Autonomous Multi-Chemistry Secondary-Use Battery Energy Storage

A Dissertation Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Mitchell Taylor Smith
December 2019

Acknowledgements

First, I would like to acknowledge my advisor, Dr. Leon Tolbert for his continued guidance and support throughout my time in graduate school. I would like to thank him for providing an environment full of opportunities for me to grow as a leader and a scientist. I would also like to acknowledge my committee members: Dr. Michael Starke, Dr. Fangxing (Fran) Li, and Dr. Matthew Mench for their time and feedback during my Ph.D. studies. Additionally, I would like to thank Dr. Starke for guiding me through the many opportunities at Oak Ridge National Laboratory (ORNL) and ensuring I found the right fit.

I would like to thank the numerous scientists and colleagues at ORNL who played a critical role in my systems design and integration research: Steven Campbell, Dr. Madhu Chinthavali, Dr. Sheng Zheng, Daniel King, Andrew Herron, Jr., and Nate Watson. I am especially indebted to Steven Campbell for the days and weeks of time spent troubleshooting thousands of lines of DSP and python code for our energy storage deployment project.

Special acknowledgement and thanks go to the engineers and scientists at Spiers New Technologies. Dirk Spiers, Bryan Schultz, John Junger, Tim Willis, and Brandon Carter all provided valuable guidance, information and materials that were pivotal to my research efforts.

I thank the Bredesen Center for Interdisciplinary Graduate Research and Education for welcoming me to the university and providing a group of students that made going through the ups and downs of graduate school a much more bearable prospect. Thank you, Dr. Lee Riedinger and Wanda Davis, for taking a chance on me and inviting me into your family.

I would also like to thank my parents for their support throughout my entire Ph.D. studies. I am eternally grateful for their help keeping me focused, driven, and happy. I owe

much of my perseverance to my sisters, who kept me following in their footsteps with gentle advice and guidance every step of the way.

Finally, I acknowledge ORNL and the United States Department of Energy, Office of Energy, Energy Storage Program for providing the funding opportunities that sponsored this work in part.

Abstract

Battery energy storage is poised to play an increasingly important role in the modern electric grid. Not only does it help with peak shaving and shifting for renewable resources like wind and solar, it can also provide a host of other ancillary grid-stabilizing services. Cost remains a limiting factor in deploying energy storage systems large enough to provide these services on the scale required by an electric utility. Secondary-use electric vehicle batteries are a source of inexpensive energy storage systems that are not yet ready for disposal but are no longer appropriate for vehicle applications. However, the wide range of manufacturers using different battery chemistries and configurations mean that integrating these batteries into a large-format system can be difficult. This work demonstrates methods for the autonomous integration and operation of a wide range of stationary energy storage battery chemistries. A fully autonomous battery characterization is paired with a novel system architecture and transactive optimization to create a system which can provide utility-scale energy services using a multitude of battery chemistries in the same system. These claims are verified using a combination of in-situ testing and a computer modelling testbed. Results are presented which demonstrate the ability of the proposed method to combine a wide range of batteries into a single system.

Table of Contents

Chapter 1 - Introduction 1
Secondary-Use Batteries 7
Purpose of this Project 9
Chapter 2 - Review of Autonomous Multi-Chemistry ES Systems 12
Energy Storage System Architectures 12
Hybrid Energy Storage Systems 15
Multi-Modular Storage Systems 16
Optimization of BESS 21
Optimization for Maximum Service 22
Increasing ES Value Through Optimization 23
Optimizing to Account for Battery Lifetime 25
Transactive Optimization using an Energy Market 30
Parameterization of Secondary-Use Batteries 34
Battery Chemistry Identification 34
Battery Capacity Measurement 39
State of Health Estimation 41
Chapter Summary 44
Chapter 3 - Autonomous Battery Parameterization 46
Method 46
Charging to 100% State of Charge 47
Determining Pulse Depth and Length 48
Pulse Discharge Parameterization Test 51
Static Discharge Test for Capacity 53
Experimental Set-Up 54
Results 59
Manual Battery Parameterization 59
Autonomous Battery Parameterization 63
Discussion and Conclusions 63
Chapter 4 – Architecture for Control, Actuation, and Optimization 69
Overview of Agent-based Control Nodes 70
Central System Controller 71
Utility Agent 72

Intelligence Agent 73
Inverter Agent 74
Transactive Optimization Agent 76
Internal Communications Agent 77
Energy Storage Module Controller 78
Internal Communications Agent 78
Converter Agent 80
Battery Management System (BMS) Agent 81
Price Signal Agent 84
Intelligence Agent 85
Summary 87
Chapter 5 - Energy Storage Integration Demonstrations 88
Direct BMS Integration 88
AMIE Residential Energy Storage 89
Lessons Learned from Direct Integration 93
Agent-Based BMS Integration 95
Agent-based RES for BMS Integration Testing 96
BMS Agent Implementation 101
BMS Testing and Results 104
Discussion and Summary 113
Chapter 6 - Distributed Control through Transactive Optimization 115
Transactive Control Overview 115
Distributed Transactive Energy Market 118
Energy Market Overview 118
Price Signal Generation 121
Optimizing the Bid Curves 123
Virtual System Testbench 126
Energy Storage Module Controller Implementation 128
Central System Controller Implementation 130
Optimization Testing Plan 132
Results and Discussion 132
Early Optimization Results 133
Universal Charge/Discharge Binary Variable Results 134
Individual Charge/Discharge Binary Variable 139

Conclusions 142
Chapter 7 – Summary, Conclusions, and Future Work 143
Summary of Completed Work 143
Future Work 146
Conclusions 148
References 150
Vita 157

List of Tables

Table 1.1 - Electrical energy storage services and their typical parameters 6
Table 1.2 - Characteristics of various EV battery packs currently available for secondary-use applications. 10
Table 2.1 - Results showing various battery capacities and terminal voltages resulting in very similar gradients for single chemistry types. Results from original authors of [41]..... 39
Table 3.1 - A list of battery configurations and nameplate capacities run to test the autonomous parameterization techniques presented in this research. 58
Table 3.2 - Test results from running manual parameterization techniques alongside the autonomous method. 64
Table 4.1 - Sample of errors that trigger a system shutdown 75
Table 4.2 – Common battery parameters required by the BMS Agent 81
Table 4.3 – Errors monitored by the BMS agent 82
Table 5.1 - Sample of cases from the FMEA of the RES 107
Table 5.2 - OSHA hazard determination table..... 108
Table 5.3 - Sample of experiments used to test fault modes of the battery system 109
Table 6.1 - Parameters for Testing Optimization Functionality 132

List of Figures

Figure 1.1 - Li-Ion pricing and forecasts from 2006 to 2030 as presented by Nykvist et al. [11].. 3
Figure 1.2 - Energy storage technologies and their capabilities [13]. 4
Figure 1.3 - LBNL cost of power interruption breakdown [14] 5
Figure 1.4 – EV and PHEV sales by the Alternative Fuels Data Center. 9
Figure 1.5 – Examples of differing battery form factors and parameters 11
Figure 2.1 – Single source system with DC/DC boost converter and grid-connected inverter 14
Figure 2.2 – Single source system grid connected through a boosting transformer 14
Figure 2.3 – Semi-active system topology with multiple storage elements 15
Figure 2.4 - The parallel-active configuration for HESS 16
Figure 2.5 - Series-style multi-modular architectures 17
Figure 2.6 - Parallel-style multi-modular architectures 19
Figure 2.7 - A reduced section of the system proposed by Tesla. 21
Figure 2.8 - Battery chemistry figures of merit normalized to fit on a scale of 0 to 100 [42]. 28
Figure 2.9 - Generalized sample of a pulse charge identification of battery chemistry. 35
Figure 2.10 - A simplified circuit diagram of the experimental set-up. 36
Figure 2.11 - An example pulse discharge test showing the change in V_{oc} and V_{loaded} with time. Data is from completed work. 37
Figure 2.12 - A close-up of two open-circuit pulses. V_a represents the time when the relay is opened and the load disconnected. V_b represents the point when the relay is closed and the load connected 37
Figure 2.13 - a graph of transient voltage difference measured from a test. P1 and P2 represent the linear area of the curve that can be used to measure the gradient. 38
Figure 2.14 - Battery parameter estimation techniques for SOH measurement [54]. 42
Figure 2.15 - A seven-node equivalent circuit model for state of health determination. 43
Figure 3.1 – High-Level process to determine battery parameters autonomously. 47
Figure 3.2 – Process for determining if a battery has reached full charge based on voltage. 48
Figure 3.3 - An example voltage gradient with the three zones of the process. 49
Figure 3.4 - Iterative method to determine depth and length of the chemistry ID pulses 50
Figure 3.5 - Flow chart outlining the process used to execute the pulse test autonomously. 52
Figure 3.6 - Circuit diagram of power components in the parameterization test-stand. 55
Figure 3.7 - The testing apparatus used to execute the pulse-discharge test on the battery. 56
Figure 3.8 - Raw terminal voltage data for a single 12 Ah lead-acid battery. 60
Figure 3.9 – The result of finding the peaks and valleys of the raw dataset. 61
Figure 3.10 – The voltage difference between the peaks and valleys shown above 61
Figure 3.11 - Battery current from the pulse-discharge test shown above. 62
Figure 3.12 - An example of NiMH cells reversing voltage. 65
Figure 4.1 – System architecture and electrical topology of the multi-chemistry BES [58] 70
Figure 4.2 - General agent structure of a node [60] 71
Figure 4.3 - Central system controller agent framework [58] 72
Figure 4.4 - Agent framework of an ESMC [58] 78
Figure 4.5 - ESMC state machine including state triggers 86
Figure 5.1 - Photos of the AMIE Second-Life Battery System 90
Figure 5.2 - Simplified electrical diagram of the PCS in AMIE. 91
Figure 5.3 - Front-view of the AMIE power conditioning container 91

Figure 5.4 - High-level diagram of the control devices in the AMIE ES 92
Figure 5.5 - AMIE ES installed at the evaluation site [74] 94
Figure 5.6 - Cell voltage and balancing data from the Orion BMS 94
Figure 5.7 - System diagram of the agent-based RES 97
Figure 5.8 - Power conditioning for the GMLC RES project 97
Figure 5.9 - Two RES units using different COTS BMS units 100
Figure 5.10 - The RES testbed used for BMS testing 101
Figure 5.11 - Example BMS agent configuration file for the Watt Mini 2.0 102
Figure 5.12 - Example pairing code between BMS and Agent variables 104
Figure 5.13 - Demonstration of two BMS agents communicating with two BMS units 105
Figure 5.14 - Example of ES full cycle at 4.5 kW 110
Figure 5.15 - Price optimized testing of the RES 112
Figure 5.16 - Deployed and tested RES using the BMS agent and agent-based inverter 113
Figure 6.1 - Traditional price curves for large generation assets [66] 116
Figure 6.2 - Asymptotical bid function increasing price as departure time nears. 117
Figure 6.3 - Example price curves for three ES units at ~50% SOC 119
Figure 6.4 - Data flow between the CSC, API, and ESMC 120
Figure 6.5 - Example present value versus state of charge. 122
Figure 6.6 - Agent-Framework and Transactive Optimization Testbed [58] 127
Figure 6.7 – Li-Ion discharge curve with linear portion highlighted. 128
Figure 6.8 – State machine for modelling converter behavior 130
Figure 6.9 – Example dispatch schedule programmed into Utility Agent 131
Figure 6.10 - Oscillations in dispatch results from multiple optimal solutions 133
Figure 6.11 - Dispatch of three units with different initial SOCs 135
Figure 6.12 - Dispatch of three units with varying power ratings 136
Figure 6.13 - Dispatch of three units with varying energy capacity 137
Figure 6.14 - Dispatch of three units with different capital costs 138
Figure 6.15 - Test 1 results with new optimization constraints 140
Figure 6.16 - Various manual dispatch tests demonstrating optimization behavior 141
Figure 7.1 - Architecture for future work to integrate with a power electronics testbed 149

Chapter 1 - Introduction

In the coming decade, energy storage technologies are expected to revolutionize the electric grid. Whether energy storage is used to support grid-scale solar or wind plants or placed in a homeowner's backyard to provide load shifting, peak shaving, and solar mitigation, energy storage will help usher in a transformation of how electricity is bought and sold. The growth in energy storage interests and deployments is coming for a host of reasons.

Global energy policy has been increasing restrictions on carbon emissions in a push to move electricity generation towards renewable resources such as wind, solar, tidal, and hydroelectric. These resources often need some form of energy storage to convert the time the energy is produced to the time the energy is consumed. In the United Kingdom, environmental objectives have targeted an 80% reduction of greenhouse gases between 1990 and 2050. The United States has similarly targeted a 26-28% reduction of carbon emissions between 2005 and 2025 with a goal of 80% reduction by 2050 [1]. Yet, electricity demand grows all the while with an expected growth in demand of 25% during the same period [2]. With these policy changes and the rapidly increasing popularity of solar and wind, renewable resources are beginning to take a significant hold on the global energy market. In 2013, global PV capacity was estimated to have reached 139 GW with a projected exponential growth [3]. Similarly, wind power's steady growth resulted in a 2013 capacity of 318 GW [4]. Yet, intermittency remains an impediment to further integration of renewables given that solar irradiance and wind are not constant year-round or even day-to-day. Energy storage enables increased renewable penetration by managing the intermittency of solar and wind. When renewable generation cannot meet demand requirements or exceeds consumer demand, energy storage is able to act as a buffer to provide or absorb energy to and from the grid. Past issues with peak energy supply has led to

questions of why more grid energy storage does not currently exist. Summertime rolling blackouts have become common in states like California where supply cannot meet peak demand on hot days [5]. As recently October of 2019, PG&E has implemented rolling blackouts for over 150,000 customers with California Edison cutting power to more than 300,000 people. The measures reduce transmission demand and are an attempt to prevent failed lines from starting wildfires [6]. Energy storage could help with intermittency problems and provide peak demand relief; especially in a state with 13.8% renewable penetration and plans to continue increasing renewables and eliminating coal and nuclear energy generation [7]. In most cases, the answer comes down to cost versus useful product lifetime.

Today, the expected time to recoup the cost of an energy storage investment often exceeds the expected life of the system. For example, older lead-acid battery technologies require low capital investment, but face problems with the number of charge/discharge operating cycles being in the range of one or two thousand. Lithium-ion (Li-ion) can cycle tens of thousands of times but are prohibitively expensive with costs typically exceeding $300 per kWh [8]. For newer systems like redox flow batteries (RFB) promising high capacities at a low $/kWh cost, utilities show hesitance to invest in what they deem untested or incomplete technologies. Cost reduction of all storage methods has been a thematic focus of most US energy storage programs since the DOE first reported on energy storage in 2003. In 2013, the DOE proposed a short-term energy storage goal for 2019 of approximately $250/kWh with a long-term goal reducing the price to $150/kWh by 2023 [9]. Sandia National Lab (Sandia) and the Electric Power Research Institute (EPRI) showed that most demonstrated energy storage technologies cannot yet meet these goals as of 2013 [10]. In 2015, the Lithium-Ion industry was evaluated by Nykvist et al. In their work, summarized in Figure 1.1, they found that price

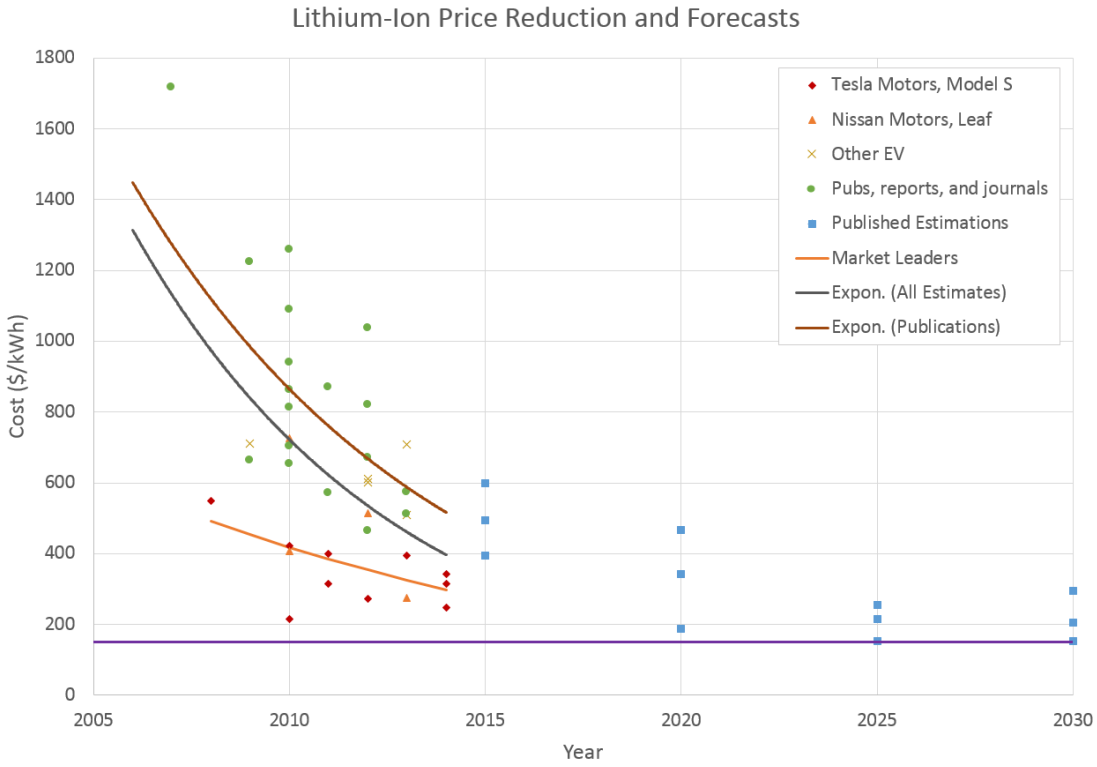


Figure 1.1 - Li-Ion pricing and forecasts from 2006 to 2030 as presented by Nykvist et al. [11]

reductions and forecasts were not predicted to reach below $150/kWh until sometime after 2030 [11]. Bloomberg New Energy Finance (BNEF) followed with another evaluation in 2017 with a report that suggested batteries could theoretically be produced in a Korean production facility in 2017 for $162/kWh and reach cost targets as low as $74/kWh by 2030. However, BNEF is careful to note that these are production costs and are not commercial pricing suggestions [12].

In meeting a reasonable ROI, cost of the energy storage is one issue, and the value achievable by performing ancillary services to either offset load or earn credit from a utility is another value stream. Energy storage systems have been demonstrated to cover a wide-reaching range of methods and applications. Figure 1.2 gives a generalization of the portfolio of grid

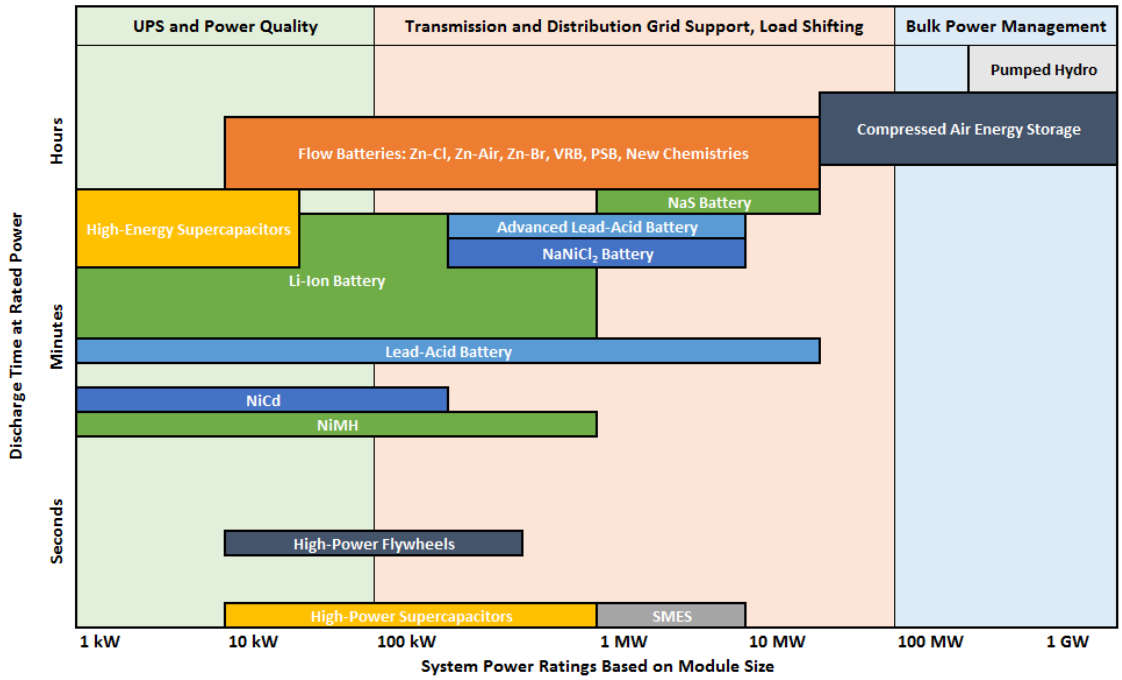


Figure 1.2 - Energy storage technologies and their capabilities [13].

applications which range in power level from 1 kW to 1 GW with best suited energy storage technologies mapped by capacity ranging from seconds to hours [13].

Providing reliability through uninterrupted power services has been the driving force behind much of the past energy storage work. Besides preventing nuisance interruptions or mitigating the need for load-shedding activities, there is a large financial incentive to provide uninterrupted power. Lawrence Berkley National Lab (LBNL) performed a survey in 2004 estimating the financial impact of short-term power interruptions on various power sectors. In their study, they demonstrated that an hour-long outage could cost an industrial customer up to $3,253. Commercial customers would be less impacted by an hour-long outage with an estimate of only $886. In Figure 1.3, the estimated financial impact of power interruptions is shown by

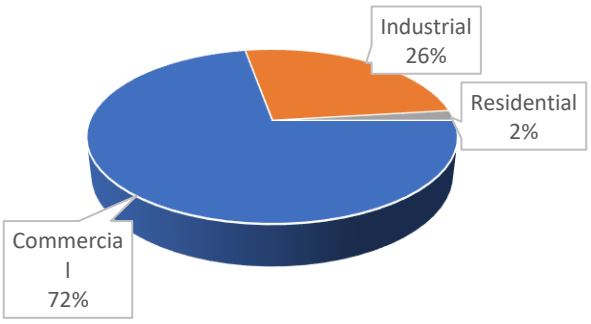


Figure 1.3 - LBNL cost of power interruption breakdown [14]

energy sector with a total yearly estimate of around $79 billion [14]. For commercial and industrial consumers, these values are easy to estimate given that power outages result in calculable losses from production downtime. For residential consumers, financial losses from a blackout are more difficult to calculate since this value largely rests on how much worth the customer attributes to having energy in the home.

To provide additional value to ESS and offset the cost, services beyond back-up power have been explored. In Table 1.1, a list of energy storage services are presented and described. These services are often difficult to provide with traditional generation resources [10]. Many of the services described in Table 1.1 have been demonstrated using various energy storage types. Some have been implemented on megawatt-scale flow batteries while others have been shown on kilowatt-scale residential energy storage devices. In 2015, Starke et al. present a "community energy storage" device capable of providing power quality, retail energy time-shifting, and demand charge management services for a small number of residential consumers. Their work demonstrated the ability of a 25 kW, 50 kWh secondary-use battery energy storage system to provide power factor correction, load-correction, load-leveling, and peak shaving operation [15]. Other demonstrations can be found in the 2013 DOE report on the utilization of ESS [16].

Table 1.1 - Electrical energy storage services and their typical parameters

Name	Typical Size	Discharge Duration	Min. Cycles per Year	Description
Bulk Energy Services				
Energy Time-Shift (Arbitrage)	1-500 MW	< 1 hour	250 +	Shift energy from peak production times to peak consumption times from PV and/or wind farms
Electric Supply Capacity	1-500 MW	2-6 hours	5-100	Scheduled and used just like a generator (hydro)
Ancillary Services				
Regulation	10-40 MW	15-60 min	250-10,000	Control-area specific maintenance of frequency
Spinning, Non-Spinning, Supplemental Reserve	10-100 MW	15-60 min	20-50	Used in place of spinning reserve generators, but with little to no losses while idling
Voltage Support	1-10 MVAR	N/A	N/A	Support system voltage by injecting or absorbing reactive energy (real power not needed)
Black Start	5-50 MW	15-60 min	10-20	Energize transmission lines to bring power plants back on line after a catastrophic grid failure
Load Following	1-100 MW	15-60 min	N/A	Similar to damping variability in wind and PV, ES can dampen variability in load as well.
Frequency Response	Variable	Seconds	N/A	Helps when freq. suddenly drops due to the sudden loss of a generator or transmission line.
Transmission Infrastructure Services				
Upgrade Deferral	10-100 MW	2-8 hours	10-50	Mitigate the need for new transmission equipment by providing small generation boosts
Congestion Relief	1-100 MW	1-4 hours	50-100	Store energy at low congestion times and use to reduce peak transmission capacity requirements
Distribution Infrastructure Services				
Upgrade Deferral	.5 -10 MW	1-4 hours	50-100	Offset peak load to help mitigate the need for new generation assets to be built.
Voltage Support	.5-10 MW	1-4 hours	50-100	Provide voltage sag support along distribution feeder when heavy loads could result in an under-voltage event
Customer Energy Management Services				
Power Quality	100kW – 10 MW	10 sec – 15 min	10-200	On-site storage monitors power quality and discharges/charges to smooth power to the load
Power Reliability	Load-dependent	Length of outage	High	Storage systems allow customers to island and self-power during system outage events
Retail Energy Time-Shift	1-1000 kW	1-6 hours	50-250	Also called “peak-shifting,” customers store energy at off-peak times and discharge at peak times
Demand Charge Management	50 kW - 10 MW	1-4 hours	50-500	Also a time-shifting function, but used to reduce energy consumption during peak demand-charge

Secondary-Use Batteries

Widespread adoption of energy storage has been limited by high capital costs and low return-on-investment (ROI) figures. For example, residential customers looking to obtain a 25 kWh battery system for home use would still be looking at a $6500 cost using the DOE's cost goal of $250/kWh for 2019. After adding in the bulk-of-system cost components including an inverter and other installation hardware, a unit capable of providing residential ancillary services and limited black-out services could end up with costs exceeding $10,000[1]. While this may be a feasible option for some consumers, this likely does not have economic merit for most consumers [15]. In a 2011 report, Oak Ridge National Laboratory (ORNL) proposed the adoption of used electric vehicle (EV) batteries as a source of inexpensive batteries for energy storage [17].

Second-life (or secondary-use) batteries for energy storage are taken from EVs or PHEVs at the end of the battery system's useful life. A secondary-use battery can be defined as any battery pack that has been used to the point where only 70-80% of the initial rated capacity remains [17]. As with solid-state lighting devices, these end of life (EOL) conditions are defined at these values because there is no sharp drop-off that signals EOL like an incandescent light-bulb burning out. Instead, a device is guaranteed to stay above an operational threshold for a given time window. For example, a new Nissan Leaf is capable of a minimum 60 miles of range for 6-10 years. While the car starts with 83 miles of range, the vehicle is guaranteed to operate over 60 until end of life is reached [18]. While this depends heavily on the usage of the car, the ongoing assumption is that the 30 kWh battery pack in the car would still maintain at least 21

[1] This assumes an inverter cost of approximately $2000, which is approximately the current cost of a 10-kW commercial solar inverter. It also assumes $1500 in skilled labor for installation.

kWh of energy storage capacity. Capacity is used as a primary metric since cycle life can be harder to predict since it can change as the use-case for the energy storage device changes.

The early adopters of EVs and PHEVs have begun to see the end of useful life. Additionally, manufacturing flaws or other warranty-replacement issues have resulted in increased availability of used batteries. As a result, the availability of these used vehicle battery packs has begun to significantly increase and will continue to increase as the sales of EVs and PHEVs continue to grow as sales data show in Figure 1.4. Cooling of the market and the dip in 2015 were caused by a significant decrease in gas prices in 2014 and 2015 [19]. The ongoing decision for most original equipment manufacturer (OEMs) is whether these systems should be destroyed, recycled, or reused. Presently many of these systems are being shelved with the hope of developing rapid reprocessing. Spiers New Technologies (SNT), a battery reprocessing plant, reports storing nearly 1500 vehicle battery packs from Nissan vehicles alone.

An evaluation of reuse led to the 2011 study, and ORNL predicted that these batteries could be sold for as little as $75-250/kWh, far below the DOE's 2019 target of $250/kWh [17]. Nissan, maker of the Leaf electric vehicle, has valued their spent battery packs at $1000 as part of their upgrade program to replace current vehicle owners' spent batteries [20]. Assuming a spent pack contains 70% of its initial 24 kWh of energy, this means that the battery could be resold by Nissan at around $60/kWh. With minimal repackaging and grading processes, this number could beat the DOE target in a much shorter time window than developing technologies. Some predictions like those in Figure 1.1 do not see commercial prices reaching $150/kWh until 2030 or later. Others, like BNEF, see the manufacturing costs for new batteries finally approaching $75 /kWh around 2030, but do not mention a possible retail cost for the same prediction [12].

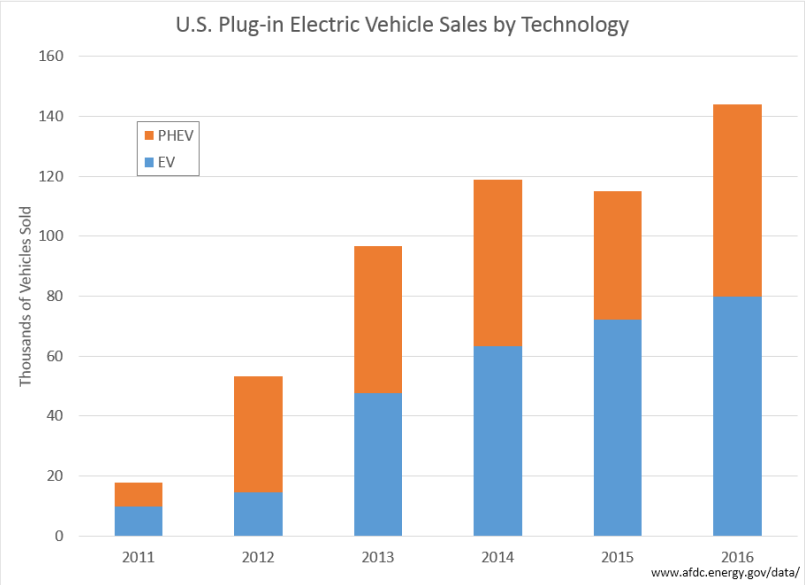


Figure 1.4 – EV and PHEV sales by the Alternative Fuels Data Center.

With utilities searching for ES technology capable of 100 kWh to 1 MWh of storage, secondary-use EV systems should not be overlooked. For example, the Leaf battery packs currently in storage at SNT come to approximately 38 MWh of storage. Assuming 25% of these units are serviceable for vehicles and another 25% are destined for recycling, this still leaves 19 MWh of inexpensive storage medium. With an appropriate hardware and software architecture for incorporating these units into an ES system, many utilities could afford to invest in ES to mitigate the growing amount of intermittent renewable energy.

Purpose of this Project

A major impediment to providing affordable energy storage is the wide variety of battery types available. As shown in Table 1.2, different vehicle manufacturers use a variety of battery chemistries, resulting in an array of pack voltages, capacities, and capabilities. Many full-electric or PHEV vehicles use various Li-ion chemistries while hybrid EVs or "e-assist" vehicles

Table 1.2 - Characteristics of various EV battery packs currently available for secondary-use applications.

Vehicle	Battery Chemistry	Rated Capacity (Ah)	Capacity (kWh)	Pack Voltage (V)	Nominal Cell Voltage (V)
Nissan Leaf	$LiMn_2O_4$	33.1	24	403.2	3.8
Chevy Volt	NMC-LMO	50	18.4	395	3.75
Fiat 500e	Li-ion	60	24	364	4.01
BMW i3	Li-ion	60	22	360	3.75
Toyota Prius	NiMH	6.5	1.31	288	7.2

utilize smaller Li-Ion packs or nickel-metal-hydride (NiMH). This poses a challenge to developing a uniform grid-interface since the source of the energy can have many different limiting factors. Additionally, these packs of different size, capacity, and voltage will not have all aged equally. Due to variations in manufacturing, individual cell aging, and total cycling, batteries of the same type could still end up behaving differently. Cell aging is often captured using the term "State of Health" (SOH). State of Health is most often defined by the measured capacity of the battery versus the initial capacity and by a measurement of the internal resistance of the battery.

Integrating vehicle battery packs into a common system for grid energy storage is a challenge. The challenge comes in the variation of voltage, capacity, and response ranges because of different chemistries and battery pack ages. Figure 1.5 shows the various battery module configurations along with the module voltage and cell configuration information which are being used throughout this work. For this work, a battery module is the smallest discrete unit that a pack can be broken down into. A cell will be defined as a sealed, prepared unit which could theoretically be separated from the whole module; though in most cases, this would prove time consuming and difficult. As an example, a Nissan Leaf battery pack contains 48 battery modules. Each module consists of four cells with two parallel groups in series (denoted 2P2S).

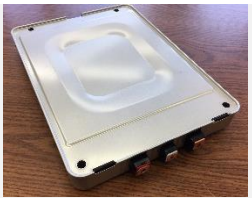
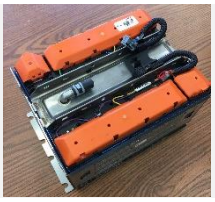
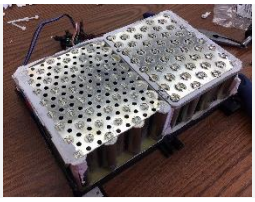
Nissan Leaf	Fiat 500e	BMW i3	Toyota Prius
			
3.3 - 4.15 V	2.8 – 4.15 V	3.1 – 4.15 V	6.5 – 9.2 V
20 Ah/cell	60 Ah/cell	0.92 Ah/cell	7 Ah/cell
2P2S	5S (or 6S)	54P2S	6S

Figure 1.5 – Examples of differing battery form factors and parameters

Effectively, this means there are 48 modules containing 96 measurable cells in a pack, but 192 pouch cells in the entire system. The goal of this work is to demonstrate an autonomous methodology for the seamless integration of a wide range of battery systems using a framework which is independent of these features.

In the following chapter, a review of medium-scale battery energy storage systems and associated technologies is presented to depict existing battery grid integration frameworks. Chapter 3 will describe the methods used to autonomously identify battery chemistry and parameters along with results from testing. Chapter 4 will present a software and hardware architecture for multi-chemistry utility-scale battery energy storage. In Chapter 5, an implementation of the agent-based software will be presented, tested, and deployed for a related project. This will be followed by a description of the transactive optimization method designed to control the BES and the results from implementation in Chapter 6. To close, Chapter 7 will present a summary of the research, conclusions, and a discussion of future related work.

Chapter 2 - Review of Autonomous Multi-Chemistry ES Systems

Battery systems composed of multiple aged batteries or different chemistry units that were also capable of self-diagnostics and chemistry identification were not present in the literature. However, the concepts behind a multiple energy source grid-connected inverter has been thoroughly explored in previous work such as Hybrid Energy Storage Systems (HESS) and other grid-scale energy storage devices. In this chapter, the architecture of grid interconnected energy storage systems is reviewed. Additionally, control and optimization methods for these types of systems are investigated. This is followed by a discussion of energy storage autonomy and how systems today monitor batteries operating in a storage system.

Energy Storage System Architectures

At their core, all energy storage systems (ESSs) consist of a power conditioning system (PCS) and a form of energy storage element. The purpose of the PCS is to manage the bidirectional power flow between the AC electric grid and energy storage medium which is usually DC. The last stage of the PCS usually consists of a power electronic system (PES) called an inverter that converts DC to AC for grid interconnection. For an inverter to be able to control bi-directional power flow, the voltage level of the dc connection must be increased above the bridge rectification voltage of the inverter. In a typical six-switch inverter for a three-phase system, the reverse conduction body diodes of the switches will act as a full-bridge rectifier if the DC voltage on the input does not exceed the natural rectification voltage of the AC connection. For a three-phase system, this voltage is calculated as:

$$V_{dc} = \frac{3 * \sqrt{3} * V_{peak}}{\pi} \; where \; V_{peak} = \sqrt{2} * V_{rms_ln} \qquad (2.1)$$

On a three phase – 480 Vac system, this means that the DC voltage must exceed 648 Vdc to prevent the body-diodes from performing uncontrolled rectification. Not accounting for voltage drops across the switches or passive components, this is also the minimum voltage required to synthesize real power only; with an inverter having both real and reactive capabilities needing a much higher DC-link voltage. Higher DC-link voltages also enable the inverter to push power onto the grid more easily.

For safety reasons, most battery stacks do not have nominal voltages high enough to push onto a distribution network or even a 480 Vac bus. In fact, most systems are designed to have DC link voltages below 600 Vdc to avoid internal isolation issues with the batteries and to reduce the need for cables with larger insulation, which occupy more volume. To achieve higher voltages, another set of converters is often used to boost the voltage. Figure 2.1 demonstrates the first way this can be achieved in a single-source system where the voltage is boosted first with a dc/dc converter and then sent through an inverter to connect to the grid. Often, the inverter coupled to the grid through a 1:1 ratio isolation transformer, but with the proper attention to grounding, this step may be eliminated. Alternatively, the power can be inverted first and then sent through a boosting AC transformer to couple to the grid as Figure 2.2 shows [21] [22]. These two system topologies represent the most basic form of energy conversion for energy storage. The switch configuration of the internal dc/dc and dc/ac converters can be changed to suit the needs of the energy source and the grid interconnection. When another energy storage source medium is added to this type of system, the system complexity increases and is the focus of the following discussions.

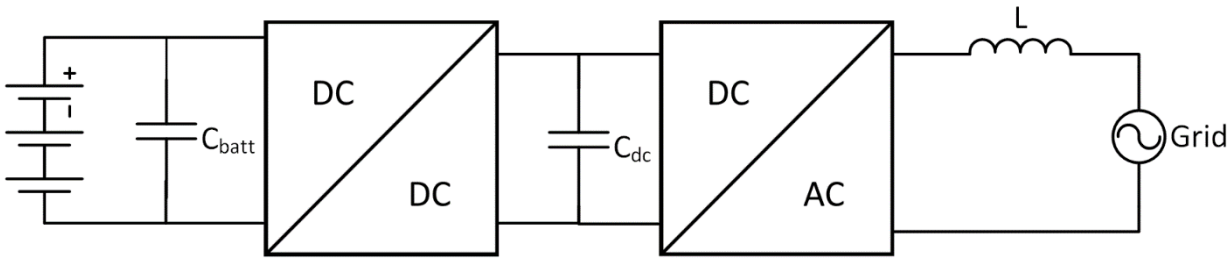


Figure 2.1 – Single source system with DC/DC boost converter and grid-connected inverter

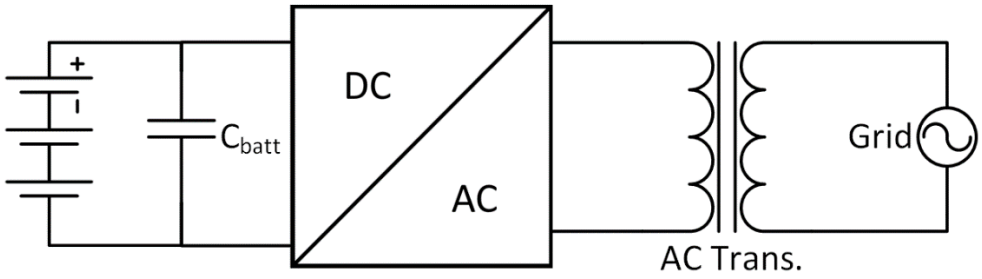


Figure 2.2 – Single source system grid connected through a boosting transformer

Hybrid Energy Storage Systems

Integrating a second energy source has been previously explored with hybrid energy storage systems. Whether integrating a supercapacitor or a superconducting magnetic energy storage (SMES), many have proposed variants on a combination of active and semi-active topologies. In both cases, the architecture follows the same basic principle outlined in Figure 2.1. By closely regulating the DC bus, an excessive amount of current or voltage ripple can be prevented from appearing on the battery. Figure 2.3 represents a semi-active architecture where a power-dense ES device regulates the DC bus voltage through a direct connection and a DC/DC converter regulates current flow to a more energy-dense ES element. In a system such as this, the power-dense ES device must be a fast-responding element such as a SMES or supercapacitor in order to reduce the presence of high-current pulses that usually degrade batteries [23] [24].

Another architecture for DC hybridization is the parallel active system shown in Figure 2.4. Here, both energy storage elements use a dc/dc converter to carefully control the bidirectional power flow to their respective element. A system such as this can be used to pass the bulk of high power pulses to either energy source depending on the present condition of each

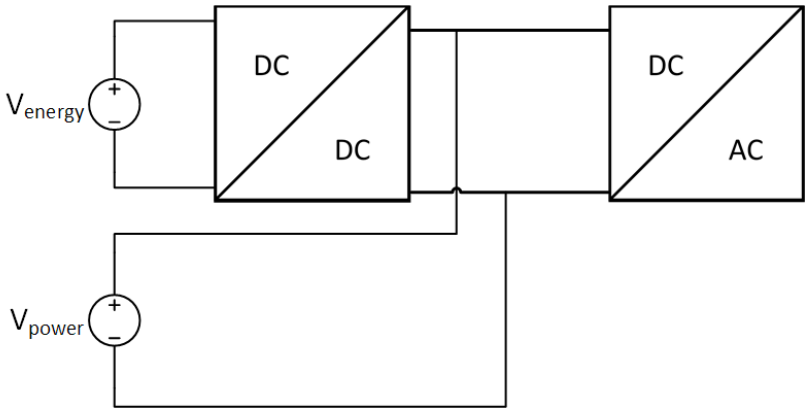


Figure 2.3 – Semi-active system topology with multiple storage elements

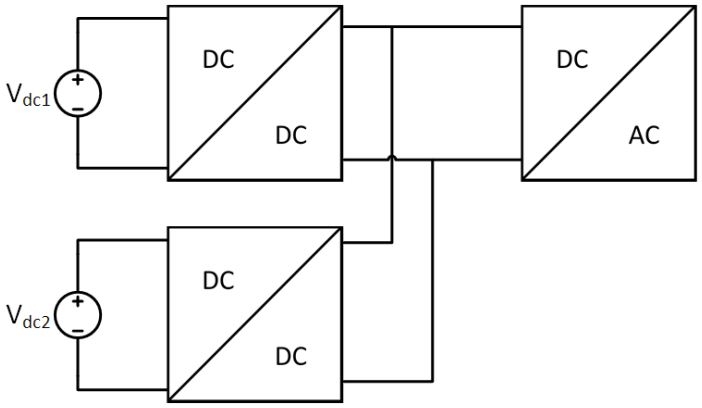


Figure 2.4 - The parallel-active configuration for HESS

ES element. The authors of [23] conclude that this topology gives the best results for a HESS, but this comes with increased cost and slightly reduced reliability due to the greater number of switches. Both architectures provide a means for integrating two energy storage mediums, but only the second can expand to incorporate a larger number of energy storage elements. Furthermore, the control systems architectures for this HESS have been limited to two sources and have not proven modular expandability.

