

BATTERY MANAGEMENT SYSTEMS

Equivalent-Circuit Methods

Gregory L. Plett

VOLUME II

Battery Management Systems

Volume II

Equivalent-Circuit Methods

For a listing of recent titles in the
Artech House Power Engineering and Power Electronics,
turn to the back of this book.

Battery Management Systems

Volume II

Equivalent-Circuit Methods

Gregory L. Plett



**ARTECH
HOUSE**

BOSTON | LONDON
artechhouse.com

Library of Congress Cataloging-in-Publication Data

A catalog record for this book is available from the U.S. Library of Congress.

British Library Cataloguing in Publication Data

A catalog record for this book is available from the British Library.

ISBN-13: 978-1-63081-027-6

Cover design by John Gomes

**© 2016 Artech House
685 Canton Street
Norwood, MA 02062**

All rights reserved. Printed and bound in the United States of America. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the publisher.

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Artech House cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

10 9 8 7 6 5 4 3 2 1

Contents

Preface ix

1	<i>Battery-Management-System Requirements</i>	1
1.1	<i>Battery-pack topology</i>	5
1.2	<i>BMS design requirements</i>	7
1.3	<i>Requirement 1a. Battery-pack sensing: Voltage</i>	8
1.4	<i>Requirement 1b. Battery-pack sensing: Temperature</i>	9
1.5	<i>Requirement 1c. Battery-pack sensing: Current</i>	11
1.6	<i>Requirement 1d. High-voltage contactor control</i>	13
1.7	<i>Requirement 1e. Isolation sensing</i>	15
1.8	<i>Requirement 1f. Thermal control</i>	18
1.9	<i>Requirement 2. Protection</i>	19
1.10	<i>Requirement 3a. Charger control</i>	20
1.11	<i>Requirement 3b. Communication via CAN bus</i>	21
1.12	<i>Requirement 3c. Log book function</i>	22
1.13	<i>Requirement 4a. State of charge estimation</i>	22
1.14	<i>Requirement 4b. Energy estimation</i>	26
1.15	<i>Requirement 4c. Power estimation</i>	27
1.16	<i>Requirement 5. Diagnostics</i>	29
1.17	<i>Where to from here?</i>	29
2	<i>Simulating Battery Packs</i>	31
2.1	<i>Modeling battery cells</i>	31
2.2	<i>Modeling approach 1: Empirical</i>	32
2.3	<i>Modeling approach 2: Physics-based</i>	36
2.4	<i>Simulating an EV</i>	40
2.5	<i>Equations for vehicle dynamics</i>	41
2.6	<i>EV simulation code</i>	46
2.7	<i>EV simulation results</i>	54
2.8	<i>Simulating constant power and voltage</i>	55
2.9	<i>Simulating battery packs</i>	59

2.10	<i>PCM simulation code</i>	62
2.11	<i>Example PCM results</i>	64
2.12	<i>SCM simulation code</i>	66
2.13	<i>Example SCM results</i>	66
2.14	<i>Where to from here?</i>	67
3	<i>Battery-State Estimation</i>	69
3.1	<i>SOC estimation</i>	70
3.2	<i>A careful definition of state of charge</i>	71
3.3	<i>Some approaches to estimate SOC</i>	74
3.4	<i>Review of random processes</i>	81
3.5	<i>Sequential probabilistic inference</i>	95
3.6	<i>The linear Kalman filter</i>	101
3.7	<i>The extended Kalman filter</i>	114
3.8	<i>Implementing an EKF using the ESC cell model</i>	121
3.9	<i>Problems with EKF, improved with sigma-point methods</i>	130
3.10	<i>The SPKF</i>	135
3.11	<i>Implementing SPKF using the ESC cell model</i>	142
3.12	<i>Real-world issues pertaining to sensors, initialization</i>	147
3.13	<i>Reduced computational complexity using bar-delta filtering</i>	151
3.14	<i>Where to from here?</i>	159
3.15	<i>Appendices: Algorithms</i>	160
4	<i>Battery Health Estimation</i>	167
4.1	<i>Need for health estimates</i>	167
4.2	<i>Negative-electrode aging</i>	171
4.3	<i>Positive-electrode aging</i>	177
4.4	<i>Sensitivity of voltage to R_0</i>	178
4.5	<i>Code to estimate R_0</i>	181
4.6	<i>Sensitivity of voltage to Q</i>	183
4.7	<i>Estimating parameters via Kalman filters</i>	184
4.8	<i>EKF parameter estimation</i>	186
4.9	<i>SPKF parameter estimation</i>	190
4.10	<i>Joint and dual estimation</i>	191
4.11	<i>Robustness and speed</i>	194
4.12	<i>Unbiased estimate of total capacity using linear regression</i>	195
4.13	<i>Weighted ordinary least squares</i>	197
4.14	<i>Weighted total least squares</i>	200
4.15	<i>Goodness of model fit</i>	204
4.16	<i>Confidence intervals</i>	205

4.17	<i>Simplified total least squares</i>	207
4.18	<i>Approximate full solution</i>	210
4.19	<i>Code to simulate the methods</i>	217
4.20	<i>Example HEV simulations</i>	220
4.21	<i>Example EV simulations</i>	224
4.22	<i>Discussion of simulations</i>	226
4.23	<i>Where to from here?</i>	227
4.24	<i>Appendices: Algorithms</i>	228
5	<i>Cell Balancing</i>	237
5.1	<i>Causes of imbalance</i>	238
5.2	<i>Not causes of imbalance</i>	240
5.3	<i>Balancer design choices</i>	242
5.4	<i>Circuits for balancing</i>	245
5.5	<i>How quickly must I balance?</i>	253
5.6	<i>Balancing simulation results</i>	259
5.7	<i>Where to from here?</i>	262
6	<i>Voltage-Based Power-Limit Estimation</i>	263
6.1	<i>Voltage-based power limits</i>	264
6.2	<i>Limits via simple cell model</i>	266
6.3	<i>Limits via full cell model</i>	269
6.4	<i>Bisection search</i>	271
6.5	<i>Where to from here?</i>	277
7	<i>Physics-Based Optimal Controls</i>	279
7.1	<i>Minimizing degradation</i>	279
7.2	<i>SEI formation and growth</i>	282
7.3	<i>SEI ROM results</i>	289
7.4	<i>Lithium plating on overcharge</i>	292
7.5	<i>Plating ROM results</i>	294
7.6	<i>Optimized power limits</i>	298
7.7	<i>Plug-in charging</i>	299
7.8	<i>Fast-charge example</i>	301
7.9	<i>Dynamic power calculation</i>	303
7.10	<i>Where to from here?</i>	306
	<i>About the Author</i>	313
	<i>Index</i>	315

Preface

This book comprises the second volume in what is planned to be a three-volume series describing battery-management systems. The first volume focused on deriving mathematical sets of equations or *models* that describe how battery cells work, inside and out. This second volume applies equivalent-circuit style models to solve problems in battery management and control. The third volume will show how physics-based models also can be used to solve problems in battery management and control, leading to better results. The intent of the series is not to be encyclopedic; rather, it is to put forward only the current best practices, with sufficient fundamental background to understand them thoroughly.

This particular volume is organized in the following way:

- Chapter 1 introduces the requirements of a battery-management system. These include sensing, control, protection, state and health estimation, and communications.
- Chapter 2 reviews equivalent-circuit models of lithium-ion cells and shows how to use them to simulate the response of a battery pack to an input stimulus.
- Chapter 3 investigates battery-cell state estimation. Nonlinear Kalman filters are shown to give very good estimates, along with dynamic error bounds that enable confident use of the estimates when computing battery-pack energy and power.
- Chapter 4 looks at state-of-health estimation. A simple method can be used to estimate cell resistance, but we find that it is more difficult to estimate cell total capacity. A regression technique based on total-least squares provides optimal unbiased results.
- Chapter 5 discusses cell balancing. It considers factors that lead to imbalance, some questions that must be addressed when designing a balancer, and some circuit options for balancing cells.
- Chapter 6 explores computation of power limits where terminal-voltage constraints are applied. The simple method from Chap. 1 is extended to work with a full equivalent-circuit cell model.

- Finally, Chap. 7 exposes the fundamental flaw with voltage-based power-limit estimates and introduces physics-based methods that can be used alongside a circuit model to give better limits.

The intended audience for this material is someone with an undergraduate degree in engineering—principally electrical or mechanical. The reader does not need to be intimately familiar with all the concepts from Volume I of this series to be able to benefit from the topics in this volume. However, the deeper understanding that can be developed by studying Volume I will add richness to the study, and will aid understanding some concepts explored in Chaps. 2, 4, and 7 of this book that would otherwise be quite opaque.

