharging. Although the methods to measure DCIR as described in this paper were tightly controlled, studies of the repeatability of the measurement outcome (on the same cells) were not carried out.

7.2 Test results from cells cycled at 8C

While I originally believed that adding air-cooling to the validation experiment cell group in chapter 6.2, this proved to be an error in experimental design. All cell groups fell below to 70% of their original capacity in less than 300 cycles, with the higher DCIR Δ cell groups degrading much faster than that. I will not further analyze these graphs due to the impracticality of using cells under such conditions. While the initial validation experiment proved promising with a cycle life of over 500 cycles to 70% of initial capacity, the forced air cooling used in the mass cell cycling experiments reduced the average temperature of the cells from about 70°C to 35°C. While the cooling effect was initially believed to be beneficial in reducing SEI layer buildup and electrolyte degradation (which happens more quickly at higher temperatures) the extremely high C rates are believed to have led to lithium plating at the reduced temperatures, which caused even more rapid degradation. If done again, I should have included forced-air cooling in the validation experiment to screen for this issue.

7.3 Effects of water cooling on cells being cycled at 4C and 8C

Adding water cooling furthered the effect of degradation due to lithium plating at below-optimal operating temperatures; the cells being cycled at 8C degraded even faster than the air-cooled cells. The water cooling also negatively impacted the cells that were being cycled at 4C, for two reasons. One, lithium plating that is similar to what the air-cooled 8C cell groups experienced. Both water-cooled 4C samples fell to less than 80% of their initial capacity within 300 cycles (the air-cooled cells being tested at 4C maintained more than 80% of their capacity at well past 1500 cycles). I believe that this is due to lithium plating arising from the sub-optimal operating temperature. Two, the capacities of both samples fell drastically past 300 cycles due to the physical growth of a layer between the negative cell terminal and the can; while distilled water was used, this growth is believed to be a salt deposition that added to the resistance of the cell. If done again, I would use a more neutral fluid such as transformer oil instead of distilled water. I would also implement a more thorough validation test of cycle life vs. cell temperature before carrying out further samples to confirm results.

7.4 Experimental validation of modeling results

As evidenced by Figure 7-2, the variance in experimental measurements was high and much of the anticipated data was lost due to unexpected lithium plating that occurred due to over-cooling of the cells. In order to validate the model in Chapter 3.4, more experimentation is required. However, the fact that the degradation curves for the individual cells appear to follow a linear pattern is promising for the model's accuracy.

One way to collect data for this modeling experiment more quickly than by DCIR Δ matching is by conducting cycle life tests at slightly different C-rates, indicative of the difference in current distribution that these cells would face in an imbalanced DCIR case.

Chapter 8 Summary and Future Work

8.1 Summary of key points:

From a manufacturing standpoint, it is important to precisely control the thickness of the anode and cathode electrode layers; if they vary from the design specification, the cell will end up performing unexpectedly (a high-power cell that has too thick of an electrode active material layer will behave like a high-energy cell and therefore have higher internal resistance).

Variations in the thickness of the electrode *within* a cell can lead to localized hot. Variations in electrode thickness *from cell to cell* can lead to an entire cell taking more of the current when placed in parallel with other cells in a battery pack.

Mismatching of capacity *or* DCIR in parallel-connected cells can lead to one cell being worn out faster, and the remaining cells taking a much higher current toward the end of the discharge cycle. This can lead to premature aging due to lithium plating if the neighboring cells are not designed to handle this abnormally high current.

Binning by DCIR values hasn't been done at the factories that I visited, including two of the larger manufacturers in China who use highly automated production lines. Binning by DCIR is beneficial for pack assemblers who are certain of the charge/discharge profiles that their cells will experience, and they can make sure that these cells are well below the C-rate at which the voltage curves begin to disperse (from cell to cell within a batch).

The low correlation between initial DCIR value and capacity fade supports the argument that capacity fade is faster within mismatched parallel-connected cell groups, and that this fade is independent of the initial DCIR values of the individual cells.

The repeatability of the high C-rate DCIR pulse test may hinder its use as a metric of cell performance. Issues such as rest periods (covered in chapter 7.1.1.2.1) of days to weeks can affect a cell's voltage during pulse discharging. Although the methods to measure DCIR as described in this paper were tightly controlled, studies of the repeatability of the measurement outcome (on the same cells) were not carried out.

If a cell's capacity degradation curve is non-linear (as a function of DOD and C-rate), there will be optimization points that make it possible to plot battery pack cycle life as a function of initial variance in cell internal resistances. This can lead to further optimization in industrial planning, where the cost to achieve those tolerances is known. Thus, cycle life can be plotted as a function of incremental manufacturing cost. The almost-linear degradation profiles of the LiFePO₄ cells that were tested in this thesis makes this less of an issue, but it could be more interesting for other lithium-ion chemistries with less linear degradation profiles.