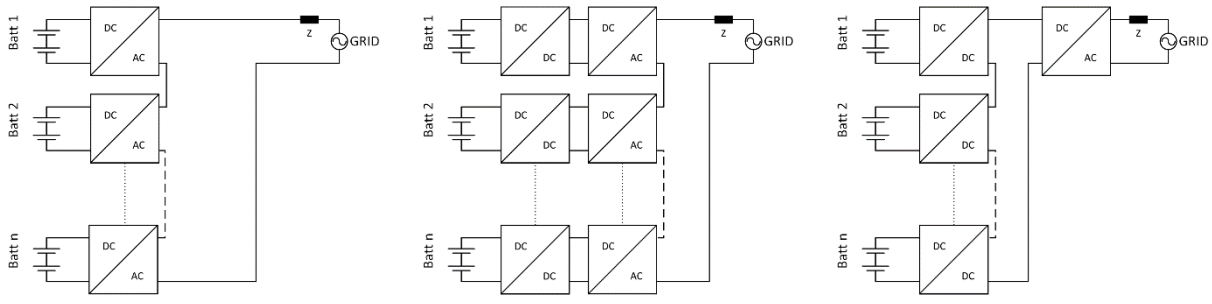
Multi-Modular Storage Systems

For the autonomous multi-chemistry BES to be easily expanded, system modularity is desired. As mentioned above, cell-level systems are being omitted from this discussion of energy storage devices since this type of architecture would be prohibitively expensive in a secondary-use type system where battery modules are already constructed. However, the architecture and terminology used to describe these systems can be applied at a higher level to a multi-chemistry system. Most of the modular architectures can be split into two groups; cascaded (or series) systems and parallel systems. Cascaded systems involve configurations like

those shown in Figure 2.5(a-c), while parallel systems are like those in Figure 2.6(a-c) and Figure 2.4 [25].

Cascaded systems follow the same basic architecture described in Figure 2.2, but differ in how the post-inverter voltage boost is achieved. Figure 2.5(a) shows a cascaded PWM converter where each converter synthesizes a part of the AC waveform. In this topology, the AC waveform is split into smaller sections and each device is only responsible for providing the power needed to synthesize one part of the waveform. As such, much lower voltages are needed since the systems are in series and the voltages stack across the synthesized waveforms. These systems can achieve very good AC output with a THD of less than 3.5% even at very high power levels [26]. Additionally, this architecture meets the criteria of keeping the independent battery stack voltages low since the converters are in series.

Figure 2.5(c) uses the idea of a series boost to gain the same voltage boosting effect. However, instead of using a dc/ac converter on each battery section, the current flow from each cell is regulated by a dc/dc converter which then feeds the DC link of a dc/ac converter. This architecture was used in 2014 for the design of a cell-level charge balancing controller [27]. For



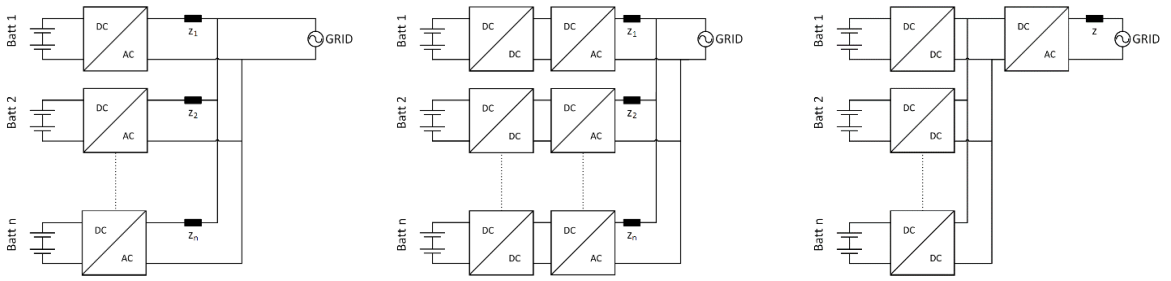
(a) Cascaded dc/ac inverters *(b) Cascaded dc/dc/ac converters* *(c) Series dc/dc with dc/ac inverter*

Figure 2.5 - Series-style multi-modular architectures.

this application the voltages per cell are in the range of 3 to 24 Vdc. The last of the series configurations, Figure 2.5(b), operates like the one shown in Figure 2.5(a), but has the added benefit of a dc/dc converter to isolate the ES element from the AC converter. Operationally the series modules in this system are like the first conversion strategy demonstrated in Figure 2.1.

Benefits to using a cascaded architecture include lower required DC voltages per pack while still maintaining the same AC output voltage. A 2014 reliability review of BES converters also calculated that series-connected systems are, on average, lower cost and more reliable to run [25]. A significant drawback to the cascaded architectures is that individual battery packs could not be charged while another pack is being discharged. The grid must be involved in all energy transactions between the various energy storage elements, which means that internal balancing cannot occur without grid energy exchange. Also, if individual packs can no longer supply the voltage or power needed to sustain their part of the AC waveform, the system would no longer be able to provide any power, even if other packs were still mostly charge. With the architecture shown in Figure 2.5(a), if a storage element fails, it can be bypassed only if the remaining pack voltages are high enough to support the AC waveform production. This means that the system can only afford a certain number of battery failures before it must be decommissioned, despite several healthy storage elements remaining. This makes this system conducive to using many identical battery systems in order to prolong the life of a newly manufactured system and less appealing when a variety of capacities and chemistries are present.

Parallel multi-modular architectures can require higher battery voltages like in Figure 2.6(a) because they must be able to synthesize the full AC bus voltage. However, systems like this are favored when micro-grid controls are already in place. For instance, in 2014 the CERTS team proposed and demonstrated a P vs. f droop-controlled energy storage micro-grid. In this



(a) parallel dc/ac inverters *(b) parallel dc/dc/ac converters* *(c) parallel dc/dc with a dc/ac inverter*

Figure 2.6 - Parallel-style multi-modular architectures

work, the authors connected several energy storage elements to an existing 480 Vac 3Ø grid and demonstrated the abilities of secondary-use batteries to respond to grid-islanding events [28]. A demonstration of this architecture was completed in Zhangbei, China as a 14 MW Li-Ion BESS. The system was split into nine separate units of varying sizes connected to a local micro-grid through nine transformers [29]. This parallel architecture was also used in the simulation of a HESS using a Li-ion BESS and SMES for smoothing the output of a wind-farm. To reduce the required DC voltage, a dc/dc converter can be added before the dc/ac inverter. Essentially, this is the same configuration as Figure 2.6(a) but using the two-stage energy conversion system that was detailed in Figure 2.1. This approach of integrating multiple energy storage elements in the style of an AC microgrid was shown to be robust but did require one controller monitoring all devices simultaneously [30]. This means there must be constant and fast communications between the grid-connected devices in order to dispatch the system in an optimal way.

As previously mentioned, the parallel structure of Figure 2.6(c) has been used for several years by HESS groups. In 2011, the authors hybridized a super-cap onto a lead-acid battery

system using this architecture. The system used three dc/dc converters to interface with the lead-acid battery system, a supercapacitor, the regenerative braking component, and a dc/ac converter for the electric motor. Through simulation, the authors found this architecture reduced the demand on the battery system by as much as 20% [31].

Tesla Motors began an attempt to patent this class of architectures in 2016, but as of January 2018, the patent has yet to be awarded. As described in the patent, the system consists of a parallel and series group of batteries connected to a DC/DC converter. Strings of two battery packs are placed in series with the center-point of the battery tied to ground. This creates a positive voltage from one stack and a negative from the other. This is done to reduce the voltage-ratings of the cables while still maintaining a positive to negative dc-link high enough to convert to 480 Vac. Multiple of these units are then placed in parallel to increase capacity. Subsequently, a plurality of these DC/DC converters are connected to a main DC/AC converter which can then interface with more energy storage units through an AC micro-grid like Figure 2.6(a). Figure 2.7 shows a reduced electrical model of the system proposed by Tesla's un-awarded patent. Each battery symbol represents an entire battery. The patent shows a system where up to sixteen 6.5 kWh battery units are placed in parallel behind a DC/DC converter [32]. This seemingly solves both the voltage issues and capacity issues that most of the above system topologies have encountered. However, placing batteries directly in series and parallel relies on the capacities and internal resistances of the units being nearly identical. In most cases, this means the battery packs should not only be new units, but also intentionally matched for ideal operating conditions. While this could be done with second-life batteries, it would be difficult since matching battery characteristics requires extensive measurements to be made of every cell

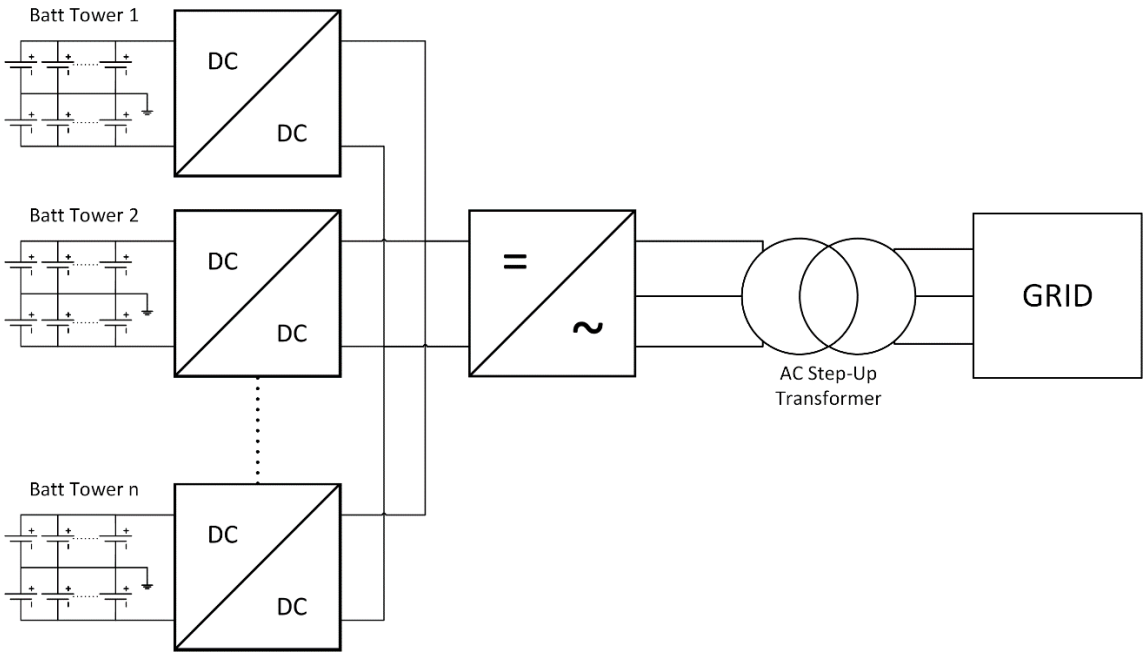


Figure 2.7 - A reduced section of the system proposed by Tesla.

within a used EV pack. This process would also significantly increase the cost associated with recycling secondary-use ES elements.

Optimization of BESS

Architectural complexity, while improving usability and efficiency, can increase until the cost of the power electronics converters far outweighs the cost of the batteries. Without an effective monitoring, optimization, and control strategy, the batteries could be dispatched in a manner that quickly degrades the system without maximizing the benefits. Previous works have implemented many different optimization strategies for controlling the energy flow in and out of batteries. Typical optimization strategies focus on providing the maximum amount of service, regardless of the monetary value. This method can be refined so that the energy storage system provides service whenever the monetary value to the owner or consumer is maximized. Additionally, considerations for battery lifetime can be added to maximize battery profits while

extending battery life. The ES optimizations presented here focus on single-element optimization. Multi-element ES optimizations are largely absent from the literature. In their place, transactive optimization methods used with microgrids will be presented. These could be adapted for use in a multi-battery energy storage system.

Optimization for Maximum Service

In Chapter 1, Table 1.1 detailed the many services that can be provided by an ESS. Among these services is demand-charge management, which can be enacted using techniques such as peak shaving and shifting. Peak shaving is the action of storing energy during off-peak times and pushing it back to the grid to reduce peak load. For instance, a 2 MW BESS system attached to a substation for electric urban rail transit, energy from regenerative braking was stored during off-peak. During peak energy-use times, the energy was discharged back into the substation to offset the energy-use of the trains. However, choosing how much energy to store during off-peak time and when to dispatch it during a peak requires an optimization to reduce the peak power draw of the train system. The optimization is then constrained to prevent over-charging or over-discharging the BESS. Adding this system optimization aided in the determination of a peak power threshold, over which the BESS would either absorb power during off-peak times or release power to the system to reduce demand during peak times [33].

$$\begin{cases} F_1 = \left(\dfrac{1 - \sum_{i=1}^{n} E_{sub,ess}(i)}{\sum_{i=1}^{n} E_{sub}(i)} \right) * 100\% \\ F_2 = \left(\dfrac{1 - \sum_{i=1}^{n} P_{sub,ess}^{peak}(i)}{\sum_{i=1}^{n} P_{sub}^{peak}(i)} \right) * 100\% \end{cases} \tag{2.2}$$

$$-P_{batt,max} \leq P_{batt-ref} \leq P_{batt,max} \tag{2.3}$$

$$\Delta U_{dis} \geq 0, \Delta U_{char} \geq 0 \tag{2.4}$$

The objective function for this optimization is shown in (2.2). Here, $E_{sub}(i)$ and $P_{sub}^{peak}(i)$ are the output energy and peak power of the substation if it did not contain energy storage. $E_{sub,ess}(i)$ and $P_{sub,ess}^{peak}(i)$ are for a station containing energy storage. By maximizing these functions, the optimization could determine how to dispatch the ESS. To prevent charging or discharging too quickly, the system is constrained by (2.3). To ensure that the correct amount of energy is being charged and discharged and the system is not expected to provide or absorb more energy than it contains, the system is also constrained by (2.4). These values (ΔU_{char}, ΔU_{dis}) are determined by a genetic algorithm in a way that maximizes energy savings and minimizes peak power at the substation [33]. While this type of optimization may result in the best service provided to the system, it does not take in to account the possibility for the ES to earn revenue for performing these ancillary services. Given that capital costs of energy storage have long been a concern, it would be better if the system could maximize the revenue generated from providing these services.

Increasing ES Value Through Optimization

Wu et al. have shown through a simulated energy storage system that the dispatch of an ES device can be optimized to capture multiple services with the goal of maximizing the revenue generated by the unit [34]. Revenue generation is dependent on the services provided and the perceived value to the utility. In this work, an optimal control strategy for BES is proposed that uses a look-ahead window to balance the service capacity of the battery on an hour-by-hour basis. Each hour, the optimization problem is solved to determine a base operating point for the battery. This base operating point is the amount of power that the unit exchanges with the AC grid. Linking these optimization windows together provides a unique challenge when dealing with revenue maximizing optimizations. Without interconnection between the windows, the

system may earn a significant revenue in one hour but not leave enough energy for the system for future use, which would decrease the overall revenue of the system. In the objective function, (2.5), used to achieve this goal, the net revenue generated from trading energy with the grid ($\lambda_k p_k$) and the balancing services ($\beta_k r_k$) provided to the grid are maximized within similar constraints to those listed in (2.3) and (2.4) [34].

$$\max_{p_k^+, p_k^-, p_k, p_k^{batt}, l_k, r_k^+, r_k^-} \sum_{k=1}^{K} [\lambda_k p_k + \beta_k^+ r_k^+ + \beta_k^- r_k^-] \tag{2.5}$$

In the first term, the hourly energy price (λ_k) is multiplied by the net power exchange (p_k) with the grid to give the net value generated from buying or selling power. The second term ($\beta_k^+ r_k^+$) gives the revenue generated by absorbing power from the grid as a service while the third term ($\beta_k^- r_k^-$) gives the revenue generated by pushing power to the grid as a service. Together, these three terms are maximized over a 24-hour window and a dispatch schedule is generated for the battery system. This objective function is appropriately constrained to limit the amount of energy that the battery can supply or absorb and the rate at which it performs these actions. A final constraint is given that sets the desired end-of-day SOC for the battery so that at the end of any given 24-hour window, the battery is not left exhausted of all energy. To verify the optimization, the algorithm was tested in simulation on a system based on a real-life 4 MW, 16 MWh flow-battery. The results showed that over a 20-year lifespan, the system could generate $26.6 million in revenue. When the capital costs of $14.8 million and operating costs of ~$5.5 million are deducted from this number, the total profit return came in at $6.5 million [34]. While this system provides promising numbers, it does not take in to account the reduction in life over time that a traditional battery system will suffer. Flow batteries experience aging like traditional storage mediums, but preventative maintenance can refresh the battery to the initial

specifications and the cost of this maintenance can be absorbed into operating costs. With traditional Li-ion batteries, the aging processes are not easy to reverse, if they are reversible at all. If the battery does not make it to 20 years, then the revenue earned may not outweigh the system capital cost.

Optimizing to Account for Battery Lifetime

Cycle life, which is a component of battery lifetime, has been a concern for BES systems and a major impediment to their widespread adoption. The optimization must be constrained to prevent overcharge or over-discharge which can significantly degrade the lifetime of the system. However, heavy cycling at high charge and discharge rates can also impact the life of the battery [35]. Works investigating EV charging stations have studied these effects extensively. In fact, EVs using "fast-charging" stations at 50 kW rates often only charge the batteries to 80% to avoid doing damage to the battery system [36] [37]. Optimizations attempting to target battery life mainly do so through two methods. First, the power variation to and from the battery is limited so the system is not constantly ramping up and down. The other method is to mathematically place a weight on cycling the battery such that the system only dispatches the battery if the benefit outweighs the negative effects on the battery. As the battery ages, the objective weight increases to make it harder to dispatch the system, thus prolonging the battery life.

Looking at how hybrid energy storage systems are managed provides an example of prolonging battery lifetime by minimizing the amplitude of power (and subsequently current) variations. For instance, batteries discharged or charged at a steady current instead of a sinusoidal or rapid step-change current have been shown to have longer lives than their counterparts. It should be noted that intermittent testing using current pulses will not cause significant degradation, but that consistent large amplitude current variations should be avoided.

In a work which paired a supercapacitor with a lithium-ion battery, the authors designed an optimization which minimized the amplitude and variation of the current requests to the battery [38]. Minimizing the current amplitude means the battery is not cycled as hard, which reduces the internal temperature of the battery. In turn, this reduces the thermal cycling of the battery. Mitigating the thermal cycling and temperatures of the battery can help prolong lifetime. Different reactions can occur in the battery as temperature increases. These "side-chain" reactions can cause permanent damage to the battery since many are irreversible once they have occurred. Some can even plate lithium out of the battery and cause internal short-circuits [39].

The authors of [38] capture this idea by creating a utility function for the battery as shown in (2.6). The battery lifetime, u_{batt}, is comprised of u_{ave}, the current amplitude, and u_{dif}, which is the current variation $\frac{dI_{batt}}{dt}$. The coefficients a and b are designed to normalize these equations to 0 when the maximum current level is reached, or the maximum current ramp rate is reached. Both of these values are user defined based on the power ripple limits placed on the system. A similar function is generated for the capacitor which aims to maximize the energy efficiency of the supercapacitor.

$$u_{batt} = w_{ave}u_{ave} + w_{dif}u_{dif} \; where:$$

$$u_{ave} = 1 - a(I_{bat} - I_{ave})^2 \tag{2.6}$$

$$u_{dif} = 1 - b(I_{bat_i} - I_{bat_{i-1}})^2$$

These utility functions can then be combined into a single objective function using a weighted-sum to balance the minimization of the system. The weighted-sum approach uses coefficients for each objective (u_{bat} and u_{cap}) which add up to 1 and are strictly between 0 and 1 to balance the goal of the optimization. For instance, if the aim of the system is to lower the

current amplitude, then the current variability must be allowed to be higher than usual. In (2.7) the objective function for the HESS is shown. The system is constrained to stay within the upper and lower current limits of the battery and supercapacitor. It is also constrained so that the currents between the supercapacitor, the battery, and the load must add to zero. This means if the system is pushing power to the load, but the supercapacitor needs to charge, then the battery must provide all the energy [38].

$$f_{min} = -w_{ave}u_{ave} - w_{dif}u_{dif} - w_{cap}u_{cap} \tag{2.7}$$

Essentially this optimization (along with other similar methods), attempts to extend the life of the battery in an ES by reducing the amount the system is cycled and the severity of the cycling regardless of the value to the battery owner. In fact, the authors of [40] developed an optimization which reduces the frequency of charging and discharging cycles and makes sure that each cycle is as shallow as possible. As the value-based optimizations demonstrated, this will reduce the overall revenue gained from the system. Barnes et al propose a system which attempts to marry these together into a planning tool that optimally selects battery chemistry and size for a future system by evaluating the net present value of a variety of proposed systems and their potential benefit over a predetermined planning period [41].

Net present value (NPV) measures the value of a project while accounting for the time value of money. By using NPV, an optimization can pursue different project options and create a plan which best suits the needs of the project (or constraints). The optimization can select different battery chemistries based on the cost, efficiency, and lifetime of the units as determined by calculated figures of merit for each chemistry (shown in Figure 2.8). NPV can also be used to determine the instantaneous value of the components of a system pro-rated over the system life.

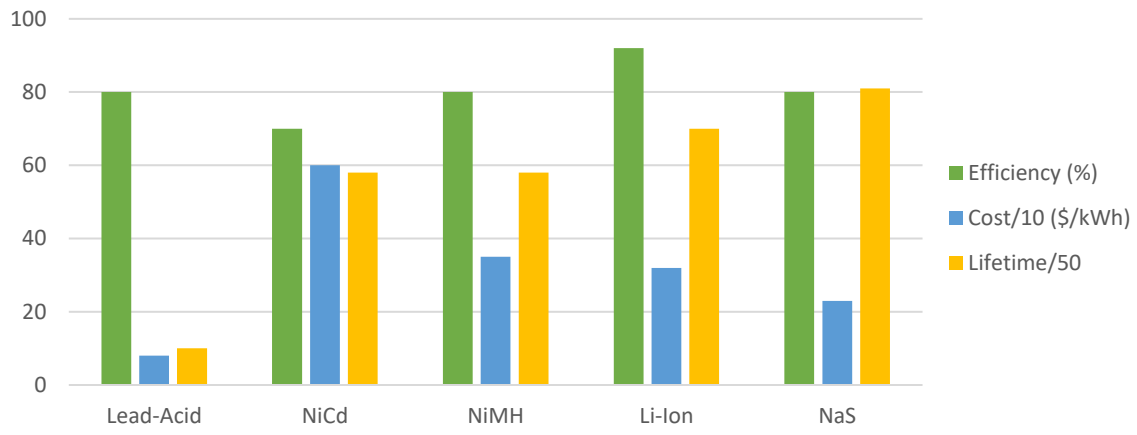


Figure 2.8 - Battery chemistry figures of merit normalized to fit on a scale of 0 to 100 [42].

By combining these with the constant-Ah model[2] for battery lifetime and comparing it to estimated time-of-use (TOU) scheduling over the lifetime of the system, the optimization can select the right battery chemistry and capacity to provide the maximum value to the selected system.

$$f = \left(c_{conv} + c_{siting} + \sum_{n=1}^{N} c_n' E_n^r\right) + r_1^{yearly} \sum_{y=1}^{Y} \frac{(1+\alpha)^{y-1}}{(1+i)^y} - c_{om} \sum_{y=1}^{Y} \frac{1}{(1+i)^y} \tag{2.8}$$

$$r_1^{yearly} = \sum_{k=1}^{K} d_k \sum_{m=1}^{M_k} (c_i + c_{k,m}^{elec}) \Delta t_{k,m} p_{k,m} \tag{2.9}$$

In the first term $\left(c_{conv} + c_{siting} + \sum_{n=1}^{N} c_n' E_n^r\right)$ of (2.8), the capital cost of the unit is estimated for each chemistry, *n*, and the rated capacity, E_n^r. The cost per unit energy is accounted for by c_n'. To account for the converter and siting costs, c_{conv} and c_{siting} are added to

[2] The constant-Ah model of a battery's life is a simplified measurement of how many cycles a battery can perform. Lifetime is measured in terms of Ah and is assumed to be roughly independent of depth of discharge.

this first term. The second term gives the annual revenue as a function of the revenue earned during the first year of operation. It is assumed this will increase over the duration of the project, *Y*, as energy costs rise. The first-year revenue is estimated using (2.9) to consider different pricing structures depending on the season. This term determines how much revenue the unit will generate over the stated project lifetime. Yearly operating and maintenance costs are assumed to decrease over the life of the project. Lifetime of the battery unit is accounted for in the constraints on the objective function. In addition to constraining the problem with round-trip efficiencies and the energy limits of the battery, the battery is monitored to ensure that as it cycles the capacity fades and it becomes unable to cycle as deeply [42]. This is accounted for in (2.10).

$$Y \sum_{k=1}^{K} d_k \sum_{m=1}^{M} \Delta t_{k,m} p_{n,k,m}^{d} \leq E_n^{f_0} E_n^{rated} \tag{2.10}$$

This novel method includes a basic model for battery life in the planning stages of the system to make sure that the battery can always provide the necessary resources over the life of the project. Instead of just extending the life of the battery, the authors give monetary value to the battery's life. This means that depending on the capital costs of the system, it could be more valuable to the system to run the battery very hard and degrade the cells quickly because more revenue can be extracted from the system. An online optimization which uses a similar valuation technique to give monetary value to an active battery's life has not yet been demonstrated in literature. Each of the three methods demonstrated above were designed to work for a single battery energy storage. Multi-element storage would need to be treated as a single element to utilitize these strategies. A single centralized optimization would be needed to collect

information about every ES element and optimize the system as one unit. Alternatively, transactive methods for optimization may be used to distribute the process.

Transactive Optimization using an Energy Market

ES optimization methods in the literature predominantly focus on the dispatch of a single unit connected to a distribution network or microgrid. Expanding these methods to optimally dispatch multiple ES units while having the unit act as a single ES would result in a complex, centralized optimization with intimate knowledge of every attached storage element. Scaling the system to include a large array of ES elements could result in an optimization which is cumbersome and does not solve very quickly. Transactive optimization techniques focus on distributing the system decision making such that individual elements bid into a central energy market with a preference for operation based upon a local optimization. System dispatch is determined by a central controller which aggregates the operating preferences from the individual units and optimally arranges these preferences to meet the external demands on the system. At a transmission and distribution level, transactive energy markets have defined grid operations for decades. Traditional grid power generation entities bid into an energy market with operating preferences scaled to an energy production cost. These bids are selected by system operators to provide power in the most inexpensive manner while also focusing on reliability.

This idea of an energy market has made its way to utilities as distributed energy resources (DERs) and intelligent loads have become more popular on distribution networks and microgrids since it alleviates the decision making load on a central controller while also limiting the amount of information that must be provided by prosumers[3] [43]. Optimally and autonomously

[3] Prosumers are owners of distributed generation and load assets capable of responding to utility dispatch signals.

dispatching these resources has been the focus of research into transactive market systems where prosumers can express their operating preferences by exchanging pricing data with a centralized controller [44]. Transactive systems have been described for multiple levels of distribution from neighborhood smart-device control, to microgrids, to networked microgrids [44] [45] [46]. This discussion will focus on two examples from the literature of transactive methods on a distribution network.

Coordinating PHEV Charging

A prime example of transactive energy control is coordinating charging between a fleet of PHEVs or EVs. Shifting or reducing peak energy demands by the chargers while meeting vehicle availability needs creates the basis of a problem widely studied since the large demands created by fleets of charging EVs can pose a threat to microgrid stability. Many authors have proposed using timeslots and priority algorithms to solve this problem. In Cramer et al, a market-based approach is used to continuously re-prioritize the charging needs of a network of PHEVs to reduce peak power consumption while also meeting the charging needs in the available window [47]. In their proposed method, each networked charger reports a "willingness to charge" to a central decision-maker in the form of a linear bid function. This function relates the preferred charging power level to the price that the charger is willing to pay to charge. Using (2.11), the bid is determined as a function of the time-to-departure, t_{dep}, and the energy needed to charge, ΔE. A caveat to this equation is that if t_{dep} is greater than 12 hours, then it is assumed to be 12 hours. This assumption can be made because the window of time for the optimization is 12 hours. Therefore, if the time to departure exceeds this window, it can just be assumed to be the length of the optimization window.

$$p_r = 1 - \left(\frac{1}{2} * \frac{\Delta t_{dep}}{12} + \frac{1}{2} * \frac{\Delta E}{E_{max}}\right) \quad (2.11)$$

Based on the present grid pricing and conditions, the central control sets a producer bid function where the willingness to deliver energy from the grid increases as the collective average price offered by the chargers increases. Essentially, this means that as chargers increase their price (through willingness to charge), the amount of power delivered by the grid increases.

The main goal of Craemer's work was to use a market based control method to reduce the necessary communications between the chargers and the central controller while using a traditional peak reduction technique and meeting the charging needs of the PHEVs. This showed promise in that the necessary data traffic was reduced by 80% while meeting the demands of the vehicles 95% of the time over a 24 hour period. It also demonstrates a simple, but effective transactive method for coordinating ES charging. Present transactive methods for EV charging lack the ability for the vehicles to push power back to the grid and earn revenue for providing grid services.

Distributed Transactive Optimization

Another example of transactive optimization comes from a study which aimed to increase internal micro-grid reliability and stability and decrease reliance on spinning reserve generation and the grid connection. The authors defined two types of resources in a micro-grid: flexible resources, such as ES, with controllable power outputs and noncontrollable resources, such as load or renewables, which produce or consume power based on a forecasted schedule. Using a 24 hour look-ahead window, the power production and consumption of non-controllable resources is forecasted using weather and behavioral models. During the window, if the noncontrollable resource is over or under producing, it contacts its neighboring resources to

either sell or buy energy to make up the difference to the forecasted value. Flexible, or elastic, neighbors bid on the opportunity to provide or absorb the excess energy.

In their proposed method, the authors fully distribute the optimization process by having each flexible resource perform an optimization which maximizes its personal welfare. Each non-controllable resource optimizes the price which it offers to other units to prevent the need to sell energy back to the grid or purchase energy from the grid (assuming this will provide the least-ideal market partner) using a neural network. A neural network is used so that a fixed resources can adjust their pricing based on past performance metrics. Meanwhile, flexible resources aim to maximize the revenue earned through energy exchange with other resources. Global constraints on the net power exchange and prices are determined by a master device monitoring the microgrid condition. These constraints are placed on all connected devices.

Having each resource optimize its own response reduces the need for large amounts of published data on the network. The only data passed between devices is price signals and power requests. However, the authors concluded that the real-time nature of the optimization required unrealistic time-sensitive networking to ensure that the system converged to a globally optimal condition. Without tracking when price signals and requests were being issued, the system could easily fall to non-optimal results [43].

Transactive optimization techniques provide an opportunity to reduce the computational workload associated with finding a globally optimal operating state for large system. However, the transactive methods for energy storage demonstrated in the literature largely focus on coordinating EV charging or on highly-complex micro-grid management strategies. For an autonomous multi-chemistry battery energy storage system, a simple transactive optimization is desired which separates optimal system dispatch from individual battery system optimization.

Using this method can provide a flexible and expandable system by distributing the optimization back to the devices, This helps de-couple the control of the energy storage elements from the utility control of the multi-element system.

Parameterization of Secondary-Use Batteries

In Chapter 1, the various battery chemistries used in electric vehicles and stationary energy storage devices were introduced. Secondary-use batteries were presented as a viable alternative to newly manufactured batteries due to the less restrictive energy and power density needs for stationary energy storage and a significantly lower cost factor. For this research, a plug-and-play interface has been proposed for integrating all types of batteries into one device. To safely and optimally manage the varying types of energy storage, the battery chemistry and capacity must be known. In almost all large-scale ES systems, both these values are already known since the unit is installed new with name-plate capacities. However, in a large system of multiple different chemistries and capacities, it would be prudent to automate these processes.

Battery Chemistry Identification

Battery identification aids in accurately measuring SOC and SOH, since battery charge and discharge curves depend heavily on the electrochemical features of each battery chemistry. The electrochemistry of the battery determines the relationship between the remaining energy content and the voltage of the battery. SOC algorithms heavily depend on the open-circuit voltage of the battery to determine a starting-point for charge-counting SOC calculations. Thus, understanding the battery chemistry in use is pivotal to estimating the battery state. Most importantly, a battery should not be charged without first understanding which chemistry is in use and the capacity and voltages of the overall pack. While battery chemistry can be inputted manually by a technician, a system to verify settings can ensure safety and decrease manual set-

up time. It can also be used to detect anomalous behavior that could indicate a bad battery pack. Previous work has focused on the design of chemistry identification algorithms for the charging of small capacity portable electronics.

A frequently used method for determining battery chemistry is pulse discharging of a system. The terminal voltage of a battery is monitored as it undergoes discharge current pulses and the voltage waveforms are recorded. As shown in the generalized sample in Figure 2.9, discharge pulsing the battery and then recording the voltage as it settles can help give a qualitative guess to battery chemistry. Voltage settling is the electrochemical reaction of the battery returning back to an equilibrium state after charging or discharging. Nickel-metal-hydride batteries exhibit fast voltage settling, while lead acid batteries settle over a much longer period, and Lithium-ion chemistries exhibit a settling time somewhere in the middle [48].

In their paper, "A Self Adaptable Intelligent Battery Charge for Portable Electronics," Abeyrante et al. show a method for quantitatively identifying a battery's chemistry using these

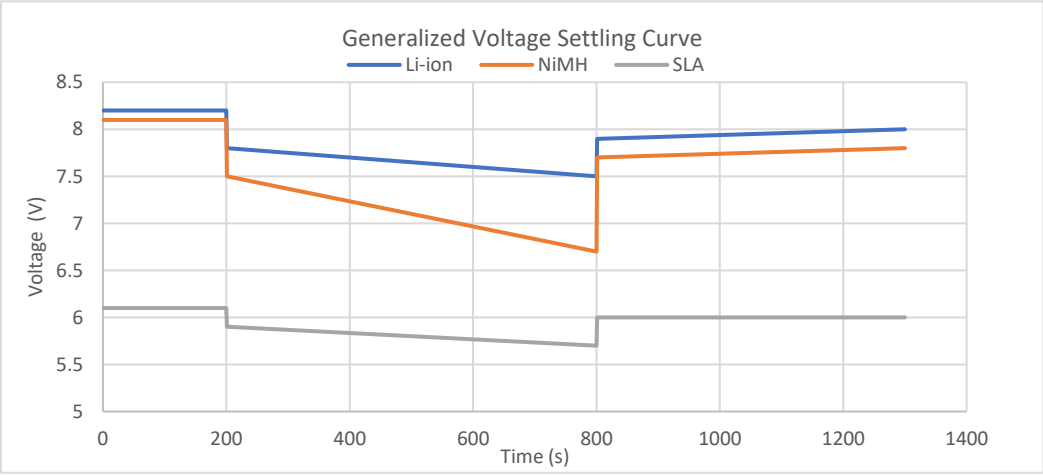


Figure 2.9 - Generalized sample of a pulse charge identification of battery chemistry.

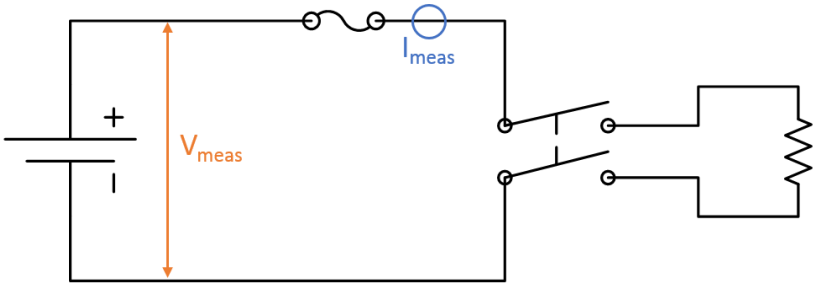


Figure 2.10 - A simplified circuit diagram of the experimental set-up.

features. In this work, a portable electronics charger is proposed which can differentiate between Li-Ion, NiMH, and SLA. Using a similar test set-up like the one shown in Figure 2.10, the authors performed a pulsed-discharge of the cells at a constant current and recorded the battery terminal voltage from a near-full charge to a near-empty charge. Results are like those show in Figure 2.11. As expected, the voltage of a battery under load drops compared to the open circuit voltage, V_{oc}. As the battery discharges, the V_{oc} and V_{loaded} voltages drop in tandem. However, as the battery nears depletion, the voltages begin to diverge. The speed and magnitude of this divergence is independent of many battery parameters, but fully related to battery chemistry [49].

The pulse discharge is performed by turning a relay on and off. As indicated in Figure 2.12, the relay is turned off at V_a, disengaging the constant current load, and turned back on at V_b, placing a load on the battery. By measuring the voltage at these points, a voltage differential can be determined through simple subtraction of V_a from V_b. When this is performed for every discharge pulse, a graph of differential voltage versus depth of discharge can be produced as shown in Figure 2.13.

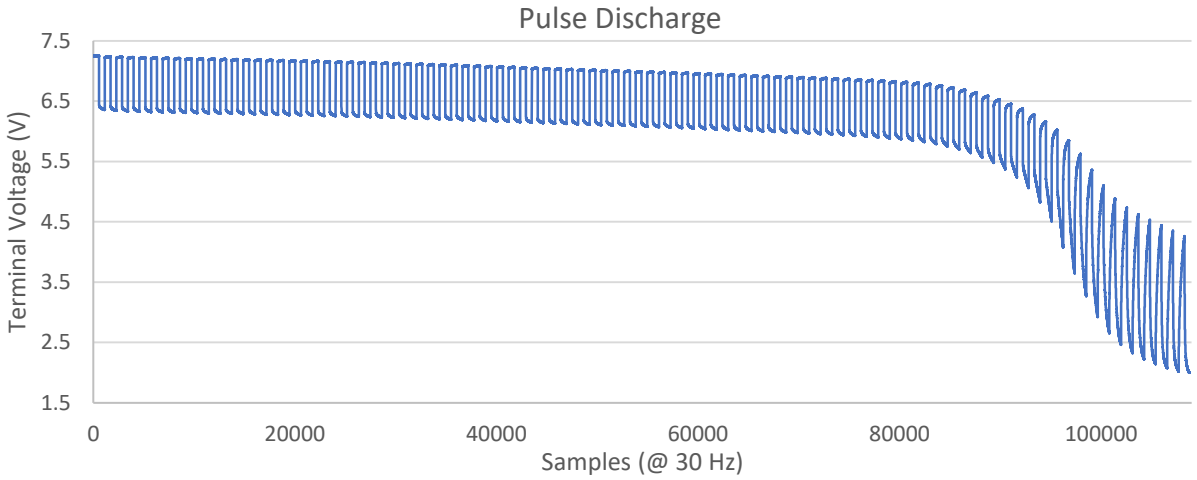


Figure 2.11 - An example pulse discharge test showing the change in V_{oc} and V_{loaded} with time. Data is from completed work.

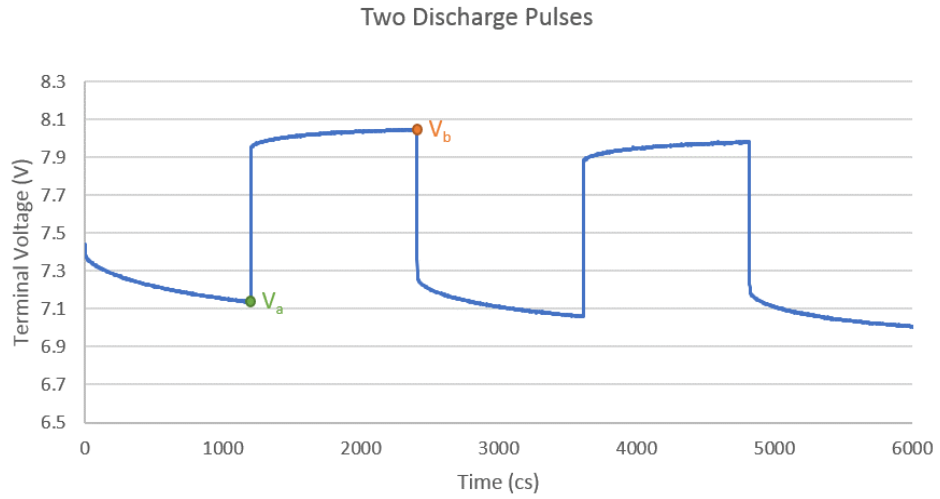


Figure 2.12 - A close-up of two open-circuit pulses. V_a represents the time when the relay is opened and the load disconnected. V_b represents the point when the relay is closed and the load connected

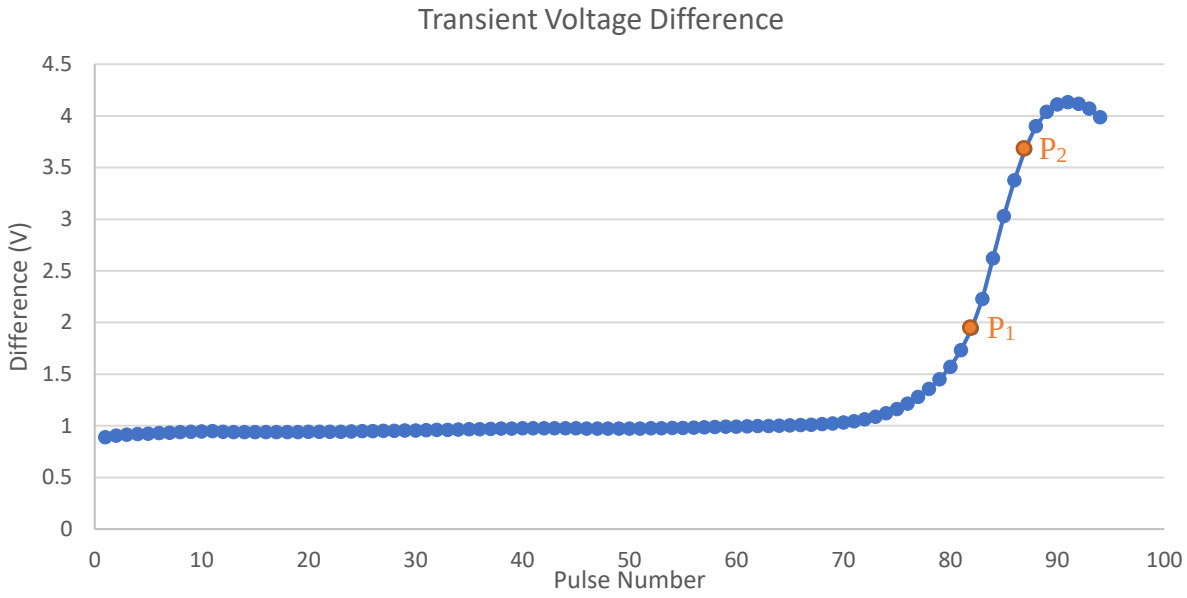


Figure 2.13 - a graph of transient voltage difference measured from a test. P1 and P2 represent the linear area of the curve that can be used to measure the gradient.