The content in this book has been taught multiple times to students of diverse backgrounds in *ECE5720: Battery Management Systems* at the University of Colorado Colorado Springs. Lecture notes and lecture videos are available at <http://mocha-java.uccs.edu/ECE5720/index.html>. As the lecture videos sometimes explain the concepts of this book in a somewhat different way, the additional perspective may be an advantage to the learner.

I am greatly indebted to a number of my colleagues and students who have supported and assisted me over the years in understanding and developing the theory and methods presented in this work. First, I would like to acknowledge Dr. Daniel Rivers, founding CEO of Compact Power, Inc. (now, LGCPI), who introduced me to this field in the first place. Without his support and encouragement, I never would have studied battery management, and this book would not exist. I would also like to thank Dr. Saeed Siavoshani for inviting me to participate as an instructor in an SAE Hybrid-Electric-Vehicle Academy some years ago. Much of the material of this book was first developed to present at that academy; it was then expanded to become the *Battery Management Systems* course just mentioned, and now it has matured to book form. Among my students, I owe a special debt of gratitude to Mr. Lukas Aldrich and Mr. Kirk Stetzel for critiquing the content and helping with many of the examples, and to Mr. Alfred Randall and Mr. Roger Perkins for their work on reduced-order models of cell degradation that are presented in Chap. 7. My colleague and friend Dr. M. Scott Trimboli has also been a great encourager of this work, as he was with Volume I.

Despite my best intentions, there are certain to be errors and confusing statements in this book. Please feel free to send me corrections and suggestions for improvements.

1

Battery-Management-System Requirements

This book investigates the proper management of battery packs, a task that requires both hardware (electronics) and software (computer program) components. The hardware elements incorporate electronic circuits to ensure the safety of the battery pack and its operator and to make measurements that include battery-cell voltages, electrical current, and temperature. The software portions monitor and coordinate the activities of the battery pack.

While we look at both hardware and software aspects in this book, we devote most of our attention to software methods or *algorithms*. These implement mathematical calculations that use measured data to estimate and summarize battery-pack present operational status and to predict its near-future performance limits. And, although most of the approaches that we will discuss can be applied to battery packs comprising cells of any chemistry, we will focus on applications involving lithium-ion battery cells.

A survey of the relevant literature uncovers many different methods for different aspects of battery management. We will explore some simple methods to introduce many of the key concepts but will devote most of our study to some more complex but also more accurate and robust approaches, which we prefer. However, we recognize that an implementation of a more complex algorithm requires more processing power—and hence a greater cost—than a simpler counterpart, so these algorithms are best suited to applications involving mission-critical systems or large battery packs comprising many battery cells where a substantial investment must be protected.

The methods and algorithms we discuss would typically be implemented by a *battery-management system* or BMS.¹ A BMS is an *embedded system*; that is, purpose-built electronics plus processing to enable a specific application. For example, Fig. 1.1 shows the electronics por-

1.1	Battery-pack topology	5
1.2	BMS design requirements.	7
1.3	Rqmt. 1a. Voltage sense	8
1.4	Rqmt. 1b. Temperature sense	9
1.5	Rqmt. 1c. Current sense.	11
1.6	Rqmt. 1d. Contactor control	13
1.7	Rqmt. 1e. Isolation sense	15
1.8	Rqmt. 1f. Thermal control	18
1.9	Rqmt. 2. Protection	19
1.10	Rqmt. 3a. Charger control	20
1.11	Rqmt. 3b. CAN comms.	21
1.12	Rqmt. 3c. Log book.	22
1.13	Rqmt. 4a. SOC estimation	22
1.14	Rqmt. 4b. Energy estimation	26
1.15	Rqmt. 4c. Power estimation	27
1.16	Rqmt. 5. Diagnostics	29
1.17	Where to from here?	29

¹ IEEE Standard 1491 defines a *battery monitoring system* as “A permanently installed system for measuring, storing, and reporting battery operating parameters.”

tion of a prototype BMS being developed for research purposes at the University of Colorado Colorado Springs.

THE PRIMARY PURPOSES of a battery-management system are:

- First and foremost, to protect the safety of the operator of the battery-powered system. The BMS must detect unsafe operating conditions and respond. This may demand disconnecting and isolating the battery pack from the load, warning the operator by some display or alert, and so forth.
- Second, to protect cells of the battery pack from damage in abuse or failure cases. This may involve active intervention under software control, or specialized electronics that can detect failures and isolate the failing components from the rest of the battery pack and from the load it powers.
- Third, to prolong the life of the battery under normal operating cases. The BMS does so by coordinating with the controller of the load it powers, advising it of dynamic limits on power that can be sourced or sunk over some short future interval that ensures that the battery pack will not be overcharged or overdischarged. It also controls the thermal-management system, ensuring that the battery pack is kept within its design operational-temperature range.
- Fourth, to maintain the battery pack in a state in which it can fulfill its functional design requirements. Thus, for example, it will not allow a battery pack to become so far discharged that it cannot deliver its rated discharge power, nor will it allow the pack to become so highly charged that it cannot receive its rated charge power at any point in time.

THERE IS A COST associated with advanced methods of battery management, so not all applications implement all features. This cost adds to the purchase price of the battery pack, so the battery-management algorithms must provide tangible value. A good rule of thumb says “your battery is ‘cheap enough’ if you can’t remember the last time you replaced it.” The idea is that replacing the inexpensive battery in something like a television remote control is not financially painful, so spending more money on an advanced remote control that makes more efficient use of the battery is probably not worth it. However, mission-critical and large battery packs represent a greater investment and motivate better battery management. If you are required to replace an expensive battery prematurely, or if you are unable to complete your mission due to battery failure because of poor management, you will remember it for a *very* long time! Another way of



Figure 1.1: An example BMS.

looking at this is to consider the cost of the battery pack in relation to the cost of a battery-pack failure. The cost of battery failure for a television remote control is not large; however, the cost of battery-pack failure for a mission-critical or large battery installation can be very high. Consequently, as this book focuses on advanced methods for management and control of battery packs, it is most relevant for applications where the cost of battery-pack failure is high and where the added cost can be justified, although the methods that we discuss are quite general.

An important application category considers vehicles having electric-drivetrain components. These vehicular applications include the following generic subcategories:

Hybrid-electric vehicles (HEVs). These vehicles have motive power provided by an electric motor plus at least one other source (e.g., a gasoline engine). A battery pack stores a small amount of energy, and the battery–motor combination is used primarily for power boost when the vehicle must accelerate, or as a power sink when the vehicle must decelerate. This enables the gasoline engine to operate at a more constant operational point comprising a combination of revolutions per minute (RPM) and torque, which can be more efficient and can also allow the vehicle to achieve the same overall peak performance requirements with a smaller engine. HEVs have essentially zero all-electric vehicle range and are never plugged in to recharge their battery pack; instead, the gasoline engine recharges the battery when extra power is available. An example HEV is the Toyota Prius, shown in the top frame of Fig. 1.2.

Plug-in hybrid-electric vehicles (PHEVs). These vehicles are similar to HEVs but have a somewhat larger battery pack and motor. They can operate in electric-only mode under some operating conditions, typically at lower speeds such as for residential or city driving. Consequently, they have some all-electric range, often on the order of 10 to 20 miles. The vehicle can be “plugged in” to the utilities grid to recharge the battery pack. Subsequently, the vehicle operates first in a *charge-depletion* mode where the majority of the traction power is taken from the battery and not from the gasoline engine. When the battery charge is depleted to some minimum allowed level, the vehicle then switches to a *charge-sustaining* mode where it operates just like a standard HEV. An example PHEV is the Ford C-MAX Energi, which is shown in the second frame in Fig. 1.2.

Extended-range electric vehicles (E-REV). E-REVs are similar to PHEVs but have somewhat larger battery pack and motor. They can operate in electric-only mode under nearly all operating conditions so long as battery power is available. Their all-electric range is also greater, often on the order of 35 or more miles. As with a PHEV,



Figure 1.2: Vehicles having electrified drivetrain. From top to bottom: Toyota Prius, Ford C-MAX Energi, Chevy Volt, and Tesla Model S.

(All photos carry Creative-Commons “Attribution 2.0 Generic (CC BY 2.0)” license, and were taken by Soknet Soknet, Kyle Harris, Kārlis Dambrāns, and Niels de Wit, respectively. Cropped from originals downloaded from Flickr.com.)

they may be plugged in to charge their battery pack and they operate in charge-depletion and charge-sustaining modes. Since many commuters do not drive more than 35 miles in a day, they will rarely exceed the all-electric range of the vehicle and so the car will operate essentially as an electric vehicle for them. But if they ever do need to drive a distance farther than the battery energy alone allows, “extended range” is provided by the gasoline engine. An example E-REV is the Chevy Volt, which is shown in the third frame in Fig. 1.2.