8.2 Future work

If I had an extra year to continue work on this project, I would expand on the following:


- Testing of DCIR Δ vs. cycle life among other chemistries that don't have a flat discharge curve like LiFePO_4 , where the interactions between cells may be mitigated in the middle of the discharge cycle instead of at the very end of it.
- Testing DCIR variance vs. cycle life vs. cell thickness with a prismatic pouch cell, for which total cell thickness measurement is feasible.
- Determining at which C-rates DCIR pulse current measurements should be taken, and how long these pulses should last, in order to develop a rigorous test procedure for binning cells based on DCIR values. Determining the repeatability of DCIR measurements with this technique is also critical for creating a practical test procedure.
- Developing an empirical model with this experimental data and expand it to multi-parallel cell packs.

Chapter 9 References





- [1] Gholam-Abbas, N., 2009, *Lithium Batteries: Science and Technology*, Springer.
- [2] Doyle, M., Fuller, T. F., and Newman, J., 1993, "MODELING OF GALVANOSTATIC CHARGE AND DISCHARGE OF THE LITHIUM POLYMER INSERTION CELL," *Journal of the Electrochemical Society*, 140(6), pp. 1526-1533.
- [3] Botte, G. G., Johnson, B. A., and White, R. E., 1999, "Influence of some design variables on the thermal behavior of a lithium-ion cell," *Journal of the Electrochemical Society*, 146(3), pp. 914-923.
- [4] Tarascon, J. M., and al, e., 1992, *Electrochim. Acta*.
- [5] Arora, P., White, R. E., and Doyle, M., 1998, "Capacity fade mechanisms and side reactions in lithium-ion batteries," *Journal of the Electrochemical Society*, 145(10), pp. 3647-3667.
- [6] Vetter, J., Novak, P., Wagner, M. R., Veit, C., Moller, K. C., Besenhard, J. O., Winter, M., Wohlfahrt-Mehrens, M., Vogler, C., and Hammouche, A., 2005, "Ageing mechanisms in lithium-ion batteries," *Journal of Power Sources*, 147(1-2), pp. 269-281.
- [7] Ratnakumar, B. S. V., Smart, M. C., and Surampudi, S., 2002, "Electrochemical impedance spectroscopy and its applications to lithium ion cells," *Seventeenth Annual Battery Conference on Applications and Advances, Proceedings*, pp. 273-277.
- [8] Christensen, J., and Newman, J., 2004, "A mathematical model for the lithium-ion negative electrode solid electrolyte interphase," *Journal of the Electrochemical Society*, 151(11), pp. A1977-A1988.
- [9] K.E. Thomas, R. M. D., J. Newman, 2002, *Advances in Lithium-Ion Batteries*, Kluwer Academic Publishers.
- [10] Verbrugge, M. W., Baker, D. R., and Koch, B. J., 2002, "Mathematical modeling of high-power-density insertion electrodes for lithium ion batteries," *Journal of Power Sources*, 110(2), pp. 295-309.
- [11] Ning, G., and Popov, B. N., 2004, "Cycle life modeling of lithium-ion batteries," *Journal of the Electrochemical Society*, 151(10), pp. A1584-A1591.
- [12] Dubarry, M., and Liaw, B. Y., 2007, "Development of a universal modeling tool for rechargeable lithium batteries," *Journal of Power Sources*, 174(2), pp. 856-860.
- [13] Phostech Lithium, I., 2012, "LifePower P2 submicron C-LiFePO4 active material," http://www.phostechlithium.com/prd_LiFePO4P2_e.php.

Chapter 10 Appendices





10.1 Combined capacity fade of all 4C air-cooled cells





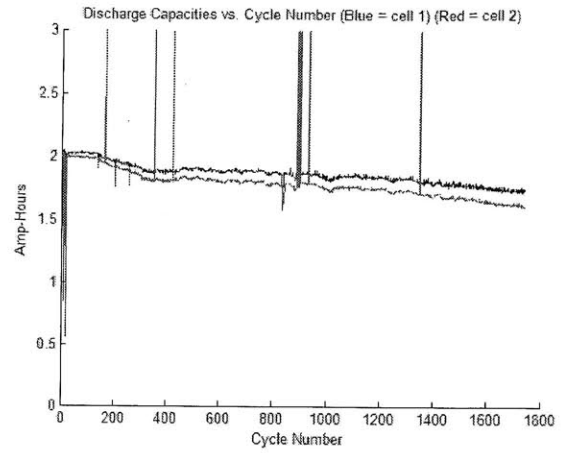

10.2 Cycling data from High DCIR Δ Cell group Sample 1 (initial DCIR Δ = 18.4%)







10.3 Cycling data from High DCIR Δ Cell group Sample 2 (initial DCIR Δ = 20.3%)







10.4 Cycling data from Medium DCIR Δ Cell group Sample 1 (initial DCIR Δ = 9.6%)



10.5 Cycling data from Medium DCIR Δ Cell group Sample 2 (initial DCIR Δ = 9.5%)



10.6 Cycling data from Low DCIR Δ Cell group Sample 1 (initial DCIR $\Delta = 0.17\%$)



10.7 Cycling data from Low DCIR Δ Cell group Sample 2 (initial DCIR Δ = 0.0%)