The authors note that the section of the curve marked P_1-P_2 in the figure is approximately linear. This gradient, which was dubbed G_{batt}, is unique for each battery chemistry regardless of the terminal voltage of the battery or the battery capacity. The gradient is calculated using (2.12) where the percent discharge is determined by estimating the energy that has left the battery versus the total capacity of the battery. It should be noted that all of this is done at a constant current. A flaw in this method is that it assumes previous knowledge of the total capacity of the battery, as indicated by the inclusion of the Q_{tot} term in the equation below [49].

$$G_{batt} = \frac{(V_1 - V_2)}{(Q_1\% - Q_2\%)} \quad where\ Q_i\% = \left[\frac{I_{DC} \times \Delta t}{Q_{tot}}\right] \tag{2.12}$$

However, G_{batt} is a simple slope between two points as demonstrated in (2.13). By making the gradient independent of capacity, it can be measured while data is being taken. Additionally, it means that no knowledge of the battery needs to be known before taking the measurement [50].

$$G = \frac{dV}{dt} = \frac{1}{2dt}(V_{i+1} - V_{i-1}) \quad (2.13)$$

In a paper containing continued work, the same authors demonstrate clearly that the gradients for each battery chemistry are unique enough that they may be sorted into bands for identification by a simple algorithm on an embedded controller. Table 2.1 shows a series of example tests run by these authors. Two problems exist with this method. First, the algorithm can deplete the battery below conventionally safe levels if an automatic shutoff system is not included. Second, previous work in the literature demonstrating this method has only done so on small-capacity, low-voltage cells intended for use in portable electronics.

Battery Capacity Measurement

Knowing the capacity of a battery system is critical to its operation. Battery capacity is the amount of energy that the system can store when charged to the top of the cell voltage range (Vmax). Capacity (C) is used to calculate SOC and can be used to determine safe discharge rates for a battery. Capacity can be easily measured through a timed discharge at a constant current. Even at a non-constant discharge rate, the integral of a current curve vs time can be found to give a capacity in amp-hours (Ah). This method is recommended for use by the US Advanced Battery Consortium (USABC) in their Battery Test Manual for Hybrid Electric Vehicles [51].

Table 2.1 - Results showing various battery capacities and terminal voltages resulting in very similar gradients for single chemistry types. Results from original authors of [50]

Battery Chemistry	Terminal Voltage (V)	Capacity (Ah)	Idc (mA)	Gbatt (V)
Lead Acid	6	4.50	600	1.10
Lead Acid	6	4.50	200	1.00
Lead Acid	6	7.00	600	0.80
NiMH	4.8	0.90	200	2.60
NiMH	3.6	0.65	200	2.25
Li-Ion	3.7	1.05	200	0.19
Li-Ion	3.6	0.85	200	0.23

Perceived battery capacity can change with the rate of discharge. A battery under heavy discharge for a short period of time may appear to have less energy capacity than one that is lightly discharged for a long period of time. This is due to the voltage drop phenomenon previously discussed. For this reason, the USABC test manual recommends performing a static discharge test at a discharge rate of C/3 from V_{max} to V_{min}. This rate is defined as the "current corresponding to the manufacturer's rated capacity for a three-hour discharge." Using this test, USABC has standardized the measurement of capacity across electric vehicles. For instance, if the battery's rated capacity is 60 Ah, the test would statically discharge the battery from its maximum operating voltage to its minimum operating voltage at a rate of 20 A. Then, the capacity could be measured by (2.14).

$$Capacity\ (Ah) = I_{DC} * \frac{\Delta t}{3600} \tag{2.14}$$

However, given an unknown battery of unknown capacity, this method should not be used since the nameplate capacity (C) of the battery is unknown. Discharge rates that are too high could damage the battery. Rates that are too low will take much longer and provide an unrealistic estimation of capacity. To solve this issue, the authors of [49], who developed the chemistry identification technique, proposed using the results of the pulsed discharge test. Based on many experiments, the authors estimated that the linearized gradient typically represents approximately 5% of the total charge of the battery. Using (2.14) from above, this can be expressed as shown in (2.15). Therefore, the capacity of the battery can be found by re-arranging the equation to look like (2.16). This process estimates the amount of energy discharged by the battery for a known period and assumes that it comprises 5% of the total energy contained within the battery. Extrapolating from there, the total capacity of the battery

can be found. This method can easily provide a false reading of capacity since it is extrapolating the total capacity from a relatively small window.

$$\Delta Q\% = 5 = \left(\frac{I_{DC} \times \Delta t}{Q_{tot}}\right) \times 100 \tag{2.15}$$

$$Q_{tot} = \left(\frac{I_{DC} \times \Delta t}{5}\right) \times 100 = (I_{DC} \times \Delta t) \times 20 \tag{2.16}$$

Other methods of "blind" capacity measurement have been proposed which similarly pulse-discharge or pulse-charge the battery and use complex learning algorithms or data-fitting techniques to guess the capacity from a short window of battery data [52] [53]. However, these are meant for in-situ measurements to determine the current SOC of a battery system. Since the USABC manual already has a static-discharge test, and the pulse discharge test is done over the entire discharge range of the battery, it seems that using the Coulomb counting method to estimate and then verify the capacity of a battery would be more reliable. Yet, a method such as this has yet to make an appearance in the literature.

State of Health Estimation

With the battery chemistry and capacity verified, the next step is determining the State of Health of the battery. This is valuable information since it can be used to provide more accurate SOC readings and will be used to prolong the battery's life through optimal cycling management. State of Health (SOH) is a nebulous term when discussing EV batteries. Two degradation parameters are most often used to quantify the health of the battery; the capacity of the battery and the power capability. Battery capacity measurement techniques were discussed in detail above. Power capability is measured by characterizing the electrical properties of a battery. Most often, the internal impedance is cited as the parameter most heavily influencing the power capability. Determining the internal impedance can be done in many ways as

illustrated in Figure 2.14. What follows is a brief discussion of methods. Although this project will not be proposing a novel health measurement, it is worth understanding how health is measured.

The top group of methods (electrochemical impedance spectroscopy) is a complex laboratory measurement technique which provides highly sensitive models of the battery parameters. So sensitive, in fact, that they can be affected by the environment the battery is tested in such as proximity to magnetic objects, temperature, etc. Active EIS uses an active signal generator to excite the battery with small amounts of sinusoidal currents at different frequencies. By recording the resulting voltage response, the impedance can be calculated at

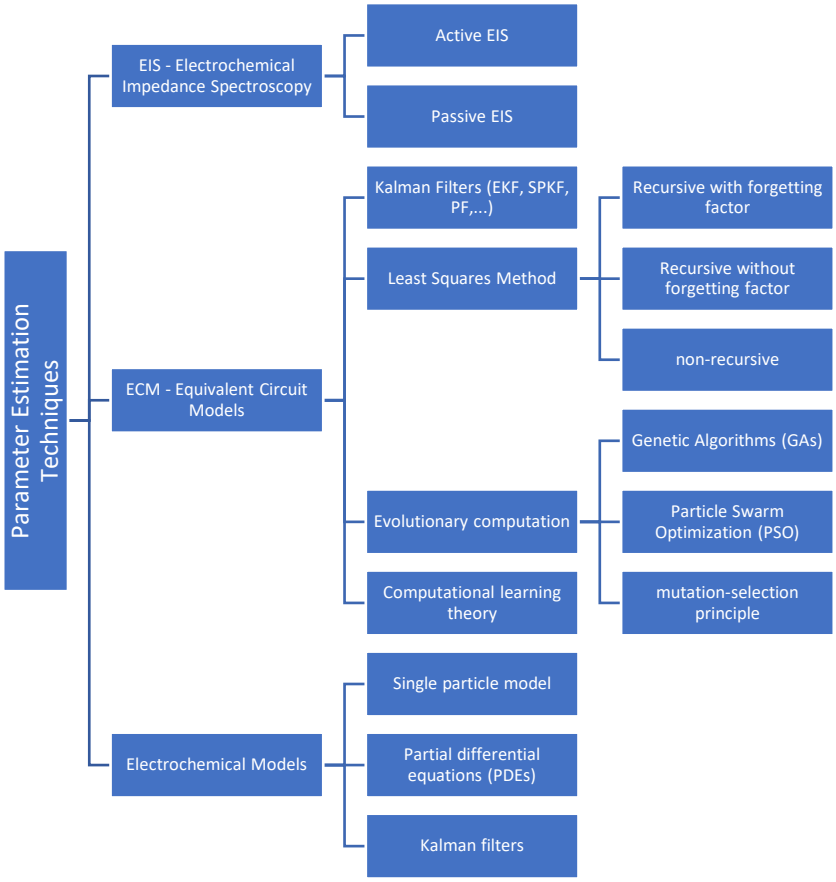


Figure 2.14 - Battery parameter estimation techniques for SOH measurement [54].

each frequency using a Fourier transform. From this, a Bode plot of the battery can then be generated. Passive EIS is performed similarly but uses the incident fluctuations of current caused by loading the battery to provide the frequencies instead of actively generating frequencies and pushing them through the battery. Both techniques require careful measurements that are tightly time-correlated and result in measurements that are closely linked to temperature, SOC, and SOH.

The middle group uses various algorithms to tune an equivalent circuit model (ECM) to match data being measured from the actual battery. This method is favored by EV manufacturers for parameterization because it can be executed on a finished battery pack without special equipment. Using an automatic battery cycler or a large bidirectional DC supply, the battery is charged and discharged for quick pulses for *t* intervals of 5-10% of the SOC of the battery. USABC recommends performing the pulses every 10% in what they call the Hybrid-Power Pulse Characterization test (HPPC) [51]. The voltage response of the battery as well as current values are recorded and used to curve-fit a multi-node ECM to the battery response. Figure 2.15 shows a typical seven-node ECM used for parameterization. Various techniques can be used to perform this curve-fitting but Kalman filters are preferred for Li-ion chemistries while

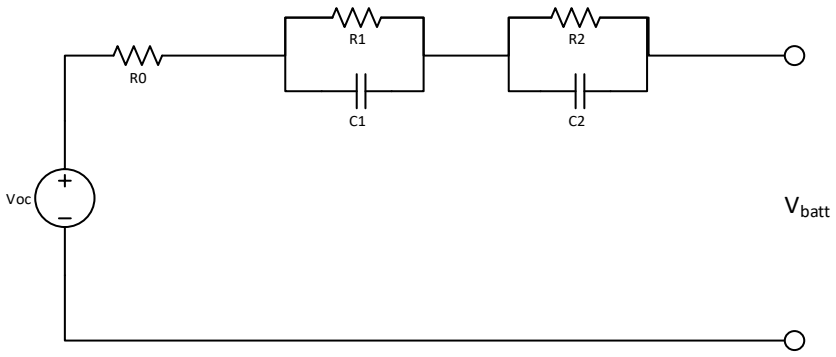


Figure 2.15 - A seven-node equivalent circuit model for state of health determination.

easier-to-use Recursive Least Squares (RLS) fitting is more often used on NiMH batteries. Kalman filters tend to provide more accurate battery models, however they are computationally more burdensome than the RLS method. Other works have proposed using genetic algorithms or machine learning to refine this process, but industry has not picked up on these methods yet [54]. Populating the ECM results in the internal impedance of the battery being well-defined.

The last group operates very similar to the ECM, but instead of populating an electrical model of the battery, data is used to populate an electro-chemical model of the battery. Like the middle methods, linear or sigma-point Kalman filters may be used to estimate the internal battery energy, conductivity of electrolyte, and SoC. These parameters are then used to draw conclusions about the internal resistance of the battery. Literature is lacking when it comes to using this method in an online battery system. Additionally, the electrochemical models are computationally more intense to run than the simple curve fitting algorithms above, making them less appealing for deployment on a battery system [55]. For this project, an ECM will be used in accordance with the recommendations from the Battery Test Manual for Hybrid Electric Vehicles. The test procedure for a HPPC has been well defined and proven to give stable results for EV batteries.

Chapter Summary

Above, the literature was reviewed for battery energy storage hardware architectures, health and life optimizations, and battery parameterization techniques, which are all critical components to developing an autonomous multi-chemistry battery energy storage system. This chapter demonstrated the lack of several key enabling technologies in the literature. First, several viable architectures have been demonstrated. Yet none of these systems incorporate multiple battery types of different ages into a single unit. Next, a review of optimization

literature showed that previous optimization objectives focus on either maximizing battery service or prolonging battery life. Currently, no methods assign a value to the life of the battery and optimize to achieve a maximum profit regardless of the impact on battery lifetime. This is critical since the most valuable use of a battery may not be to extend its life indefinitely. Additionally, transactive methods for optimization were explored as an alternative to the centralized single-ES methods commonly found in the literature. Finally, chemistry identification and parameterization techniques were reviewed. Many methods and techniques were found, yet none of them focused on automating the process. Instead, most techniques were focused on taking measurements for analytical studies of batteries or for future implementation into portable electronics. Missing from the literature was an autonomous method for chemistry identification and parameterization. In the next chapter, an overview will be given of an autonomous multi-chemistry battery energy storage system which fills in the missing components from literature.

Chapter 3 - Autonomous Battery Parameterization

Battery chemistry identification and parameterization is a feature for an advanced system containing mixed battery chemistries. Understanding the chemistry of an energy storage component helps parameterize the current and voltage characteristics. In Chapter 2, the processes for chemistry identification, capacity determination, and impedance measurements demonstrated in the literature were described. However, previous work in this area falls short of autonomous identification on an unknown high capacity, multi-cell battery. Specifically, the autonomous method developed and demonstrated in this work affirms that pulse testing can be scaled from small cells, to strings of cells in series and parallel.

This chapter describes the automation process developed to determine the chemistry and capacity of an unknown battery pack. The goal of this research is to create an autonomous program that tests an unknown battery and produces reliable chemistry identification and capacity measurements. Results are presented which demonstrate the ability of the system to accurately parameterize the battery given a fully executed test. A discussion of the results including errors will close the chapter.

Method

Building on methods developed by the USABC and Ayoub et al., the autonomous parametrization process shown in Figure 3.1 was developed and demonstrated to work on high-capacity and multi-cell battery systems [49]. To begin, the battery is charged to approximately 100% SOC by monitoring the terminal voltage. This is followed by executing a test to determine the pulse length and current (depth) necessary for a fast, but safe chemistry identification. This is followed by running the chemistry identification test, which produces pulse data that is then

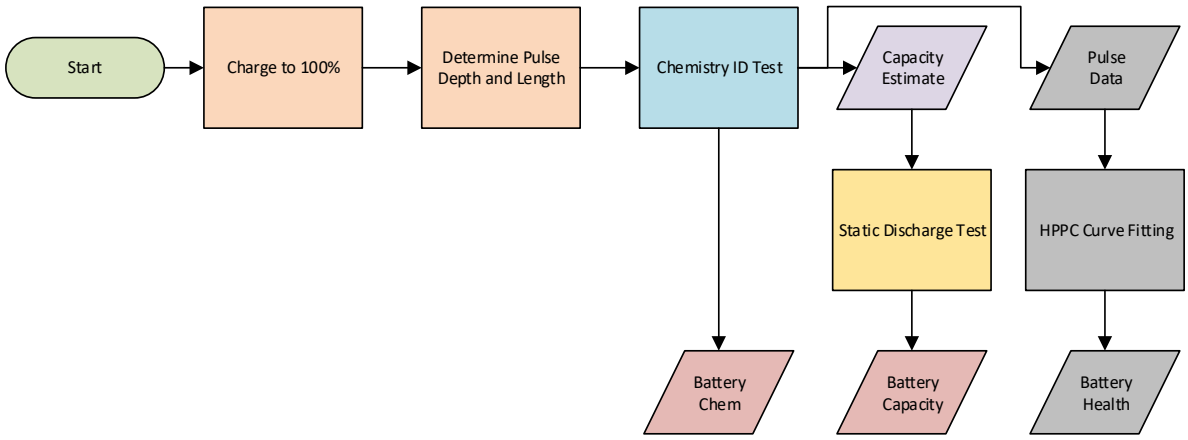


Figure 3.1 – High-Level process to determine battery parameters autonomously.

used to determine the battery chemistry and an estimated capacity. Pulse data can also be used to perform hybrid power pulse characterization (HPPC) curve fitting as described by USABC to estimate internal resistance and parasitic battery losses. For this work, the HPPC data are saved but not used to determine battery health. Instead, battery chemistry and capacity were determined to be enough to perform a system-level optimization. The decision to omit this data will be further discussed in the conclusions section as well as Chapter 4 when the architecture of the distributed optimization is presented. This section describes each of these test sub-components in further detail.

Charging to 100% State of Charge

In Figure 3.2, a flow diagram is presented which shows the program logic used to determine when the battery under test has achieved a near 100% charge. Since the battery chemistry is as-yet unknown, a 10 A charging current is used since this value is much less than the maximum allowable charging current for most EV battery systems [52]. Terminal voltage is measured at five second intervals and two timesteps are compared. When the voltage difference

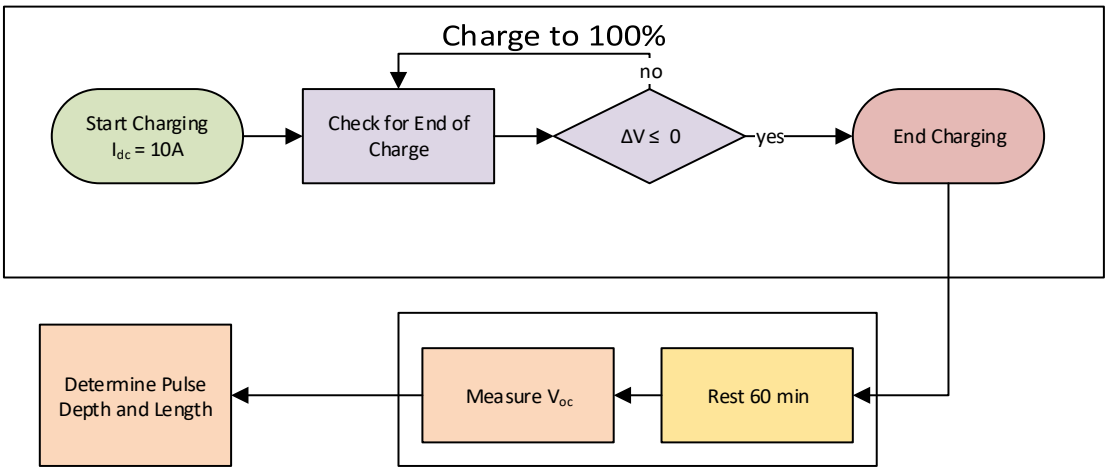


Figure 3.2 – Process for determining if a battery has reached full charge based on voltage.

between two timesteps becomes zero or negative, charging is stopped, and the system rests for 60 minutes. This rest time is used to allow any thermal effects from charging to dissipate and allow the open-circuit voltage of the battery to settle. Lithium-ion batteries do not require this rest time other than to prolong life, but NiMH cells will demonstrate widely varying capacities if not rested post charging. Once this period has elapsed, the open-circuit voltage (V_{oc}) is measured and is passed to the next process to determine the pulse depth and length. The voltage of the highest cell is also noted at this point.

Determining Pulse Depth and Length

As the battery is discharged under a constant current load, the open-circuit voltage falls at a slower rate than the loaded voltage (V_{load}). Mapping these voltage differences across the discharge range of a battery produces a unique voltage gradient per chemistry that is independent of capacity, terminal voltage, temperature, and discharge rate. Breaking down the voltage gradient graph into component parts is important to understanding how the pulse depth and length can be determined autonomously. Figure 3.3 shows an example gradient. First, the

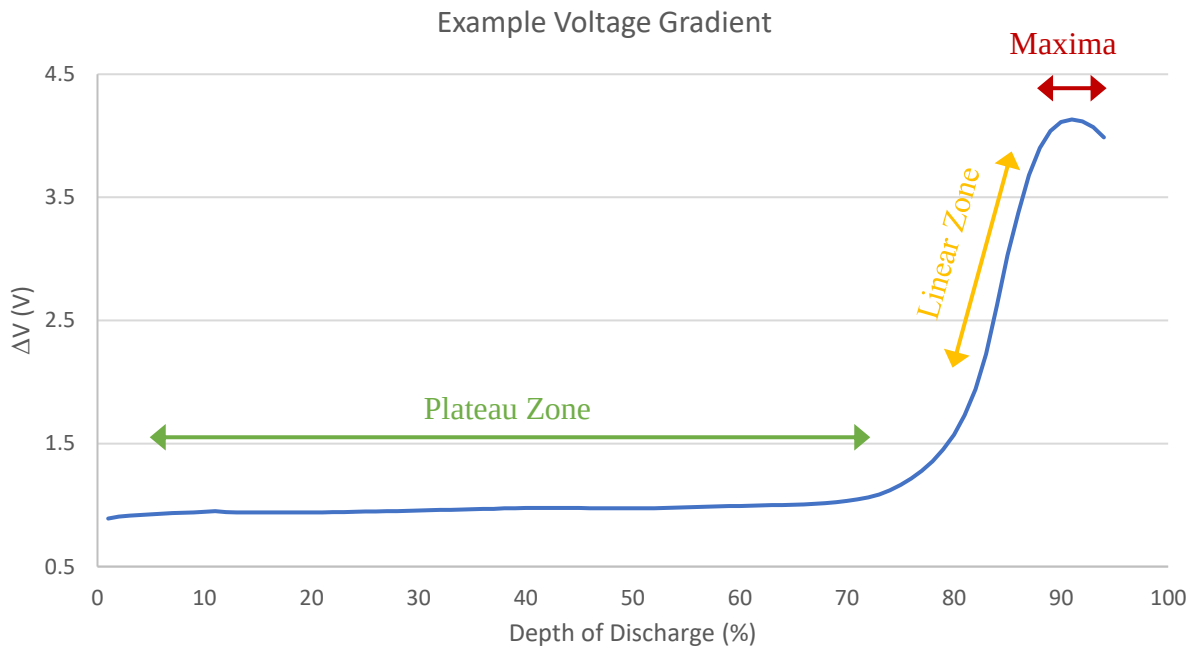


Figure 3.3 - An example voltage gradient with the three zones of the process.

"plateau" period of the graph represents the normal, near linear discharge of the battery where the V_{oc} and V_{load} fall at nearly the same rate. As the battery nears empty, the gradient rises into the linear zone where the V_{load} decreases more quickly than V_{oc}. During the peak, the battery is nearly depleted the gradient reaches a local maximum and begins to decrease. This indicates that one or more cells in the string is fully depleted and testing should stop.

For this identification method to work, the current pulses used to determine the battery chemistry must have consistent length and be discharged at a constant current, but since the discharge current is independent, it can be increased to run the pulse discharge test more quickly. The higher the discharge current, the more energy will be discharged with each pulse and the greater the different in terminal voltage at the end of each unloaded pulse. From measuring the ΔV values across batteries from a variety of OEMs at a variety of pulse depths, a convenient rule

of thumb was derived qualitatively. The quickest test results came from setting the pulse depth to a current such that the difference between the V_{oc} and V_{load} value (known as ΔV_g) was approximately 0.1% of the pack V_{oc} measured while the battery was full charged. Setting the current lower than this extended the test length unnecessarily. Using a current that gave a larger ΔV often meant that finding a linear section of the gradient was difficult because of a lack of data points off which to base the calculation.

The value of 0.1% was found by evaluating nine early data runs on lithium-ion Leaf cells and noting the peak difference during the plateau part of the voltage gradient graph. Three data sets gave between 3 and 5 linear data points on the chemistry identification part of the gradient curve. Together, the ΔV during the plateau period for each of these runs averaged 0.12% of the V_{oc} for the attached cells. This was reduced to 0.1% to ensure that a minimum of 3 data points are found in the linear period. The procedure shown in Figure 3.4 was then developed to find the pulse depth and breadth to execute the fastest test for an unknown battery.

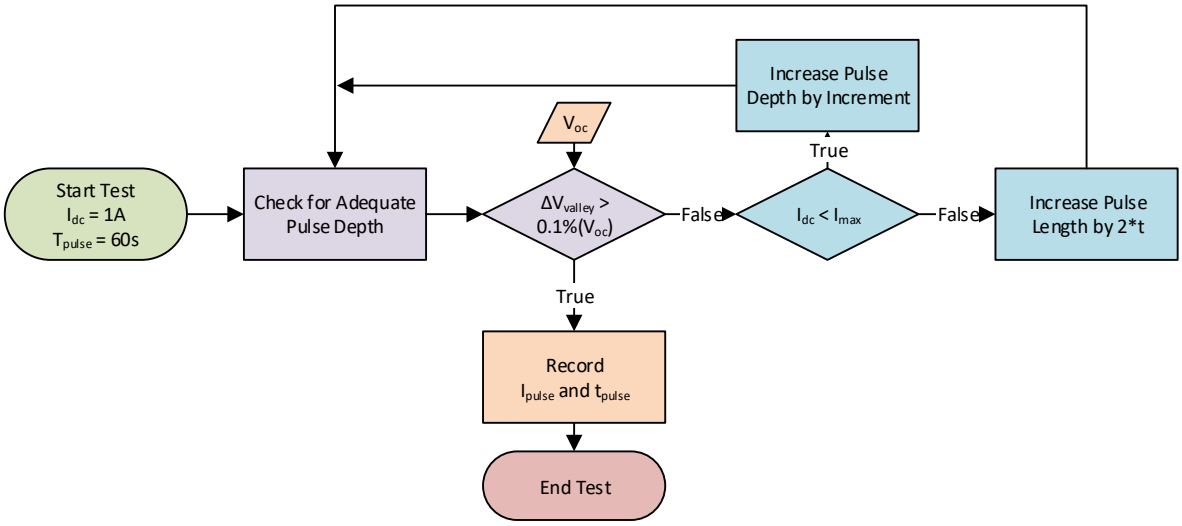


Figure 3.4 - Iterative method to determine depth and length of the chemistry ID pulses

To begin, the pulse is set to a low current and short duration. A single pulse cycle is executed and the ΔV_g analyzed versus 0.1% of V_{oc}. If this condition has not been met, the pulse depth is increased in 1 A increments and re-executed given that it is under a pre-specified current safety limit. While this safety limit can be arbitrary, 20A was chosen for this algorithm because this value would not be overly taxing on most present vehicle batteries. Additionally, it was within the agreed-upon safety limitations of the experimental set-up described in the next section. Should the current limit be reached before the voltage difference satisfies the depth requirement, the pulse length is doubled, and the test re-executed. Once the pulse length and depth have been determined, this information is recorded and used to then execute the pulse discharge test.

Pulse Discharge Parameterization Test

Previous research in the literature executed the pulse discharge test over the entire discharge range of the battery. Ending the test was signaled by the battery no longer providing the constant current required. However, as Figure 3.3 shows, the linear zone of the voltage difference curve used to identify the battery occurs before the graph of ΔV experiences a local maximum; which is a salient feature that can be used to terminate the pulse test. To shorten the test time and prevent discharging the battery to nearly 0% SOC, a differential monitor was set up using the method shown in Figure 3.5.

The test begins with a charge-to-full using the method described earlier to ensure that the process of determining pulse depth and breadth has not brought the battery too close to the linear zone. Pulse discharge testing begins after the battery has rested and uses the parameters determined in the previous step. For each pulse step, the voltage gradient, V_g, is calculated using (3.1) with the completion of each cycle from V_{oc} to V_{load}.

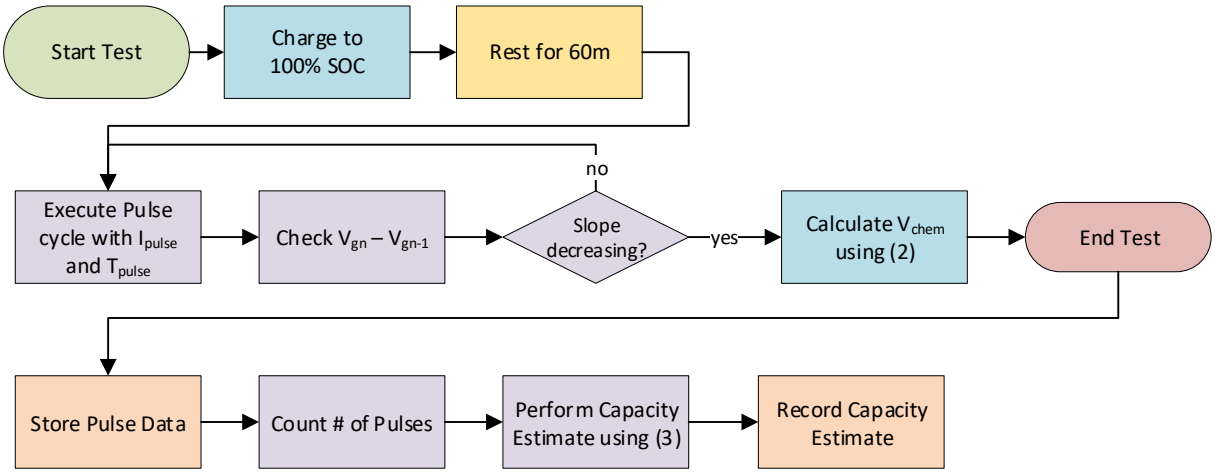


Figure 3.5 - Flow chart outlining the process used to execute the pulse test autonomously.

$$V_g = V_{oc} - V_{load} \tag{3.1}$$

Testing ceases when the measured V_g stops increasing and begins to decrease. This is detected by taking the difference between the current V_g and previous V_g. This difference will keep increasing until the end of the linear zone. To verify that the data is nearing a true maximum, three consecutive differences must indicate that V_g is decreasing within a +/- 0.2 V margin of error to cease the test. These steps were taken to ensure that noise in the data did not cause a test to prematurely stop. The slope of the linear zone is then calculated using (3.2) where *n* is the pulse count of the final cycle which was executed. This calculation disposes of the three consecutive decreasing measurements and produces a gradient based on the last three measured points, which are assumed to be nearly linear. This assumption was demonstrated to be valid by comparing the results of data sets that were analyzed by hand versus identification performed autonomously and getting proper battery identifications.

$$V_{chem} = \frac{V_{g_{n-3}} - V_{g_{n-5}}}{2} \quad (3.2)$$

From here, the chemistry is determined via a look-up table. Bands for gradient identification were determined qualitatively based on testing various OEM batteries in many different series and parallel configurations. These values were also compared to those found by Ayoub et al in the literature and found to be consistent even from single, low capacity cells to large multi-cell arrangements [49]. Based on the consistency of the bands, a look-up table is deemed the best option for chemistry ID. If a battery's gradient falls outside of these ranges or the data is otherwise imperfect, the battery is flagged for manual inspection and re-testing. Pulse data is then passed to the static discharge test to obtain an estimate of battery capacity.

Static Discharge Test for Capacity

In Chapter 2, USABC's static capacity test is described, providing a method for verifying the nameplate capacity of a known battery. Using an unknown battery with this test would be difficult since the discharge rate for the test is dependent on the capacity of the battery; specifically 1/3 of the nameplate capacity. To automate this process and allow for an unknown battery to be tested, the capacity is first estimated by measuring the energy discharged during the pulse-discharge test. The capacity estimate is found using (3.3), which counts the number of loaded pulses executed during the pulse discharge test and multiplies by the pulse length in seconds to determine an "on" time. The total time is then multiplied by the current to get an estimate of system capacity in ampere-seconds (As). Finally, the value is divided by 3600 to convert the value to the standard capacity unit of ampere-hours (Ah).

$$Q_{est} = \frac{N_{load} * T_{pulse} * I_{pulse}}{3600} \quad (3.3)$$

Since the battery was left discharged at the end of the previous test, the pack and minimum cell voltages are noted and the system is then charged back to full using the method from Figure 3.2. The the battery is discharged to the minimum pack voltage noted from the previous test at a constant rate of 1/3 * Q_{est} per the USABC Battery Manual method. Like in (3.3), the capacity is calculated by multiplying the time it took to discharge the entire battery by the rate at which the discharge was performed. This is shown in (3.4). To verify the accuracy of the measurement, the battery is cycled from full to empty two more times, with each of the measured capacities logged. If these capacities are within +/- 5% of one another, then the measurements are averaged and presented as the system capacity. If not, the system is flagged for further testing.

$$Q = \frac{T_{discharge} * I_{discharge}}{3600} \qquad (3.4)$$

In order to verify that the methods described above produce the desired results, a battery parameterization test stand was designed and constructed. The following section provides a description of the test setup used to validate the autonomous parameterization techniques.

Experimental Set-Up

To validate the method above, a battery parameterization test-stand was constructed. The test-stand uses a DC power supply for charging, a current-controlled load for discharging, and a data acquisition (DAQ) system to take and log the data. Figure 3.6 shows an electrical schematic for the test-stand. The battery under test can be electrically isolated from both the charging circuit and the discharging circuit by keeping both switch sets open. Fuses in the circuit are to protect both the battery and the programmable devices from any over-current situations from

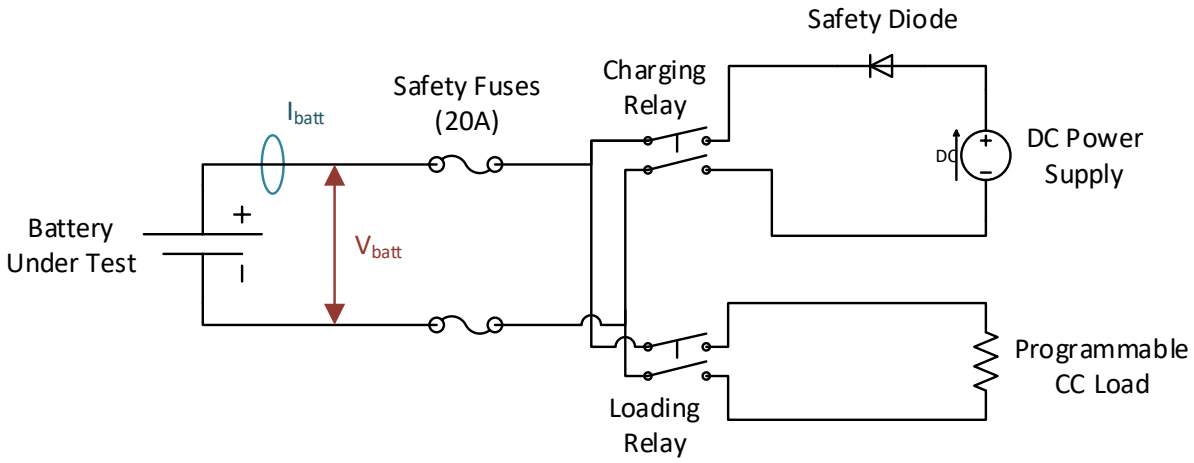


Figure 3.6 - Circuit diagram of power components in the parameterization test-stand.

faulty wiring or a backwards connection. The safety diode on the charging circuit prevents the battery from accidentally injecting current into the supply. The DC voltage setpoint on the power supply was set based on the measured terminal voltage.

Figure 3.7 shows a picture of the constructed system mounted into a 19" rack for ease of access to components. Omron MY2N DPDT relays were used as the charging and loading relays. These were chosen since they were rated to handle up to 60Vdc and 20A. Current sensing was performed by LEM HX 20-P sensors, which was selected because they have a selectable measurement range which could be set to +/- 20A. Voltage measurements were taken directly by a National Instruments (NI) 9220 synchronous analog input module as part of the NI cRIO 9068 that was used for data acquisition and relay control. While the NI 9220 can sample data at a minimum of 2000 Hz, the necessary data resolution was only needed as sub-second and data was down sampled to10 Hz. The voltage produced by the LEM HX 20-P was also sampled by this module.

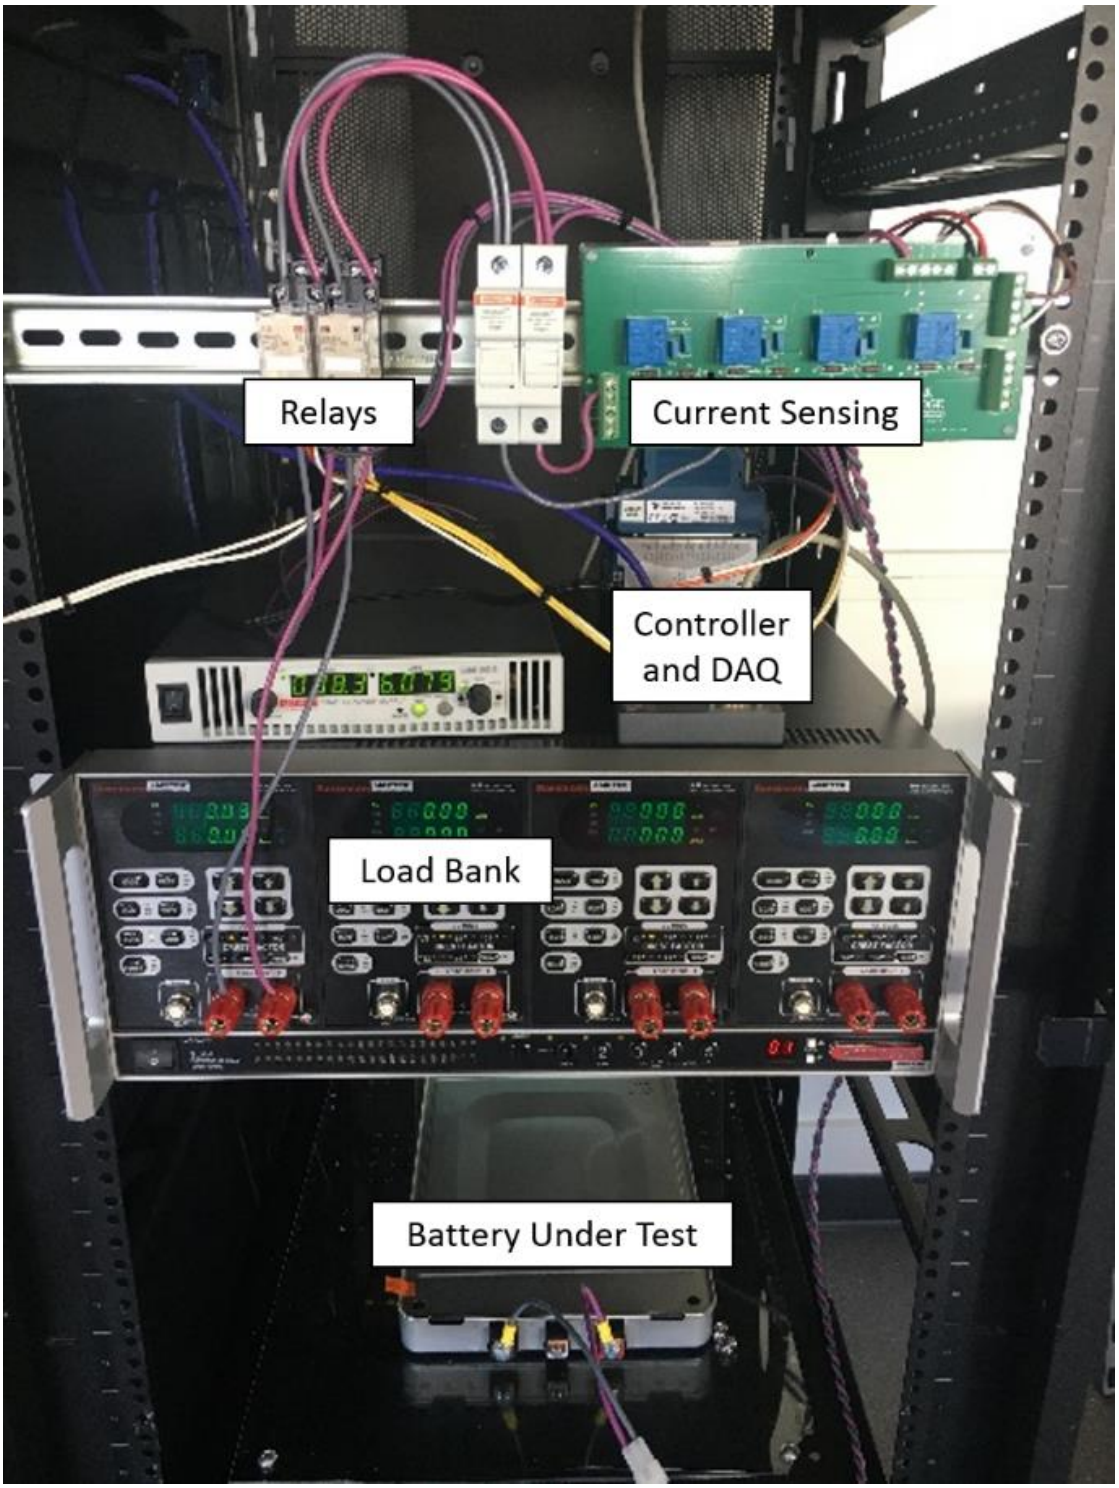


Figure 3.7 - The testing apparatus used to execute the pulse-discharge test on the battery.

Calibration of the measurement devices was performed by replacing the battery under test with a DC power supply and closing the loading relay. Voltage measurements were confirmed using a calibrated multimeter, and the NI module needed no scaling. Calibration of the current sensor was performed by setting the programmable load to a known value between 0 and 10 A in 2 A increments. The measurement was confirmed on a calibrated multimeter. The calculated scale and offset were then used to turn the signal measured by the cRIO to a current measurement accurate to within +/- 0.2 A.