Electric vehicles (EVs), also known as *battery-electric vehicles (BEVs)*. For these vehicles, the battery–motor combination provides the only source of motive power. There is no gasoline engine. Consequently, the vehicle design is much simpler than for any of the flavors of hybrid, but the vehicle range is limited by the amount of useable energy that can be stored by the battery. The size of the battery pack then becomes a critical design variable that is factored into the economic optimizations performed when designing the vehicle. Some commercial EVs have range of less than 100 miles, but others have range more than 300 miles between recharges. An example EV is the Tesla Model S, which is shown in the bottom frame in Fig. 1.2.

ALL OF THESE VEHICLE TYPES employ battery packs that are “large,” “high voltage,” and “high current.” There are some distinctions in design, which we will detail when necessary. However, the commonalities of their battery systems are more significant than their differences; so when distinctions aren’t important, we refer to any member of the class of vehicles having electric-drivetrain components as *xEV*. As *xEVs* represent a very important category of application requiring large battery packs, many of the examples presented in this book will be described in terms of an *xEV* implementation.

However, we do note that there are other significant and growing application domains requiring large battery packs. These include some that support and supplement the standard utilities electrical grid. For example, a large battery bank may be used for *grid-storage* of energy when opportunistic generation exceeds demand, such as from a solar or wind farm. The stored energy can be supplied back to the grid at a later time when the primary generation power is unavailable. Or, a battery pack may be used for *grid-backup* to provide energy to a load during grid power outages. For example, large, mobile grid-backup systems can be installed in the trailer of a semitruck and driven to neighborhoods to restore power temporarily during an emergency or during scheduled grid maintenance. Also, large battery packs are becoming more common in *frequency regulation* applications, where the battery pack operates as small-scale grid storage to make use of short-term surplus generation to charge the pack and in turn

provide short-term power to the grid when insufficient generation is available.

These applications—and others—are very different externally, and do place some different requirements on the battery-management-system design. However, they all tend to share the same kinds of algorithmic requirements, which is the primary focus of this book. Therefore, most of what is presented herein will be of interest to any battery-pack controls designer.

1.1 Battery-pack topology

Electrical *power* is computed as electric current multiplied by voltage: that is, $p = iv$. Therefore, high-power battery packs must deliver either high current, high voltage, or both. To achieve a design requirement on maximum power, the battery-pack engineer must make decisions regarding the topology of the pack. What should be its voltage range and peak current?

The chemistry of an individual cell fixes its voltage range; so for high-voltage packs we must connect multiple cells in series. The pack voltage is then the sum of the individual cell voltages. Roughly speaking (i.e., if all cell voltages are assumed equal), we have that $v_{\text{pack}} = N_s \times v_{\text{cell}}$, where the battery-pack topology wires N_s cells in series.

Cell construction places limits on the maximum current that the cell is designed to sustain; so for high-current packs we must wire cells in parallel. The pack current is then the sum of the currents going through the individual cells that are wired in parallel. Roughly speaking (i.e., if all of these currents are assumed equal), we have that $i_{\text{pack}} = N_p \times i_{\text{cell}}$, where the battery-pack topology wires N_p cells in parallel.

The tradeoff regarding how many cells to wire in parallel versus in series in order to achieve a battery-pack maximum-power design requirement is generally determined by economic and safety factors. Groups of cells termed *modules* are usually designed to have maximum voltage less than 50 V for safety,² and battery-pack total voltage is generally kept to less than 600 V because power electronics that operate at higher voltages are presently very expensive. Within these voltage ranges, when designing to meet a power-demand requirement, higher voltages tend to be preferred over higher currents to minimize resistive power losses in both the wiring that is internal to the battery pack and to the load (computed as $i_{\text{pack}}^2 \times R$). This also allows using wire having smaller diameter (and lower cost) as power loss scales with the square of current versus only linearly with resistance.

² There is an evolving discussion regarding what dc voltage range is considered to be safe and what is considered to be hazardous. The 2012 edition of the *NFPA 70E: Standard for Electrical Safety in the Workplace* listed a dc shock hazard limit of 50 V dc, but this was raised to 100 V dc in the 2015 edition. Many designers still adhere to the 50 V limit to minimize risk.




Figure 1.3: Two different approaches to modularizing 300 cells in a high-power battery pack: PCM versus SCM.

BATTERY-PACK DESIGN is often *modular*. That is, a design is first created and optimized for a module comprising a small group of cells; then, modules are connected in series and/or parallel to achieve the overall battery-pack design objectives. This allows reusing a fixed-module design for many different possible application scenarios and minimizes *nonrecurring engineering (NRE)* costs.

Modules may be constructed with any number of cells wired in series and parallel, and we will explore some of the tradeoffs in Chap. 2. Here, we note only that two extreme cases are shown in Fig. 1.3. The *parallel-cell-module (PCM)* approach wires cells in parallel to make modules, then wires modules in series to create a battery pack. The battery pack drawn in the top of the figure configures 300 cells by wiring three cells in parallel to make a PCM, and then wiring 100 PCMs in series to make the pack. The *series-cell-module (SCM)* approach instead wires cells in series to make modules, then wires modules in parallel to make a battery pack. The battery pack drawn in the bottom of the figure configures the same 300 cells by wiring 100 cells in series to make an SCM, and then wiring three SCMs in parallel to make the pack.³

We can design battery packs and battery-management system for either approach. Most often, however, we use something in between these extremes. For example, a 3P6S module has eighteen cells configured with three cells wired in parallel and six cells in series. The module power and energy are then both approximately eighteen times that of a single cell (but not exactly, as we shall discover later).

³ If all cells are identical, then these two battery packs provide the same energy, power, and so forth. However, if there are differences between the cells, the packs can operate in quite different ways. We explore this more in Chap. 2, where we see how to simulate battery packs having different topologies.

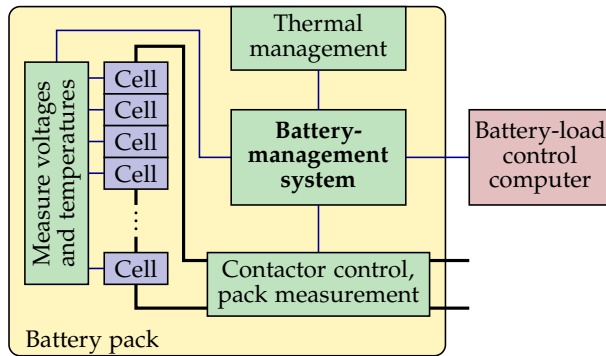


Figure 1.4: The context of the BMS inside the battery pack.

1.2 BMS design requirements

The BMS is interconnected with all major components in the battery pack, as illustrated in Fig. 1.4. These include the *battery stack*, comprising all of the cells, the sensing and control electronics, and at least some part of the thermal-management system.

The electronics may be implemented as a single monolithic design, or there may be separate elements having different functionality distributed throughout the battery pack. As an example, the drawing shows cell-voltage and temperature-sensing elements as distinct from pack sensing and control and from overall supervisory battery management, all within the confines of an overall battery pack, which is drawn as the yellow shaded box. In an xEV, the battery-load control computer is the overall vehicle controller; in other applications, the corresponding controller would be interfaced.

Regardless of electronics topology or battery application, BMS functional requirements can be broken down into five overall categories:

1. *Sensing and high-voltage control*: The BMS must measure cell voltages, module temperatures, and battery-pack current. It must also detect isolation faults and control the contactors and the thermal-management system.
2. *Protection*: It must include electronics and logic to protect the operator of the battery-powered system and the battery pack itself against overcharge, overdischarge, overcurrent, cell short circuits, and extreme temperatures.
3. *Interface*: The BMS must communicate regularly with the application that the battery pack powers, reporting available energy and power and other indicators of battery-pack status. Further, it must record unusual error or abuse events in permanent memory for technician diagnostics via occasional on-demand download.
4. *Performance management*: It must be able to estimate state of charge

(SOC) for all the cells of the battery pack, compute battery-pack available energy and power limits, and balance cells in the battery pack.

5. *Diagnostics*: Finally, it must be able to estimate state-of-health (SOH), including detecting abuse, and may be required to estimate state-of-life (SOL) of the battery cells and pack.

In this chapter, we present a high-level overview of these five requirements. Later chapters will develop performance-management and diagnostic topics in detail. In particular, Chaps. 3 through 7 focus on SOC estimation, SOH estimation, balancing, and power-limits estimation in more depth.

As we proceed through the remainder of this chapter, the section titles will refer to requirement numbers from the above list. For example, all section titles beginning with “Requirement 1” have to do with sensing and high-voltage control, and so forth.

1.3 Requirement 1a. Battery-pack sensing: Voltage

In a battery pack comprising lithium-ion cells, *all* individual cell voltages must be measured.⁴ Cell terminal voltage can provide a measure of the relative balance of cells in the battery pack, and is a critical input to most SOC and SOH estimation algorithms.

Out-of-bounds cell voltage is also an indicator of a number of serious lifetime and safety issues. For example, overcharging a lithium-ion cell can initiate unwanted internal chemical side reactions that degrade the cell. Overdischarging can precipitate a sequence of events leading to a cell short-circuiting. In extreme cases, either of these can cause *thermal runaway*, where heat generated by side reaction or short circuit accelerates the failure mechanisms via natural positive feedback, ultimately resulting in a battery fire and/or explosion. Due to various reasons discussed in Chap. 5, we cannot assume that all cells automatically have the same voltage, and so we cannot skip measuring any voltage in a lithium-ion battery pack.

A number of silicon vendors have come to the aid of the BMS electronics designer by designing integrated-circuit (IC) “chipsets” that can assist with cell-voltage measurement.⁵ These are low-cost devices with little or no general-purpose processing capabilities on-chip. They implement the difficult task of measuring analog voltages with high accuracy, high common-mode rejection, and fast response in high-electromagnetic-interference (EMI), high-temperature, and high-vibration environments, such as encountered in a vehicular application. A single IC can be used to measure the voltages of a

⁴ If cells are wired in parallel, they will share the same terminal voltage, and so voltages of parallel cells cannot be distinguished. We could more correctly say something like “all voltages of individual parallel groups of cells must be measured.”

⁵ See, for example, Andrea, D., *Battery Management Systems for Large Lithium Ion Battery Packs*, Artech House, 2010, which has quite a bit more detail on this topic, and on BMS electronics design in general.

number of cells connected in series, and two ICs can often be placed in parallel for redundant fault-tolerant designs.

One example is the popular LTC6803 part, designed in Colorado Springs by Linear Technology Corporation. A functional block diagram of this IC is shown in Fig. 1.5. A single LTC6803 can monitor the voltages of up to 12 cells wired in series in a module. An internal analog multiplexer connects cells to a 12-bit delta-sigma analog-to-digital converter, and a single “read” command can trigger the measurements of all 12 voltages in rapid sequence. Up to 10 ICs can be daisy-chained together to monitor up to 120 cells in a battery pack and electrically isolated communications are built into the design.

It supports either internal (slow) or external (faster) dissipative cell-balancing circuitry and up to four temperature measurements per module (in addition to the internal IC die temperature). It can be powered by the module of cells itself or from an external source.

When selecting a voltage-monitoring IC for a BMS electronics design, some of its specifications to consider and compare to design requirements for the battery pack include how many cells can each IC monitor? How many cells in total can be monitored? Does the IC support either passive or active balancing? What is the measurement resolution and accuracy? How many temperature measurements will it support? How many wires are required to communicate between ICs (if multiple ICs are required in a design)? And, what is the chipset availability and per-cell cost?

1.4 Requirement 1b. Battery-pack sensing: Temperature

As we have already explored in detail in Volume I of this series, lithium-ion cell dynamics are strong functions of temperature. For example, resistance increases and chemical processes tend to slow down at cold temperatures. So, we must know cell temperature to be able to predict near-future cell performance. Further, degradation mechanisms are also temperature-dependent: generally, high temperatures accelerate degradation rates, although charging a battery pack at low temperatures can also lead to premature failure via lithium-plating. Also, a BMS must be able to sense temperature to control the thermal-management system to maintain battery-pack temperatures in a safe region.

Ideally, we would measure each cell’s *internal* temperature. However, cells are not (mass-)produced with temperature sensors built in, so we must rely on measurements of cell *external* temperatures. In addition, as there is an expense per sensor, we would like to minimize the total number of sensors required. With an accurate pack thermal model, we can place a limited number of temperature sensors




Figure 1.5: An example IC for battery-stack voltage measurements. (Figure used with permission of Linear Technology Corporation.)

external to one or more cells per module and then estimate internal temperatures of all cells. This is usually sufficient.

But, how do we measure temperature? To measure any simple physical quantity electronically, we must first represent that quantity as a voltage signal and then measure the voltage via an analog-to-digital converter circuit. There are two primary approaches to doing so to measure temperature.

The first is to use a *thermocouple*, which is a device comprising two dissimilar metals in contact with each other that acts as a miniature battery. A very small voltage is produced when the temperature of the thermocouple is different from a reference temperature at another part of the measurement circuit, and this voltage depends on magnitude of the temperature difference. The thermocouple voltage can be amplified and measured and temperature can be computed from this measurement. A design challenge when using thermocouples is that the reference temperature must be independently known or measured, which probably makes thermocouples best suited for laboratory testing and not for production BMS designs.

The second method uses a *thermistor*, and is probably better-suited for use in commercial products. All resistors have value that varies somewhat with temperature but are usually designed to minimize this variation. A thermistor, on the other hand, is designed to maximize and to capitalize on temperature variation. Negative-temperature-coefficient (NTC) thermistors have resistance that varies inversely with temperature, and positive-temperature-coefficient (PTC) thermistors have resistance that varies directly with temperature. If we can measure thermistor resistance, we can then infer temperature.

To measure resistance, we can use a voltage-divider circuit, such as shown in Fig. 1.6. In the circuit, the top resistor R_1 has resistance that does not vary appreciably with temperature, but the lower resistor $R_{\text{thermistor}}$ has value that is designed to vary significantly with temperature. We compute overall current as

$$i = \frac{v}{R_1 + R_{\text{thermistor}}}.$$

Then, we note that the measured voltage is $v_{\text{thermistor}} = iR_{\text{thermistor}}$ or

$$v_{\text{thermistor}} = \frac{R_{\text{thermistor}}}{R_1 + R_{\text{thermistor}}}v.$$

The value of R_1 is designed to limit power loss through the measurement circuit but at the same time provide a useful measurement range for $v_{\text{thermistor}}$.

If we measure $v_{\text{thermistor}}$ and know the circuit-design parameters,




Figure 1.6: Voltage-divider circuit.

we can compute

$$R_{\text{thermistor}} = \frac{v_{\text{thermistor}}}{v - v_{\text{thermistor}}} R_1. \quad (1.1)$$

The thermistor data sheet will give an expression relating $R_{\text{thermistor}}$ to temperature. For example, a representative relationship computes

$$R_{\text{thermistor}} = R_0 \exp \left(\beta \left(\frac{1}{273.15 + T} - \frac{1}{273.15 + T_0} \right) \right), \quad (1.2)$$

where the temperature being measured is denoted as T and nominal resistance at reference temperature T_0 is denoted as R_0 ; temperatures are converted from celsius to kelvin by adding 273.15, and β is a device parameter. The top frame of Fig. 1.7 plots thermistor resistance for an NTC device having $R_0 = 100 \text{ k}\Omega$ at $T_0 = 25 \text{ }^\circ\text{C}$ and $\beta = 4282$.

If this device were placed in the lower leg of a voltage divider having a 5 V source and $R_1 = 100 \text{ k}\Omega$, we would measure the thermistor voltage $v_{\text{thermistor}}$ as a function of temperature that is plotted in the middle frame of Fig. 1.7. We can then compute thermistor resistance via Eq. (1.1) and temperature via solving Eq. (1.2) for that value of resistance. For efficient computation in an embedded BMS, the overall relationship between measured voltage and temperature can be precomputed and stored in a *lookup table (LUT)*. The result for this example is plotted in the bottom frame of Fig. 1.7.

1.5 Requirement 1c. Battery-pack sensing: Current

A BMS must measure battery-pack current. In part, this is to detect and log abuse conditions and to ensure safety. However, battery current is also a critical input to most SOC and SOH estimation algorithms.

There are two basic electronics elements that can be used in a circuit to sense current: current-shunts and Hall-effect sensors.

A *current-shunt* is a low-value (e.g., $0.1 \text{ m}\Omega$) high-precision resistor placed in series with the battery pack, usually at the negative terminal. The voltage drop v_{shunt} across the shunt resistance R_{shunt} is measured using a standard analog-to-digital converter, and current is computed as $i = v_{\text{shunt}}/R_{\text{shunt}}$. Since the shunt resistance must be small (to avoid large power losses due to $i^2 R_{\text{shunt}}$ heating), the voltage drop across the shunt will be small as well. So, the voltage is usually amplified before sensing and the calculation for current is adjusted accordingly.