Table 3.1 shows the battery testing configurations that were used to verify the autonomous parametrization process described above in the Methods section. Two chemistries of Li-ion were used from a first generation Nissan Leaf and a BMW i5. Nickle-metal hydride batteries from a second generation Toyota Prius were also tested alongside a handful of various SLA batteries. Nissan Leaf batteries were chosen due to their abundance in the secondary-use space. Despite the BMW i5 batteries being less prevalent in the market, they were included in tests since the internal battery format differs significantly from that of a Nissan Leaf. Fiat 500e batteries were available for testing but could not be tested beyond a single module due to laboratory regulations limiting the available energy and the physical configuration of the module.

As a control, each battery configuration was tested manually with data analysis occurring after the test, and then it was tested again using the autonomous procedure. The values for manual parameterization versus autonomous were then compared to each other and to the nameplate values given by the battery supplier. In the next section, results for these tests are presented and analyzed in depth.

Table 3.1 - A list of battery configurations and nameplate capacities run to test the autonomous parameterization techniques presented in this research.

Vehicle Model	Chemistry	Configuration (Modules)	Min. Voltage (V)	Max Voltage (V)	Capacity (Ah)
Nissan Leaf	Li-Ion (Set 1)	Single[+]	5.6	8.3	38.9
		2 Series	11.2	16.6	38.9
		2 Parallel[+]	5.6	8.3	77.8
		2P2S	11.2	16.6	77.8
		4 Series	22.4	33.2	38.9
	Li-Ion (Set 2)	Single	5.6	8.3	37.6
		2 Series	11.2	16.6	37.6
		2 Parallel	5.6	8.3	75.2
		2P2S	11.2	16.6	75.2
BMW i5	Li-Ion	Single	5.6	8.4	99*
		2 Series	11.2	16.6	99*
		2 Parallel	5.6	8.4	198*
Fiat 500e	Li-Ion (Set 1)	Single Module (5 Series)	14.5	20.75	60*
	Li-Ion (Set 2)	Single Module (5 Series)	14.5	20.75	60*
Toyota Prius	NiMH (Set 1)	Single[+]	3.5	7.2	6.8
		2 Series	7	14.4	6.8
		2 Parallel[+]	3.5	7.2	13.6
	NiMH (Set 2)	Single	3.5	7.2	5.6
		2 Series	7	14.4	5.6
		2 Parallel	3.5	7.2	11.2
	NiMH (Set 3)	Single	3.5	7.2	0.8
		2 Series	7	14.4	0.8
		2 Parallel	3.5	7.2	1.6
		4 Series	14	28.8	0.8
		4 Parallel	3.5	7.2	3.2
N/A	Lead-Acid (Set 1)	Single[+]	11.9	12.7	12
	Lead-Acid (Set 2)	Single	11.9	12.7	4.5*
		2 Series	23.8	25.4	4.5*
		2 Parallel	11.9	12.7	9*

* Indicates a nameplate capacity based on a new battery. Actual present capacity is unknown.
[+] Results from this test were published in [56]

Results

For each battery configuration shown in Table 3.1, a manual test was run where identification and parameterization was performed using data analysis tools in NI Diadem. This section will demonstrate the manual process for identifying the battery chemistry and capacity. These values are then presented alongside of battery parameters obtained using the autonomous method presented above.

Manual Battery Parameterization

Determining the chemistry identification gradient manually requires analyzing each pulse cycle to find the peak voltage during unloaded recovery as well as the lowest voltage (valley) during battery loading. National Instruments Diadem software was used to perform these operations since the data was recorded as TDMS databases and the software excels at handling large datasets. Diadem software provides tools to calculate peaks and valleys in a dataset given a search window width and either a maximum value or minimum value based on whether valleys or peaks are desired. As an example, the dataset shown below in Figure 3.8 was analyzed manually and the results follow. Data shown is for a 4.5 Ah lead-acid battery discharged with 1 A pulses. Each pulse was 200 seconds in length, with the battery terminal voltage measured at a down-sampled frequency of 30 Hz[4].

Peaks were found using a window width of 4500 samples (~ 150 seconds) and a minimum threshold of 5 V. Valleys were found using the same Diadem tool in valley mode with a window width of 4500 samples and a maximum threshold of 12.5 V. The tool identifies the window location of each peak and valley, making it simple to align the results and ensure each

[4] Many manual tests were run at various data collection frequencies before the rate of 10 Hz was decided upon for the final test configuration. Data collection speeds showed no signs of affecting the resulting data.

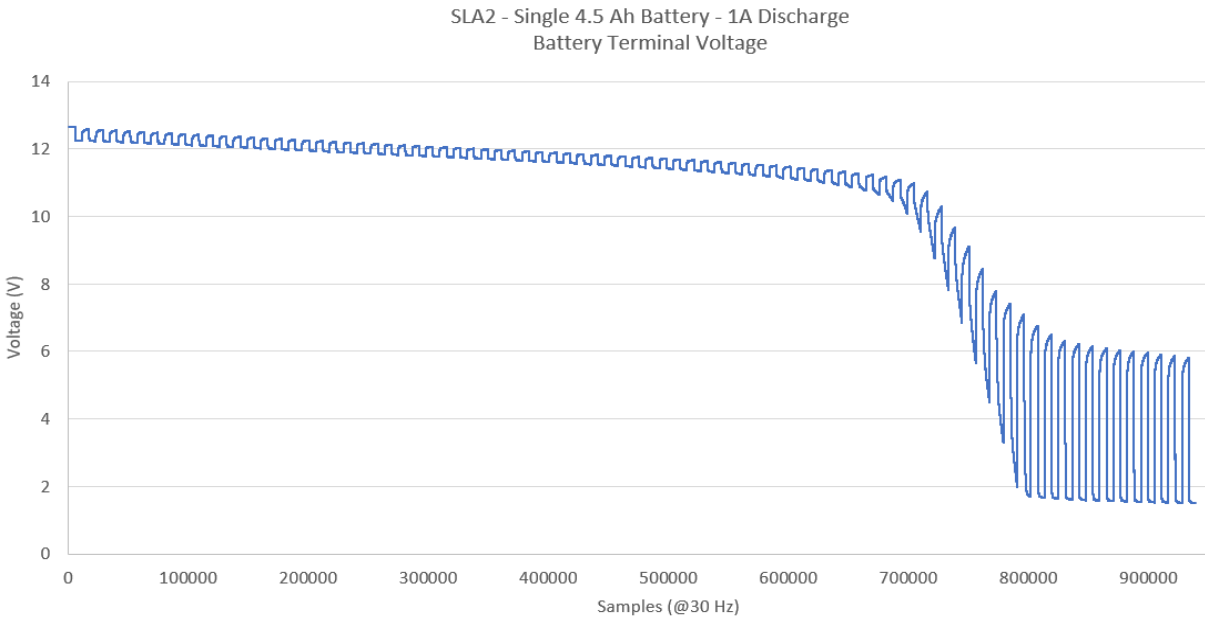


Figure 3.8 - Raw terminal voltage data for a single 12 Ah lead-acid battery.

peak and valley align with their corresponding value from the same pulse. Figure 3.9 shows the results. As expected, the shape follows that of the outline of the raw dataset.

From here, each valley is subtracted from the corresponding peak from the same pulse to give a voltage difference plot as shown in Figure 3.10. From this chart, the linear zone can then be approximated and used to determine the voltage gradient for identification. For this battery the points chosen to calculate the slope and obtain the identification gradient are highlighted in orange. Based on the values given, the identification gradient for this battery is 0.75 V, which places it in the range of a lead-acid battery. Further manual identification results are presented in the following subsection alongside the results from autonomous identification tests.

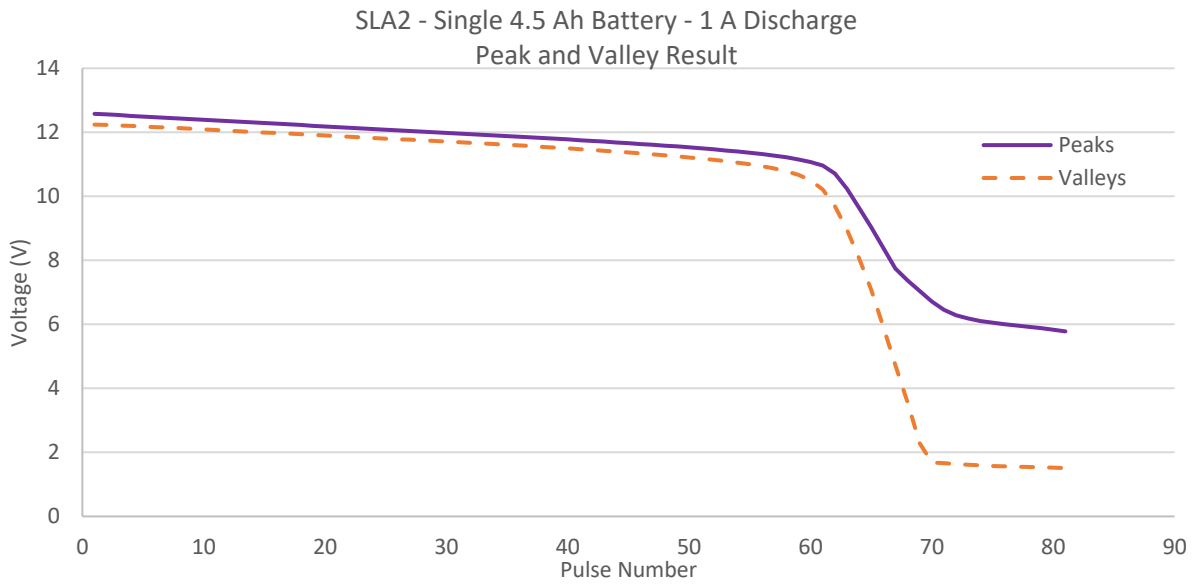


Figure 3.9 – The result of finding the peaks and valleys of the raw dataset.

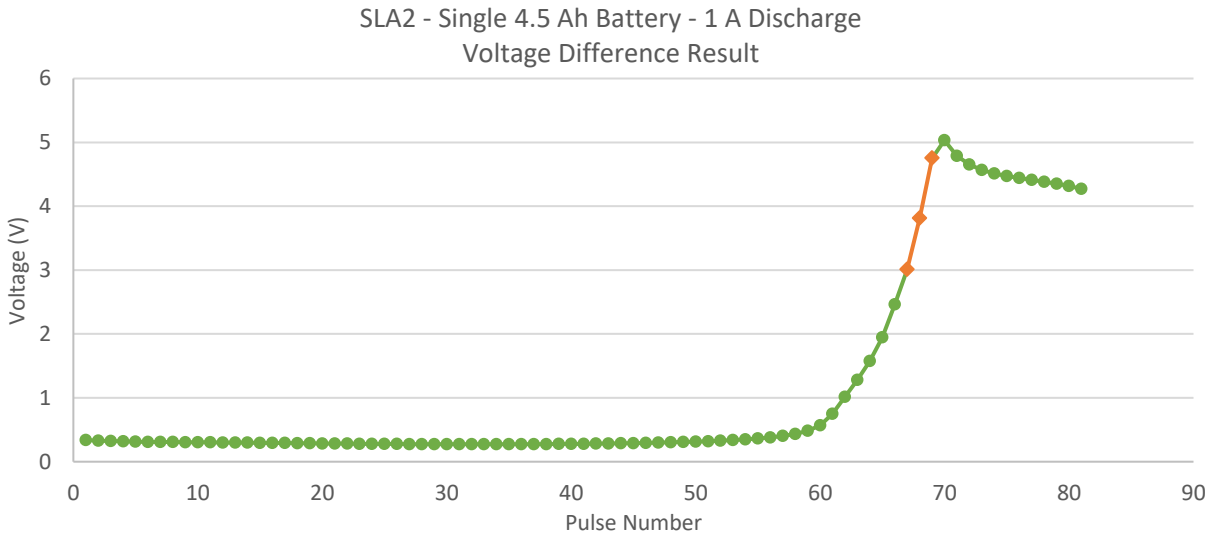


Figure 3.10 – The voltage difference between the peaks and valleys shown above

The final step in the manual parameterization process is to calculate the estimated capacity and compare it to the autonomous parametrization data. Since the capacity determination is done using a modified version of USABCs static discharge test, the same autonomous capacity test was used to verify the estimation on both sets of test results. Figure 3.11 shows how the current pulses remained consistent until right after the linear portion of the voltage difference curve. The current tapers at the end of testing as the battery nears 100% DOD. For the capacity estimation, only pulses that sustained the full 1A current for the entire 200s pulse cycle were chosen; resulting in a pulse count of 71. Using (3.3) from above, this results in an estimate capacity of 3.94 Ah. This is approximately 87% of the nameplate capacity of the battery. However, it should be noted that these lead-acid batteries were recovered from a gently used uninterruptible power supply (UPS), so they may have experienced some health degradation before testing.

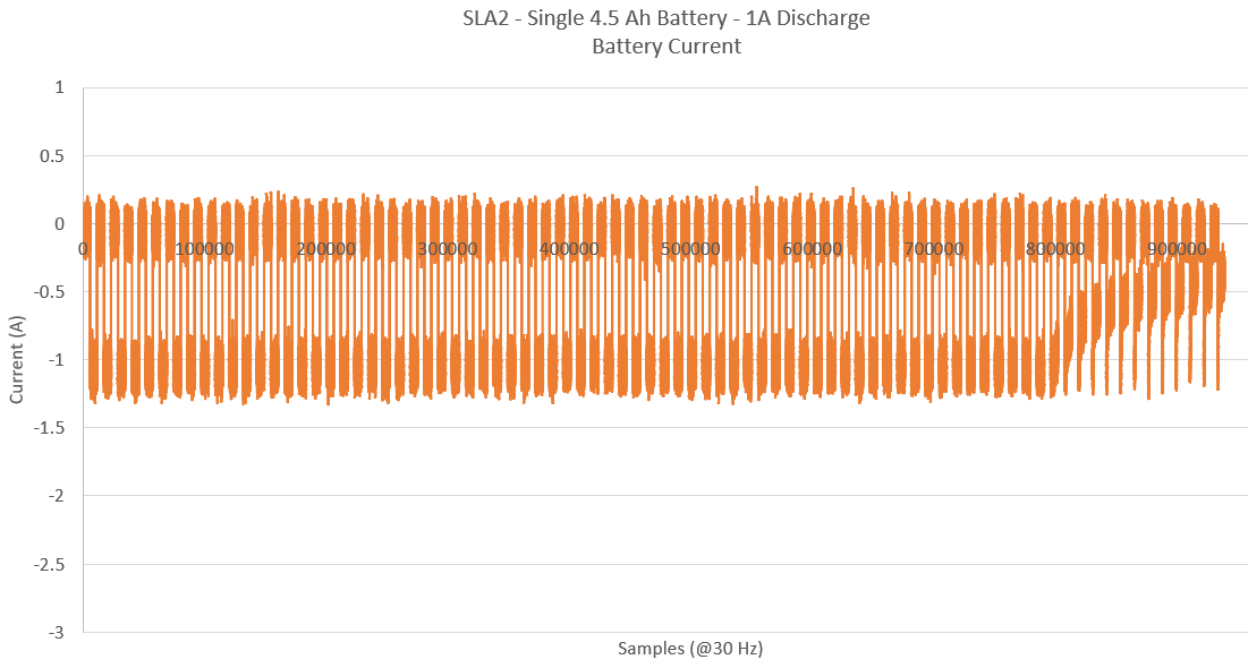


Figure 3.11 - Battery current from the pulse-discharge test shown above.

Autonomous Battery Parameterization

Results from the autonomous battery testing are presented in Table 3.2 alongside their counterparts from manual testing. Autonomous testing followed the procedure outlined in the Methods section of this chapter. In the case of a failed test or misidentified battery, it was repeated until a successful test was obtained, or the test failed three times.

Discussion and Conclusions

Analysis of the voltage gradients for identification above shows that despite the various test configurations and differences from the manual testing values, the autonomous battery parameterization process still results in clear battery chemistry identification bands. Li-ion can be defined as a gradient between 0 and 0.6 V, lead-acid as between 0.7 and 1.8 V, and NiMH as anything between 1.9 and 3.5 V. These bands worked well for identification if the pulse test could fully execute. However, sometimes the tests would fail repeatedly resulting in the system rejecting the battery for autonomous parameterization.

Failed tests are defined as ones which reach near 100% depth of discharge (DoD) without ever having three consecutive pulses with decreasing slope to trigger the end of the test. When the battery nears 100% DoD, it will be unable to provide the requested current which indicates to the decision-making tree that the battery is fully discharged, but not identified. Most often, this occurred when a string of series-cells became significantly unbalanced from one another. Figure 3.12 shows an example of one test where this occurred and the resulting voltage difference curve. Each jagged slope is the result of one cell of the series reaching empty before the next. Most often, this occurred with the NiMH Toyota Prius batteries.

Table 3.2 - Test results from running manual parameterization techniques alongside the autonomous method.

Chemistry	Config. (Modules)	Manual G_{batt} (V)	Auto. G_{batt} (V)	Manual Q_{est} (Ah)	Auto. Q_{est} (Ah)	$Q_{nameplate}$ (Ah)
Leaf Li-Ion (Set 1)	Single[+]	0.20	0.18	36.2	35.4	38.9
	2 Series	0.32	0.33	34.4	33.1	38.9
	2 Parallel[+]	0.17	0.21	71.4	69.9	77.8
	2P2S	0.22	0.24	72.7	70.0	77.8
	4 Series	0.11	0.14	33.3	33.6	38.9
Leaf Li-Ion (Set 2)	Single	0.59	0.42	38.0	38.1	37.6
	2 Series	0.47	0.26	37.7	36.5	37.6
	2 Parallel	0.41	0.40	74.1	72.2	75.2
	2P2S	0.38	0.49	69.9	68.9	75.2
BMW Li-Ion	Single	0.40	0.43	80.2	80.0	99*
	2 Series	0.29	0.32	80.6	80.1	99*
	2 Parallel	0.27	0.27	159	157	198*
Fiat Li-Ion (Set 1)	Single Module (5 Series)	0.18	0.20	44.3	43.3	60*
Fiat Li-Ion (Set 2)	Single Module (5 Series)	0.09	0.12	48.5	48.5	60*
Prius NiMH (Set 1)	Single[+]	2.13	2.28	6.62	6.48	6.5
	2 Series	2.60	2.71	6.55	6.54	6.5
	2 Parallel[+]	2.54	2.78	12.8	12.1	13.6
Prius NiMH (Set 2)	Single	2.13	2.28	1.87	1.63	2.5
	2 Series	2.25	2.30	2.33	2.21	2.5
	2 Parallel	2.44	2.45	4.99	5.01	5.0
Prius NiMH (Set 3)	Single	2.41	2.40	0.83	0.77	0.8
	2 Series	2.38	2.40	0.78	0.77	0.8
	2 Parallel	2.02	2.22	1.53	1.59	1.6
	4 Series	Failed to execute				0.8
	4 Parallel	2.55	2.63	3.05	3.02	3.2
Lead-Acid (Set 1)	Single[+]	0.89	1.21	11.88	11.87	12
Lead-Acid (Set 2)	Single	0.75	0.81	3.94	4.11	4.5*
	2 Series	1.11	1.15	3.95	4.07	4.5*
	2 Parallel	1.08	1.22	8.16	8.13	9*

* Indicates a nameplate capacity based on a new battery. Actual present capacity is unknown.

[+] Results from this test were published in [56]

One explanation of this phenomenon is because each Prius NiMH battery is six cells in series with no balancing, regulation, or bypass resistors. When single cells drain faster than their counterparts, they reach a point when the voltage behavior collapses, and the cell essentially becomes a resistor in series with the rest of the battery. This is evidenced in Figure 3.12 where the terminal voltage falls rapidly, stabilizes, and then falls again. In this case, two cells stopped producing a voltage, though the current demand was still being met by the remaining cells. As the cell recovers, the terminal voltage experiences strange surges as the individual cells recover as well. However, this behavior is not good for the battery and is indicative of cells with uneven aging. This means that as the battery modules age, any unbalance can be exacerbated since the module has no way of accounting for the issue. This hypothesis is supported by data that showed older, low capacity modules had a higher rate of incomplete pulse-testing than the newer, less

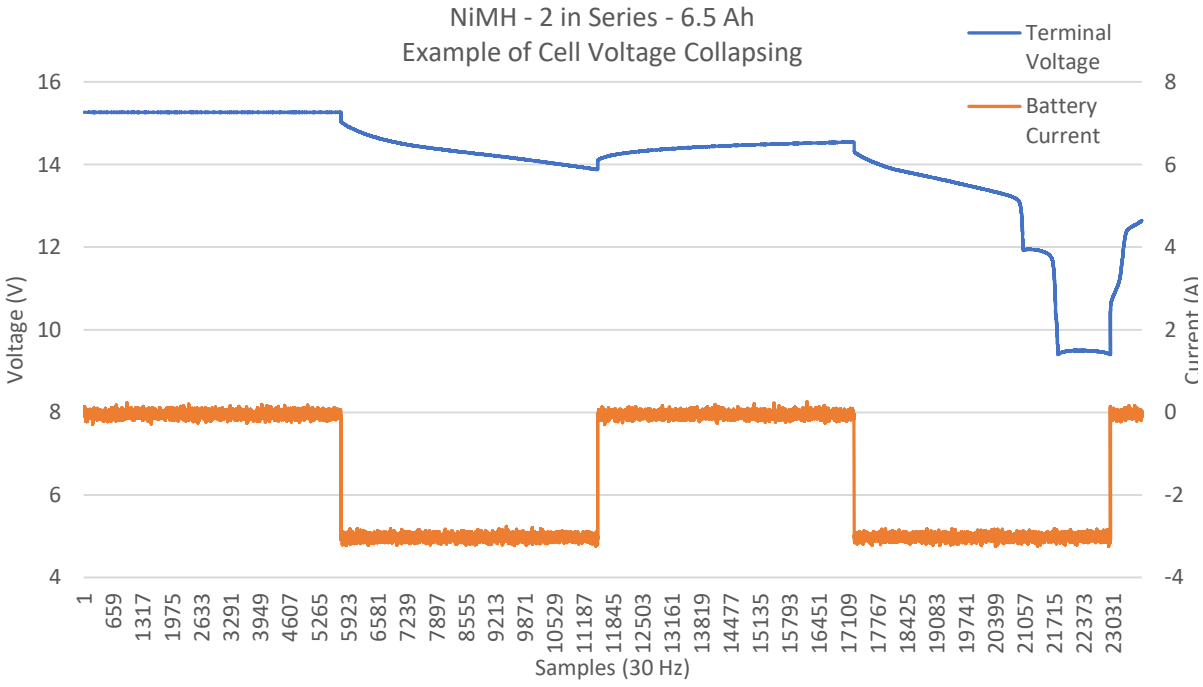


Figure 3.12 - An example of NiMH cells reversing voltage.

aged modules. While the lack of a battery management system (BMS) may seem like an oversight, for new NiMH batteries, self-balancing is all that is required.

Further investigation of Table 3.2 shows that the autonomously determined chemistry identification gradients were on average about 5-10% higher than the ones determined manually. Since this resulted in a +5-10% offset, the ability of the autonomous identification algorithm was not affected. However, it is helpful to understand why this occurred. Some difference between hand-calculated values and automatically generated values is to be expected. The consistently higher nature can most likely be attributed to the fact that the identification algorithm is using the last three values with the highest slope between them before the difference curve reached the local maximum. During manual identification, it was easier to pick out a larger linear area of the curve, which would end up averaging the slope slightly lower.

Capacity estimations between the autonomous and manual methods were also consistently different. When performed autonomously, the capacity measurements typically resulted in a slightly lower capacity estimate. This can be attributed to the fact that the tests ended earlier than the manual tests, resulting in fewer full-current pulses to be counted towards the capacity estimation. As examination of Table 3.2 reveals, these estimates were still within 20% on average, which means that the impact on the static-discharge test is minimal. Changes to the measured capacity of the battery are going to appear when the discharge rate is significantly different between measurements. Modifying the discharge rate by tenths of amps will have less impact on the capacity measurement of the battery.

Given that the autonomous method produced differing results for capacity from the manual tests, the accuracy of the tests may come into question. While the results were slightly different, they still resulted in consistent identification of the battery chemistry and an

approximation of the usable capacity. It is important to differentiate that the capacity measurement presented here is the usable capacity of the battery versus the nameplate capacity of the battery. When batteries are used in an actual system, often the entire capacity is not used in order to avoid damaging the cells. For instance, in the residential energy storage system discussed in Chapter 7, the battery is only allowed to charge to 80% SOC and only discharged as low as 20% SOC. The operable range inside of these values is the true usable capacity of the battery without the voltage getting too high or too low and possibly resulting in severe battery degradation. Having a measurement of the usable capacity of the system means it can operate more safely within the limits of the system.

Finally, while the pulse data was recorded and could be used to perform a Hybrid Power Pulse Characterization (HPPC) test on the battery modules as initially discussed, it was decided that it would not be necessary for this work. HPPC can provide a moderately accurate measurement of the internal resistance of the battery system at various frequencies, giving good model of the battery behavior. However, without an exhaustive look-up table of previous tests to measure these resistances against, these values become little more than an approximation of the internal system losses. For that reason, it was decided to log the pulse data for the tests but not execute the HPPC. Companies like SNT and OEMs log thousands of module tests to get a statistical approximation of measured internal resistance versus capacity to approximate where a battery falls in its lifecycle. Unfortunately, this data is often proprietary since it gives valuable performance data for the battery chemistries being used.

While health data is important to the architecture described in the following chapter, capacity and chemistry play a more pivotal role in determining the dispatch value of a battery in the multi-chemistry system. In the future, the health results from an HPPC could be used to

more accurately approximate the remaining life cycles on a battery and give a slightly better valuation for the connected unit. For now, a significant value increase has already been achieved by limiting the amount of time a technician must interact with an energy storage unit to get it up and running. An autonomous battery parameterization technique has been demonstrated which allows for an unknown battery to be installed and characterized with minimal technician interaction. Using this information to optimally dispatch a multi-chemistry system using a distributed control and actuation architecture is the focus of the following chapters.

Chapter 4 – Architecture for Control, Actuation, and Optimization

Chapter 2 described common architectures and electrical topologies for large-scale energy storage. Present methods often increase system voltage by making large series strings of batteries. The system energy capacity is then increased by placing these high-voltage strings in parallel. Examples of this method can be found in Tesla's patent for their large-scale energy storage system or from Mercedes Benz's large-format energy storage [32] [57]. When using secondary-use batteries, creating large groups of parallel batteries for a high capacity system can be difficult because of the various types of EV batteries. To circumvent this issue a system architecture is proposed where battery units of various capacities and voltages are intertied through parallel DC/DC converters to allow carefully controlled dispatch of the different energy storage elements. The topology chosen for the multi-chemistry system is shown in Figure 4.1 and functions as a DC microgrid with a point of common coupling (PCC) to the grid through a high-power DC/AC inverter.

Each DC/DC node of the system contains an energy storage element with management and monitoring system, a power electronics converter (PEC), and an energy storage module controller (ESMC). The ESMC provides a modular, agent-based interface for interpreting and executing controls between each component of the converter node and the system-level controls. Communications between the ESMC and the central system controller are handled using an application programming interface (API). Energy is added and removed from the utility-scale energy storage through the DC/AC inverter node. Contained within this node is the PEC and the central system controller (CSC), which manages the system-wide operational state as well as the status of the locally connected inverter. The CSC creates and maintains a transactive energy market with each of the connected ESMCs as a 'customer' in order to optimally dispatch the

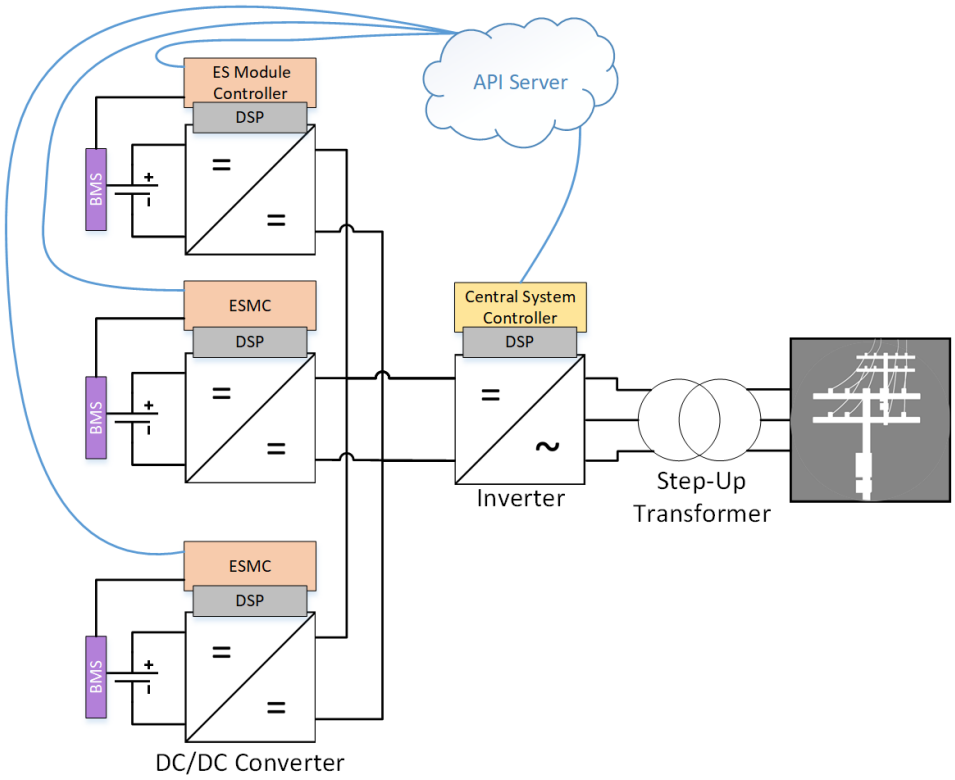


Figure 4.1 – System architecture and electrical topology of the multi-chemistry BES [58]

system and obtain the maximum value over the life of the connected energy storage elements. Like the ESMC, the CSC uses a modular, agent-based approach to monitoring and managing the various locally and remotely connected devices. Agent-based interfaces with power electronics systems have previously been explored in the literature [58] [59]. This chapter will present the architecture and control methods used in the CSC and ESMC to achieve distributed control.

Overview of Agent-based Control Nodes

An agent-based architecture was chosen for the control nodes due to their modularity and novel method for integrating a wide range of physical components. In an agent framework, separate applications running on the same computational device interchange data with a shared message bus. Agents can be added or removed as necessary to interface with different devices

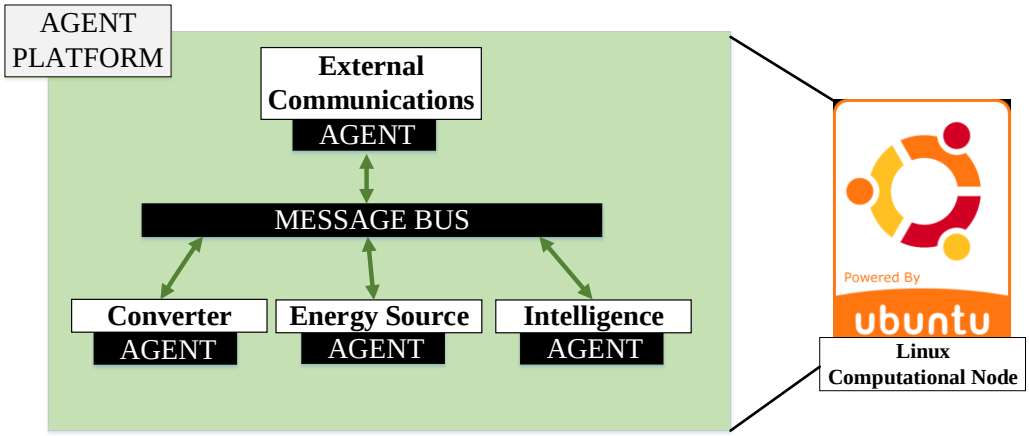


Figure 4.2 - General agent structure of a node [60]

that may be attached to the control node. Figure 4.2 shows the basic building block used to develop a generic converter node. Each node requires an energy source agent, a power electronics converter interface agent, a communications agent, and an intelligence agent to coordinate the actions of the other agents. Depending on the application of the node, extra agents are needed to generate pricing signals or provide specific communications abilities. Using an agent framework allows these items to be added to the system without completely reconfiguring the communication structure.

Central System Controller

Co-located with the DC/AC inverter that is serving as the PCC with the grid, the central system controller (CSC) serves as the coordinating controller for the transactive architecture as well as the intelligence for the inverter. A map of the agent framework used in the CSC is shown in Figure 4.3. Each agent in the framework serves a distinct purpose and runs as a separate application with constant access to any information published to the internal message bus of the

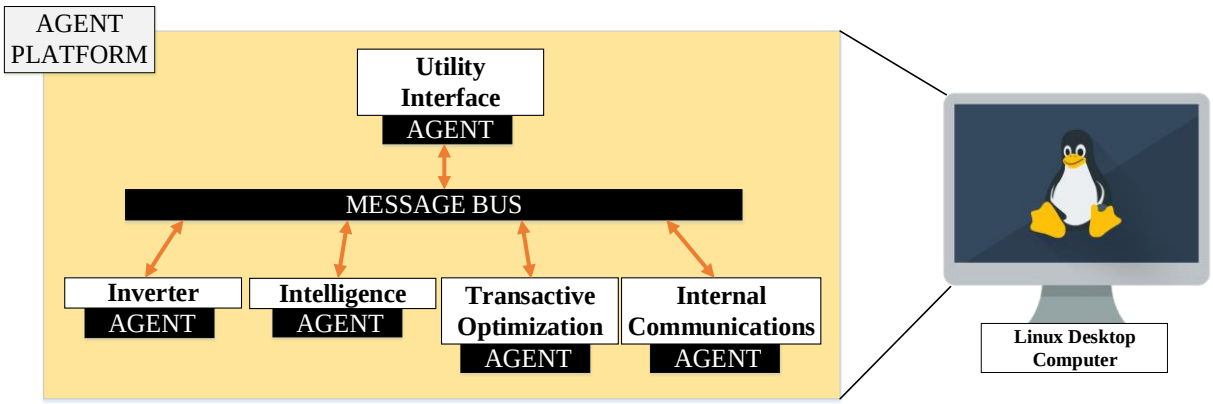


Figure 4.3 - Central system controller agent framework [58]

computer. The following provides a description of the purpose and basic operation of each agent in the CSC.

Utility Agent

The utility interface agent acts as the system contact to the outside world. It can be adapted to use the communications protocol of the utility's choice, provided a system set-point is communicated at least every 5 minutes. Data about system status, present energy availability, and power level are collected by the utility interface from the message bus and made available to the utility by this agent as well. For this research, the utility agent provides a central graphical user interface (GUI) to interact with and view the status of the entire ESS. The utility interface agent also allows different use-cases to be loaded for testing the system response under a variety of dispatch scenarios. This is done by specifying a "dispatch-script" for the system to follow over time.

Each script divides a 48-hour window into five-minute windows of desired positive (charging) or negative (discharging) power dispatch. Five-minute intervals were chosen since this provides plenty of time for the transactive optimization to solve, while also not requiring the

system state change too rapidly. For example; a single 50 kW, 150 kWh battery[5] discharging at full power could only discharge approximately 4.17 kWh of energy in a given 5-minute interval or about 2.78% of the battery capacity. This means that if a battery is nearing its charge and discharge state of charge (SOC) limits, the system will not be able to exceed those limits by more than three percentage points. It is common for the power limit of the battery to be de-rated towards the SOC limits as well, so the maximum energy transfer could even be less than 3% [60].

Intelligence Agent

Set-point and system state change requests are passed from the utility interface to the intelligence agent for processing. Possible system states include standby, start-up, on-grid normal, shutdown, error, and lockout. Based on information collected from the inverter agent and the energy storage modules, the intelligence agent makes decisions about whether it is possible and safe to energize the system. When the agent receives a start-up command and has no system-wide errors, it changes states to begin operation.

Upon initialization, the intelligence agent locks-out the system and requires that a "Clear Errors" message be received from the utility in order to move the system to the standby state. This is done to ensure that after a service interruption, the system does not restart in an unsafe manner by responding to old commands. After clearing system errors and receiving the start-up command, the intelligence agent performs a series of checks to make sure that executing the start-up sequence will not damage any sub-systems. An energization sequence is then executed. If the start-up is performed successfully, the agent moves to the "on-grid normal" state and

[5] A typical utility-scale battery energy storage system would contain 3+ hours of energy at the maximum rated power.

signals the inverter agent to start normal operation. Once in the normal state, the system monitors for any errors or a shutdown command. When one of these triggers occurs, the system either quickly shuts down (error), or signals for the system to shut down by tapering power transfer, and then executing a shut-down sequence.

Should an error occur resulting in a lockout, the system must be reset from this state by the utility with a "Clear Errors" message to the CSC. The intelligence agent would then move to the standby state and wait for the next start-up command while watching for possible system errors. If the system shuts down normally, then the agent will automatically move to standby once shutdown is completed.

Inverter Agent

The intelligence agent helps coordinate the actions of the inverter and the rest of the energy storage system. When the intelligence sends a command to the inverter agent, it processes the information and sends the appropriate commands over a communications protocol such as TCP/IP, UDP, or ModBus to the inverter controller. In order to control the inverter and energize the system successfully, the inverter agent must be able to control a contactor on the DC side of the inverter as well as a contactor on the grid side of the inverter. The states of these contactors, as well as the present state of the power stage must be communicated to the inverter agent along with voltage, current, power, and frequency measurements. This information is used by the inverter agent and intelligence agent to determine if the ESS should be in an operational state.

The inverter agent monitors the incoming data from the power stage for several errors shown in Table 4.1. Note that these errors occur on a timescale of tens to hundreds of milliseconds rather than microseconds. The inverter agent still relies on the signal processor of

Table 4.1 - Sample of errors that trigger a system shutdown

Name	Tolerance	Description
AC Over/Under Voltage	+/- 5%	An RMS voltage measurement (5 cycle average) arrives that is too low or too high, causing a safe shut-down
Inverter Communications Timeout	10 seconds	Communications is lost with the high-speed inverter controller. No up-to-date data means the system must shut down by removing DC voltage
DC Over/Under Voltage	950-1050 Vdc	The DC voltage measurement (16 ms average) of the bus goes outside of the safe operating range.

the inverter to identify transient problems with AC or DC voltage and shut down appropriately. However, if the system drifts towards an unsafe operating point, the inverter agent can shut the system down using a standard de-energization procedure instead of just aborting energy transfer mid-operation. For example, if the temperature of the power stage approaches a limit, the inverter can ramp the power-setpoint to zero and then de-activate the system.

As an example, the residential ESS inverter agent allowed the AC voltage to drift +/- 5% of the nominal grid voltage. If an over or under-voltage was sustained for at least 4 cycles (~68 ms), then the intelligence agent would execute a normal system shutdown. However, should a more severe voltage event occur and take the system +/- 10% of nominal, then the digital signal processor (DSP) would deactivate the power stage and open the protection contactors after at most a quarter cycle (~4.2 ms). DC voltage protection actions can happen even faster on the DSP since the measurement is not a periodic signal.

For this research, the inverter agent was designed to start-up and behave as if a 100 kW power stage were being controlled. This power rating is re-configurable through a configuration file. Instead of executing communications with an outside device, the inverter agent contains a simple model of a power electronics converter that assumes that the system will remain under

normal operating conditions. The assumptions are that the DC link will remain fixed at 1000 Vdc and the grid connection will remain steady at 480 Vac.

While these assumptions do not cover transient converter interactions and behaviors, they provide a basis for testing the operation of the communications systems and control algorithms. That means the data being relayed by the model to the inverter agent is more important than testing real-world transients. However, real-world testing of an agent-based system will be discussed more in Chapter 5 with a residential agent-based ESS. To verify function error detection, manual input was added to cause DC and AC voltage and current deviations outside of the configured limits.

Transactive Optimization Agent

Dispatch of each ES element in the multi-chemistry system is determined by the transactive optimization agent. This agent requires a price curve and present status from each of the attached energy storage elements. Based on the status of each attached system (lockout, standby, or normal), the transactive optimization agent determines whether the unit should be included in the dispatch scheduling of the system. This way, a unit with an error will be left out of the system optimization. The transactive optimization agent also uses this to recalculate the dynamic total system capacity and the present available capacity (SOC) of the larger ESS. These values are communicated with the utility agent so that the end-user is aware of any system problems impacting capacity.

Every five minutes, an optimization uses the price curves, statuses, and the utility dispatch request to maximize the obtainable value from the system. It is important to note that the optimization is only using the price curve data from each ESMC to make the determination of the dispatch level for the attached energy storage unit. This simplifies the centralized

optimization, while allowing each battery to manipulate its price curve to obtain its desired dispatch value. Instead of the optimization having a series of weighted objectives to balance battery life, value, and usage for each unit, each battery determines the best balance of life and use by assigning it value. The central optimization dispatches the system to minimize the cost to operate the battery system.

Internal Communications Agent

The last agent in the CSC is the internal communications agent that allows it to communicate with the RESTful API server that provides inter-device communication. The server, represented in Figure 4.1 as a cloud-based device, is an application programming interface (API) that provides a publish/subscribe based messaging bus similar to that used internally in each computer. The internal communications agent subscribes to any data posted with a prescribed heading on the API and downloads that data at regular intervals. If a device disappears from the API or ceases to update, the internal communications places the unit in an error state which removes the device from the optimization. When a new device appears on the API, the internal communications agent creates a new ESMC object on the internal message bus to relay the received data to the other agents. Each agent using this data automatically retrieves the data it needs. For example, the transactive optimization agent will retrieve the price signal from the server and incorporate the new device into the next optimization cycle.

When the transactive optimization solves, the internal communications agent posts the results of the optimization alongside the requested dispatch for each ESMC to the API server. Results are published as either “optimal”, “infeasible”, or “standby.” The first two states are only produced when the system is running successfully, and the optimization is running on

regular intervals. The Standby State is used when the inverter is not energized, and the system is waiting for a dispatch command from the utility agent.

Energy Storage Module Controller

Each DC/DC converter in the system has an agent-based energy storage module controller (ESMC) paired with the high-speed converter controller. Like the CSC, the ESMC coordinates the actions of the DC/DC converter, the attached energy storage element, and regulates the data interaction between the ESMCs and the CSC. It also creates and manages the role that the energy storage element plays in the transactive architecture by manipulating a price signal that is sent to the CSC. Figure 4.4 illustrates a map of the agent framework developed to universally work with a range of energy storage devices [58].