Fig. 1.8 shows a photograph of a current shunt (top frame) and a block diagram for current sensing using a current shunt (bottom frame). Examining the current shunt itself in more detail, note that

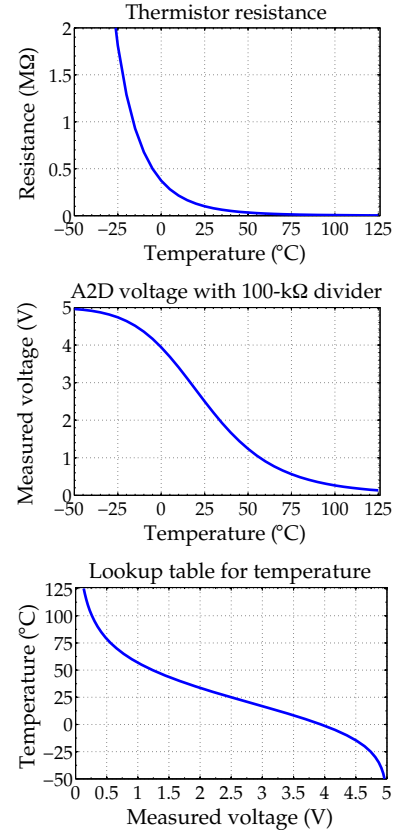


Figure 1.7: Thermistor usage example.

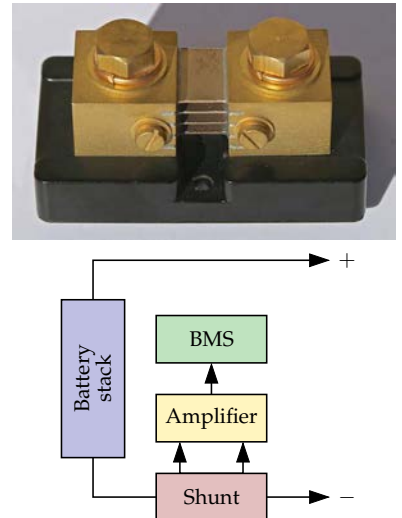


Figure 1.8: Sensing current with a current shunt.

there are four connection terminals. The large terminals on the top of the device are for connecting to the main battery-pack current-carrying wires: one side is connected to the negative terminal of the battery stack, and the other side is connected to the output negative terminal of the battery pack. The entire pack current then passes through the parallel plates that form the calibrated resistance in the center of the shunt, between these large terminals. Specifically, the resistance between the two smaller screw terminals is calibrated, and the sensing leads are connected to these smaller terminals.⁶

The primary advantage of current shunts (over Hall-effect sensors, described next) is that they have no offset at zero current, regardless of temperature. This is of vital importance if one is updating an SOC estimate by integrating current into and out of a battery pack (this is known as *coulomb counting*). However, the amplification and measurement electronics can still introduce an offset, which must then be calibrated out of every measurement.

A disadvantage when using current shunts is that they must usually be electrically isolated from the main BMS circuitry (e.g., in an automotive application, the BMS is powered via an external 12 V supply that must be isolated from the high-voltage battery). This, plus amplification circuitry, adds complexity to the design. The resistance of the current shunt also changes with temperature, so temperature should be measured and resistance calibrated. In addition, the shunt itself introduces some energy losses, and the heat that is generated must be dissipated via the thermal management system.

IF A COIL is wrapped around a primary current-carrying conductor, the electromagnetic field produced by the conductor induces a secondary current in the coil. *Hall-effect sensors* measure this induced current to infer the primary current. Fig. 1.9 shows a photograph of a Hall-effect sensor (top frame) and a block diagram for current sensing using a Hall-effect sensor (bottom frame). The main battery-pack current-carrying wire passes through the oval opening in the center of the sensor—no direct electrical connection is made between the sensor and the high-voltage battery pack. This yields the distinct advantage that Hall-effect sensors are automatically isolated electrically from the high-voltage battery and so no special isolation circuitry is needed.

However, because Hall-effect sensors operate based on electromagnetic principles, magnetic hysteresis is inherent in the device. Feedback signal-conditioning circuitry can be devised to guard against this hysteresis, and some sensors come prepackaged with such circuitry. Even so, Hall-effect sensors suffer from at least some measurement offset at zero current, which drifts over time and which changes

⁶ This scheme is called a *Kelvin connection* and enables *four-wire voltage measurement*. Since essentially zero current is drawn by the voltage-sensing electronics, there is negligible voltage drop across the resistance of the smaller terminals, and current can be calculated as stated earlier. However, if one were to (mistakenly) connect the voltage-sensing wires to the larger screw terminals, the voltage drop of the large battery-pack current passing through the uncalibrated resistance of the large screw terminals would significantly degrade the accuracy of the current calculation.




Figure 1.9: Sensing current with a Hall-effect sensor.

with temperature. Even if this offset is “zeroed out” in software (by subtracting an initial measurement) at some initial temperature during a BMS startup routine (when it is known that zero current is flowing since contactors are not yet closed), this ad hoc calibration does not correct for time- and temperature-varying drift in the bias. As the bias plays havoc with a number of BMS algorithms, some kind of compensation is necessary. For example, if the BMS ever knows for certain that zero current is flowing, the bias can be “zeroed out” at those times. Some HEV applications allow for this.

1.6 Requirement 1d: High-voltage contactor control

For safety reasons, as described in more detail in Sect. 1.7, high-voltage battery packs are designed to be isolated electrically from chassis ground. If someone were to touch chassis ground with one hand and either terminal of the battery pack with the other, the electrical isolation should be sufficient such that he or she is completely safe. This will not be true if there is an *isolation fault* or *ground fault*, which is why we investigate how to detect such faults in Sect. 1.7.

For similar safety concerns, the battery pack internal high-voltage bus is completely disconnected from the load at both battery-pack external terminals when not in use. This requires two high-current capable relays known as *contactors*. Contactors used in battery packs are normally-open devices, so if the BMS loses power for any reason, the terminals of the contactor are de-energized and the battery pack is disconnected from the load.

Battery loads are often capacitive. For example, in a vehicle application, the motor-driving circuitry requires large capacitors to filter the transients caused by switching power to the motor-drive coils. This causes a complication during the startup sequence that connects the battery pack to its load. Supposing that the capacitive load is in a discharged state, if both contactors were to be closed simultaneously, a huge spike of current would attempt to flow instantly, potentially welding the contactors closed or blowing a fuse.

So, a third *precharge contactor* is used. The overall startup process is depicted in Fig. 1.10. In frame (a), the pack is initially disconnected from the load. All contactors are open and the thick red lines—indicating the power path—extend from the battery stack only to the internal side of the contactors.⁷ The negative contactor is activated first, as illustrated in frame (b), where the thick blue lines indicate the low-voltage signal paths being activated. This connects the “-” terminal of the battery stack to the “-” terminal of the load.

Next, the precharge contactor is activated (frame (c)). The precharge resistor limits current flow and allows the battery pack to charge up

⁷ Incidentally, the figure also shows that a fuse is usually inserted in middle of the battery stack. If the fuse is blown for some reason, this limits the maximum voltage encountered by the service technician to at most half of the overall battery-stack voltage.