Internal Communications Agent

Like the CSC, an internal communications agent acts as the interface to other devices in the system. Using the same RESTful API server, data is published along with an identification code so that the CSC can uniquely identify each attached ESMC. This agent published on a regular five second interval whether the data on the message bus has been updated or not. If no

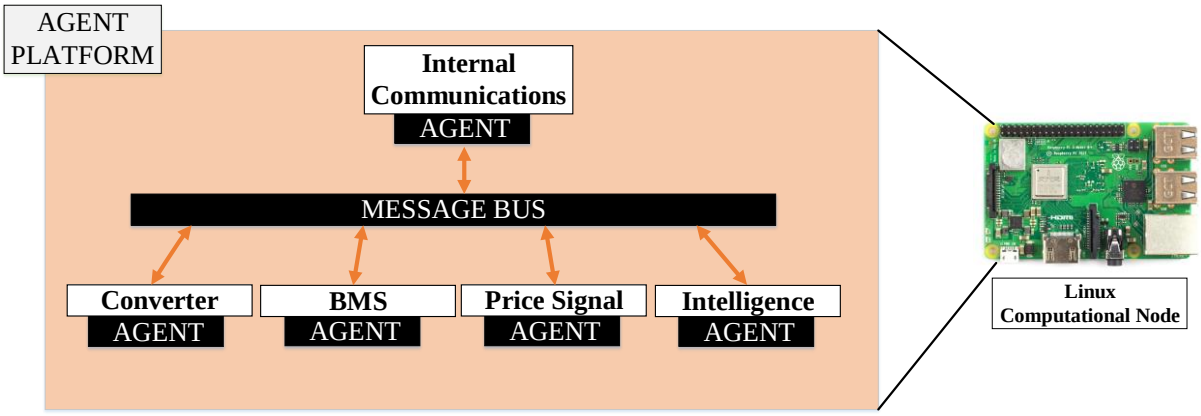


Figure 4.4 - Agent framework of an ESMC [58]

data is present for a value, an empty string is published in its place. To avoid creating multiple copies of data on the server or overwriting an existing ESMC with a duplicate name, the internal communications agent first polls the API server to check if an ESMC object exists with its identification code. If the server responds with an internal "404 – not found" error, then the agent sends a request to the server to create a new ESMC object at that address. If the ESMC has not attempted to post data before but receives anything other than a 404 error, then the internal communications agent posts an error to the message bus to inform the intelligence agent that there is a duplicate device on the system. This is done to prevent one device from over-writing data from another just because they happen to have the same identification code.

Each time the communications agent posts new data to the API server, it also checks for new commands from the CSC. If data is present and the timestamp of the data has updated since the last set of data, the then agent processes the data packet and posts the dispatch information from the CSC to the internal message bus for other agents. If data is present, but the timestamp has not updated, the agent increments an error timer. Should the error timer reach 6 consecutive failures, the agent assumes that the CSC has lost communications with the API server and stopped updating. An error is posted to the internal message bus alerting the intelligence agent to a loss of communications. Unless the ESMC losses communications with the API server, a case where no data is present should not occur. However, in the case of missing data, an error is sent to the intelligence agent so the system can safely shut down. Monitoring the inter-device communications enables the distributed nature of the controls. While high-speed data transfer is not necessary between devices, reliable and consistent data access allows safe operation of the devices.

Converter Agent

Like the inverter agent in the CSC, the converter agent provides an interface with the DC/DC converter that is controlling power flow in and out of the energy storage element. When commands are received from the ESMC intelligence agent, the converter agent processes the command and then sends the appropriate command via the pre-selected communications protocol. The power stage control device (DSP, FPGA, microcontroller) that is sending PWM signals to the power stage must communicate voltage, current, and power measurements to the converter agent. This information, along with input and output contactor states are received, processed for errors such as over/under voltage and over/under current, and then passed to the intelligence agent to determine the operational state of the system. This is done so that if a different DC/DC converter is used, the interface may be changed without having to change the overall control mechanism found in the intelligence agent.

In this work, the converter agent does not communicate with an external converter controller. Instead, the agent is configured to start-up and behave as a 50 kW power stage. The power limits of the DC/DC converter can be altered using a configuration file as necessary to demonstrate a variety of power stages. This is important later for testing the transactive optimization system where testing a variety of different power level converters will help determine the effectiveness of the optimization. A simple model of the converter is used in this agent which assumes that the DC input will vary between 370 and 440 Vdc while the output of the converter will remain fixed at 1000 Vdc. The converter is also assumed to have a stable power ramp rate of 25 kW/s such that any change in power set-point can be successfully executed in a single 5 second optimization time step if necessary. As discussed in the CSC

inverter agent, these assumptions are made so that the communications and optimization infrastructure may be tested.

Battery Management System (BMS) Agent

Interfacing with the wide range of battery management systems available from original equipment manufacturers (OEMs) and commercial off-the-shelf (COTS) suppliers requires the use of a BMS agent to manage a common data interface for the intelligence agent. To determine which data and control parameters are needed to properly control a BMS, several COTS BMS units were analyzed for the data that they publish or that are required for safe operation. Table 4.2 provides an overview of these 13 critical parameters. Other data is typically available from the BMS but may not be necessary for the BMS agent to operate properly. BMS units from OEM developers for EV batteries were excluded from this analysis. These developers (LG

Table 4.2 – Common battery parameters required by the BMS Agent

Name	Unit	Parameter	Description
Vbatt	V	Battery Pack Voltage	Voltage of the battery pack measured from the most positive terminal to the most negative terminal of the battery.
Ibatt	A	Pack Current	Current flow measured on the positive side of the battery. Positive is discharging, negative is charging current
SOC	%	State of Charge	An estimate of the total remaining energy in the battery, calculated by the BMS's internal algorithm
Cell_V_{min}	V	Minimum Cell Voltage	Cell voltage of the cell with the lowest voltage
Cell_V_{max}	V	Maximum Cell Voltage	Cell voltage of the cell with the highest voltage
Cell_V_{avg}	V	Average Cell Voltage	Average cell voltage across the entire pack
$Temp_{max}$	˚C	Highest Pack Temperature	Highest measured temperature in the container across all thermistors (cell or ambient)
$Temp_{min}$	˚C	Lowest Pack Temperature	Lowest measured temperature in the battery across all thermistors (cell or ambient)
$Temp_{avg}$	˚C	Average Pack Temperature	Average across all thermistors in the battery (cell or ambient)
DCL	A	Discharge Current Limit	Maximum discharge current before the BMS faults and opens to isolate the battery for protection.
CCL	A	Charge Current Limit	Maximum charge current before the BMS faults and opens to isolate the battery for protection.
Status	N/A	Activation Status	Present status of the system. Activated/deactivated
Control	N/A	Start/Stop Control	Control channel to intertie or disconnect the battery

Chem, Panasonic) often do not allow third-party interfaces with their communications bus in order to protect the proprietary state-estimation and charge-balancing algorithms implemented on the BMS. However, given knowledge of the communications protocol and a list of data registers, any OEM BMS could be added later as long as the 13 parameters identified by this work are included.

The values shown above in Table 4.2 represent the control parameters that the BMS agent monitors to maintain safe operation of the energy storage element. Like the intelligence agent, a state engine is used to monitor the present battery conditions and raise an error if the conditions reach an untenable state. Should an error occur, the BMS agent immediately sends a shutdown command to the BMS while also publishing an error to the message bus so that the ESMC converter and intelligence agents can execute an emergency shutdown. The error engine in the BMS agent monitors for the following conditions shown in Table 4.3. The bounds on these conditions are set via a configuration file since they can change from battery to battery based on size and cell configuration.

Table 4.3 – Errors monitored by the BMS agent

Error	Limits	Result	Description
Operational SOC Limit Reached	Op_SOC_{min}, Op_SOC_{max}	Once error is sent to Intel. agent, waits for shut-down command from Intel.	Limits set ~5% inside of the actual SOC cut-off of the BMS to prevent a "hard" shutdown
SOC Limit Reached	SOC_{min}, SOC_{max}	Immediately sends the shut-down command to the BMS	Set just inside the BMS limits (~1%) and mainly used to prevent the ESS from starting energy transfer
Over/Under DC Input Voltage	DC_in_{min}, DC_in_{max}	Immediately sends shut-down command to the BMS	Attempts to protect the battery from voltages outside of the normal charge/discharge range
Over/Under Pack Temperature	T_pack_{avg}	Once error is sent to Intel. agent, waits for shut-down command from Intel.	Ensures that the battery temperature does not get close to the limits
Over/Under Cell Temperature	T_Cell_{min}, T_Cell_{max}	Once error is sent to Intel. agent, waits for shut-down command from Intel.	Checks that any individual cell is not approaching the thermal limits. Executes a "safe" shutdown

For most systems, monitoring SOC provides a good estimate of the operational state of the battery. However, knowing information such as cell voltage, and pack voltage allows for the BMS error engine to make more informed decisions about the operational state in extreme conditions, such as the sudden failure of individual cells, unexpected transient voltages, or when operating near temperature boundaries. While the BMS monitors these parameters as well and will protect the battery when necessary, it will also perform these shutdown actions abruptly and without first issuing a warning[6]. Having the agent monitor these values using custom conditions set within the bounds of the BMS allows the agent to recognize unsafe operating conditions earlier and safely shut down the system instead of abruptly opening the electrical connection with the energy storage element. This can prevent cascading failures in the system that may result from the sudden transients that would occur.

For integration with the architecture described in this project, a simple battery model was implemented in the BMS agent in place of communications with an actual BMS. Using a linear charge counting equation (4.1),the model adjusts the present battery SOC at a rate that is proportional to the current power output of the system. Based on the direction of the power flow and the length of the model timestep, energy is either added or subtracted from the present energy contents of the battery. This also allows for accelerated testing by adjusting the model time-step (Δt) to represent longer amounts of time. For normal operation, this parameter would be equivalent to the model loop execution time.

$$SOC_n = SOC_{n-1} + \frac{P \times \Delta t}{Q_{total}} \tag{4.1}$$

[6] In general, the COTS BMS units evaluated do not provide a warning signal. However, the Nuvation HV BMS can provide a warning signal. Since this feature is not common, it is assumed no warning will be given by the BMS.

In (4.1), Q_{total} is the energy capacity of the battery in watt-hours (Wh), P is the present power dispatch in watts (W), and n is the current iteration of the program. It is important to note that for this model, a positive power flow represents a charging current, while a negative power flow represents a discharging current. While this is the reverse of the direction chosen by most COTS BMS units, these signs are most commonly used by the utility industry to define ES power flow. Additionally, it should be noted that an efficiency of 100% is assumed to simplify the battery model. This is done because battery efficiencies are non-linear and would need to be included in the model as a piece-wise function.

As the SOC approaches the extremes of the operational range (Op SOC), the power capability is de-rated by adjusting the current limit on the battery. For example, as the SOC approaches within 5% of the minimum SOC value the maximum discharge power limit will fall linearly until it reaches 0 W when the SOC is at minimum. This is part of the model because all the COTS BMS units studied for this work included a similar power de-rating mechanism as the batteries approach the SOC limits. This helps prevent the unit exceeding SOC limits if the timing of the communications encounters an error, since the unit will automatically begin to power down. Many BMS's also de-rate the power limits of the system based on temperature readings. However, since a nominal temperature of 23 ˚C is assumed for each battery, this is not a necessary part of the model.

Price Signal Agent

Battery system control often requires knowledge of most, if not all, of the 13 parameters described above. However, the distributed transactive architecture used for this work requires each energy storage device to reduce these 13 variables to a simple price signal consisting of 3 parts. The first part is the linear price signal. The second part is the power limits in both the

charge and discharge directions. The last part is the present status of the system which only includes an active/standby/error status and the present power output of the unit. The price signal agent of the ESMC is where this simplification of parameters occurs.

The price signal agent retrieves the needed information on the battery and the converter status from the internal message bus of the ESMC. Using this information, the three components of the price signal are determined and pushed back to the message bus. A new price signal is generated every two seconds and pushed to the message bus, where it is sent to the CSC by the communications agent. Regardless of whether the system is presently charging, discharging, or in standby, the price signal is always calculated for both the charge and discharge directions. This is done to ensure that the CSC has the most up-to-date system price curve at all times in case the dispatch request from the utility to the CSC shifts suddenly.

The price signal is sent to the communications agent via the VOLTTRON message bus, where it is packaged with other system information as well as identifying messages into a single JSON string. Using JSON provides a simplified way to package and interpret data between the communication agents in the ESMC and CSC.

Intelligence Agent

Once the CSC has negotiated a dispatch plan for the system, the power set-point is received by the communications agent and posted to the intelligence agent. Like the intelligence agent in the CSC, the primary function of the ESMC intelligence agent is to monitor data from the ESMC agents for errors and to provide a safe start-up and shutdown sequence for both

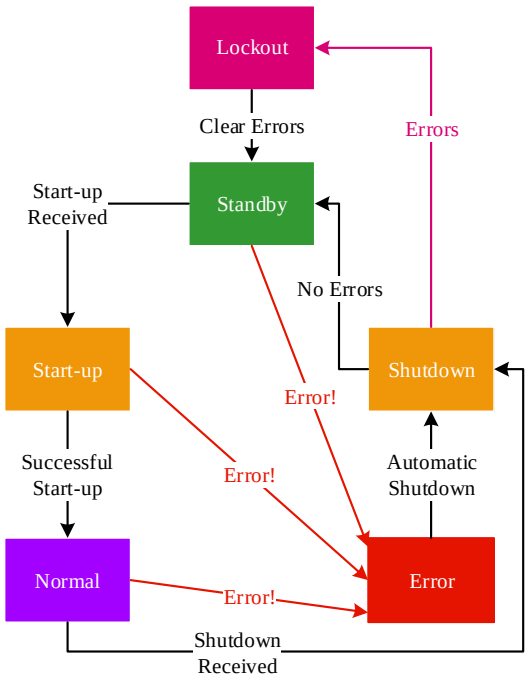


Figure 4.5 - ESMC state machine including state triggers

normal and fault scenarios. To track the system state, a state machine is used which allows the system to be in standby, start-up, normal, shut-down, error, or lockout as shown in Figure 4.5.

Upon powering-on and receiving the first commands from the CSC, the intelligence agent will begin in "lockout" and require that a "clear errors" command is sent. This is a safeguard so that the system does not come online during the middle of system operations and jump to an infeasible state causing damage. Once the ESS is in standby, if a valid power set-point is received from the CSC then the intelligence agent will begin a start-up sequence. This involves communicating with the BMS and converter agents to energize the converter in a safe manner based on measurements from these agents. Upon successful start-up the system transitions to the "normal" state. Once in the normal state, the intelligence agent monitors the message bus for errors from the other agents as well as communications errors.

Summary

The parallel DC/DC topology proposed for the multi-chemistry battery energy storage system requires a series of duplicate energy storage modules. However, flexible integration with a wide range of hardware is desired, which necessitates modular software. Using an agent-based architecture for the system-level control of the energy-storage modules allows for integration with multiple types of hardware. This means duplicate software can be used on several identical ESMCs while communicating with a range of BMS units and converters, reducing system integration costs by simplifying the communications between the PECs and the other system devices. In the following chapter, a residential energy storage deployment using an agent-based architecture will be presented and discussed.

Chapter 5 - Energy Storage Integration Demonstrations

A significant hurdle currently facing second-life battery systems is integration cost. Battery systems must be graded and integrated with a new BMS system before they can be sold as a reliable second-life system. The ability to integrate a secondary-use system with a variety of PECs without changing the BMS or communications method is important to reducing integration cost. In general, a flexible unit capable of working with many PEC platforms provides more value than a unit that is tightly integrated with one or two PEC manufacturers.

In Chapter 4, a BMS agent was described which provides a universal interface between a converter intelligence agent and the BMS device which is measuring and collecting data from the battery system. It was determined that the battery could be safely monitored and controlled using a short list of critical variables such as SOC, voltage, and current. To demonstrate the feasibility of this method, this chapter will present two residential energy storage deployments using three different battery devices. For each battery device, the method used to interface between the BMS and system controller will be described and evaluated. A successful demonstration of the agent-based BMS will allow energy storage elements with different BMS units to interface with an agent-based inverter without significant changes to the software operating the energy storage system. This is an important cost-reducing measure that will enable more widespread use of secondary-use batteries.

Direct BMS Integration

Traditionally, communications between the battery management system and inverter of an energy storage system are tightly integrated between the devices. The device (DSP, FPGA) controlling the high-speed switching of the inverter also communicates with the BMS and

manages start-up and shut-down sequencing for the ESS. This means that energy storage solutions are often created to only work with a limited number of power conditioning systems.

For example, the Tesla Powerwall residential ESS is only sold with a supported "Backup Gateway" from Tesla which provides the power management and conditioning for the attached batteries [61]. Similarly, the RESU system from LG Chem has a limited number of PCS devices with which it is compatible for residential energy use [62]. As system size increases, the level of integration required between the PCS and the energy storage element tends to increase as well. Therefore, switching to a different power stage from a separate manufacturer becomes more difficult.

AMIE Residential Energy Storage

Based on the industry model, a residential energy storage (RES) system was created for the Advanced Manufactured Integrated Energy (AMIE) demonstration at ORNL. In this system, a selection of batteries from two FIAT 500E electric vehicles were assembled into an air-cooled battery with a commercial off-the-shelf (COTS) BMS as shown in Figure 5.1(a). A stack of 104 cells from the FIAT battery systems was used since the nominal cell voltage at 50% SOC is ~3.95 Vdc. The number of batteries was chosen such that the nominal pack voltage is ~410 Vdc, meaning that an additional DC/DC converter would not be needed before the DC/AC inverter (as discussed in Chapter 2.3). Assuming a capacity near 60 Ah based on the original specifications and low mileage on these EV packs, this gives the assembled pack a capacity of ~ 24 kWh. Data acquisition and state estimation for the pack is provided by the Orion BMS shown in Figure 5.1(b), which measures pack and cell voltages, current, and temperatures. Based on this information, the BMS provides an SOC and available capacity estimate. It uses the controller area network (CAN) communications protocol for transmitting this data and receiving



(a) Array of FIAT 500E batteries



(b) The BMS in the battery container

Figure 5.1 - Photos of the AMIE Second-Life Battery System

commands. The commercial-off-the-shelf (COTS) Orion BMS from Ewert Energy Systems was chosen based on recommendations from the battery supplier, Speirs New Technologies (SNT).

Power conditioning for the AMIE RES is performed in a second container and consists of an AgileSwitch 30 kW power stage, capacitors for DC filtering, inductors for AC filtering, and a National Instruments cRIO 9039. This FPGA-based device has been configured with the necessary analog and digital control modules to perform the inverter controls as well as high-level device coordination duties. Contactors to isolate the power stage from the battery and the grid are also included in order to allow the PCS to protect itself in the event of a fault on either the grid or battery side. Figure 5.2 shows a simplified electrical diagram of the system. Control of the contactors, as well as the closed-loop inverter controls providing PWM signals to the power stage is handled by the cRIO 9039 device located inside the metal cage seen in the front-view of the inverter container shown in Figure 5.3. Communications to the BMS are also handled by the cRIO controller in addition to externally interfacing with the home energy

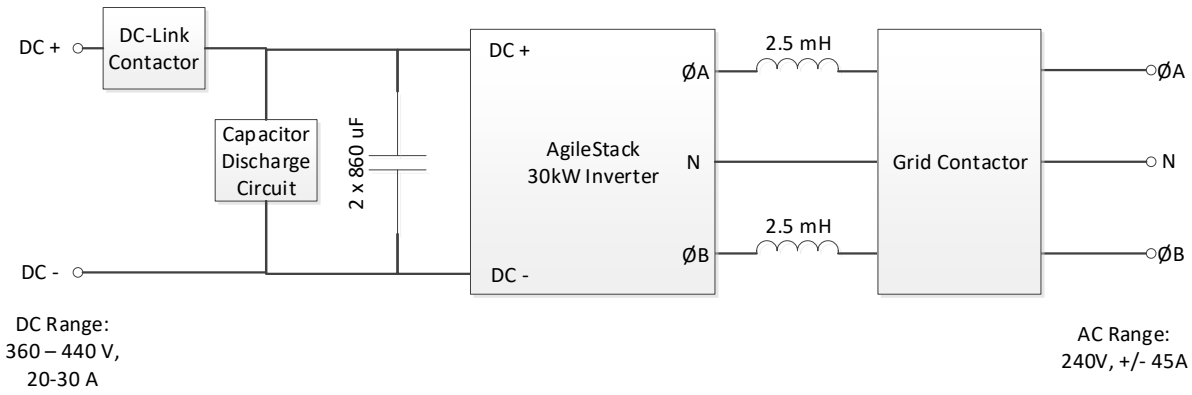


Figure 5.2 - Simplified electrical diagram of the PCS in AMIE.

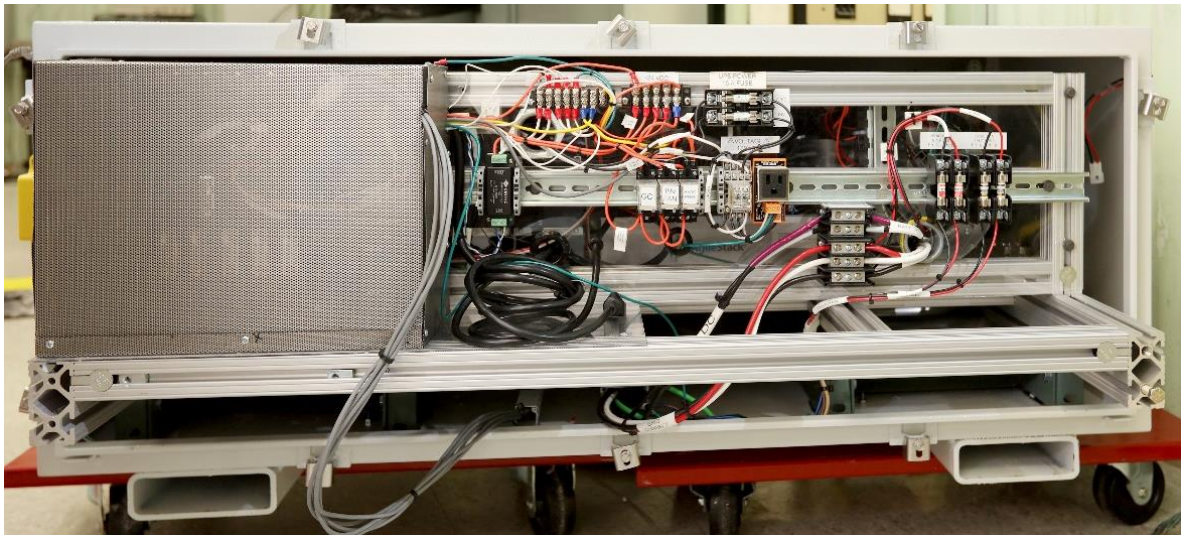


Figure 5.3 - Front-view of the AMIE power conditioning container

management system (HEMS) that handles dispatching the battery. The power stage, filters, and other high-voltage electrical components are located behind plexiglass to minimize exposure to hazardous voltage.

Figure 5.4 shows a high-level diagram of the AMIE energy storage system. Communications interconnections are shown in the diagram to specify how the system communicates between devices. Note that each device must communicate with the field programmable gate array (FPGA) on the cRIO, which runs a single piece of code at high speed. When integrating the battery and power stage together, CAN decoding for interpreting data from the BMS had to be built directly into the high-speed inverter controller. The CAN addresses were known and specifically monitored by the FPGA on the cRIO. When new data arrives, it is internally published to the real-time (RT) computer co-located on the cRIO-9039. The real-time layer of the computer monitors the data from the BMS as well as its own sensors for error

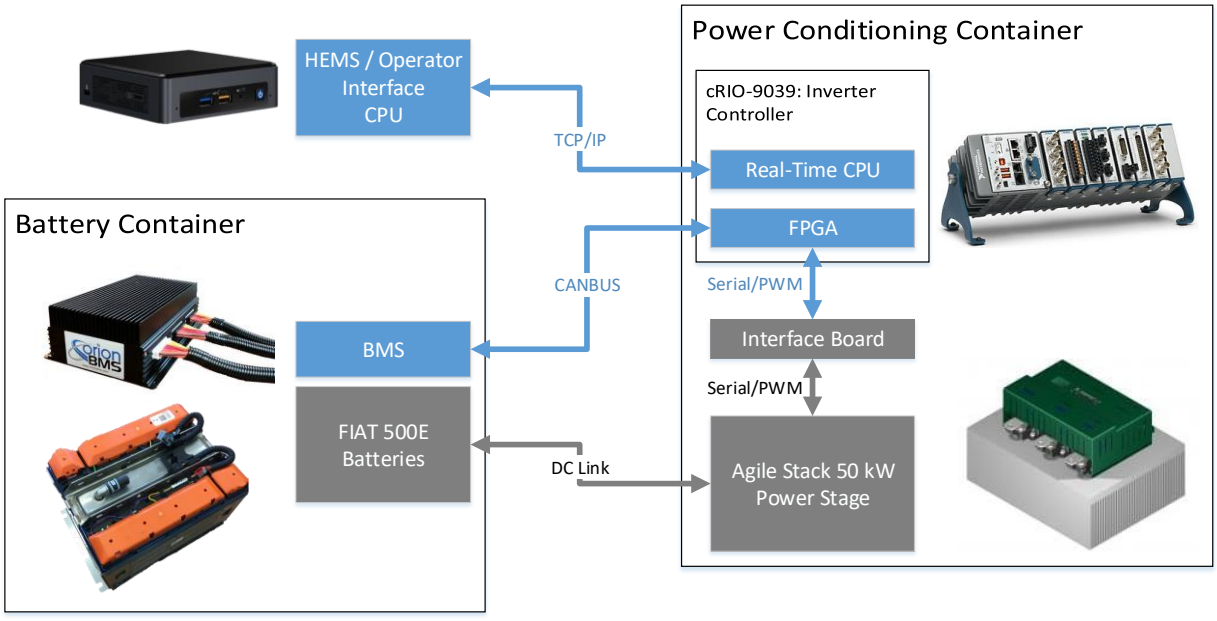


Figure 5.4 - High-level diagram of the control devices in the AMIE ES

conditions in the battery and the PCS. The FPGA also provides the closed-loop controls that execute the switching on the power stage.

When the battery needs to be replaced, the BMS would have to be the same as the current configuration for the integration to be seamless. Should the new BMS have data located at different CAN addresses, the FPGA code (which also executes the closed-loop inverter controls) would have to be updated and re-verified as functional. If the communications protocol of the BMS were to be updated to ModBus or RS-232, hardware would have to change as well since CAN required the installation of a special module in the cRIO-9039. All these steps represent significant integration time increases.

Lessons Learned from Direct Integration

The secondary-use residential energy storage system described above was successfully operated as part of the AMIE demonstration. Figure 5.5 shows the unit installed under the "back porch" of the residence. The battery unit functioned as expected for several demonstrations, but over time experienced significant voltage deviations (shown in Figure 5.6) that the BMS could not overcome with traditional balancing. This occurred because the batteries had been integrated into the system without being graded or tested meaning that significant differences in battery health were present that were untested. In Chapter 1, the importance of matching battery health and capacity as part of the secondary-use integration process was discussed. Over time, the differences in health meant that weaker cells were more heavily taxed by charge-discharge cycling and degraded quickly, exacerbating the cell voltage disparities. At first, faulty battery modules were replaced, but this meant replacing five or six cells every time an individual cell went bad because of the FIAT module configuration. Eventually, replacing



Figure 5.5 - AMIE ES installed at the evaluation site [74]

Live Cell Data

Highest Resistance:	0.70	[001]	**Highest Cell Volt:**	4.055	[086]	**Pack SOC:**	95.0%	○ Internal Resistances
Lowest Resistance:	0.70	[001]	**Lowest Cell Volt:**	3.841	[045]	**Pack Current:**	0.0A	◉ Live Cell Voltages
Avg Cell Resistance:	0.70		**Avg Cell Voltage:**	3.916		**Pack Voltage:**	407.2V	○ Open Cell Voltage
Delta Cell Resistance:	0.00		**Delta Cell Voltage:**	0.214		**Current Limits:**	+198A / -33A	
						Relay Status:	Discharge:OFF, Charge:OFF, Safety:ON	

Cell 1	Cell 2	Cell 3	Cell 4	Cell 5	Cell 6	Cell 7	Cell 8	Cell 9	Cell 10
3.971	3.974	3.972	3.970	3.977	3.973	3.977	3.979	3.977	3.974
3.976	3.982	3.972	3.982	3.979	3.976	3.982	3.979	3.976	3.978
3.973	3.972	3.978	3.984	3.969	3.974	3.971	3.975	3.975	3.974
3.969	3.976	3.975	3.975	3.974	3.968	3.849	3.851	3.848	3.849
3.850	3.851	3.846	3.846	3.841	3.850	3.842	3.856	3.848	3.851
3.861	3.849	3.850	3.859	3.850	3.849	3.848	3.848	3.849	3.848
3.846	3.850	3.848	3.848	3.850	3.848	3.848	3.849	3.849	3.849
3.847	3.849	3.971	3.973	3.980	3.975	3.980	3.980	3.850	3.848
3.848	3.849	3.848	3.849	4.042	4.055	4.053	4.048	4.051	3.973
3.973	3.980	3.976	3.976	0	0	3.851	3.850	3.850	3.850
3.852	3.850	3.852	3.850	3.852	3.850	0	0	0	0

Figure 5.6 - Cell voltage and balancing data from the Orion BMS

modules became too time consuming to keep the unit running and replacing the battery unit seemed to be the more viable solution.

When investigating options to replace the batteries of the ESS with a commercial storage solution from SNT or another vendor, it became apparent that the tight integration between the BMS and the inverter was a detriment. Changing the number of cells in the system or re-addressing where data appears on the CAN protocol requires software modifications to the base code. If a different communications protocol were to be used, changes to the hardware could be required which would result in costly configuration changes. For example, the cRIO platform used in the AMIE ES demonstration costs approximately $10,000 with the CAN module contributing $700 to the cost.

To reduce integration time, improve system flexibility, and reduce cost, three features are needed. First, to allow the system to interact with a range of BMS options, a program communicating with the BMS should run separately and asynchronously with the power stage controls. This allows programming alterations to be made without impacting the inverter controls. Next, changing the communications protocol should not require significant hardware changes to the system that increase cost. Using a platform which requires specialized hardware interfaces contributes to cost and increases integration time. Finally, changing the BMS that is controlling the battery should not result in having to change any safety features of the power conditioning system. Re-certifying that all the safety systems in the PCS operate with a new battery system adds significant time to the integration process.

Agent-Based BMS Integration

Based on the lessons learned from the AMIE ES integration project, an agent-based energy storage was developed for the Grid Modernization Laboratory Consortium (GMLC)

residential energy storage project [59] [60]. The goal of this project was the creation of a low-cost residential energy storage solution which provides significant financial benefit to the homeowner when coupled with a residential rate structure that charges for on and off-peak energy usage. Like the design of the CSC and ESMC agent frameworks, the GMLC residential energy storage uses an agent platform to support integration with various hardware to lower the overall system cost and integration time. Chapter 4 provides a detailed description of an agent-based architecture for integrating a wide range of PECs and BMS units. The following is a description of the agent-based system constructed for the GMLC project and its use as a testbed for verifying the principles of the BMS agent with real battery systems.

Agent-based RES for BMS Integration Testing

Figure 5.7 shows an overview of the agent-based system used in the project as well as the hardware interfaces. The integration and interoperability of a variety of BMS units was tested using the PCS and agent platform developed for the GMLC as a testbed. Each component of the energy storage system shown below was designed, constructed, and tested with the assistance of the Electric Energy Systems Integration team at ORNL.

Power Conditioning System

Due to the lack of commercially available 10 kW bidirectional power stages, the ORNL power electronics team designed and constructed a custom unit for this project shown in Figure 5.8. It consists of a 240 Vac single phase power stage using IGBTs with an L-C input filter and a large inductor on the AC output. With forced-air cooling, the power stage is capable of more than 25 kW but has been derated to 10 kW for the purposes of this project. Current and voltage measurements are read by a Texas Instruments 28377 digital signal processor (DSP) which also provides closed-loop current control and generates PWM signals. An interface board between

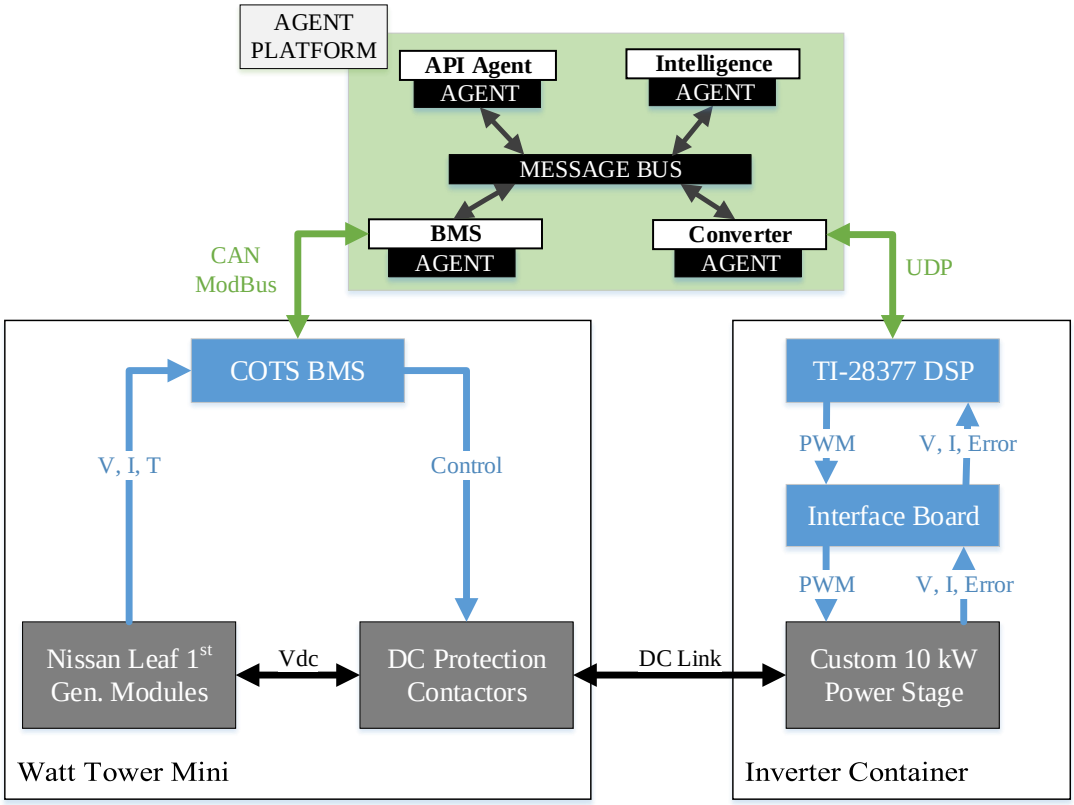


Figure 5.7 - System diagram of the agent-based RES

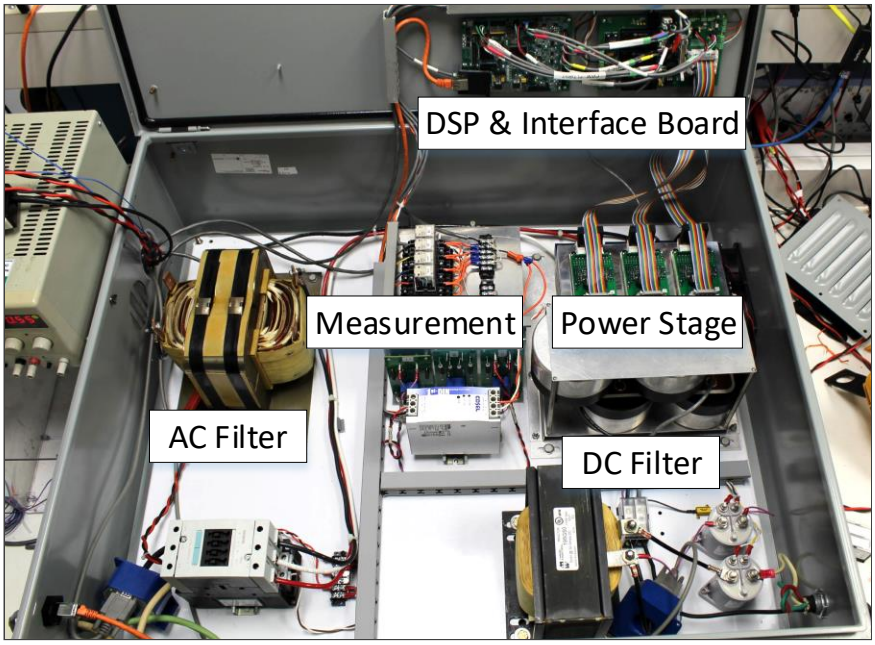


Figure 5.8 - Power conditioning for the GMLC RES project

the power stage and the DSP is used to condition and isolate the measurement and PWM signals. This is done to properly scale the signals so that a variety of power stages can be used instead of just the custom unit designed for this project. The DSP communicates with the converter agent of the agent-framework using the User Datagram Protocol (UDP) for fast data communications. As Figure 5.7 shows, this then communicates with the agent platform where operational decisions are made.

Agent Platform

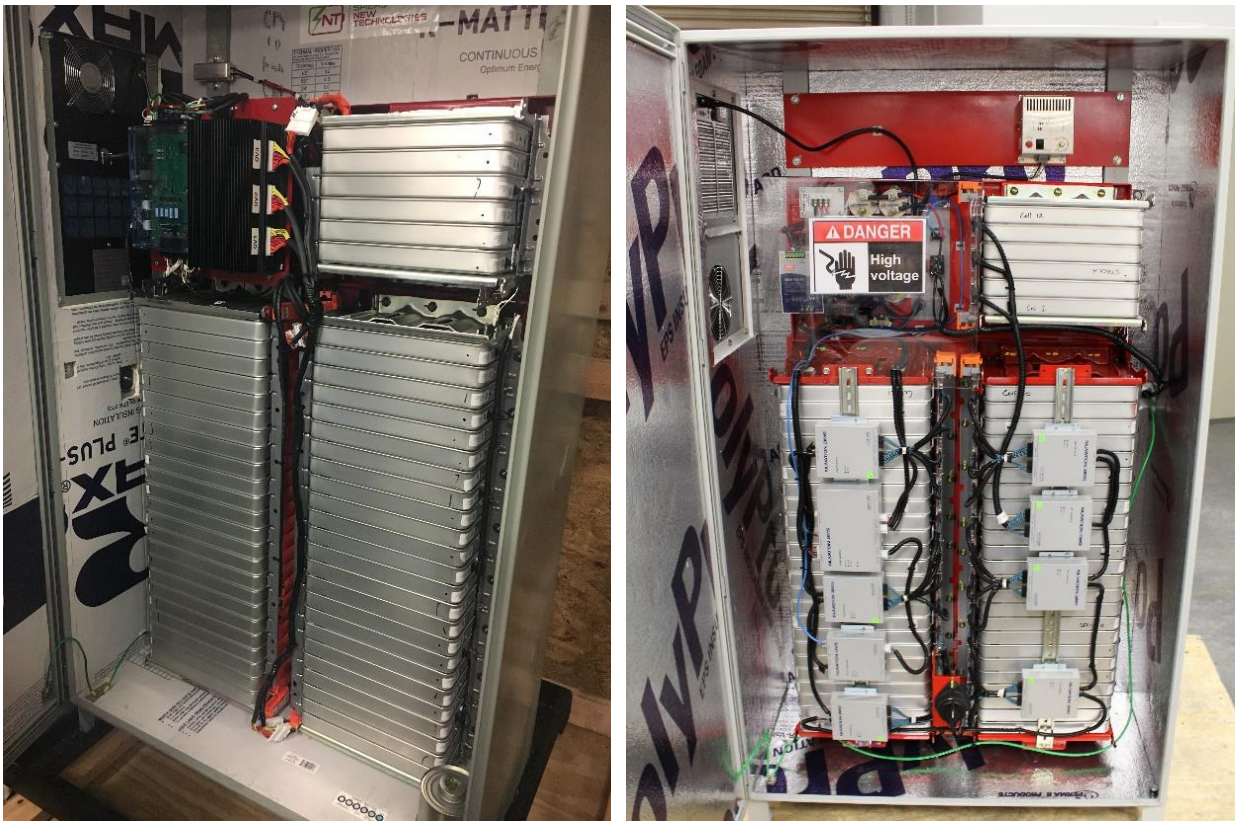
The agent platform is similar in design to the Energy Storage Module Controller (ESMC) from Chapter 4 and auto-launches upon powerup of the small form factor computer running Ubuntu Linux version 16.4.10. Key differences between the frameworks include the lack of a pricing agent and the use of communications with actual hardware for both the converter and BMS agents instead of hardware models. A pricing agent is not needed since the dispatch of this unit is determined externally by a remotely located master controller and sent to the unit through the API agent. This API agent collects operational data on the RES, which it sends to the master controller for monitoring, data storage, and control. When new commands are posted by the master controller, the API agent downloads the operating mode and dispatch command and updates the control state of the intelligence agent accordingly. While this is like the intra-system communications described for the multi-chemistry BES, it includes extra layers of security and handshaking between the API agent and the server to ensure that the communications cannot be maliciously interrupted or intercepted. Like the intelligence agent in the ESMC, the intelligence agent in the RES platform coordinates the start-up, normal, and shutdown actions of the other agents. The final agent shown in Figure 5.7 is the BMS agent, which is identical to the one

outlined for the ESMC in Chapter 4, except that it has been configured to also interact with the energy storage systems described below and not just a battery model.

Secondary-Use Residential Energy Storage Systems

Two secondary-use battery units were tested during the GMLC RES project. Both units use 54 Nissan Leaf Generation 1 modules for a total measurable cell count of 108. Like the AMIE ES, this gives both units a nominal pack voltage at 50% SOC of ~410 Vdc. Both units were certified by SNT to contain 38 Ah which gives a useable storage capacity of 14 kWh. These specifications were primarily chosen based on battery module availability from SNT and the desired pack voltage to avoid the use of a separate DC/DC converter to boost the voltage.