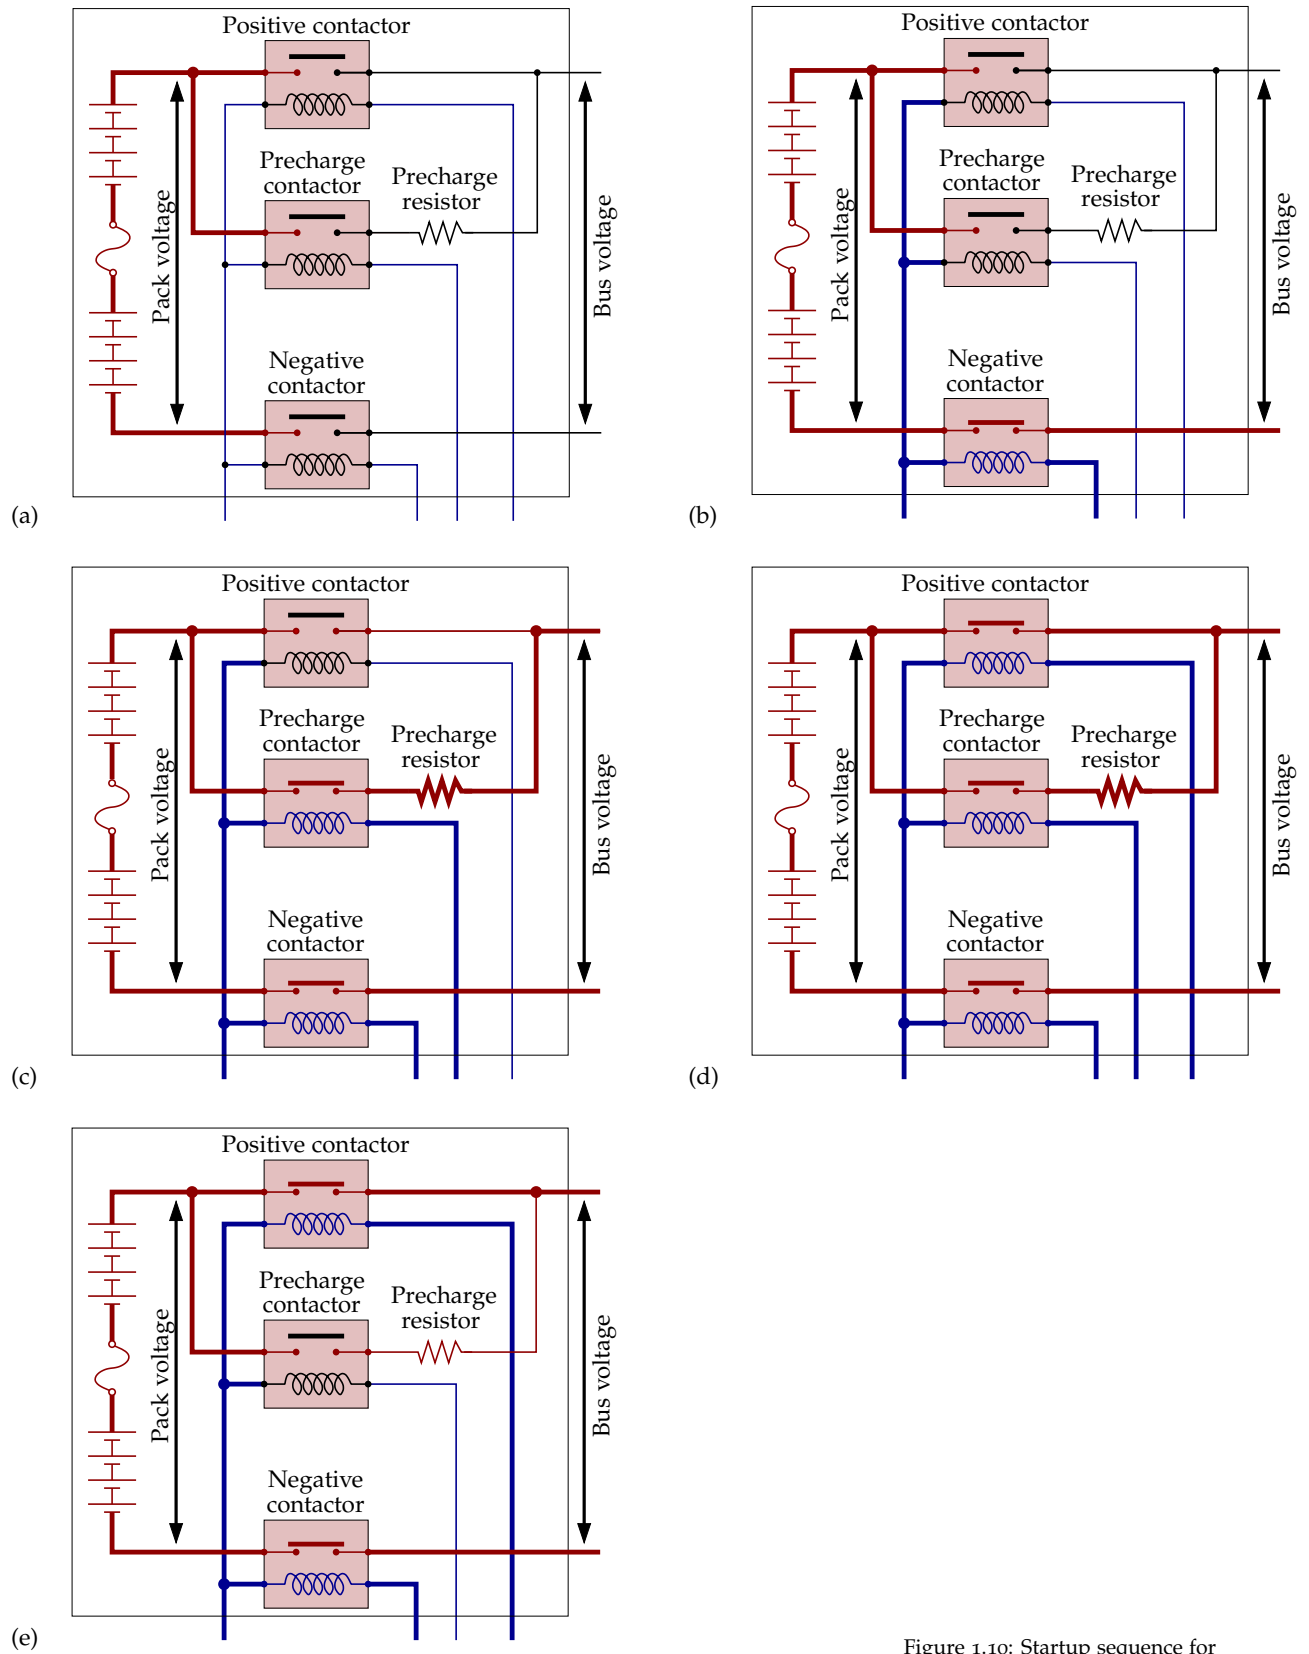


Figure 1.10: Startup sequence for connecting a battery pack to its load.

the capacitive load at a controlled and safe rate. Often, the precharge-resistor temperature is monitored—if it becomes too high the load may have a short-circuit fault: The startup process is aborted and the pack is disconnected from the load. Similarly, the bus and battery-stack voltages are monitored. If these voltages don't converge to the same neighborhood after a specified interval, the load may have a short-circuit fault: The pack disconnects.⁸

Assuming that bus and battery-stack voltages have become “close enough” “quickly enough,” the BMS then closes the positive contactor (frame (d)). This connects the “+” terminal of the battery stack to the “+” terminal of the load. The load is now connected directly to the battery stack through a low-resistance path. The precharge contactor is then opened (frame (e)).

The procedure to follow when the battery pack is being shut down is not as clear. Since the capacitive load and the battery voltage have already come into equilibrium, the danger of contactor fusing/welding is not as severe. As long as the inductance of the load is small, simply opening the main contactors is probably sufficient. (If not, the startup procedure could be used in reverse to drain the inductive energy prior to complete disconnect of the battery pack.)

1.7 Requirement 1e. Isolation sensing

In a standard automobile, the negative terminal of the 12 V lead-acid battery is connected directly to the vehicle frame or *chassis*. This saves money on wiring as only the positive voltage requires separate wires to distribute power throughout the vehicle. It also does not pose a significant safety risk as it is unlikely that 12 V across a person's body would cause harm.⁹

However, if someone were to touch both terminals of a high-voltage battery pack, injury or even death could occur. Therefore, battery-pack designers must take great care to minimize the likelihood of this happening. So, all wiring is fully insulated and at no point in the circuit is either high-voltage battery terminal connected to the highly exposed vehicle chassis.

This greatly enhances safety. However, if the vehicle were to be involved in an accident, or if vibration or any other kind of wear breaks through the insulation on the high-voltage wiring, it is possible that one of the terminals of the battery pack could come into contact with the vehicle chassis through a low-resistance path. This then poses a safety hazard, which we must detect in order to warn the vehicle operator of the danger.

The Federal Motor Carrier Safety Administration dictates a procedure for detecting loss of isolation.¹⁰ The premise is that it should be

⁸ Measuring battery-stack and bus voltages requires high-impedance voltage dividers and isolation circuitry. Bus voltage can instead be inferred from stack voltage and pack current: when pack current approaches zero, we can conclude that bus voltage has approached stack voltage.

⁹ Recall the discussion from footnote 2 on what dc voltages are considered hazardous.

¹⁰ These are documented in Federal Motor Vehicle Safety Standard FMVSS 305 and Society of Automotive Engineers standard SAE J1766.

safe for someone to touch either terminal of the battery pack and the chassis ground at the same time. The metric states that isolation is considered sufficient if less than 2 mA of current flows when connecting chassis ground to either the positive or negative terminal of the battery pack via a direct short circuit.

We derive this procedure with reference to Fig. 1.11. The battery pack should not be connected in any way to the chassis. In the circuit, both R_1 and R_2 should be infinite. However, it is possible that isolation has been lost. In that case, either R_1 or R_2 will become less than infinite. We call the lesser of R_1 and R_2 the *isolation resistance* R_i , and it is our primary objective here to determine thus resistance. According to the safety criterion, R_i must be greater than $v_b/0.002$ A or $R_i > 500v_b$, where v_b is the battery voltage.

For the BMS to sense whether the pack is sufficiently isolated from the chassis, it must somehow determine R_i . To do so, we measure v_1 and v_2 using a high-impedance analog-to-digital measurement circuit, itself having impedance greater than $10\text{ M}\Omega$.¹¹ Inserting the sensor to make these measurements breaks strict isolation, but the high impedance of the sensor ensures that it will not violate the 2 mA current limit by itself.

Redrawing the circuit in Fig. 1.12, we see that R_1 and R_2 form a voltage divider across the total battery voltage. As we wish to find the smaller of the two resistances we solve for R_1 if $v_2 > v_1$; otherwise we solve for R_2 . Note also that the current through both resistors in this circuit must be equal in order to satisfy Kirchhoff's current law. So, $v_1/R_1 = v_2/R_2$. We will use this identity shortly.

1.7.1 Potential isolation fault on negative side: Find R_1

If there is a fault on the negative side of the battery pack, then $v_2 > v_1$. So, if we measure $v_2 > v_1$, we wish to solve for R_1 and then check to see whether it is sufficiently large to avoid an isolation fault.

To find R_1 , we insert a known (large) resistance R_0 between the positive terminal of the battery and chassis ground via a transistor switch, as shown in Fig. 1.13. This again breaks strict isolation, but not enough to worry about if R_0 is "big enough" (i.e., $R_0 \gg 500V_b$).

We now measure v'_2 . By Kirchhoff's current law, the sum of currents through R_0 and R_2 must equal the current through R_1 , so

$$\frac{v_b - v'_2}{R_1} = \frac{v'_2}{R_2} + \frac{v'_2}{R_0}.$$

Substituting $v_b = v_1 + v_2$ and $R_2 = R_1(v_2/v_1)$ gives

$$\frac{(v_1 + v_2) - v'_2}{R_1} = \frac{v'_2}{R_2} + \frac{v'_2}{R_0}$$




Figure 1.11: Step 1 of measuring isolation.

¹¹ Note the polarity of the voltmeters—both v_1 and v_2 are positive.




Figure 1.12: Step 1 of measuring isolation, redrawn.




Figure 1.13: Step 2 of measuring isolation if fault is on the negative side.

$$= \frac{v'_2(v_1/v_2)}{R_1} + \frac{v'_2}{R_0}.$$

Finally, solving for R_1 , we find

$$\begin{aligned} \frac{(v_1 + v_2) - v'_2 - v'_2(v_1/v_2)}{R_1} &= \frac{v'_2}{R_0} \\ R_1 &= \frac{R_0}{v'_2}(v_1 + v_2 - v'_2 - v'_2(v_1/v_2)) \\ R_1 &= \frac{R_0}{v'_2} \left(1 + \frac{v_1}{v_2}\right) (v_2 - v'_2). \end{aligned} \quad (1.3)$$

In summary, if $v_2 > v_1$, there is a possible isolation fault on the negative side of the battery. We then measure v'_2 and compute $R_i = R_1$ from v_1 , v_2 , and v'_2 using Eq. (1.3). Isolation is deemed sufficient if $R_i > 500v_b$.

1.7.2 Potential isolation fault on positive side: Find R_2

The procedure is similar if $v_1 > v_2$, except that now we want to find $R_i = R_2$. We insert a known large resistance R_0 between the negative terminal of the battery and chassis ground using a transistor switch, as shown in Fig. 1.14, and measure v'_1 .

Again, by Kirchoff's current law,

$$\frac{v_b - v'_1}{R_2} = \frac{v'_1}{R_1} + \frac{v'_1}{R_0}.$$

Substituting $v_b = v_1 + v_2$ and $R_1 = R_2(v_1/v_2)$ gives

$$\begin{aligned} \frac{v_1 + v_2 - v'_1}{R_2} &= \frac{v'_1(v_2/v_1)}{R_2} + \frac{v'_1}{R_0} \\ \frac{v_1 + v_2 - v'_1 - v'_1(v_2/v_1)}{R_2} &= \frac{v'_1}{R_0}. \end{aligned}$$

Solving for R_2 , we find

$$\begin{aligned} R_2 &= \frac{R_0}{v'_1}(v_1 + v_2 - v'_1 - v'_1(v_2/v_1)) \\ R_2 &= \frac{R_0}{v'_1} \left(1 + \frac{v_2}{v_1}\right) (v_1 - v'_1). \end{aligned} \quad (1.4)$$

In summary, if the potential isolation fault is on the positive side of the battery, we measure v_1 , v_2 , and v'_1 , and compute $R_i = R_2$ using Eq. (1.4). Isolation is deemed sufficient if $R_i > 500v_b$.

IT IS ALWAYS WISE to design with fault-tolerance and built-in test circuitry in mind. In the case of the isolation-sense requirement, we can test the function of the circuitry without changing the circuit




Figure 1.14: Step 2 of measuring isolation if fault is on the positive side.

design from what we have discussed already. We simply activate the transistor switches to insert resistance R_0 between both battery-pack terminals and chassis ground at the same time. We measure v'_1 and v'_2 : they should both be equal to $v_b/2$ to within the tolerance of the resistors used for R_0 .

1.8 Requirement 1f. Thermal control

In this book, we do not go into detailed thermal management and control strategies. However, they are very important aspects of BMS design as they address some significant battery-pack longevity and safety concerns.¹²

In general, lithium-ion cells last longest if they are maintained in a temperature band from about 10 °C to 40 °C.¹³ Some of the degradation mechanisms that are accelerated by temperature effects are discussed qualitatively in Chap. 4.

As we saw in Volume I of this series, models of heat generation in a battery cell can be complicated. However, irreversible and joule heating terms tend to dominate (especially over short timescales with random-looking input current), so heat generation can be crudely approximated as i^2R where i is cell current and R is a thermal resistance. Therefore, when the battery pack is sourcing or sinking a large amount of current per cell, heat generation is high.

In an EV, while the total battery-pack current is high, the relative C-rate per cell is low as the battery pack must have high total capacity for acceptable vehicle driving range. Internal heat generation is also therefore low. So, air cooling may be sufficient.

However, in an HEV, C-rates are high. Therefore, liquid cooling may be necessary. In any xEV application, careful coupled thermal/electrical simulation and analysis of battery-pack operation should be conducted in the design phase to determine the correct size and type of thermal management system.

While it is common (even universal) for a thermal-management system to be required to cool cells, it is rare for one to consider heating cells. Even if a vehicle were started when the ambient temperature is quite low, the i^2R self-heating tends to bring the battery pack to a reasonable operating temperature fairly quickly. Some plug-in vehicles, however, do heat a too-cold battery pack using grid power while plugged in, but to do so using the pack's own power when not plugged in would deplete the available energy and therefore driving range, so there is little reason to do so.

¹² See, for example, Santhanagopalan, S., Smith, K., Neubauer, J., Kim, G-H, Keyser, M., and Pesaran, A., *Design and Analysis of Large Lithium-Ion Battery Systems*, Artech House, 2015, which has much more in-depth discussion of battery-pack thermal requirements.

¹³ A good rule of thumb is "if you are comfortable at a certain temperature, the battery pack is also 'comfortable' at that temperature."

1.9 Requirement 2. Protection

Referring back to the list of BMS design requirements enumerated in Sect. 1.2, we have now completed our discussion of Requirement 1: Sensing and high-voltage control. We now quickly consider Requirement 2: Protection.

Energy storage in any form is potentially hazardous. If energy were to be released in an uncontrolled way, there could be catastrophic consequences. Because of long experience, we have grown accustomed to the risks associated with gasoline-powered vehicles, and understand very well how to minimize the likelihood and impact of those hazards.

We are still learning how to control the risks associated with energy storage in high-capacity battery packs. The electronics and software of the BMS are integral parts of an overall risk-management strategy. They must provide monitoring and control to protect cells from out-of-tolerance ambient operating conditions, and to protect the user from consequences of a battery failure. There are large challenges involved. For example, if a cell develops an internal short circuit, hundreds of amperes can develop in microseconds. Therefore, protection circuitry must act *very* quickly to isolate the fault; otherwise, the cell may go into thermal runaway, resulting in a battery-pack fire or explosion. Safety management for high-capacity battery packs is a developing field, as conventional methods using circuit elements such as fuses tend to act too slowly, and new devices and methodologies are needed.¹⁴

In a battery pack, protection must address the following undesirable events or conditions: Excessive current during charging or discharging, short circuit, overvoltage or undervoltage, high ambient temperature or overheating, loss of isolation, and abuse. Whenever possible, fallback protection paths should be implemented.