The first unit, named the Watt Mini 1.5, is shown in Figure 5.9(a) and uses the same Orion BMS used in the AMIE ES. It is supplemented with a custom Energy Storage System Control Board (ESS-CB) built by SNT to monitor and control the DC contactor states and provide a pre-charging circuit. Pre-charging is performed in most ES units to protect the PCS and batteries from excessive in-rush currents when the DC link is energized. The Orion BMS and ESS-CB uses the CAN communications over RS-485. The second unit, named the Watt Mini 2.0, uses a modular BMS from Nuvation Energy Systems called the Nuvation High Voltage (HV) BMS. This ES, shown in Figure 5.9(b), can communicate using the CAN protocol. However, Modbus-TCP was chosen for increased communications reliability in an environment prone to electromagnetic interference (EMI) from the PCS. SNT also transitioned to this BMS because it includes DC contactor control, eliminating the need for the custom ESS-CB. It also has Underwriters' Laboratory (UL) approval, which is necessary for the final residential deployment.



(a) (b)

Figure 5.9 - Two RES units using different COTS BMS units

Testbed Implementation

The agent platform, PCS, and battery systems described above provided the components for testing the ability to quickly integrate and change BMS units. However, this still requires a grid connection to test bidirectional energy transfer. Since verifying the state-machine operation of the BMS agent requires intentionally faulting the system, a safer option is to use a four-quadrant grid emulation tool like the Ametek RS90 instead of an actual grid connection. The implemented testbed is shown in Figure 5.10. Since the PCS of this system was designed for 240 Vac single-phase operation, the grid emulator was programmed to behave like this type of grid connection. Later iterations of the testbed included metering of the AC connection for 3rd party measurement and verification, but data presented in the testing and results section below

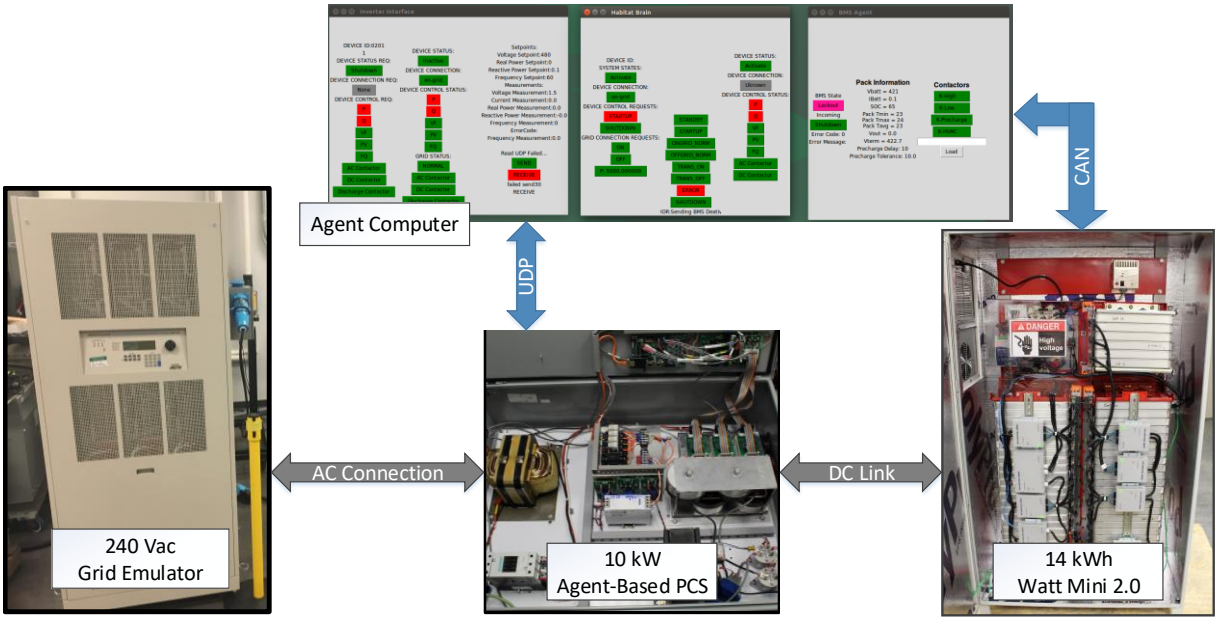


Figure 5.10 - The RES testbed used for BMS testing

was collected by the inverter and verified using a Tektronix Oscilloscope paired with a Yokogawa Power Analyzer.

BMS Agent Implementation

During the discussion of the ESMC BMS Agent in Chapter 4, the details of a state machine for controlling the start-up, operation, and shut-down procedures are described. The same agent and state machine are used to communicate with and control the BMS hardware used in the two batteries shown in Figure 5.9(a-b). Using the BMS agent to interface with real hardware requires switching the source of the battery data from the internal battery model to the communications protocol and addresses of a chosen system. This is done in a configuration file that is read when the BMS agent is first launched during the commissioning of the RES system. Figure 5.11 shows an example configuration file which has been set-up to launch the BMS agent to communicate with the Nuvation High Voltage BMS via Modbus-TCP. Specifically, this file

```
{
        "agentId": "BMS",
        "BMS_Manufacturer": "Nuvation",
        "BMS_Model": "MESA",
        "Communication_Medium": "Modbus TCP",
        "Communication_Settings":{
                "CAN_Baud_Rate": "250 kbps",
                "Modbus_IP": "192.168.53.75",
                "Modbus_Port": "502"
        },
        "Pack_Data":{
                "DC_Voltage_Max": "443",
                "DC_Voltage_Min": "356",
                "Temperature_Max": "45",
                "Temperature_Min": "0",
                "SOC_Max": "95",
                "SOC_Min": "5",
                "Capacity":"16000",
                "PowerRating":"10000"},
        "Operational":{
                "SOC_Max": "80",
                "SOC_Min": "20"
        }
}
```


Figure 5.11 - Example BMS agent configuration file for the Watt Mini 2.0

has been configured to communicate with the Watt Mini 2.0 in the testbed described above. While the BMS agent operates, this information is used to inform decisions about the locations and methods used to read data from the battery and where commands are sent to start and stop power transfer from the Watt Mini 2.0.

For example, when the BMS agent first starts the core function that reads data from a real BMS or the battery model, the configuration file is checked. If it is a known BMS, then the agent attempts to open communications with the device. If opening the communications channel fails, or the type of BMS is not yet supported by the BMS agent, then an error state is triggered, and an error message is displayed by the agent. In the case of the Nuvation Energy BMS, the communications core opens a Modbus-TCP connection to the IP address specified in the configuration file and begins reading the BMS data at regular intervals. Similarly, when the intelligence agent of the RES signals the BMS agent to transition to “start-up,” the state-engine

performs BMS-dependent actions. After checking which BMS is being used, the agent sends the proper start-up command over Modbus-TCP or CANBUS and then waits for the acknowledgement command to be read back. If completed successfully, the BMS agent moves to the "normal" state. If the BMS and battery system fail to acknowledge the command and energize, then the process is attempted two more times before the BMS agent transitions to an "error" state and signals the RES intelligence agent that a problem has occurred with the battery system.

To integrate a new BMS, the communications protocol (data addresses and types) must be added to the BMS agent under the communications core as a new BMS type. The user manual for the Nuvation HV BMS specifies the communications addresses for each of the 13 parameters identified in Chapter 4. Each piece of necessary data is matched with data type, address, and any scaling factors required to retrieve the properly formatted data. This information is then paired with the variable that is used in the decision making process within the agent. An example of the portion of code which provides this functionality for the Nuvation BMS using Modbus-TCP is shown in Figure 5.12. The internal agent variable (in this case *self.BMS.Status.Measurements.DC_Voltage*) is matched with the result of retrieving certain data from the BMS. In this case, the MESA-ESS[7] standard is being used, which provides a base address for certain data sets followed by an offset to retrieve specific pieces of data. For CAN data frames, slightly more processing is involved, but the retrieved information is stored in the same variable name. This means that a single BMS agent can only communicate with one BMS at a time, but multiple BMS agents can be launched in parallel as necessary.

[7] MESA was formerly an acronym when it was part of the Sun-Spec Alliance, but is now the name of the organization spearheading open standards for energy storage communications [73].

```
elif self.Communication_Medium == 'Modbus TCP':
    self.MESAOneBase = 40070
    self.MESATwoBase = 40094
    self.MESAThreeBase = 40116
    try:
        self.BMS.Status.Measurements.DC_Voltage = str(self.ModbusRead(self.MESATwoBase+11, 'uint16') * self.ModbusRead(self.MESATwoBase+19, 'sunsf'))
        self.BMS.Status.Measurements.DC_Current = str(self.ModbusRead(self.MESAThreeBase+11, 'int16') * self.ModbusRead(self.MESAThreeBase+16, 'sunsf'))
        self.BMS.Status.Measurements.CCL = str(self.ModbusRead(self.MESATwoBase+12, 'int16') * self.ModbusRead(self.MESATwoBase+20, 'sunsf'))
        self.BMS.Status.Measurements.DCL = str(self.ModbusRead(self.MESATwoBase+13, 'int16') * self.ModbusRead(self.MESATwoBase+20, 'sunsf'))
        self.BMS.Status.Measurements.Pack_Temps_Max = str(self.ModbusRead(self.MESAThreeBase+7, 'int16') * self.ModbusRead(self.MESAThreeBase+15, 'sunsf'
        self.BMS.Status.Measurements.Pack_Temps_Min = str(self.ModbusRead(self.MESAThreeBase+9, 'int16') * self.ModbusRead(self.MESAThreeBase+15, 'sunsf'
        self.BMS.Status.Measurements.Pack_Temps_Avg = str((int(self.BMS.Status.Measurements.Pack_Temps_Max) + int(self.BMS.Status.Measurements.Pack_Temps
        self.BMS.Status.Measurements.SOC = str(self.ModbusRead(self.MESAOneBase+11, 'uint16') * self.ModbusRead(self.MESAOneBase+22, 'sunsf'))
        self.BMS.Status.Measurements.Cell_Vmax = str(self.ModbusRead(self.MESAThreeBase+3, 'uint16') * self.ModbusRead(self.MESAThreeBase+14, 'sunsf'))
        self.BMS.Status.Measurements.Cell_Vmin = str(self.ModbusRead(self.MESAThreeBase+5, 'uint16') * self.ModbusRead(self.MESAThreeBase+14, 'sunsf'))
        self.BMS.Status.Measurements.Cell_Vavg = 'PlaceHolder'
        self.BMS.Status.Measurements.ActiveState = str(self.ModbusRead(self.MESATwoBase+16, 'uint16'))
```

Figure 5.12 - Example pairing code between BMS and Agent variables

Once the process of matching data to BMS agent variables is completed, switching between different BMS hardware or the battery model only requires changing the model and communication mode specified in configuration file. An additional benefit to this method allows the SOC, voltage, and capacity limits to be set without significant programmatic changes as well, so that batteries of various sizes and configurations can all be used with this interface.

BMS Testing and Results

Preliminary testing demonstrated the ability of the BMS agent to communicate with the two different battery units used in the RES testbed described above, two identical BMS agents were launched simultaneously, but configured to communicate with either the Orion BMS via CAN or the Nuvation BMS via Modbus-TCP. In Figure 5.13, the two BMS agents running side-by-side can be seen. Comparing the two agents, there are no noticeable differences in the data available to the agent-based RES system. From the perspective of the intelligence agent of the RES, there is no difference between the batteries except for slight changes in nameplate capacity, with the Watt Mini 1.5 being listed as 16 kWh instead of 14 kWh in order to account for slight miscalculations by the BMS in estimating SOC. Each unit can be treated identically during the regular operational sequences.

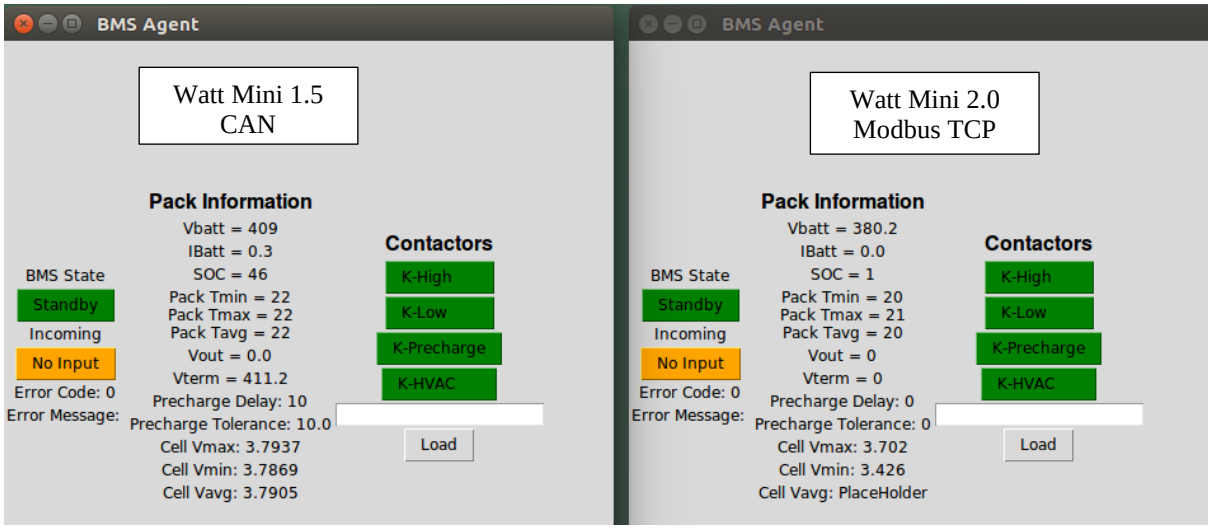


Figure 5.13 - Demonstration of two BMS agents communicating with two BMS units

Both the Watt Mini 1.5 and 2.0 systems communicated successfully with the agents with minor errors to the data structures resulting in misread values at first. However, these were easily corrected since both battery units provided other means to verify that the values being communicated were correct. One problem encountered on the Watt Mini 2.0 was the SOC estimation of the pack. At a pack voltage of 380 Vdc, the state estimation algorithm estimated a 1% state of charge. This meant that the agent would error upon launch, lockout, and refuse to come out of an error state. To avoid this problem, the minimum allowable SOC had to be lowered in the BMS configuration file to 0%. This allows the system to launch and come out of the error state so that it can at least begin to charge.

Watt Mini 2.0 System Failure

While the Watt Mini 2.0 successfully integrated with the BMS agent, power transfer tests were never successful. The Watt Mini 2.0 uses a modular BMS from Nuvation Energy which has a proprietary communications bus between the various data acquisition modules and the master controller. If too many errors are accumulated on this bus during operation, the BMS will

open the battery contactors to interrupt energy transfer since it assumes the battery is experiencing possible problems. While power transfer was not successful, these failed tests did help verify the error behavior of the BMS agent's state engine. When the Nuvation would abort power transfer, the BMS agent recognized that, despite being in the "Normal" state where they should be closed, the K-High and K-Low contactors were open. This immediately sends the BMS agent to the "Error" state. Upon receiving an error from the BMS agent, the remainder of the agent-platform from the testbed performs the necessary actions to quickly and safely shut down the system. This includes having the PCS open both DC and AC contactors and informing the master controller that an error has occurred in the battery via the API agent.

Fault Mode Effects Analysis (FMEA) Testing

Developing a robust agent-based system for deployment means that each failure point must be reinforced to prevent damaging or catastrophic failures from occurring. As discussed above, system testing involves intentionally and unintentionally faulting the system to observe the behavior and inspecting the results. Performing a Fault Mode Effects Analysis (FMEA) on the BMS agent system involves analyzing each of these failure points and determining the frequency and severity at which they can occur. Before significant power transfer testing or deployment, this analysis is used to develop a series of tests to stress the system and ensure that safety measures operate and prevent dangerous faults from occurring.

In Table 5.1, an example of the FMEA performed on the interconnection between the Watt Mini 1.5 and the agent-framework is given. As the chart shows, each failure point is given a severity and frequency grade, which in turn denotes the overall fault severity as a color. This is based off of an Occupational Health and Safety Administration (OSHA) chart show in Table 5.2.

Table 5.1 - Sample of cases from the FMEA of the RES

Failure	Causes	Consequences	F	S	Risk	Safeguards
HVDC Unplugged (while live)	User error, cable stress	DC link broken, possible arc in connector	U	L	UL	Inv agent detects loss of voltage, opens DC contactor, shuts BMS agent down (opens battery contactors)
HVDC cable severed (DC short circuit)	Accident, cable stress	DC link broken, exposed HVDC, short-circuit current of battery is 2.5 kA	R	H	RH	Inv. Agent detects loss of voltage, aborts power transfer. Battery HVDC is fused at 60A, fuses blow, eliminating HVDC
CAN cable becomes disconnected (while live)	Accident, cable stress, user error	BMS and BMS agent lose communications, cannot send signal to shut down battery	R	L	RL	BMS agent detects loss of comms, enters error state. Intelligence agent aborts power transfer, opens contactors, will not restart.
CAN cable becomes disconnected (while idle)	Accident, cable stress, user error	BMS and BMS agent lose communications, cannot send start-up signal.	R	L	RL	BMS agent detects loss of comms, enters error state. Intelligence agent cannot execute start-up sequence, system will not energize
SOC value outside safe range	Battery drift, parasitics, accident	Battery could overcharge or undercharge, creating permanent damage to the battery	U	M	UM	BMS agent faults, issuing shut-down command to both the BMS and the intelligence agent. System de-energizes
DC Voltage outside safe range (battery V)	Grid event, battery problem	Battery could overcharge or undercharge, creating permanent damage to the battery	U	M	UM	BMS agent faults, issuing shut-down command to both the BMS and the intelligence agent. System de-energizes

Safeguards taken to prevent this failure method, as well as the software actions that are automatically executed are shown in the final column of Table 5.1. While the Frequency and Severity ratings are qualitative values, they help provide insight to how dangerous a fault is perceived to be.

Results from the FMEA are used to add protections or safety mechanisms if the risk factor falls in the orange or red areas of Table 5.2. They are also used to develop tests to ensure that the safety mechanisms respond as expected in the case of a fault. For example, the first entry in Table 5.1 was tested by pressing the emergency stop (E-Stop) on the Watt Mini 1.5 while the power conditioning system was discharging the system at 1 kW. Pressing the E-Stop

Table 5.2 - OSHA hazard determination table

Freq. / Risk	Rare (R)	Unlikely (U)	Occasional (O)	Frequent (F)
Very Low (V)	VR	VU	VO	VF
Low (L)	LR	LU	LO	LF
Medium (M)	MR	MU	MO	MF
High (H)	HR	HU	HO	HF

breaks the DC connection between the battery and the inverter, thus emulating a cable break between the devices. This test was repeated while the system was in standby (no power, de-energized), idle (no power, energized), and finally while the system was charging and discharging at a rate of 1 kW. Table 5.3 shows a sample of tests that were performed on the RES test system to verify fault mode operation before it was deployed. Tests included on this chart were for faults which effect the battery and BMS agent. Additional testing was performed to verify the error behavior of the power conditioning system and API connectivity.

Table 5.3 - Sample of experiments used to test fault modes of the battery system

Test ID #	Fault Under Test	Safe Faulting Procedure	Expected Result	Pass/Fail
1.1	CAN loss while system in standby	Before system energization, remove CAN cable. Attempt system start-up.	System should attempt start-up. When contactors do not close. Start-up fails	Pass
1.2	CAN Loss while system is idle	Attempt system start-up. Upon energization, remove CAN cable	Start-up will complete successfully. 30 s after cable is remove, system should error and shutdown	Pass
1.3	CAN loss while system is charging at 1 kW	Attempt system start-up. Send -1 kW set-point. Once power is stable, remove CAN cable	Start-up will complete successfully. 30 s after cable is remove, system should error and shutdown	Pass
1.4	CAN loss while system is discharging at 1 kW	Attempt system start-up. Send 1 kW set-point. Once power is stable, remove CAN cable	Start-up will complete successfully. 30 s after cable is remove, system should error and shutdown	Pass
2.1	Removal/severing of HVDC cable while system in standby	Before system energization, press battery E-stop. Attempt system start-up.	System should attempt start-up. When no DC-link voltage is detected. Start-up fails	Pass
2.2	Removal/severing of HVDC cable while system in standby	Attempt system start-up, upon energization, press battery E-Stop.	Start-up will complete successfully. When E-stop is pressed, BMS agent will fault and shut-down the system	Pass
2.3	Removal/severing of HVDC cable while system is charging at 1 KW	Attempt system start-up Send -1 kW set-point. Once power is stable, press battery E-Stop	Upon E-Stop, PCS should fault in ~4ms with a “DC overvoltage” error and shutdown safely DC link should not be greater than 500 Vdc on oscilloscope at any point during fault	Pass
2.4	Removal/severing of HVDC cable while system is discharging at 1 KW	Attempt system start-up Send 1 kW set-point. Once power is stable, press battery E-Stop	Upon E-Stop, PCS should fault in ~4ms with a “DC undervoltage” error and shutdown safely	Pass

Power Transfer Testing

To verify the successful integration of the BMS agent with the rest of the agent-based energy storage platform, short and long-term testing was performed on the RES testbed. The first tests consisted of charge cycling the Watt Mini 1.5 battery at various rates to confirm the stability of the agents and their interactions. One such test is shown in Figure 5.14 where the battery was fully charge cycled. The jumps in SOC are where the BMS performed a "drift-point adjustment" to the SOC calculation based on its internal algorithm.

Several times during this series of tests, the grid emulator experienced internal errors resulting in the loss of the grid as measured by the PCS. This inadvertently ended up testing the error capability of the RES system in the same way the failures on the Watt Mini 2.0 tested the

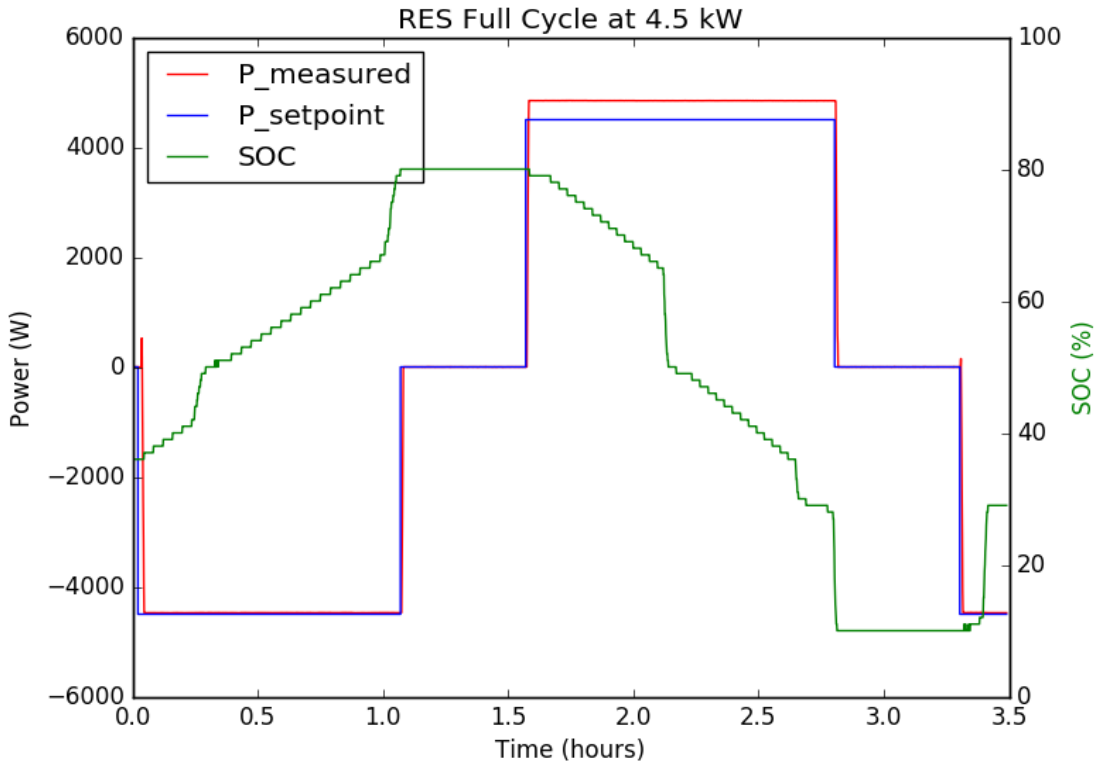


Figure 5.14 - Example of ES full cycle at 4.5 kW

error engine of the BMS agent. When the grid voltage drops to zero, the DSP first measures the loss of the grid via its AC measurement and immediately halts power transfer and electrically isolates the PCS by opening the AC and DC contactors on either side of the power stage. An error is simultaneously reported to the converter agent which enters an error state preventing energy transfer from being accidentally restarted before the error is cleared. Upon receipt an error state, the intelligence agent instructs the BMS agent to send the shut-down command to the Watt Mini 1.5 thus completing the error process and safely isolating the system. When the unit was being prepared for deployment, this failure mode was tested several more times and the speed at which the shutdown occurred was measured. On average, the DSP isolated the PCS within 16 ms and the BMS had completely shut down within 4 seconds as measured by an oscilloscope and stopwatch respectively.

Another long-term energy transfer test was to allow the RES to run continuously with the set-point being dispatched remotely over a server from a central optimization service. This test demonstrated that the energy storage can cycle and respond to a continuous stream of power setpoints over a long period of time in a stable manner. Results from one such test are shown in Figure 5.15, where the unit ran autonomously for approximately 8 hours. In this case, the goal of the optimization was to dispatch the RES in such a manner that it would earn the most revenue based on the given utility price signal (shown by a dotted black line).

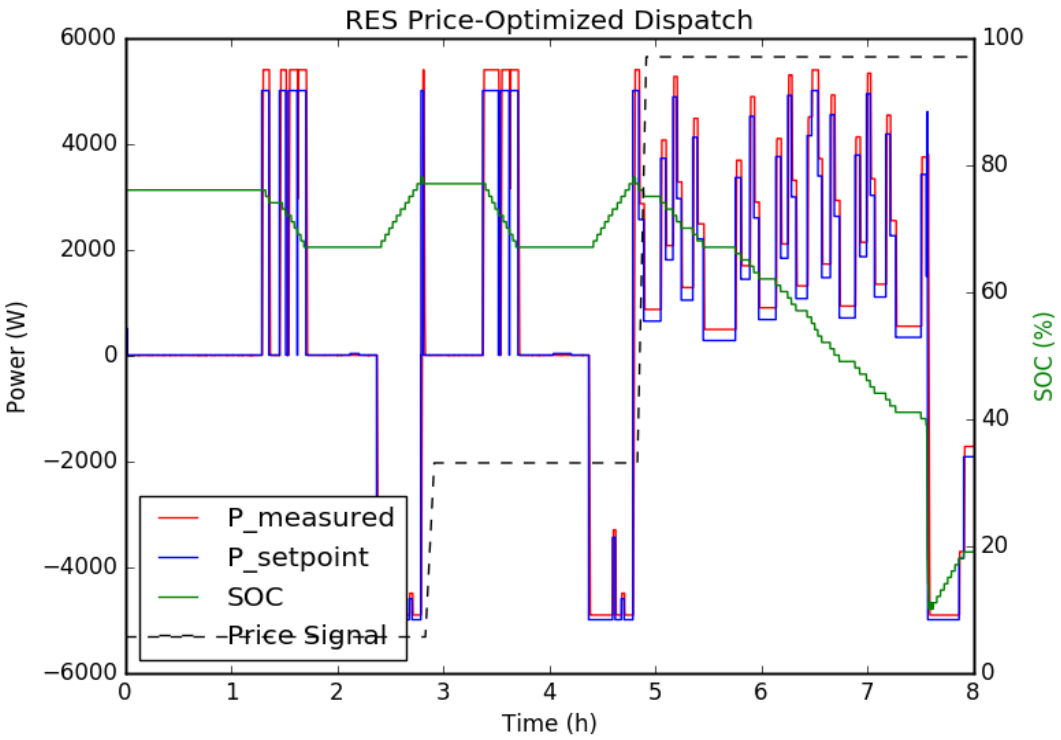


Figure 5.15 - Price optimized testing of the RES

Using a combination of these long-term tests, short-term pulse testing, and fault mode testing, the RES testbed using the BMS agent with the Watt Mini 1.5 was eventually determined to be stable for deployment. The agent-based PCS, battery unit, and agent platform were deployed at a residence in Bear Creek, NC for further testing in a real-world and unpredictable environment. Figure 5.16 shows the deployed Watt Mini 1.5 and PCS on a concrete pad outside of the home. The computer running the agent platform is located inside the PCS container and uses the homeowner's internet service to connect to the API server and master dispatch controller located remotely at SNT in Oklahoma City, OK. A full description of the deployed system, as well as testing results from running net-zero services on the house can be found in [63].



Figure 5.16 - Deployed and tested RES using the BMS agent and agent-based inverter

Discussion and Summary

Using an agent-based architecture to integrate a variety of BMS units significantly increases the speed at which the energy storage element of the residential storage systems could be exchanged. In the first generation produced for the AMIE demonstration platform, having the BMS communications hard coded directly into the inverter controls provided a convenient method for ensuring safe operations of the unit by having the inverter directly monitor and control the battery system. However, the reduced ability to swap for a different battery system from another manufacturer or use different communications protocols prevented this method from being useful on future projects.

Using the agent framework on the GMLC RES project aided in the on-time completion and deployment of the unit. Late in the project, the battery manufacturer changed the BMS being used to monitor and control the system since it eliminated the need for custom manufactured parts that did not carry Underwriters Laboratory (UL) approval. As described above, integration of the new BMS with a different communications protocol took very little time. When the Watt Mini 2.0 battery system experienced power transfer issues shortly before the planned deployment date, switching back to the first unit using the known reliable BMS was a matter of changing the configuration file in order to continue testing and deployment with the Watt Mini 1.5.

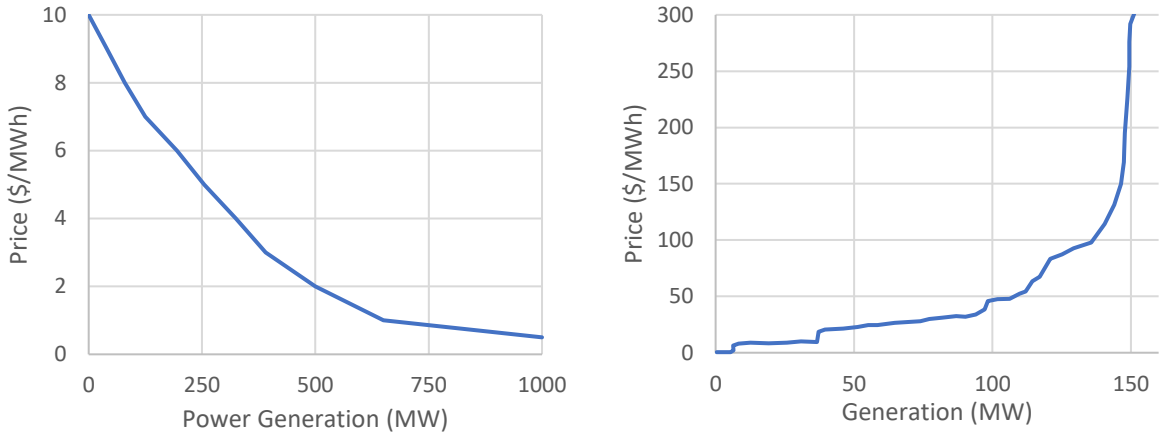
This chapter compares two energy storage projects using different methods to integrate incoming data from battery systems. Use of the agent-based framework provided a robust yet flexible method for ensuring that an energy storage system could reliably communicate with many different BMS options. While these specific examples only contain one energy storage element running through a single-stage power conditioning system, it provides a proof-of-concept for agent-based systems. Since the control software is independent of the size and configuration of the energy storage and converter systems, it can be assumed that these same frameworks can work on a larger scale for the multi-chemistry system described in Chapter 4.

Chapter 6 - Distributed Control through Transactive Optimization

In Chapter 4, the agent-based architecture for controlling the multi-chemistry battery energy storage was described with a focus on the interactions between agents and the control systems. With an agent framework in place that has been demonstrated to be reliable and robust, the next step is to control a multi-converter system as a single unit. A review of the literature in Chapter 2 showed that the traditional method for controlling an energy storage unit with multiple storage elements involves controlling all the converters as a single system. Just like the AMIE RES unit from Chapter 5, this reduces the flexibility and modularity of the system since it requires a tight integration between the converters and a central high-speed controller. This chapter proposes a transactive optimization that allows a central controller to distribute control responsibility to any number of attached energy storage elements within the agent framework. Demonstration of the transactive optimization will be done using the battery and converter models within the agent framework discussed in Chapter 4. Finally, the chapter will close with a discussion of the results.

Transactive Control Overview

Transactive systems are frequently discussed in the literature, from traditional systems used in transmission and distribution, to newer applications proposed to regulate and schedule electric vehicle charging [64] [65]. Most systems follow a producer-consumer model where energy producers bid into the market with a cost ($/kWh) versus power (kW) curve. For large generation assets, such as a nuclear power station, power is produced most efficiently and at lower per-kW costs when the asset is providing nearly 100% of available power. Thus, a bid is produced much like the one shown in Figure 6.1(a), where the per kilowatt cost of energy becomes cheaper as the plant is asked to produce more energy. Bids from multiple suppliers are



(a) – Generation asset bid curve

(b) – Aggregate cost to meet demand

Figure 6.1 - Traditional price curves for large generation assets [66]

then stacked based on price which produces an aggregate price curve like the one shown in Figure 6.1(b), which is based off of a 2008 generation stack from PJM [66]. Prices shown in the latter increase because generation assets which cost more to activate are called upon when the power demand is higher. The steps seen in the curve are when different producers take over. For instance, the first small area is large renewable farms, followed by nuclear energy, then coal assets, and finally natural gas. This asset order has since changed as natural gas becomes cheaper. The grid manager (usually an Independent System Operator or Real Time Operator) then compares this aggregate price curve to the present demand from the loads on the grid and decides which producing entities to buy energy from to minimize the overall cost.

Utilities and others in the literature have been developing similar transactive systems in which energy consumers such as residential or industrial customers bid-in to the utility with a "willingness to buy" price curve specifically for use with electric vehicles [47] [67]. To spread the impact of electric-vehicle charging, the authors developed a transactive method for sequencing the charging of EVs and PHEVs. Instead of all vehicles charging as soon as they are

connected, they must bid into a market with a "willingness to buy" asymptotical bid function like the one shown in Figure 6.2. The vehicle estimates a time-of-next-departure, at which point the battery must be fully charged. As the time-to-departure (Δt_{dep}) shrinks, the price the EV is willing to buy electricity increases. This occurs asymptotically until the vehicle can no longer achieve a full charge. Each EV submits this bid function to a central "producer" whose willingness to produce energy increases linearly with price. This way, each EV or PHEV will only begin to charge as it approaches the time-of-departure and will charge at higher powers as the departure window closes.

Still other systems have been proposed and implemented which use an optimized price curve based on forecasted energy usage. These forecasts are used to predict magnitude and timing of energy demand, which can then be used to develop a 24 to 48-hour price schedule. If a home energy management system (HEMS) is in use that receives a price signal, it can be used to drive down usage during peak times and shift that usage. Traditionally, energy storage systems

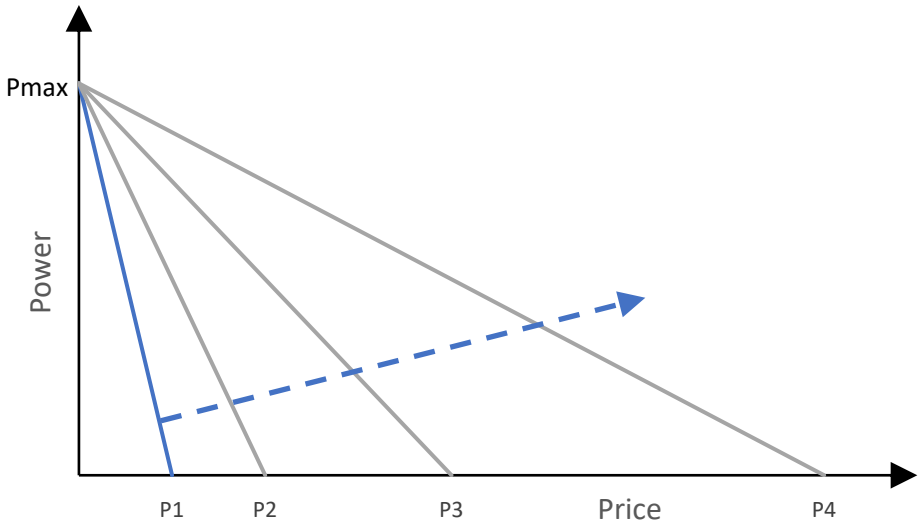


Figure 6.2 - Asymptotical bid function increasing price as departure time nears.

use an implementation of this last method. Energy is stored during "off-peak" price periods and then used to offset load consumption during peak periods. This was demonstrated by the Community Energy Storage (CES) project as well as the GMLC-RES project presented in the previous chapter. Manipulating the price signal provides a pseudo-transactive method for controlling the time and magnitude of net energy consumption from the grid.

For this project, a combination of the first two methods is developed where the ES simultaneously submits bids exhibiting different prices for both a willingness to buy (charge) and price for the remaining energy in the system (discharge). The following section describes the transactive method, price determination algorithm, and an overview of how it is incorporated into the agent framework from Chapter 4.

Distributed Transactive Energy Market

An important aspect to the transactive energy market developed for this project is that energy has a different value during charging and discharging. Both values are calculated as a function of the "remaining charge capacity" of the individual battery system and determined locally by the energy storage. In other methods, the value of the energy being purchased or sold by an energy storage is determined by the utility. The proposed method internalizes the price signals and builds an internal energy market, which means that the value of energy in the battery systems is independent of the external market.

Energy Market Overview

Creating the internal market starts with having each ESMC evaluate the economic value for battery performance based on local BMS and converter data and use this data to develop a bid curve. An example of a bid curve for three energy storage systems is presented in Figure 6.3.

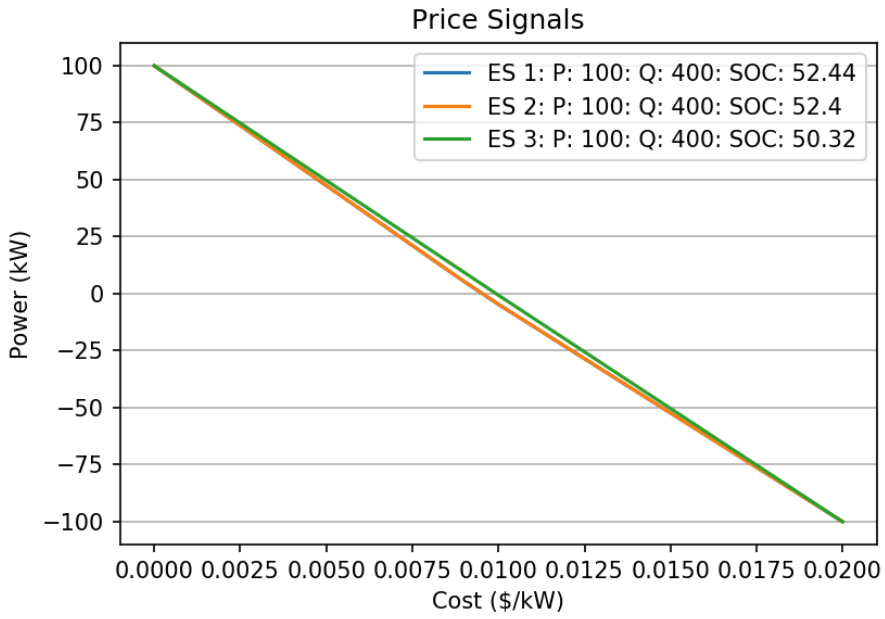


Figure 6.3 - Example price curves for three ES units at ~50% SOC

Each bid curve is shared with the CSC and processed based on the utility dispatch request to generate an optimized dispatch for the entire ESS. In this work, it is assumed that dispatch power requests are sent by a master controlling entity such as a utility distribution energy resource management system (DERMS) or microgrid controller. Based on previous integration experience, this is a good assumption since adding an interface to manipulate other utility signals into a power setpoint is trivial. However, other approaches to generate a schedule for the ESS system can be used and developed within the utility interface agent. The utility interface can also be developed to accommodate any number of different communication protocols.

The price curve is posted to the REpresentational State Transfer (RESTful) API along with a system status indicating whether the attached ESMC is presently active. Once a dispatch request is posted to the CSC message bus from the utility interface agent, the optimization agent retrieves the compiled price signals from the transactive communication agent. A value-based

optimization is performed resulting in a power dispatch level for each attached ESMC. The resulting dispatch signals for each energy storage element are then posted to the API server for retrieval by the ESMCs. The ESMCs act on these signals to create a change in power level or the activation or deactivation of the attached ES. This action, determined by the intelligence, is then translated into setpoints and control commands which are then pushed out to the relevant agents.

Power setpoints are pushed to the DC/DC converters which ramp to the setpoint according to their capabilities. Meanwhile, the DC/AC inverter monitors the DC bus voltage and adjusts power flow to and from the grid to keep the DC link within a set voltage range. This method for power dispatch and DC bus regulation has been well demonstrated for multi-level converters [68] [69]. Finally, measurements are collected by the BMS and converter agents and the process is repeated until the system encounters an error or is otherwise taken offline. Figure 6.4 shows the data movement described above between the ESMC, API, and CSC devices.