For example, consider Fig. 1.15, which illustrates one way to think about designing protection mechanisms for a battery pack where current and temperature are the input variables. In the figure, the area shaded red corresponds to the cell-manufacturer-specified operating region where cells will most likely be subject to permanent damage. Anywhere else is considered to be “okay,” but we need a margin of error in the design; hence, we generally design protection to limit the cell’s operating conditions to smaller “safe” region, shown in green. Safety devices are then specified to constrain cells to the safe region leaving a safety margin, shown in white. Note that to traverse from the green region to the red region, two protection systems must fail. This redundancy helps provide a robust protection scheme.

Fig. 1.16 is a similar diagram for a preliminary protection design

¹⁴ If a single cell goes into thermal runaway, there may be no way to stop it if, for instance, there is an internal short in the cell caused by a impurity in the cell material. Opening the main contactor may not stop the event. However, the overall battery-pack design should be such that thermal runaway in one cell will not cascade to other cells and such that all gases are vented. This objective can be achieved by mechanically isolating cells. A clever fusing approach can also help: see Kim, G-H, Smith, K, Ireland, J, and Pesaran, A., “Fail-safe design for large capacity lithium-ion battery systems,” *Journal of Power Sources*, 210, 2012, pp. 243–253.




Figure 1.15: Designing current/temperature protection.

where voltage and temperature are the input variables. In some cases, it is possible to cross only one boundary to move from the safe zone to the failure zone. To complete the design, additional protection elements should be added.

Examples of protection devices include thermal fuses (which open the contactor when temperature is above some limiting temperature), conventional fuses (which sever an electrical connection if too much current flows for too long, but may not operate quickly enough for some kinds of faults), and electronic fault detection. For the latter, the BMS continuously senses voltage, current, and temperature and takes design-specified action if a fault is detected. In Fig. 1.16 we also notice that the plug-in charger can independently monitor battery-pack voltage when the pack is being charged and automatically stop charging if it detects that the overall battery-pack voltage is too high. Similarly, the load controller can monitor voltage and stop utilizing the battery pack if the voltage becomes out of bounds.¹⁵

1.10 Requirement 3a. Charger control

BMS requirement 3 considers communications interfacing between the battery management system and the application being powered by the battery pack. In an xEV, that application is the vehicle itself, which is managed via the vehicle control computer (which fills the role of the battery-load control computer in Fig. 1.4).

The first communication interface we consider is to the battery-pack charger. Battery packs in xEV can receive charge in two ways. All xEV packs experience *random charging*, where charge is delivered in unpredictable patterns; for example, via energy recovery because of regenerative braking. The BMS controls the maximum allowed level of random charging by continuously informing the vehicle of the present operating charge-power limit that can be supported by the battery pack. This is discussed more in Chap. 6. In addition, EV, PHEV, and E-REV have plug-in modes so can support *plug-in charging*. This is a carefully regulated process where charge power is provided from the electric utilities grid, and much finer control over the charging protocol can be implemented. Often, some sequence of constant-power steps (of decreasing magnitude) are applied until the battery pack is considered to be fully charged. Battery-pack equalization, as studied in Chap. 5, is also frequently performed during plug-in charging.

The speed with which a modern battery pack can be charged depends far less on the battery itself than it does on the utility service feeding the charger. A quick back-of-envelope example should illustrate this. Consider a small passenger vehicle: depending on

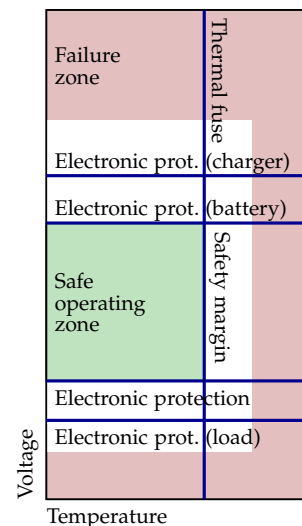


Figure 1.16: Designing current/temperature protection.

¹⁵ Fault-detection and reaction strategies are outside the principal scope of this book. Some references that include a deeper discussion of these topics include: Weicker, P., *A Systems Approach to Lithium-Ion Battery Management*, Artech House, 2014, and Santhanagopalan, S., Smith, K., Neubauer, J., Kim, G-H, Keyser, M., and Pesaran, A., *Design and Analysis of Large Lithium-Ion Battery Systems*, Artech House, 2015.

aerodynamics and so forth, the rate of energy usage is on the order of 200 Wh mile^{-1} to 300 Wh mile^{-1} . For a 300-mile range, the battery pack must have between 60 kWh to 90 kWh useable capacity. If we were to charge this battery pack in 3 min, roughly the time required to fill a passenger-vehicle gasoline tank, the utilities service would need to provide power at a rate of 1.8 MW, continuously! Current mass-produced battery technologies cannot withstand such a fast charge but neither can residential utilities.

Approaching the problem from the opposite point of view, consider a standard domestic service of 110 V at 15 A, or about 1.5 kW. This is typical of most residential wall outlets. At this rate, it would take between 40 h and 60 h to fully charge the battery pack, which is not acceptable. However, consider an upgraded domestic service of 220 V at 30 A, or 6.6 kW. This is typical of an outlet to supply an electric stove or an electric clothes dryer. At this rate, it would take between 10 h and 15 h to fully charge the battery pack. This is actually quite a good compromise for daily driving, as the car is rarely driven its full range and often has overnight to recharge. So-called “fast charge” would be needed only on extended-length trips of greater distance than the vehicle range, and while stopping 10 or 15 min to charge a vehicle battery pack is longer than the wait with which the consumer is accustomed, it is still within reasonable limits if required only infrequently.

1.11 Requirement 3b. Communication via CAN bus

The second topic we consider under the heading of communications is the protocol used for the communications. This will, of course, vary with the application, but automotive applications use the *control-area network (CAN)* protocol almost exclusively for on-board vehicle messaging. CAN is designed to provide robust communications in the very harsh automotive operating environments with high levels of electrical noise.

CAN has an electrical specification and a packet protocol. Electrically, it comprises a differential two-wire serial bus designed to network intelligent sensors and actuators. Messaging can operate at two rates: high-priority messages are sent at a higher baud rate, while low-priority messages are sent at a lower rate. High-rate messaging is used for critical operations such as engine management, vehicle stability, and motion control, while low-rate messaging is used for simple switching and control of lighting, windows, mirror adjustments, and instrument displays, etc.

The communication protocol defines the following: The method of addressing the devices connected to the bus; transmission speed and




Figure 1.17: A CAN 2.0 data frame.

priority settings; transmission sequence; error detection and handling; and control signals. Data frames, such as illustrated in Fig. 1.17, are transmitted sequentially over the bus. In the frame, there is a 1-bit start-of-frame (SOF) marker, an 11- or 29-bit address for the intended recipient of the message, a 1-bit remote-transmit request (RTR) flag indicating whether this is a data frame or a remote frame, a 6-bit control field, up to eight bytes of message to be transmitted, a 16-bit cyclical-redundancy-check (CRC) for transmission error detection, a 2-bit acknowledgment (ACK) field, and a 7-bit end-of-frame (EOF) marker.¹⁶

1.12 Requirement 3c. Log book function

A third requirement that we list under the communication category is to provide a log-book function as well as a means of accessing these data. For warranty and diagnostic purposes, the BMS must store a log of atypical and abuse events. This log should include a record of the type of abuse (e.g., out-of-tolerance voltage, current, or temperature) and the duration and magnitude of the abuse. At a minimum, a vehicle technician should be able to gain access to this log for vehicle-diagnostics, warranty, and post-crash analysis purposes.

It is also necessary to be able to store persistent diagnostics regarding, for example, the number of charge/discharge cycles completed, the battery-state and SOH (battery-model parameter values) estimates for every cell at the end of each driving cycle, and more. These data are stored in a history chip (e.g., flash memory).

1.13 Requirement 4a. State of charge estimation

Referring back to Sect. 1.2, BMS design requirement 4 has to do with performance management. Under this umbrella, we introduce SOC, energy, and power estimation in this chapter. Later, in Chaps. 3 and 6, we will explore some of these in much greater depth.

¹⁶ Fortunately, many automotive-grade microcontrollers have built-in CAN communications hardware that take care of the details involving the protocol, leaving the BMS designer free to design messages and responses to meet the requirements of the application.