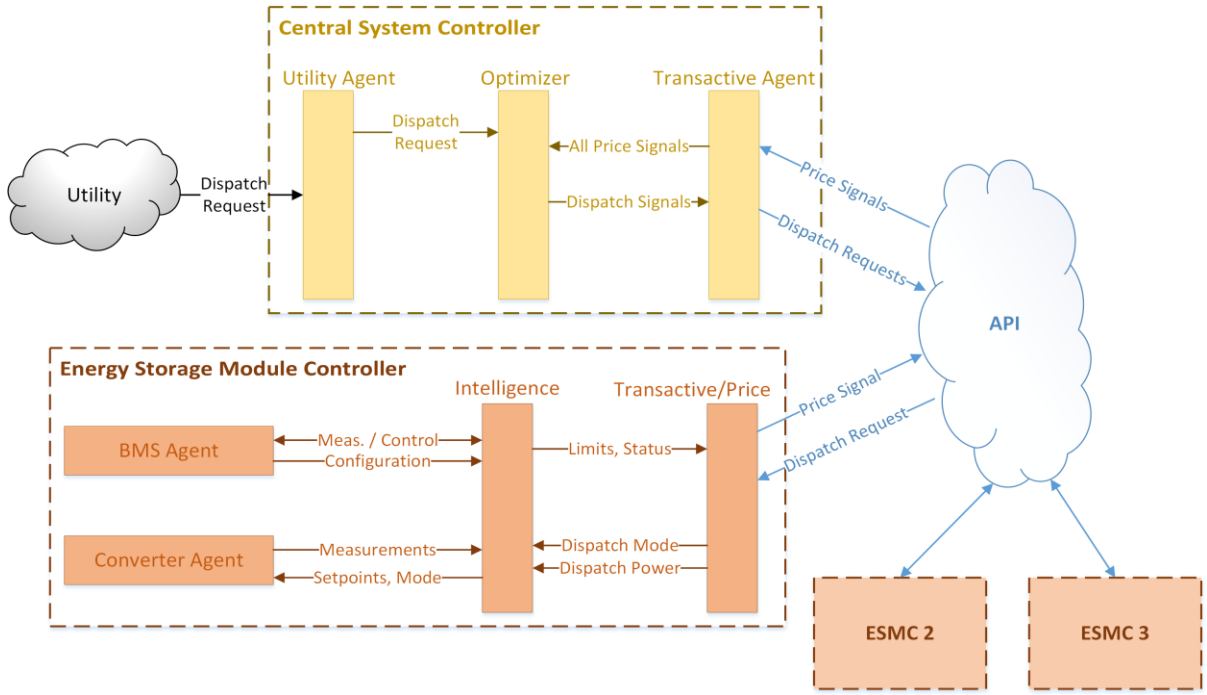


Figure 6.4 - Data flow between the CSC, API, and ESMC

Price Signal Generation

For determination of the bid curve, the ESMC calculates the net present value (NPV) of the lifetime energy production of the attached battery system. The NPV calculation uses the capital cost of the system (capex) over the lifetime energy of the system as shown in (6.1) [41]. This calculation is used as a baseline to determine the value of each kilowatt-hour of energy in the battery over the entire life of the battery. While this value can shift with the battery usage and the cycle-life, it is a good approximation for determining price. Weighting variables can be added later to account for other system factors such as charge/discharge rate or capacity fade.

$$NPV\left(\frac{\$}{kWh}\right) = \frac{Capital\ Cost\ (\$)}{Lifetime\ (cycles) \times Capacity\ (kWh)} \qquad (6.1)$$

In a multi-chemistry and multi-aged system, other factors can also weight the system behavior, such as round-trip efficiency and power capability. To capture these other factors, a traditional corner-price is extended in the negative-y direction and calculated as a two-quadrant corner-price for the battery system. Much like a traditional corner price, each ES has a linear price which is based on the maximum power capability and the present value of the charge-capacity of the system [47]. Power limits on the y-axis are based on the maximum charging (positive) and discharging (negative) rates of the system as reported by the BMS. The present system value (PSV – γ), is determined using (6.2), which calculates the discharged energy content as a function of the state of charge (SOC) of the battery system and assigns a value based on the NPV calculation.

$$\gamma_{pres.} = (NPV) \times \left(1 - \frac{SOC}{100\%}\right) \qquad (6.2)$$

The value calculated in (6.2) represents the available charge capacity of the system rather than the remaining energy capacity. This approach was chosen to capture positive and negative power flow as charging and discharging respectively. Furthermore, the objective is to reduce the availability of the battery system as the SOC reaches min and max values (high SOC battery unit is available at low cost for discharging, and at low SOC, the unit is not available as costs are high.) An example is shown in Figure 6.5. This allows the CSC to stack the energy resources based on cost optimization to maximize value during charging and minimize value during discharging, which should provide the best balance of total system performance.

The bid curve is calculated using the PSV as the x-intercept and the maximum power limits as fixed points at the minimum and maximum energy values. Charging and discharging

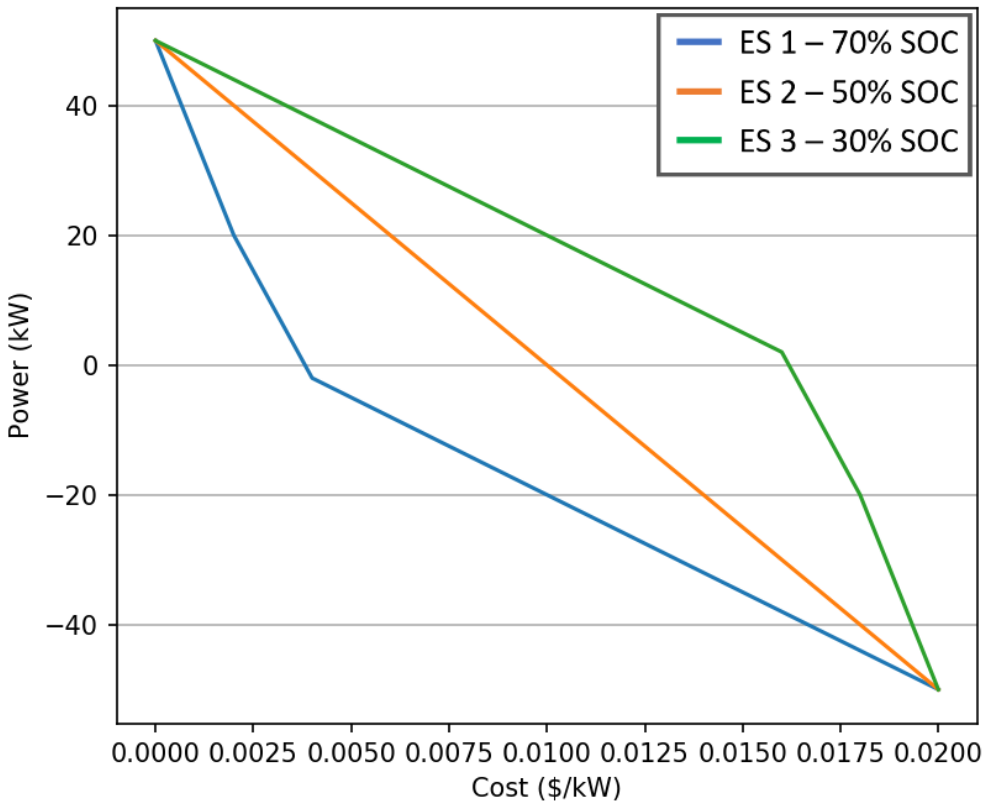


Figure 6.5 - Example present value versus state of charge.

bid equations are shown in (6.3) and (6.4) respectively where γ_{min} is the minimum value of (6.2) and γ_{max} is the maximum system value of (6.2), and λ_c and λ_d are the bid cost values. P_{cmax} and P_{dmax} are the maximum charge and discharge power limits. Since the bid curves are linear and follow a $y=mx+b$ fit, the communication of the information from EMSC to CSC is limited to the m and b variables for both the charge and discharge directions. Power limits obtained from the converter and BMS are communicated as well to constrain the optimization.

$$P_c = \frac{-P_{cmax}}{(\gamma - \gamma_{min})}(\lambda_c) + \frac{P_{cmax} * \gamma}{(\gamma - \gamma_{min})} \tag{6.3}$$

$$P_d = \frac{-P_{dmax}}{(\gamma_{max} - \gamma)}(\lambda_d) + \frac{-P_{dmax} * \gamma}{(\gamma_{max} - \gamma)} \tag{6.4}$$

Optimizing the Bid Curves

Even though the costs generated for the bid curves are only internal to the energy storage system, optimization methods used to perform value-based optimizations apply to the bid curves produced. For this linear optimization, the objective for optimizing the reported bid curves is to minimize the cost of operating the ESS in the next time interval. The objective function contains both the proposed charge (λ_c) and discharge (λ_d) costs of all units as represented in (6.5).

$$\min[\lambda_c + \lambda_d] \tag{6.5}$$

To constrain the optimization problem, the bid curves, requested utility dispatch, power limits, and power flow direction are considered. First, the cost to charge and cost to discharge are rewritten as a function of the bid curves received from N energy storage units. These are represented as the summations shown in (6.6) for charging and (6.7) for discharging. A binary charge/discharge variable has been added a well, represented by b_c and b_d respectively.

$$\lambda_c = \sum_{n=1}^{N} \left[-\left(\frac{\gamma_n - \gamma_{min_n}}{P_{cmax_n}} \right) \left(P_{c_n} - b_c \left(\frac{P_{cmax_n} * \gamma_n}{\gamma_n - \gamma_{min_n}} \right) \right) \right] \quad (6.6)$$

$$\lambda_d = \sum_{n=1}^{N} \left[-\left(\frac{\gamma_{max_n} - \gamma_n}{P_{dmax_n}} \right) \left(P_{d_n} - b_d \left(\frac{P_{dmax_n} * \gamma_n}{\gamma_{max_n} - \gamma_n} \right) \right) \right] \quad (6.7)$$

Optimization Constraints

To ensure that the power dispatch of all units reaches the requested dispatch from the utility, a constraint is added to the optimization given as (6.8). This specifies that the sum of the charging and discharging (negative power flow) dispatches must be equivalent to the $P_{dispatch}$. This equation gives the appearance that units could simultaneously charge and discharge if the dispatch power is reached. However, this is prevented by the binary variable constraint discussed below.

$$P_{dispatch} = \sum_{n=1}^{N} (P_{c_n} - P_{d_n}) \quad (6.8)$$

To ensure that each energy storage unit remains within the power limits, the optimization is constrained in both the charge and discharge direction by the power limit sent by each ESMC with the bid curve. The minimum and maximum charge and discharge limits for the battery system are given as (6.9) and (6.10) respectively.

$$0 \leq P_{c_n} \leq b_c * P_{cmax_n} \quad (6.9)$$

$$0 \leq P_{d_n} \leq b_d * P_{dmax_n} \quad (6.10)$$

The binary charge/discharge variable, b_c/b_d, is included to drive the power limit constraint to zero for the unused direction of power flow. For example, if a unit is charging, it cannot be discharging, therefore $0 \leq P_d \leq 0$ limiting the discharge power to 0 kW while the unit charges. All units are assumed to be operating simultaneously, hence the use of a universal

binary variable. When combined with (6.11), this ensures that all operating ESMCs are dispatched either to charge or discharge simultaneously. This is done to prevent the system from charging one battery while discharging another unit since this is inherently less efficient when converter losses are considered.

$$b_c + b_d \leq 1 \tag{6.11}$$

Constraint Modification

Results from running this optimization demonstrated that the universal binary variable had a negative impact on total system capacity since it resulted in the optimization being infeasible once any single energy storage unit reached either the top or bottom of its SOC range. These results will be discussed later in this chapter. To account for this problem, individual charge/discharge variables are used to constrain the energy storage units as shown in (6.11).

This is used in conjunction with universal binary variables, β_c and β_d, to ensure that if a given energy storage unit is dispatched, then the power flow is in the same direction as the other active units. This is done using (6.12) and (6.13) to sum the individual binary variables and ensure that they only sum to the number of ES units, *N*. When combined with the universal binary constraint shown in (6.14), this results in individual units being able to activate or deactivate while keeping unidirectional power flow among multiple units.

$$N \times \beta_c \geq \sum_{n=1}^{N} b_{c_n} \tag{6.12}$$

$$N \times \beta_d \geq \sum_{n=1}^{N} b_{d_n} \tag{6.13}$$

$$\beta_c + \beta_d \leq 1 \tag{6.14}$$

Optimization Output

The result of the optimization is a set of dispatch power setpoints for the battery systems. These are distributed through communication with the API as described above. Once distributed, the optimization agent waits 30 seconds, then re-runs the process with the updated price curves from the ESMCs. If an infeasible result is found, the optimization obtains the latest price curve data from the server and is immediately rerun. If no feasible solution is found within 3 attempts, a power setpoint of zero is posted to the API for all ESMCs. Repeated dispatch commands of zero cause the ESMC units to deactivate and deenergize to an idle state.

This solution was chosen because it could be used to help trigger the utility dispatch used to test the charge and discharge cycling of the optimization. When the utility agent measured zero power output from the unit, it would flip the sign of the dispatch command to have the unit start dispatching in the opposite direction. However, as discussed in the conclusions of this chapter, other solutions are possible. For example, individual units at their limits could shut down, removing themselves from the optimization and allowing the remaining system to operate. While other solutions are available, this one was chosen for the simplicity and to help trigger dispatch cycling.

Virtual System Testbench

Verification of the transactive optimization method described above was performed on a networked testbed using the agent framework described in Chapter 4. Instead of communicating with hardware for the BMS, DC/DC converters, and central inverter, these individual components were modelled in software within their respective agents. Their behavior was approximated using simple models since plans for this work include integrating the transactive control system with a more complex hardware-in-the-loop testbed. Using a virtual testbench

allows for the acceleration of the time constant as well. Each component in the testbed runs 10x faster than it would normally, with the optimization being run every 3 seconds. The hardware used in the testbed is shown in Figure 6.6.

The networked testbed uses four Raspberry Pi 3 B+ single-board computers running Raspbian Linux in conjunction with a laptop computer running Ubuntu Linux version 16.04 LTS. The Linux operating system is preferred on the single-board computers for its low overhead and simple integration with the Python 2.7 code language and the VOLTTRON agent environment [70]. Ubuntu Linux was chosen on the laptop for similar reasons including the pre-existing integration with the VOLTTRON agent system. Three of the single-board computers are used as energy storage module controllers with the fourth being used to host the RESTful API server. The central system controller and utility dispatch schedule are run on the laptop

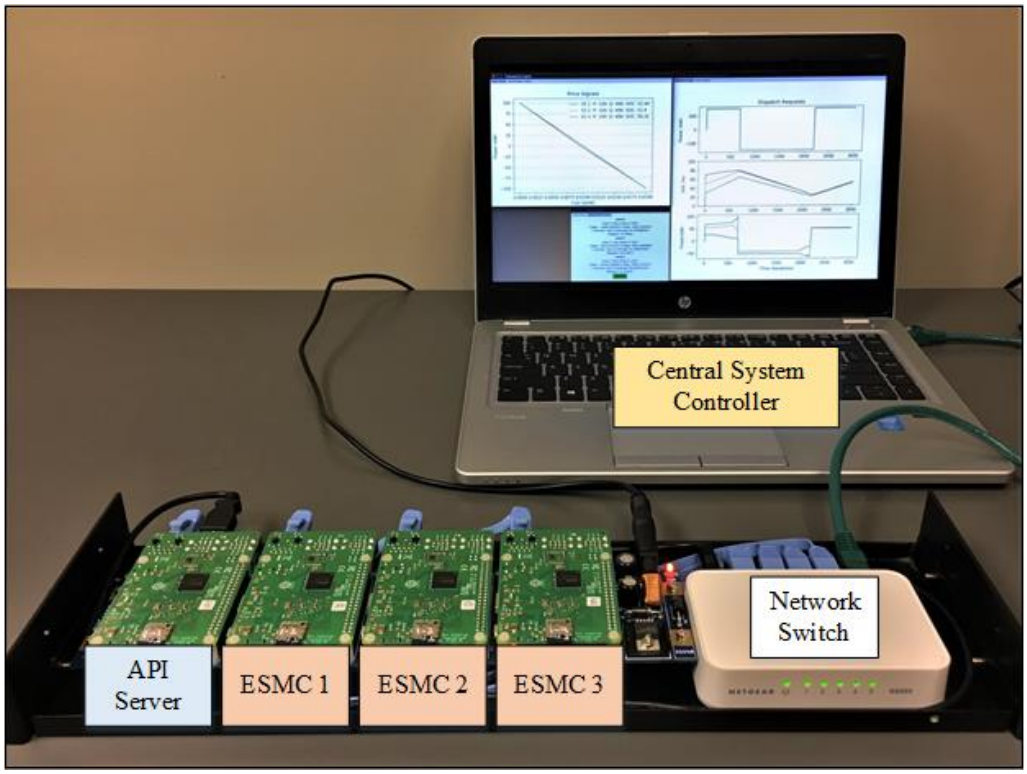


Figure 6.6 - Agent-Framework and Transactive Optimization Testbed [58]

computer because the optimization solving software could not be installed on the ARM-style processor found in a Raspberry Pi.

Energy Storage Module Controller Implementation

Each of the three Raspberry Pi units running the ESMC software were assigned battery parameters in accordance with the test plan described below. These parameters, which include capacity, cycle life, power limit, and cost define the linear battery model that is used in the BMS agent of the ESMC. Rather than connecting to a real BMS or using a complex battery model, a linear battery model is utilized in the BMS agent. As is common in real battery systems, the SOC is limited to a range of 20-80%. As Figure 6.7 shows, standard lithium-ion battery behavior over this range is nearly linear, indicating that a linear battery model will accurately represent the system behavior.

To capture this behavior, a charge-counting approximation is used which relates the rate of charge or discharge (current) to the SOC as shown in (6.15) where *k* is the iteration number,

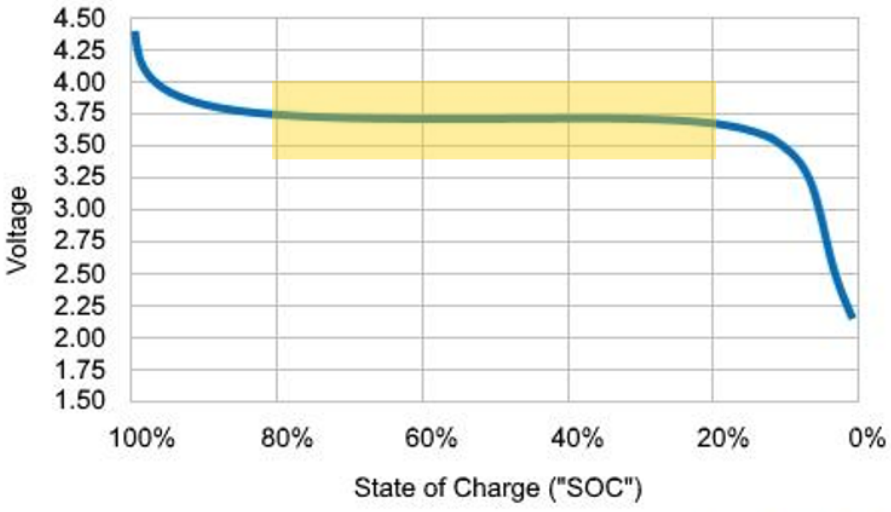


Figure 6.7 – Li-Ion discharge curve with linear portion highlighted.

Δt is the iteration length in seconds, and $E_{capacity}$ is the battery energy capacity. Using this equation, the SOC will rise and fall at a rate that is proportional to the power output level.

$$SOC_k = SOC_{k-1} + \frac{P_{dispatch} * \Delta t}{E_{capacity}} \tag{6.15}$$

In addition to this model, dynamic power limits are used at either end of the SOC range to represent a common behavior in battery management systems, which is to limit power flow as the system nears the upper and lower charge limits. These dynamic limits are represented as a series of 'if' statements in the code, which reduce the power capability of the system as the energy storage unit nears the SOC limits. These can best be represented by the piecewise equations shown in (6.16) for charging and (6.17) for discharging. The upper and lower limits of 80% were chosen since they are extremely common in battery systems to prevent overcharge and over-discharge, which can both be damaging to the battery. However, these limits are not hard-coded into the BMS agent, instead being set externally from a configuration file. Therefore, if a certain battery allows for a wider range of operation, these limits can be changed to reflect that. The same is true for the 25% and 75% values, which are instead calculated as 5% away from the upper and lower limits.

$$P_{c_{limit}} = \begin{cases} 0 & if\ SOC \geq 80\% \\ \dfrac{80\% - SOC}{5} * P_{c_{max}} & if\ 75\% \leq SOC < 80\% \\ P_{c_{max}} & if\ SOC < 75\% \end{cases} \tag{6.16}$$

$$P_{d_{limit}} = \begin{cases} 0 & if\ SOC \leq 20\% \\ \dfrac{SOC - 20\%}{5} * P_{d_{max}} & if\ 20\% < SOC \leq 25\% \\ P_{d_{max}} & if\ SOC\ > 25\% \end{cases} \tag{6.17}$$

The converter agent contains a model which represents the local bidirectional converter as a state engine as depicted in Figure 6.8 and operates as an ideal DC/DC converter. Any power setpoint is processed quickly, and the device is assumed to have a stable ramp-rate of 15 kW/s.

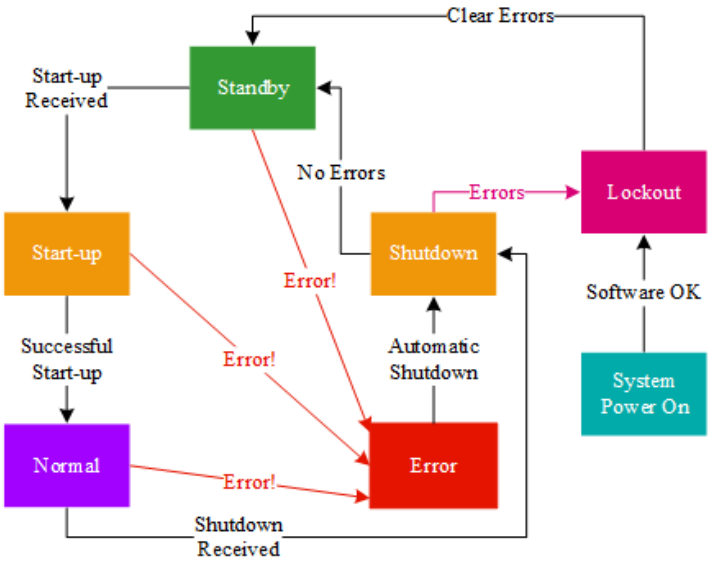


Figure 6.8 – State machine for modelling converter behavior

Based on presently available 100 kW converters, these assumptions can be made with confidence [71]. The converter model uses the state engine to respond to all intelligence agent commands to simulate real system behavior. Additionally, it contains the ability to fault the converter as necessary to test protection capabilities. Such protections include under and over-voltage protections to prevent equipment damage in the event of transient events on the battery side or inverter side of the converter.

Central System Controller Implementation

Just as the ESMC contained models of the battery system and DC/DC converter, the CSC interacts with an inverter model and a simulated utility. A similar state-engine to the one shown in Figure 6.8 is utilized by the inverter agent to model the system behavior and respond to all intelligence agent commands. One difference is that the power throughput of the inverter is dependent on the sum of the power being supplied or absorbed by the individual DC/DC converters. This information is read from the API server by the intelligence agent and reported

to the inverter. Using this method best simulates one proposed method for the physical system implementation that is proposed in Chapter 7 under Future Work. Rather than the inverter operating at a fixed setpoint, it uses a "DC-bus droop" based control to monitor the DC bus and adjust power flow accordingly to maintain a fixed bus voltage. For this research, the assumption is made that this method from literature is effective and results in stable system behavior [68] [69].

Instead of interfacing with a separate dispatch program, the utility agent contains a script-reading program to generate utility dispatch commands like the one shown in Figure 6.9. Two control modes are allowed by the utility agent. The first is user-generated dispatch where a user may input dispatch commands, and the second is scripted mode. When in scripted mode, the utility agent follows a pre-programmed dispatch schedule. For this system, a positive power dispatch indicates charging and a negative power dispatch indicates discharging.

Optimization is performed using the PuLP library for Python 2.7 within the optimization agent. PuLP provides an easy-to-use interface to the COIN-OR "Branch and Cut" (CBC) open-source solver. Other Python MILP librarys and solving tools have since been tested, but PuLP provided the simplest interface with reliable results from the CBC solver.

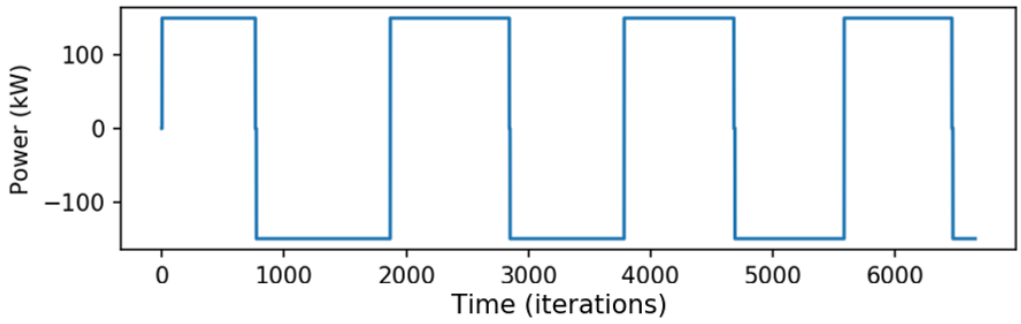


Figure 6.9 – Example dispatch schedule programmed into Utility Agent

Optimization Testing Plan

The virtual system testbench provides insight into whether the interactions between the various agents, devices, and models are behaving as expected. Testing the transactive element of the multi-chemistry battery energy storage system is done by modifying the various elements that contribute to the price-signal generation and system optimization. Each test in the plan shown in Table 6.1 changes one fundamental part of the transactive equation over three separate energy storage components. The values used are approximations for real-life systems based on investigations into the size of currently available systems and ongoing energy storage deployment projects with utilities. Cost estimations were made based on a $50/kWh capacity estimation and a fixed inverter cost. Cycle life was chosen based on new batteries.

Results and Discussion

Using the optimization testing plan in conjunction with the virtual system testbench, simulations of the optimization and system behavior were generated for each test. In the following section, a failed first attempt at optimization will be discussed that was tested using the same methods. This will be followed by a presentation of the results using a single universal charge/discharge binary variable. The section will close with results from an optimization that

Table 6.1 - Parameters for Testing Optimization Functionality

Test	Initial SOC (%)			Power Rating (kW)			Energy Rating (kWh)			Capex ($)		
1	ES-1	ES-2	ES-3	100			400			80k		
	70	50	30									
2	50			ES-1	ES-2	ES-3	400			80k		
				100	75	50						
3	50			100			ES-1	ES-2	ES-3	80k		
							400	300	200			
4	50			100			400			ES-1	ES-2	ES-3
										80k	60k	40k

was run under the modified constraints that were developed to allow the system to run individual energy storage units as necessary.

Early Optimization Results

First attempts at the transactive optimization system resulted in heavy system oscillations like the version of Test 2 shown in Figure 6.10. Oscillations like these can occur when the optimization finds multiple optimal results and just picks one. In the next cycle, the optimization may pick a much different dispatch for the units.

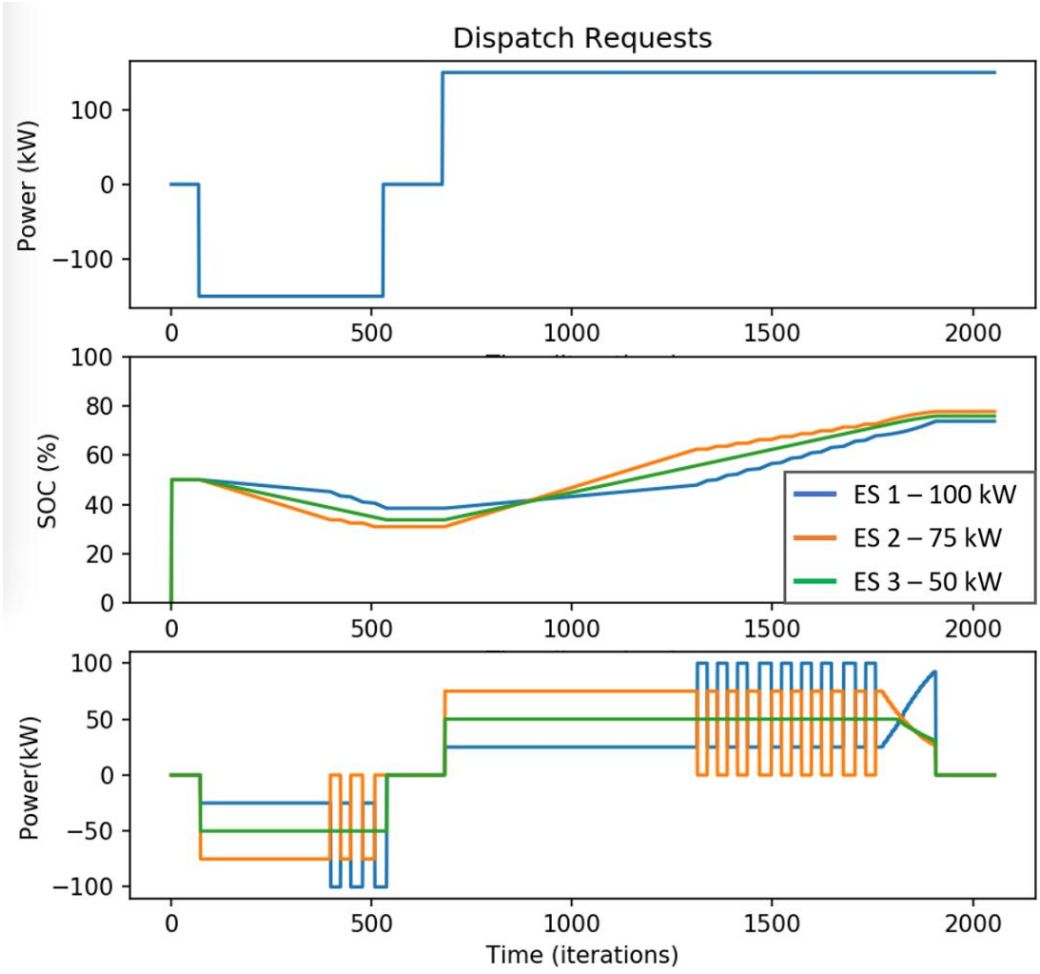


Figure 6.10 - Oscillations in dispatch results from multiple optimal solutions

In particular, the root cause of these results was improper programming of the cost variable in the optimization. Part of what makes a transactive market work is that the chosen cost at which energy is purchased or sold is the same across all the bidding elements. In this first implementation of the optimization algorithm, cost could differ from unit to unit because of a programming error. This meant that instead of "stacking" the energy from various units, the optimization could pick and choose whatever operating points satisfied the optimization for the individual energy storage units. As shown above in Figure 6.10, this results in the optimization oscillating between multiple possible solutions.

Universal Charge/Discharge Binary Variable Results

Following the test plan outline above in Table 6.1, the optimization utilizing the universal charge/discharge binary variable was tested on the virtual system testbench[8]. In the first test, three identical units are configured with SOCs at 70%, 50% and 30%. The results shown in Figure 6.11 demonstrate that despite different initial starting conditions, the battery systems converge to similar SOCs at which point the dispatch is balanced across the three units. During charging, the optimization charges the lowest SOC unit first since the price signal offers the best energy value. Discharging mode sees the highest SOC unit dispatched at the highest level since it offers the best value. This behavior is expected for three identical units with different initial SOC measurements. The slight dispatch deviations at the SOC extremes are caused by the dynamic power limits implemented on the battery model.

[8] Results from these tests were published in the IEEE Energy Conversion Congress and Exposition (ECCE) conference proceedings in October 2019 [ECCE CITE].

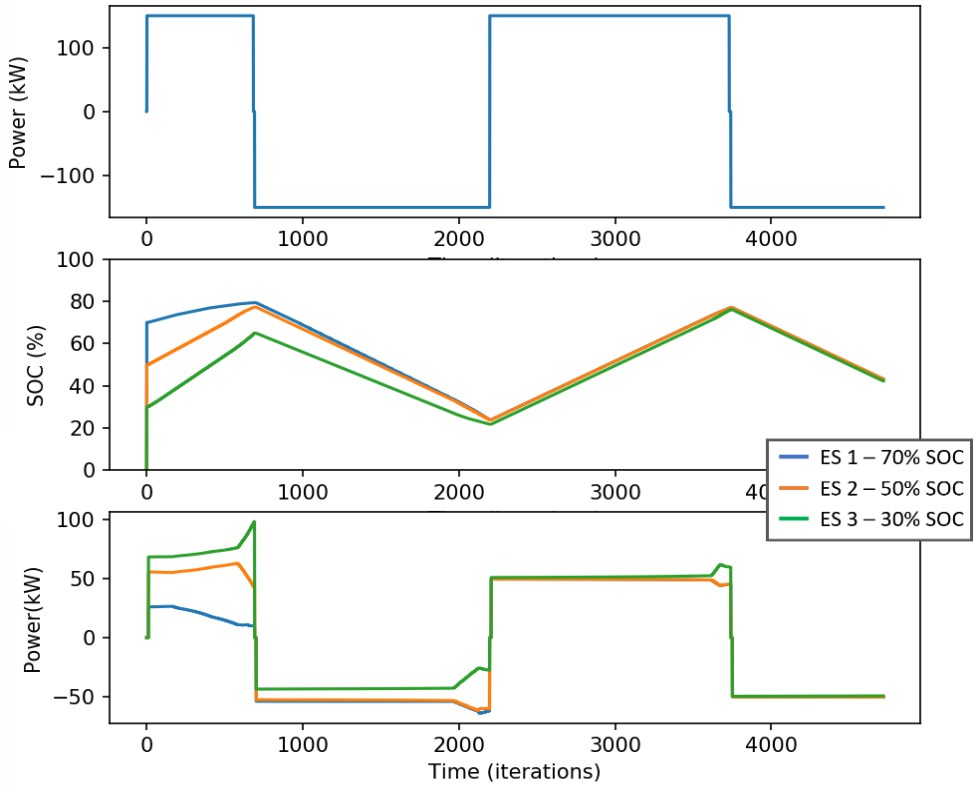


Figure 6.11 - Dispatch of three units with different initial SOCs

Next, the energy storage units are reconfigured such that they have three different power capabilities including 100 kW, 75 kW, and 50 kW. Each unit is started at 50% SOC and charging commences at the optimization-determined dispatch levels. Results from this test are shown in Figure 6.12. The energy storage units participated proportionally to the power ratings. Since each battery system has the same capex, the larger power rating unit will have a lower apparent energy value and is thus utilized more. This leads more quickly to a depleted system, which can be observed in the results. As the dynamic power limit lowers the output on the first unit, the remaining two battery systems ramp up to compensate. One issue that arises from this result is that the system will have significant unused capacity within the other two energy storage units as a result of the optimization stopping once the least expensive unit reaches its SOC limits.

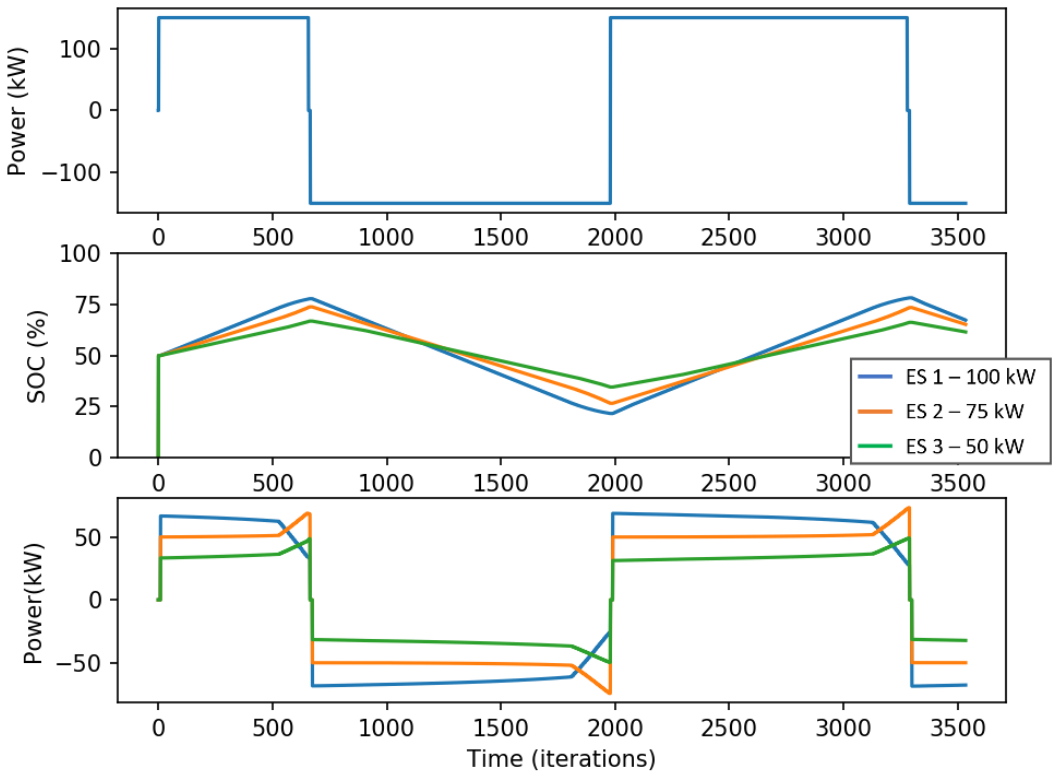


Figure 6.12 - Dispatch of three units with varying power ratings

In test 3, the three energy storage units are given differing energy storage capacities of 200 kWh, 300 kWh, and 400 kWh. As the testing table shows, the power ratings and capex are set to be consistent across the units and the SOC is reset to 50% for all units. Figure 6.13 shows the results from the virtual system testing. Based on the NPV calculation, the unit with the east capacity ends up with the highest value energy when compared to units with equivalent capital cost. The inverse is true for units with high capacities, which end up with lower NPV. Because the optimization favors charging units with higher value first and discharging lower value units first, the lowest capacity unit is charged at the highest rate and discharged at the lowest rate. Conversely, the highest capacity unit is charged at the lowest rate but takes the biggest share of the dispatch during discharging.

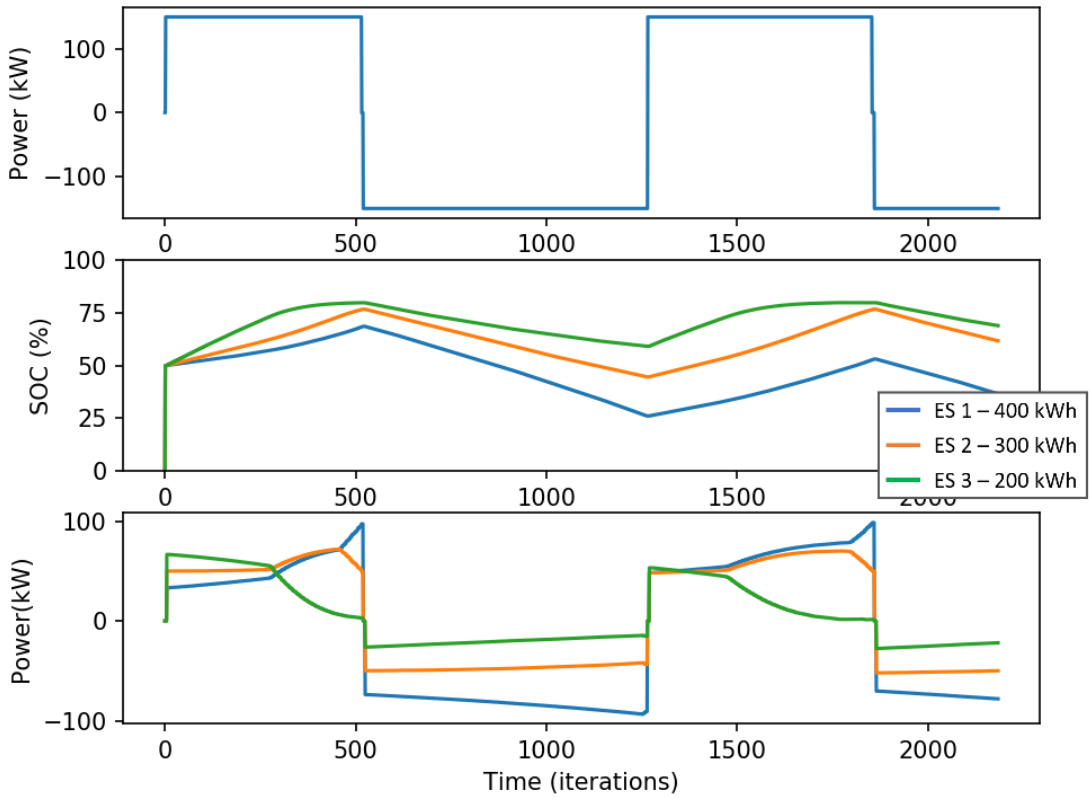


Figure 6.13 - Dispatch of three units with varying energy capacity

Finally, the capital costs for three units were changed to $40k, $60k, and $80k. These values are based off a battery cost estimation of $100/kWh, $150/kWh, and $200/kWh respectively for a 400 kWh system. These numbers assume that the cost of the BMS components are included in the ES cost. Consequently, these numbers tend to be low but will demonstrate the relationship between energy storage units with different capex values as perceived by the optimization.

Results from the fourth test with varying capital costs are shown in Figure 6.14. When all other NPV contributing factors are equal, higher capex results in a higher system valuation than an identical system of lower capital cost. This results in similar behavior from the

optimization as the previous test. The unit with the most expensive capex has the highest value, resulting in the unit being chosen first to charge and last to discharge. Meanwhile, the system with the lowest capex ends up dominating the discharge dispatch and not receiving much energy during charging periods. Over time, this causes the SOCs of multiple units to diverge, which the high value unit remaining at a high SOC and the low value unit discharging to a low SOC. As Figure 6.14 shows, this test was run for several extra cycles to ensure that the optimization makes the same decisions over multiple charge/discharge cycles. This divergence of units can be seen as a positive for lifetime management, but a negative for total system performance. Since

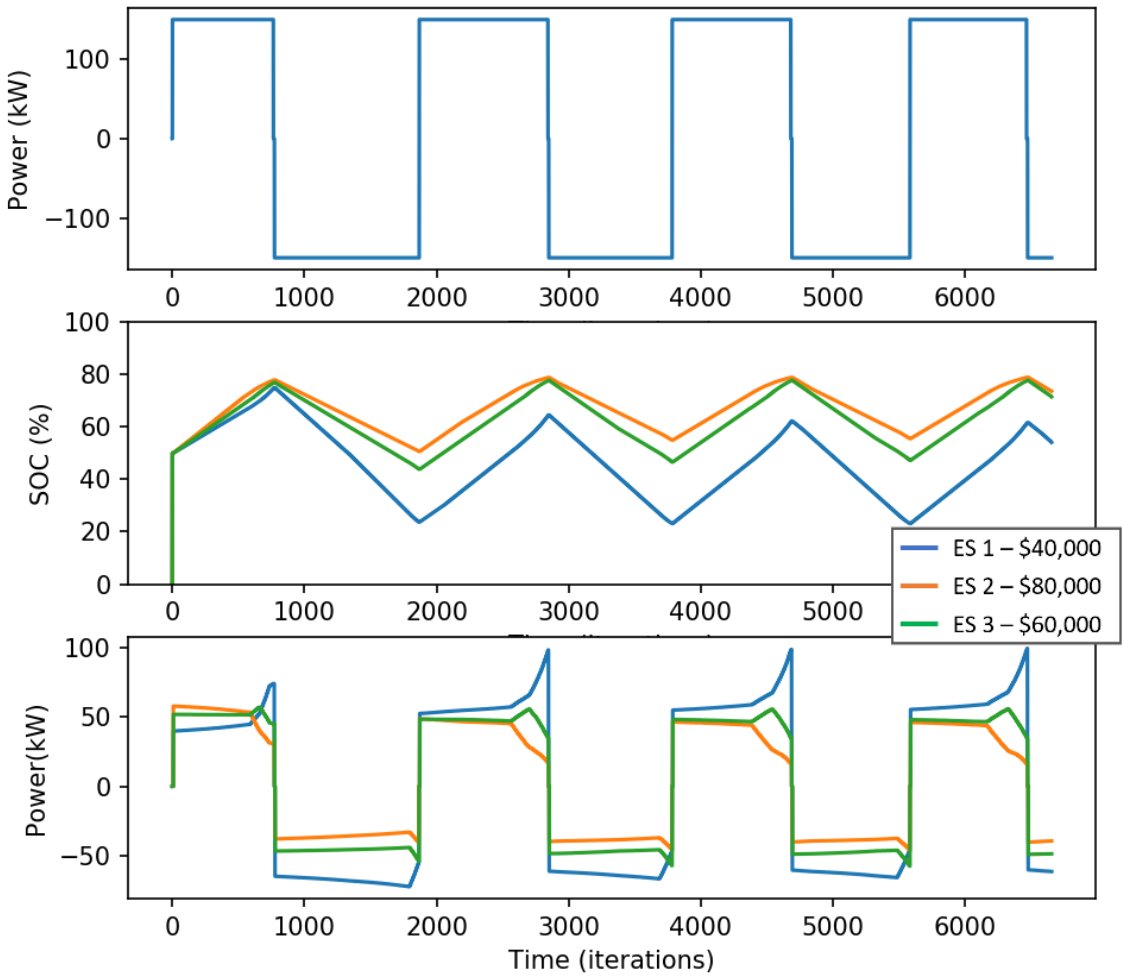


Figure 6.14 - Dispatch of three units with different capital costs

the use of the higher value unit is limited, the battery will experience an overall greater lifetime. However, since neither of the other two units can cycle through their entire functional range because of this higher value unit, the overall system performance is negatively impacted.

A single system-wide binary variable that limits charging and discharging proves problematic in the last two tests, causing the system SOC to diverge over time and limit the perceived system capacity. One way to note this is the scale of the x-axis of Figure 6.12 versus Figure 6.13. In test 1, the discharge time at 150 kW is approximately 1250 iterations. However, in the last two tests the discharge time is only 800 iterations. This ~30% reduction in capacity is due to the optimization being unable to discharge the system when any single ES unit reaches a minimum SOC limit. The problem becomes infeasible and power transfer ceases.

Individual Charge/Discharge Binary Variable

Discussed above, one possible solution to this problem is allowing the optimization to "switch" ES units on and off, if they are transferring power in the same direction as the other connected devices. To test this method, the first test from the Table 6.1 was run once more on the virtual system testbed. The results using the 150 kW utility dispatch system, shown in Figure 6.15, are nearly identical to the results from the universal binary. Under normal conditions where all units are within their SOC ranges, this result is expected. However, when one unit exceeds the SOC range, rather than turning the unit "off" as the individual binary variables allow, the solution becomes infeasible. This is indicative of a problem with the mathematical formulation of the optimization constraints.

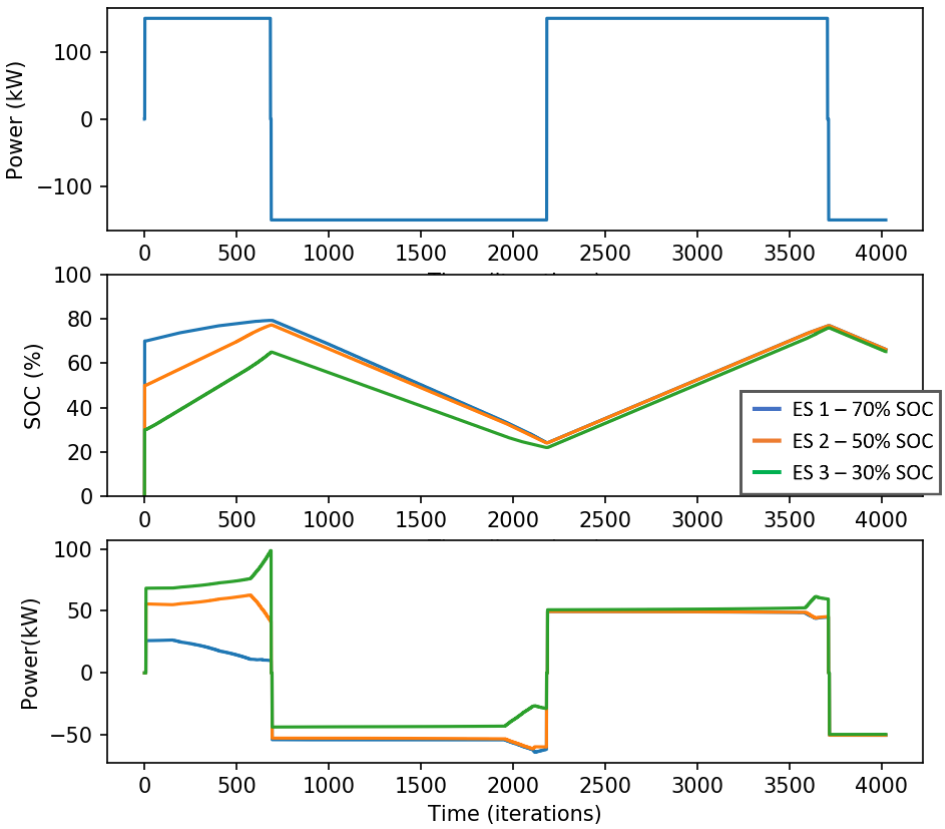


Figure 6.15 - Test 1 results with new optimization constraints

Another test was run on the individual binary variables optimization to verify that the problem was with the optimization constraints. This time, using the same system parameter values from Test 1 above, the utility dispatch request was manually adjusted to 100 kW to observe the system response. As shown in Figure 6.16 during the 0 to 500 time interval, dispatching the system at 100 kW causes the optimization to shut off two systems in order to carry the dispatch burden with just a single inverter. However, the solution appears incorrect since the first unit chosen to charge has the highest SOC. When it becomes incapable of charging, the logical unit with the lowest SOC takes over. However, when unit #3's SOC matches or exceeds the SOC of #2 (same price curve), the optimization does not choose to share the burden between the ES devices.

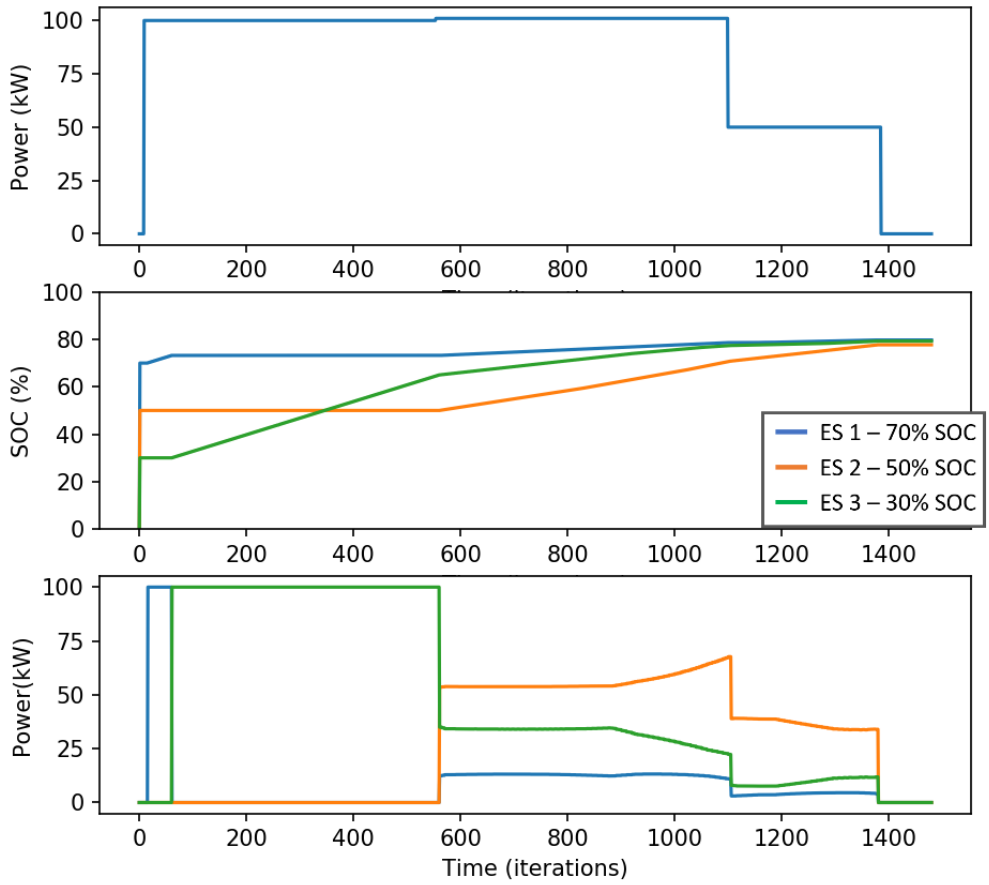


Figure 6.16 - Various manual dispatch tests demonstrating optimization behavior

After 550 iterations the setpoint was increased to 101 kW and the optimization chose a solution where the units shared the dispatch equally but used all three ES units instead of choosing to turn off the first unit (near maximum SOC). These issues point to a unique problem in the formulation of the value-based optimization. However, one way to avoid the problem is to use the ESMCs themselves to remove units close to maximum SOC from the optimization equation. Since the optimization is dynamic and based on the number of bids present on the API server, having an ESMC cease sending a bid function will remove the ES unit from the optimization. While this solution is not ideal, it does allow for an extended operational range for the energy storage system as a whole.

Conclusions

Distributing system control aids in achieving one of the goals for this research, which is flexible and modular hardware and software. In Chapter 5, the agent framework used to integrate with a wide range of battery hardware was presented, demonstrated, and shown to be effective in both an energy storage testbed and a project deployment. This chapter presented and demonstrated, through a virtual system, a reconfigurable software platform capable of integration with a wide range of battery technologies. The same BMS agent demonstrated in Chapter 5 was upgraded to include a battery model and re-used for this research on the virtual system testbed.

A crucial component of the software platform is the transactive energy market which enables distributing control back to the individual energy storage modules. Chapter 2 showed that other multi-chemistry systems investigated in the literature use centralized control that is tightly integrated with each energy storage component. In this system, limiting data and control information transmitted between the ESMCs and central system controller allows control decisions to be made more quickly, even if they do not provide a solution that is 100% optimal. At the scale of the virtual system testbench described here, limiting data-flow may not seem critical. However, growing this system to utility-scale energy storage requires dozens of elements. At these larger levels, a decentralized, lightweight control system is desired.

Results from the virtual testbench demonstrate a value-based optimization capable of dispatching the energy storage system in a logical and effective manner using limited data and by sending few commands. The flexibility of using the system was demonstrated with the multiple changes to the optimization that could be made without affecting the software on the ESMCs. In Chapter 7, future work to improve the optimization further will be proposed. Additionally, a hardware-in-the-loop upgrade to the virtual system testbed will be discussed.

Chapter 7 – Summary, Conclusions, and Future Work

Summary of Completed Work

Integrating a wide range of batteries starts with the ability to characterize the various types of batteries presently used in both stationary and vehicle applications. Previous work in the literature developed chemistry identification and capacity measurement techniques designed for use with small-capacity (<15 Wh) batteries for portable electronics. The chemistry identification techniques were designed for use with portable electronics chargers to allow a single charger to charge a wide array of devices. Improvements to capacity measurements are ongoing as they provide a crucial way to characterize new types of batteries as they are developed.

Expanding these techniques and automating the characterization process is the first contribution of this work. Through the development of a battery characterization testbench and subsequent testing, the chemistry identification techniques previously described in the literature were adapted and scaled for use in multi-cell systems. Chapter 3 presents results demonstrating that high-capacity (320 – 1290 Wh) multi-cell strings could be properly identified using the autonomous algorithms developed for this research. Furthermore, the same data used for identification can be reused as a capacity estimate to aid in the measurement of the battery system capacity. Rather than relying on technician input, an autonomous battery chemistry identification and capacity measurement has been shown to successfully measure these battery parameters. Data was taken for a wide range of EV, HEV, and PHEV battery types at various ages and states of health in order to verify that the autonomous system is reliable. Knowing the capacity, voltage, and battery type is crucial to controlling single battery systems. Integrating

multiple battery systems together into a utility-scale grid-connected energy storage provides the next challenge.

Previous multi-element battery work in the literature focused around the integration of one or two batteries into a system with supercapacitors or other fast energy storage components. These methods focus on rapid, centralized control for use with busses, trains, or electric power substations. These configurations were shown to not be modular or easily expandable to larger systems. Other work demonstrated large-scale battery energy storage using new Li-Ion batteries in large arrays of series and parallel systems. Through parallel DC/DC converters, these large arrays of batteries are interfaced with a single inverter to the grid. While this method provides the modularity desired, grouping batteries in large parallel/series arrays can be time-consuming since they must be age, capacity, and voltage matched; which is not ideal for the use of second-life batteries. However, this idea can be adapted through hardware and software design to be an ideal system for integration with secondary-use batteries. Utility-scale battery energy storage architectures paired with a decentralized communications and control architecture form the hardware backbone of the agent-framework proposed in this work.

In the literature, systems using the parallel DC/DC converter architecture centralize the control of all converters in one place, typically co-located with the central inverter. Developing and demonstrating a decentralized, flexible agent-based system using transactive optimization was the second significant contribution of this research. Each DC/DC converter in the system uses an agent framework locally situated to simplify communications between the battery management system, the converter, and the central control. The flexibility provided by using the agent framework allows the system to integrate with a wide range of different batteries from many different manufacturers. Using the residential energy storage test system, the agent-based

integration between the various components was stress-tested through a series of power tests paired with fault mode effects testing. To further demonstrate the software flexibility, two different battery management systems were integrated with the residential energy storage system using the same BMS agent. These tests were deemed successful since errors could be quickly detected and the system protected from dangerous fault conditions. With local control of each DC/DC converter demonstrated as well as a framework for the control of the central DC/AC inverter, the final step is to control and dispatch multiple energy storage units together as a single system.

In the literature, large energy storage systems are most commonly controlled as single elements due in large part to the fact that using large amounts of new batteries allows this type of system dispatch to be the most logical. However, when using batteries of different capacities, power capabilities, lifetimes, and costs, it is important that the units are dispatched in a manner to maximize their potential value over the course of their lives. This work developed a transactive energy market methodology that allows each energy storage to decide the local value of the energy contained in the battery. These values, which are SOC, power, capacity, cost, and lifetime dependent, are communicated to a central optimization as a bid into the transactive market. Using a value-based linear optimization, the dispatch of the units is then decided by the central system controller and sent back to the units for execution of the dispatch commands.

Demonstrating the transactive energy market on the agent-based framework for the energy storage system was done on a virtual system testbed developed using similar agents to the ones used on the residential energy storage testbed. Results from testing the optimization in several different scenarios show that the system can optimize the system as expected using only the transactive data provided by each energy storage. This is a significant change in energy

storage optimization from previous methods described in the literature which rely on having detailed data about each battery system being used. The results also demonstrated that the proposed decentralized software architecture and transactive system can be used to dispatch the system to meet a proposed dispatch signal from a utility. Through the multiple iterations of central optimization and price-signal generation, the flexibility of the proposed software system, for integration on a large scale across dozens of energy storage elements, was shown to be effective.

Future Work

While this work provides numerous advancements towards a modular and flexible multi-chemistry utility-scale battery energy storage system, it is also clear that there are many opportunities to continue this work and expand upon these results. This section will provide several opportunities for continuing this research including the integration of the control system with a hardware-in-the-loop control testbed.

Health evaluation of secondary-use batteries was briefly mentioned in Chapter 3during the discussion of battery parameterization. It is assumed that capacity provides an indicator of the health of the battery, therefore further explorations into determining State of Health (SOH) could be omitted. This decision was made, in part, because most State-of-Health algorithms compare a set of internal resistance measurements of the battery to a large data-set of measured batteries to determine the approximate State-of-Health. However, extending the lifetime of the secondary-use systems could be done more effectively if the remaining cycle-life could be estimated alongside the capacity of the attached system. Working with EV manufacturers to understand the life-cycle of a given battery could provide insights allowing secondary-use batteries to last even longer or be designated for use in different types of grid applications.

Ongoing research at Dalhousie University shows that the grid use-case significantly impacts the remaining cycle-life of a secondary-use battery [72]. Therefore, having insight to the remaining health of the battery could provide the system control with better information to dispatch the system in a way that increases the value obtained from the second-life battery systems.

Increasing system value can also be done through improvements to the transactive optimization system. Presently, the optimization obtains a snapshot of data about the batteries through a price curve, receives an instantaneous dispatch command from the utility, and best decides the system dispatch based on these parameters. The result of this optimization will therefore be locally optimal for the system, yet it does not consider the best way to dispatch the system over time. Future work for the transactive optimization includes modifying the transactive market to function on a 24 or 48 hour "look-ahead" schedule based on the utility dispatch requirements over the same time period. This would mean that the bid curves would need to be translated so that they also accounted for how long the energy storage would be at a certain power level. Essentially, the energy storage would send a 3-dimensional bid function where energy value changes based on the power level of the dispatch and the duration of the desired dispatch at that level. Furthermore, with knowledge of the battery health, these bid curves could be modified so that certain types of batteries are chosen for different dispatch commands, such as frequency regulation, PV smoothing, or bulk energy dispatch, based on the type of dispatch for which the battery is best suited.

Finally, an opportunity exists to upgrade the virtual system testbench to provide a more accurate model of the system interactions of the multi-chemistry system. The virtual system testbed presented in Chapter 6 provides an environment in which to test the distributed control architecture and transactive optimization. Converter and BMS models are used in the energy

storage module controllers to emulate the average behavior of the DC/DC converter and the battery for each ESMC. Similarly, the central system controller contains a state machine which models the average behavior of the DC/AC inverter. In this work, the virtual system testbench provides the necessary infrastructure to verify the communications and control architecture. However, this leaves an opportunity to integrate the system with real DC/DC and inverter controllers to see how these high-level controls handle transient events and other power electronics phenomena. One method to execute these tests is to integrate the hardware controllers with a power electronics simulation tool like the OPAL-RT real time digital simulator. This will provide an accurate platform to test if the asynchronous control method proposed above can work with real converters in a physical system without the dangers associated with working on high-voltage batteries at high power. A suggestion for the system implementation can be found in Figure 7.1.

Conclusions

Energy storage technologies are poised to revolutionize the electric grid. However, high costs versus system lifetime and new technology adoption hesitance prevent utilities from implementing energy storage at scale. Secondary-use batteries provide a promising source of inexpensive energy storage as first generation EVs, HEVs, and PHEVs reach the end of their expected lives. However, differing form factors, voltages, and capacities provide an obstacle to rapid integration of these batteries to large-scale energy storage systems. Present utility-scale BES systems utilize large counts of new batteries for easy parallelization to increase system capacity. These methods cannot be replicated with secondary-use systems because of the differing ages and battery chemistries.

To overcome these obstacles and create a modular and expandable utility-scale energy storage system, a system is proposed which uses software and hardware systems to increase modularity while decreasing complexity, enabling the use of a wider range of batteries including secondary-use units, new batteries, and even flow batteries. These novel systems were demonstrated using a mixture of hardware testbenches and virtual software systems. Results presented demonstrate that the autonomous battery parameterization, software architecture, and transactive optimization can work together to create a unified system for the integration of many different batteries.

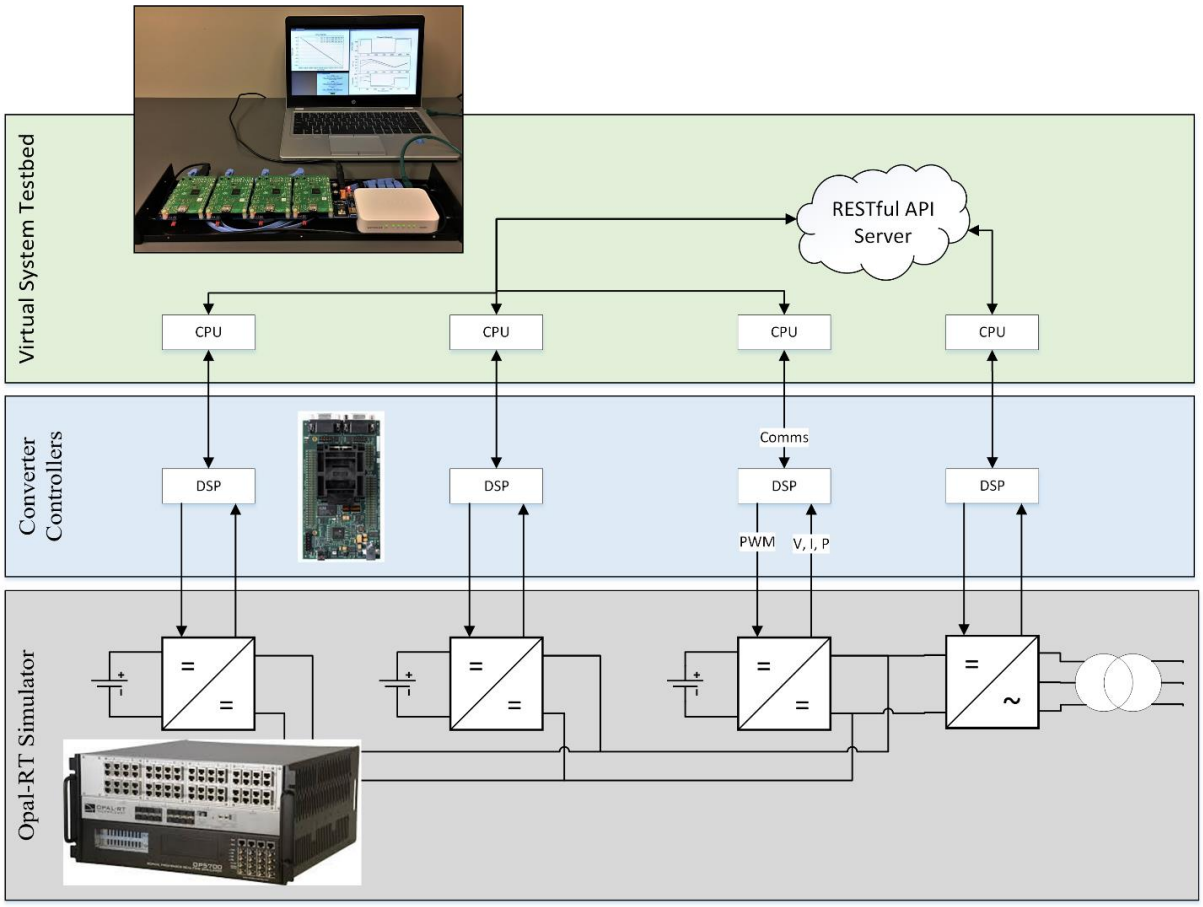


Figure 7.1 - Architecture for future work to integrate with a power electronics testbed

References

[1] G. Masson, S. Orlandi and M. Rekinger, "Global Market Outlook for Photovoltaics 2014-2018," European Photovoltaic Industry Association, 2014.

[2] U. E. I. Administration, "Annual Energy Outlook 2019 with projections to 2050," U.S. Department of Energy, Washington, D.C., January 2019.

[3] G. Masson, "PVPS Report - Snapshot of Global PV 1992-2013," International Energy Agency, 2014.

[4] Global Wind Energy Council, "Global Wind Report - Annual Market Update 2013," Brussels, 2014.

[5] R. Smith, "Southern Caifornia May See Summer Blackouts as Gas Leak's Effects Liner," The Wall Street Journal, 16 June 2016. [Online]. Available: https://www.wsj.com/articles/southern-california-may-see-summer-blackouts-as-gas-leaks-effects-linger-1466368338?mg=id-wsj. [Accessed 28 February 2019].

[6] E. Newburger, "CNBC," NBC Universal, 23 October 2019. [Online]. Available: https://www.cnbc.com/2019/10/23/pge-rebuked-over-imposing-blackouts-in-california-to-reduce-fire-risk.html. [Accessed 13 November 2019].

[7] Institute for Energy Research, "California's Self-Inflicted Energy Shortage," 27 June 2016. [Online]. Available: https://www.instituteforenergyresearch.org/regulation/california-electricity-shortage-issues/. [Accessed 28 February 2019].

[8] J. Mossoba, M. Kromer, P. Faill, S. Katz, B. Boroway, S. Nichols, L. Casey, D. Maksimovic, J. Traube and F. Lu, "Analysis of Solar Irradiance Intermittency Mitigation Using Contact DC Voltage PC and EV Battery Storage," in *Transportation Electrification Conference and Expo*, Dearborn, MI, USA, 2012.

[9] A. Brissette, A. Hoke, J. Traube, F. Lu and D. Maksimovic, "Study on the Effect of Solar Irradiance Intermittency Mitigation on Electric Vehicle Battery Lifetime," in *Technologies for Sustainability*, Portland, OR, USA, 2013.

[10] J. Traube, F. Lu, D. Maksimovic, J. Mossoba, M. Kromer, P. Faill, S. Katz, B. Borowy, S. Nichols and L. Casey, "Mitigation of Solar Irrandiance Intermittency in Photovoltaic Power Systems with Integrated Electric-Vehicle Charging Functionaloty," *Transactions on Power Electronics*, vol. 28, no. 6, pp. 3058-3067, June 2013.

[11] B. Nykvist and M. Nilsson, "Rapidly falling costs of battery packs for electric vehicles," *Nature Climate Change*, vol. 5, pp. 329-332, April 2015.

[12] C. Curry, *Lithium-ion Battery Costs and Market*, Bloomberg New Energy Finance, 2017.

[13] A. A. Akhil, G. Huff, A. B. Currier, B. C. Kaun, D. M. Rastler, S. B. Chen, A. L. Cotter, D. T. Bradshaw and W. D. Gauntlett, "DOE/EPRI Electricity Storage Handbook in Collaboration with NRECA," Sandia National Laboratory, Livermore, CA, 2015.

[14] K. H. LaCommare and J. H. Eto, "Understanding the Cost of Power Interruptions to U.S. Electricity Customers," Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA, 2004.

[15] M. Starke, P. Irminger, B. Ollis, G. Andrews, O. Onar, P. Karlson, S. Thambiappah, P. Valencia, S. Massin, A. Goodson and P. Rosenfeld, "Community Energy Stroage with Secondary Use EV/PHEV Batteries," Oak Ridge National Laboratory, 2015.

[16] I. Gyuk, M. Johnson, M. Vetrano, K. Lynn, W. Parks, R. Handa, L. Kannberg, S. Hearne, K. Waldrip and R. Braccio, "Grid Energy Storage," US Department of Energy, Dec 2013.

[17] C. K. Narula, R. Martinez, O. Onar, M. R. Starke and G. Andrews, "Economic Analysis of Deploying Used Batteries in Power Systems," Oak Ridge National Laboratory, 2011.

[18] "www.fueleconomy.gov," U.S. Department of Energy, 2015. [Online]. Available: https://www.fueleconomy.gov/feg/Find.do?action=sbs&id=34918&id=34699&id=33558&id=32154. [Accessed 23 March 2018].

[19] Vehicle Technologies Office, "Alternative Fuels Data Center," U.S. DOE Energy Efficiency and Renewable Energy, May 2017. [Online]. Available: https://www.afdc.energy.gov/data/10567. [Accessed 23 March 2018].

[20] B. Brockman, "My Nissan Leaf," My Electric Vehicle Forums, 21 October 2013. [Online]. Available: http://www.mynissanleaf.com/viewtopic.php?f=4&t=17168. [Accessed 23 March 2018].

[21] I. Serban and C. Marinescu, "A Look at the Role and Main Topologies of Battery Energy Storage Systems for Integration in Autonomous Microgrids," in *Optimization of Electrical and Electronic Equipment*, Brasov, Romania, 2010.

[22] S. Ponnaluri, G. O. Linhofer and J. K. Steinke, "Comparison of single and two stage topologies for interface of BESS or fuel cell system using the ABB standard power electronics building blocks," in *Power Electronics and Applications*, Dresden, Germany, 2005.

[23] A. Kuperman and I. Aharon, "Battery-Ultracapacitor hybrids for pulsed current loads: A Review," *Renewable and Sustainable Energy Reviews*, vol. 15, pp. 981-922, 2011.

[24] A. Narvaez, C. Cortez and C. L. Trujillo, "Comparative analysis of topologies for the interconnection of Batteries and Supercapacitors in a Hybrid Energy Storage System," in *Power Electronics for Distributed Generation Systems*, Florianopolis, 2017.

[25] N. Mukherjee and D. Strickland, "Second Life Battery Energy Storage Systems: Converter Topology and Redundancy Selection," in *Power Electronics, Machines, and Drives*, Manchester, UK, 2014.

[26] L. Maharjan, S. Inoue, H. Akagi and J. Asakura, "State-of-Charge (SOC)-Balancing Control of a Battery Energy Storage System Based on a Cascade PWM Converter," *Transactions on Power Electronics*, vol. 24, no. 6, pp. 1628-1636, 2009.

[27] W. Huang and J. A. Abu Qahouq, "Distributed Battery Energy Storage System Architecture with Energy Sharing Control for Charge Balancing," in *Applied Power Electronics Conference and Exposition*, Fort Worth, TX, USA, 2014.

[28] P. J. Hart, P. J. Kollmeyer, L. W. Juang, R. H. Lasseter and T. M. Jahns, "Modeling of Second-Life Batteries for Use ina CERTS Microgrid," in *Power and Energy Conference at Illinois*, Champaign, IL, USA, 2014.

[29] X. Li, D. Hui, M. Xu, L. Wang, G. Guo and L. Zhang, "Integration and Energy Management of Large-scale Lithium-ion Battery Energy Storage Station," in *Electrical Machines and Systems*, Sapporo, Japan, 2012.

[30] X. Lin and Y. Lei, "Coordinated Control Strategies for SMES-Battery Hybrid Energy Storage Systems," *IEEE Access*, vol. 5, pp. 23452-23465, 2017.

[31] S. Butterbach, B. Vultureescu, C. Forgex, G. Coquery and G. Friedrich, "Lead-acid battery model for hybrid energy storage," in *Vehicle Power and Propulsion Conference*, Chicago, IL, USA, 2011.

[32] M. Ozbek, B. Pierquet and A. D. Bagliano, "Scalable and Flexible Cell-Based Energy Storage System". United States 30 March 2017.

[33] T. Guo, Z. Yang, F. Lin and S. Xiong, "Optimization of Peak Load Shifting Control Strategy for Battery Energy Storage System Used in Urban Rail Transit," in *Industrial Electronics Society*, Beijing, China, 2017.

[34] D. Wu, C. Jin, P. Balducci and M. Kintner-Meyer, "An Energy Storage Assessment: Using Optimal Control Strategies to Capture Multiple Services," in *Power and Energy Society General Meeting*, Denver, CO, USA, 2015.

[35] W. Jing, C. H. Lai, M. L. D. Wong and W. S. H. Wong, "Smart Hybrid Energy Storage for Stand-alone PV Microgrid: Optimization of Battery Lifespan through Dynamic Power Allocation," in *Asia-Pacific Power and Energy Engineering Conference*, 2015.

[36] T. Jiang, G. Putrus, Z. Gao, M. Conti and S. McDonald, "Optimal Charging Strategy for EVs with Batteries at Different States of Health," in *Power Engineering, Energy, and Electrical Drives*, Istanbul, Turkey, 2013.

[37] A. Hoke, A. Brissette, D. Maksimovic, D. Kelly and A. Pratt, "Maximizing Lithium Ion Vehicle Battery Life Through Optimized Partial Charging," in *Innovative Smart Grid Technologies*, Washington, DC, USA, 2013.

[38] H. Yin, C. Zhao, M. Li and C. Ma, "Optimization Based Energy Control for Battery/Supercapacitor Hybrid Energy Storage Systems," in *Industrial Electronics Society*, Vienna, Austria, 2013.

[39] I. Buchmann, "BU-502: Discharging at High and Low Temperatures," Battery University, 10 July 2019. [Online]. Available: https://batteryuniversity.com/learn/article/discharging_at_high_and_low_temperatures. [Accessed 14 November 2019].

[40] L. Jinwei, S. Jianhui, S. Yong, M. Meiqin, Y. Xiangzhen and L. Ning, "Microgrid Secondary Frequency Control Method Based on Life Optimization of Energy Storage Units," in *Power Electronics and Application Conference and Exposition*, Shanghai, China, 2014.

[41] A. K. Barnes, J. C. Balda, S. O. Geurin and A. Escobar-Mejia, "Optimal Battery Chemistry, Capacity Selection, Charge/Discharge Schedule, and Lifetime of Energy Storage under Time-of-Use Pricing," in *IEEE Innovative Smart Grid Technologies (Europe)*, Manchester, UK, 2011.

[42] A. K. Barnes, J. C. Balda, S. O. Geurin and A. Escobar-Mejia, "Optimal Battery Chemistry, Capacity Selection, Charge/Discharge Schedule, and Lifetime of Energy Storage under Time-of-Use Pricing," in *Innovative Smart Grid Technologies (Europe)*, Manchester, UK, 2011.

[43] X. Le, S. Chen, Z. Yan, J. Ping, X. Zhao, N. Zhang and J. Xi, "Enabling a Transactive Distribution System via Real-Time Distributed Optimization," in *Transactions on Smart Grid*, 2018.

[44] W. Liu, J. Zhan and C. Y. Chung, "A Novel Transactive Energy Control Mecahnism for Collaborative Networked Microgrids," in *Transactions on Power Systems*, 2018.

[45] Y. Amanbek, Y. Tabarak, H. K. Nunna and S. Doolla, "Decentralized Transactive Energy Management System for Distribution Systems with Prosumer Microgrids," in *19th International Carpathian Control Conference*, Szilvasvarad, Hungary, 2018.

[46] J. Li, Z. Wu, S. Zhou, H. Fu and X.-P. Zhang, "Aggregator Service for PV and Battery Energy Storage Systems of Residential Building," *CSEE Journal of Power and Energy Systems*, vol. 1, no. 4, pp. 3-11, December 2015.

[47] K. D. Craemer and G. Deconinck, "Balancing Trade-offs in Coordinated PHEV Charging with Continuous Market-based Control," in *IEEE PES ISGT Europe*, Berlin, 2012.

[48] M. Ragsdale, J. Brunet and B. Fahimi, "A Novel Battery Identification Method Based on Pattern Recognition," in *Vehicle Power and Propulsion Conference*, Harbin, China, 2008.

[49] S. G. Abeyrante, S. S. Kuruppu, C. M. Divarathne, D. K. Handapangoda, S. Senevirathne and U. I. Dayaratne, "A Self Adaptable Intelligent Battery Charger for Portable Electronics," in *Industrial Electronics*, Cambridge, UK, 2008.

[50] M. Navarante, R. Koswatta and S. Abeyratne, "A Simple Battery Chemistry Identification and Implementation Technique for a Self Adaptable Charger," in *Industrial and Information Systems*, Sri Lanka, Sri Lanka, 2009.

[51] Idaho National Laboratory, "Battery Test Manual for Electric Vehicles," U.S. DOE Vehicle Technologies Program, 2015.

[52] B. Balasingam, G. V. Avvari, B. Pattipati, K. Pattipati and Y. Bar-Shalom, "Robust Battery Fuel Gauge Algorithm Development, Part 2: Online Battery-Capacity Estimation," in *Renewable Energy Research and Applications*, Milwaukee, WI, USA, 2014.

[53] R. R. Richardson, C. R. Birkl, M. A. Osborne and D. A. Howey, "Gaussian Process Regression for In-situ Capacity Estimation of Lithium-ion Batteries," *Transactions on Industrial Informatics*, vol. PP, no. 99, pp. 1-12, 2018.

[54] C. Fleischer, W. Waag, H.-M. Heyn and D. U. Sauer, "On-line adaptive battery impedance parameter and state estimation considering physical principles in reduced order equivalent circuit battery models - Part 1. Requirements, critical review of methods and modeling," *Journal of Power Sources*, no. 260, pp. 276-291, 20114.

[55] C. Fleischer, W. Waag, H.-M. Heyn and D. U. Sauer, "On-line adaptive battery impedance parameter and state estimation considering physical principles in reduced order equivalent circuit battery models: Part 2. Parameter and state estimation," *Journal of Power Sources*, no. 262, pp. 457-482, 2014.

[56] M. T. Smith, M. R. Starke, A. Herron, Jr, D. J. King and L. M. Tolbert, "Autonomous Battery Identification and Health Parameterization," in *ISA Power and Industry Division Symposium*, Knoxville, 2018.

[57] Mercedes-Benz, "Maximum Performance with Mercedes-Benz Energy Storage," May 2018. [Online]. Available: https://www.mercedes-benz.com/old/wp-content/uploads/sites/41/2018/05/MBE_FL_Salesfolder_grid_6s_EN_DE_Ansicht.pdf. [Accessed February 2019].

[58] M. T. Smith, M. R. Starke, L. M. Tolbert and M. Chinthavali, "Architecture for Utility-Scale Multi-Chemistry Battery Energy Storage," in *IEEE Energy Conversion Congress and Exposition*, Baltimore, MD, 2019.

[59] M. Starke, R. Zeng, S. Zheng, M. Smith, M. Chinthavali and Z. Wang, "A Multi-Agent System Concept for Rapid Energy Storage Development," in *IEEE Innovative Smart Grid Technologies*, Washington, DC, 2019.

[60] M. Starke, R. Zeng, S. Zheng, M. Smith, S. Campbell, B. Dean and M. Chinthavali, "Residential (Secondary-Use) Energy Storage System with Moular Software and Hardware Power Electronic Interfaces," in *IEEE Energy Conversion Congress and Exposition*, Baltimore, 2019.

[61] Tesla, Inc, "Tesla Backup Gateway 2 Datasheet," February 2019. [Online]. Available: https://www.tesla.com/sites/default/files/pdfs/powerwall/Backup_Gateway_2_Datasheet_en_GB.pdf. [Accessed May 2019].

[62] LG Chem, "LG ESS Product Information," [Online]. Available: https://www.lgesspartner.com/front/normal/en/product/productInfo.dev. [Accessed August 2019].

[63] M. Starke, M. Chinthavali, C. Winstead, T. Kuruganti, Y. Xue, R. Zeng, S. Zheng, s. Campbell and C. Thomas, "Networked Control and Optimization for Widescale Integration of Power Electronic Devices in Residential Homes," in *IEEE Energy Conversion Congress and Exposition*, Baltimore, 2019.

[64] X. Le, S. Chen, Z. Yan, J. Ping, X. Zhao, N. Zhang and J. Xi, "Enabling a Transactive Distribution System via Real-Time Distributed Optimization," *IEEE Transactions on Smart Grid*, vol. Early Access, pp. 1-11, September 2018.

[65] M. Ji, P. Zhang and Y. Cheng, "Distributed Microgrid Energy Optimization Using Transactive Control and Heuristic Strategy," in *ISSS PES General Meeting*, Chicago, July 2017.

[66] B. Posner, "The Fundamentals of Electricity Markets," The Pennsylvania State University, 2018. [Online]. Available: https://www.e-education.psu.edu/ebf200/node/151. [Accessed September 2019].

[67] S. Weckx, R. D'Hulst, B. Claessens and J. Driesen, "Multiagent Charging of Electric Vehicles Respecting Distribution Transformer Loading and Voltage Limits," *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2857-67, November 2014.

[68] S. Rivera, B. Wu, S. Kouro, H. Wang and D. Zhang, "Cascaded H-bridge Multilevel Converter Topology and Three-phase Balance Control for Large Scale Photvoltaic Systems," in *IEEE International Symposium on Power Electronics for Distributed Generation Systems*, 2012.

[69] Y. Yu, G. Konstantinou, B. Hredzak and V. G. Agelidis, "Power Balance of Cascaded H-Bridge Multilevel Converters for Large-Scale Photovoltaic Integration," *IEEE Transactions on Power Electronics,* vol. 31, no. 1, pp. 292-303, 2015.

[70] J. Haack, B. Akyol, N. Tenney, B. Carpenter, R. Pratt and T. Carroll, "VOLTTRON: An Agent Platform for Integrating Electric Vehicles and Smart Grid," in *International Conference on Connected Vehicles adn Exposition*, Las Vegas, NV, December 2013.

[71] Dynapower Company, LLC, "MPS-125 EHV: Robust 125 kW Energy Storage Inverter," 2019. [Online]. Available: https://www.dynapower.com/products/energy-storage-solutions/behind-the-meter-energy-storage/mps-125-ehv/#1476296979221-2f3a4312-f993. [Accessed October 2019].

[72] B. Thompson, *Repurposing Electric Vehicle Batteries in a Mixed Array for Grid Storage,* Halifax, Nova Scotia: Dalhousie University, April 2018.

[73] MESA Standards Alliance, "MESA Standards Home," Wordpress, 2014. [Online]. [Accessed December 2018].

[74] R. Jackson, S. Curran, P. Chambon, B. Post, L. Love, R. Wagner, B. Ozpineci, M. Chinthavali, M. Starke, J. Green, L. Tryggestad and B. Lee, "Overview of the Oak Ridge National Laboratory Advanced Manufacturing Integrated Energy Demonstration Project: Case Study of Additive Manufacturing as a Tool to Enable Rapid Innovation in Integrated Energy Systems," in *ASME International Mechanical Engieering Congress and Exposition, Vol 14*.

Vita

Mitchell T. Smith was born in West Chester, PA to Deborah and Roger Smith. His adolescence education was through the West Chester Area School District and he graduated high school in 2008. The following fall, Mitchell enrolled at Ursinus College in Collegeville, PA. While attending, he conducted undergraduate research with Dr. Thomas Carroll in magneto-optical trapping, with Dr. Nicholas Scoville in advanced graph isomorphisms, and with Dr. Christopher Cellucci in electronics for portable health equipment. He graduated in May of 2012 with a B.S. degree in Physics. In August of 2012, Mitchell began his Ph.D. studies in secondary-use energy storage at the University of Tennessee (Knoxville).